IBM Netcool Operations Insight
Version 1 Release 6

Integration Guide

SC27-8601-11

Note

Before using this information and the product it supports, read the information in Appendix B, “Notices,” on page 577.

This edition applies to version 1.6.0.1 of IBM® Netcool® Operations Insight® (product number 5725-Q09) and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Chapter 1. Solution overview

Read about key concepts and capabilities of IBM Netcool Operations Insight.

Related reference

Release notes

IBM Netcool Operations Insight V1.6.0.1 is available. Compatibility, installation, and other getting-started issues are addressed in these release notes.

What's new

Netcool Operations Insight V1.6.0.1 includes a range of new features and functions.

This description of new features and functions is also available in the Appendix A, “Release notes,” on page 559.

New product features and functions in V1.6.0.1

1.6.0.1

Updated product versions in V1.6.0.1

The Netcool Operations Insight V1.6.0 solution includes features delivered by the fix packs and fix pack extensions of the products and versions listed in the following topic:

“V1.6.0.1 Product and component version matrix” on page 13

The products are available for download from Passport Advantage® and Fix Central.

More information: “Products and components on premises” on page 9

New features in V1.6.0.1

The following features and functions are available in the Netcool Operations Insight V1.6.0.1 product:

Operations Management on IBM Cloud Private

The Operations Management on IBM Cloud Private solution provides the following new features and functions.

Cloud Native Analytics self-monitoring

A self-monitoring policy can be enabled to provide assurance that Cloud Native Analytics is processing events. For more information, see “Cloud Native Analytics Service Monitoring” on page 354.

Topology Analytics

View topological context for your Cloud Native Analytics events, where there is an associated resource. For more information, see “Topology Analytics on IBM Cloud Private” on page 356.

Connecting an on-premises Object Server to Operations Management on IBM Cloud Private

After you successfully deploy Operations Management on IBM Cloud Private, you can connect to an existing on-premises installation to create an event feed between your on-premises and cloud installations. For more information, see “Connecting an on-premises IBM Tivoli Netcool/OMNIbus ObjectServer to Operations Management on IBM Cloud Private” on page 196.

Upgrade

You can upgrade from V1.6.0 to V1.6.0.1. For more information, see “Upgrading Operations Management on IBM Cloud Private from 1.6.0 to 1.6.0.1” on page 153. For information about rolling back the upgrade, see “Rolling back Operations Management on IBM Cloud Private from V1.6.0.1 to V1.6.0” on page 161.

Storage

Network-based storage options such as Network File System (NFS) and GlusterFS are not supported. vSphere or local storage are the currently supported storage classes. For information, see “Storage” on page 104.

New product features and functions in V1.6.0

Updated product versions in V1.6.0

The Netcool Operations Insight V1.6.0 solution includes features delivered by the fix packs and fix pack extensions of the products and versions listed in the following topic:

“V1.6.0 Product and component version matrix” on page 17

The products are available for download from Passport Advantage and Fix Central.

More information: “Products and components on premises” on page 9

New features in V1.6.0

The following features and functions are available in the Netcool Operations Insight V1.6.0 product.

Operations Management on IBM Cloud Private

The Operations Management on IBM Cloud Private solution provides the following new features and functions.

Cloud Native Analytics

Cloud Native Analytics formulates event grouping policies by identifying seasonal and temporal patterns of events within a monitored environment. You can automatically deploy these policies, or you can manually select which policies are deployed. Events are consolidated into incidents, from which you can drill down into event details and time-lines, and see the seasonal, temporal and scope-based groups to which events belongs. For more information, see “Cloud Native Analytics on IBM Cloud Private” on page 342.

More configurable deployment options

Increased configuration options allow you to configure passwords, secrets and pod access, or to allow Operations Management on IBM Cloud Private to configure these for you during installation. For more information, see “Preparing for installation on IBM Cloud Private” on page 98.

Red Hat OpenShift support

Added support for Red Hat OpenShift. The Operations Management on IBM Cloud Private solution can be deployed on IBM Cloud Private with OpenShift. When you install IBM Cloud Private with OpenShift, IBM Cloud Private provides the IBM Cloud Private experience, management, and operations for applications and uses OpenShift's Kubernetes and Docker registry that is already installed by Red Hat. For more information, see IBM Cloud Private documentation: IBM Cloud Private with OpenShift.

Increased Security Assurance

Installation of Operations Management on IBM Cloud Private must now be performed with user-defined passwords and secrets or with randomly generated passwords and secrets, instead of with default passwords. For more information, see “Configuring passwords and secrets” on page 111.

Netcool/Impact scalability

The number of Netcool/Impact core server pods can be increased or decreased while Operations Management on IBM Cloud Private is running. For more information, see “Scaling the Netcool/Impact service” on page 141.

Certification Levels

Operations Management on IBM Cloud Private now has RedHat image certification, and IBM CloudPak level 2 certification, with improved resiliency, scaling and security. For more information, see https://www.ibm.com/support/knowledgecenter/en/SSBS6K_3.2.0/app_center/cloud_paks_over.html

Operations Management for Operations Insight

The base Netcool Operations Insight solution provides the following new features and functions.

Event source integration

IBM Cloud Event Management can be configured to forward public or private cloud events to Netcool Operations Insight , where they will appear in the Event Viewer. Netcool Operations Insight events can also be configured to display in IBM Cloud Event Management. For more information, see "Connecting event sources to your IBM Netcool Operations Insight on premises deployment" on page 189

Deployment modes

The base Netcool Operations Insight solution, Operations Management can be deployed in either of the following modes: on premises, or private cloud.

On premises

In this mode, all of the Netcool Operations Insight products and components are installed onto servers and use the native computing resources of those servers.

Private cloud

You can install the base Netcool Operations Insight solution, Operations Management in a private cloud, using IBM Cloud Private. In the Private cloud mode, distinct features within Netcool Operations Insight products and components, such as probes, the ObjectServer, Web GUI, and Operations Analytics - Log Analysis, run within containers, and communication between these containers is managed and orchestrated using IBM Cloud Private running on a Kubernetes cluster. For more information about deploying on IBM Cloud Private, see "Deploying on IBM Cloud Private" on page 41.

Deploying Operations Management in private cloud mode provides key benefits, including the following:

- Rapid installation and upgrade of Operations Management versions.
- Ability to start and stop Operations Management features by starting and stopping Docker containers.

It is not possible to upgrade from a version of base Netcool Operations Insight in on-premises mode to base Netcool Operations Insight in private cloud mode. However, it will be possible to install Netcool Operations Insight in private cloud mode, and subsequently migrate some or all of your data across from your Netcool Operations Insight on-premises system to Netcool Operations Insight in private cloud mode.

Note: Netcool Operations Insight in private cloud mode is limited in terms of the architecture that you are able to deploy and this in turn limits the event throughput from Netcool/OMNIbus to Operations Analytics - Log Analysis. For more details on these limitations, see "Deployment considerations for Operations Management in a private cloud" on page 42.

Products and components on IBM Cloud Private

It is possible to install the base Netcool Operations Insight solution, Operations Management, within a private cloud, by using IBM Cloud Private. Use this topic to understand the architecture of a deployment of Operations Management on IBM Cloud Private.

The Operations Management cluster is made up of a set of virtual machines that serve as nodes within the cluster. There is one master node, one management node, and the remaining virtual machines serve as worker nodes, where the Kubernetes pods and containers, which are known as workloads, are deployed. There are one or more containers within each pod.

The following figure shows the container architecture of a deployment of Operations Management on IBM Cloud Private.

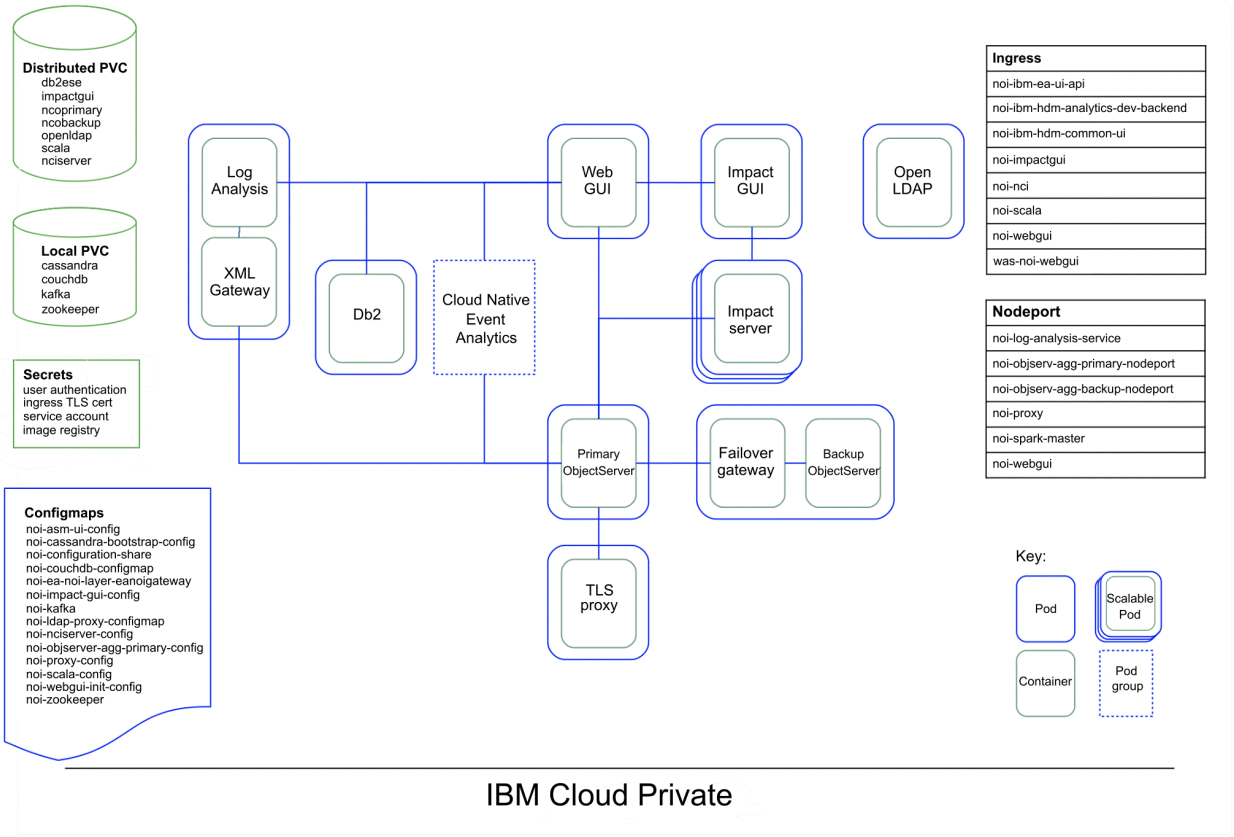

Figure 1. Architecture of a deployment of Operations Management on IBM Cloud Private

Some of the elements of the diagram are described here:

IBM Cloud Private

This is the underlying IBM Cloud Private system on which the Kubernetes cluster is deployed. IBM Cloud Pak supports the Operations Management workload, together with other deployed workloads. For Operations Management the cluster is made up of a minimum number of virtual machines, which are deployed as a master node (including management, proxy, and boot functions), and worker nodes within the cluster, together with a local storage file system.

For more information, see the following IBM Cloud Private documentation links:

- IBM Knowledge Center https://www.ibm.com/support/knowledgecenter/SSBS6K/product_welcome_cloud_private.html
- IBM Cloud Private System Administrator's Guide IBM Redbooks® http://www.redbooks.ibm.com/redbooks.nsf/RedbookAbstracts/sg248440.html?Open
- IBM Cloud Private Application Developer's Guide IBM Redbooks http://www.redbooks.ibm.com/redbooks.nsf/RedbookAbstracts/sg248441.html?Open

Operations Management cluster

The diagram displays an IBM Cloud Private installation, where the Operations Management cluster is deployed as containerized Operations Management applications within pods. The cluster is made up of a set of virtual machines that serve as nodes within the cluster. There is one master node, one management node, and the remaining virtual machines serve as worker nodes, where the Kubernetes pods and containers, which are known as workloads, are deployed. You can also create namespaces within your cluster. This enables multiple independent installations of Operations Management within the cluster, with each installation deployed in a separate namespace.

Interaction with the cluster is managed by the master node, as follows:

- Administration of the cluster is performed by connecting to the master node, either with the IBM Cloud Private GUI, or with the Kubernetes command-line interface, `kubectl`.
- Users log in to applications provided by the pods and containers within the cluster, such as Web GUI and Operations Analytics - Log Analysis, by opening a web browser and navigating to a URL made up of the master node hostname and the port number used by the relevant application.

Kubernetes manages how pods are deployed across the worker nodes and orchestrates communication between the pods. Pods are only deployed on worker nodes that meet the minimum resource requirements that are specified for that pod. Kubernetes uses a mechanism called affinity to ensure that pods that must be deployed on different worker nodes are deployed correctly. For example, it ensures that the primary ObjectServer container is deployed on a different worker node to the backup ObjectServer container.

Storage within the cluster
Storage within the cluster is provided by local Persistent Volumes. Distributed shared storage can be provided by vSphere. Currently no other distributed storage technologies are supported.

Persistent Volume Claims
Scalable network file systems within the cluster that provide storage to the pods in the cluster on demand. Local storage and vSphere storage are supported.

Pod and containers in the Operations Management cluster
Discrete Operations Management applications are deployed as containers within the cluster. One or more containers are deployed within pods. Where containers need to be tightly coupled, for example, the backup ObjectServer container and its associated bidirectional gateway container, they are deployed together within a pod to enable shared context, the same IP address and port space.

The pods and containers in the cluster are listed in the following table.

Table 1. Pods and containers in the cluster

Component or Capability	Application	Container	Pod
Netcool/OMNIbus	Primary ObjectServer	`ncoprimary`	`ncoprimary`
	Backup ObjectServer	`ncobackup-agg-b`	`ncobackup`
	Bidirectional gateway	`ncobackup-agg-gate`	
Netcool/Impact	Primary Netcool/Impact core server	`nciserver-0`	`nciserver-0`
	Backup Netcool/Impact core server	`nciserver-1`	`nciserver-1`
	Netcool/Impact GUI server	`impactgui`	`impactgui`
Db2® container	Db2 database	`db2ese`	`db2ese`
Event Search	Operations Analytics - Log Analysis	`unity`	`scala`
	Gateway for Message Bus	`gateway`	
Dashboard Application Services Hub GUIs	Dashboard Application Services Hub	`webgui`	`webgui`
LDAP	LDAP proxy server	`openldap`	`openldap`
Proxy	Proxy	`proxy`	`proxy`
Cloud Native Analytics	Cassandra	`cassandra`	`cassandra`
Cloud Native Analytics		`cassandra-change-super-user-post-install`	`cassandra`

Table 1. Pods and containers in the cluster (continued)			
Component or Capability	**Application**	**Container**	**Pod**
Cloud Native Analytics	CouchDB	`couchdb`	`couchdb`
Cloud Native Analytics	Cloud Native Analytics Action service	`ea-noi-layer-eanoiactionservice`	`ea-noi-layer-eanoiactionservice`
Cloud Native Analytics	Cloud Native Analytics Gateway	`ea-noi-layer-eanoigateway`	`ea-noi-layer-eanoigateway`
Cloud Native Analytics	Cloud Native Analytics UI API	`ea-ui-api-graphql`	`ea-ui-api-graphql`
Cloud Native Analytics	Cloud Native Analytics Collator service	`ibm-hdm-analytics-dev-collater-aggregationservice`	`ibm-hdm-analytics-dev-collater-aggregationservice`
Cloud Native Analytics	Cloud Native Analytics Deduplication service	`ibm-hdm-analytics-dev-dedup-aggregationservice`	`ibm-hdm-analytics-dev-dedup-aggregationservice`
Cloud Native Analytics	Cloud Native Analytics Normalization service	`ibm-hdm-analytics-dev-normalizer-aggregationservice`	`ibm-hdm-analytics-dev-normalizer-aggregationservice`
Cloud Native Analytics	Cloud Native Analytics Archiving service	`ibm-hdm-analytics-dev-archivingservice`	`ibm-hdm-analytics-dev-archivingservice`
Cloud Native Analytics	Cloud Native Analytics Event Query service	`ibm-hdm-analytics-dev-eventsqueryservice`	`ibm-hdm-analytics-dev-eventsqueryservice`
Cloud Native Analytics	Cloud Native Analytics Inference service	`ibm-hdm-analytics-dev-inferenceservice`	`ibm-hdm-analytics-dev-inferenceservice`
Cloud Native Analytics	Cloud Native Analytics Ingestion service	`ibm-hdm-analytics-dev-ingestionservice`	`ibm-hdm-analytics-dev-ingestionservice`
Cloud Native Analytics	Cloud Native Analytics Policy Registry service	`ibm-hdm-analytics-dev-policyregistryservice`	`ibm-hdm-analytics-dev-policyregistryservice`
Cloud Native Analytics	Cloud Native Analytics service monitor	`ibm-hdm-analytics-dev-servicemonitorservice`	`ibm-hdm-analytics-dev-servicemonitorservice`
Cloud Native Analytics	Cloud Native Analytics setup	`ibm-hdm-analytics-dev-setup`	`ibm-hdm-analytics-dev-setup`
Cloud Native Analytics	Cloud Native Analytics trainer	`ibm-hdm-analytics-dev-trainer`	`ibm-hdm-analytics-dev-trainer`
Cloud Native Analytics	Cloud Native Analytics UI server	`ibm-hdm-common-ui-uiserver`	`ibm-hdm-common-ui-uiserver`
Cloud Native Analytics	Kafka	`kafka`	`kafka`
Cloud Native Analytics	Redis sentinel	`redis-sentinel`	`redis-sentinel`

Table 1. Pods and containers in the cluster (continued)			
Component or Capability	**Application**	**Container**	**Pod**
Cloud Native Analytics	Redis server	`redis-server`	`redis-server`
Cloud Native Analytics	Spark master	`spark-master`	`spark-master`
Cloud Native Analytics	Spark slave	`spark-slave`	`spark-slave`
Cloud Native Analytics	Zookeeper	`zookeeper`	`zookeeper`
1.6.0.1 Cloud Native Analytics and Agile Service Manager topology analytics integration	Agile Service Manager Normaliser Mirror maker	`ibm-ea-asm-normalizer-mirrormaker`	`ibm-ea-asm-normalizer-mirrormaker`
1.6.0.1 Cloud Native Analytics and Agile Service Manager topology analytics integration	Agile Service Manager Normaliser Stream	`ibm-ea-asm-normalizer-normalizerstreams`	`ibm-ea-asm-normalizer-normalizerstreams`

Primary ObjectServer
This pod is made up of the Aggregation ObjectServer - primary container (`ncoprimary`). Containerized probes and on-premises probes connect to this container to send alert information by using the external node port for the pod.

Backup ObjectServer
This pod is made up of the Aggregation ObjectServer - backup container (`ncobackup-agg-b`) and a bidirectional gateway container (`ncobackup-agg-gate`). As in a traditional Netcool/OMNIbus ObjectServer pair this ObjectServer provides failover if the ObjectServer in the primary ObjectServer pod fails.

Primary Netcool/Impact core server
This pod is made up of the Netcool/Impact core server - primary container (`nciserver-0`). This container provides standard Netcool/Impact core server functionality.

Backup Netcool/Impact core server
This pod is made up of the Netcool/Impact core server - secondary container (`nciserver-1`). As in a traditional Netcool/Impact core server pair this core server provides failover if the core server in the primary Netcool/Impact core server pod fails.

Db2 database
This pod is made up of the Db2 container (`db2ese`), and is used by `webgui`.

Netcool/Impact GUI server
This pod is made up of the Netcool/Impact GUI server - secondary container (`impactgui`). This container provides standard Netcool/Impact core server functionality.

Operations Analytics - Log Analysis
This pod is made up of the Operations Analytics - Log Analysis container (`unity`) and the Gateway for Message Bus container (`gateway`). The Gateway for Message Bus container sends the event data using the Accelerated Event Notification (AEN) client to the Operations Analytics - Log Analysis container, where the event data is indexed.

Dashboard Application Services Hub
This pod is made up of the Web GUI container (`webgui`), which implements all of the GUIs used in Operations Management:

- Web GUI Event Viewer
- Event Analytics GUIs
- Event Search dashboards and GUIs

LDAP proxy server
By default LDAP proxy functionality is provided in the LDAP proxy container (`openldap`), which includes default LDAP configuration, predefined users, and single sign-on capability for those users. All of the other pods communicate with the LDAP proxy pod to support this default functionality. You can configure the LDAP container to change the LDAP server that the pods in the cluster connect to.

Cloud Native Analytics Action service
This service is responsible for updating the ObjectServer with the findings from the Cloud Native Analytics Inference service and the Cloud Native Analytics Collator service. It updates and enriches entries in the ObjectServer with correlation, seasonal and other data.

Cloud Native Analytics Gateway
This service sends ObjectServer *alerts.status* insertions to the Cloud Native Analytics Ingestion service, for archiving and processing by the Cloud Native Analytics Inference Service.

Cloud Native Analytics UI API
Internal API service used by the UI.

Cloud Native Analytics Collator service
This service generates supergroups where groups have one or more common events. It also generates metrics on the groups and supergroups that are applied to individual events.

Cloud Native Analytics Deduplication service
This service de-duplicates event actions that are received from the Inference Service into single entries. These are then used to update the ObjectServer, and by the Cloud Native Analytics Collator service.

Cloud Native Analytics Normalization service
This service takes the output of the Cloud Native Analytics Collator service from Kafka, and posts the items to a REST endpoint. Currently, the endpoint is hosted by the Cloud Native Analytics NOI Action service but it can be any service that supports the API endpoint.

Cloud Native Analytics Archiving service
This service receives events that are published by the Cloud Native Analytics Ingestion service on the internal Cloud Native Analytics Kafka events topic, and creates occurrence and instance records in the underlying Cassandra database. This event data is used for training algorithms, and by the Cloud Native Analytics UI to query events that appear in correlation groups of seasonal enrichment, by using the Cloud Native Analytics Event Query service.

Cloud Native Analytics Event Query service
This service provides a REST API for querying the Cassandra database to find NOI event instances.

Cloud Native Analytics Inference service
This service receives events and then applies relevant policies to infer relevant correlations and enrichment actions for events.

Cloud Native Analytics Ingestion service
The Ingestion service provides a REST API for the ingestion of events in NOI format into the Cloud Native Analytics backend. Events are validated and converted into an internal Cloud Native Analytics event representation before they are published onto the internal Cloud Native Analytics Kafka events topic.

Cloud Native Analytics Policy Registry service
This service provides a REST API for managing and querying policies that have been created by an algorithm or a user for correlation or event enrichment. It is used by the Cloud Native Analytics Inference service to fetch the set of policies that are related to a set of event IDs, and by the Cloud Native Analytics UI for fetching details on specific policies that have run.

Cloud Native Analytics Service Monitor service
This service provides a single REST API that can be queried to get the current service health status of all other services in the deployment. The service has a list of deployed service endpoints, and queries each one and then returns a consolidated service view to the caller.

Cloud Native Analytics Setup service
The setup job is created and run as part of the deployment of NOI. It performs the initial setup and creation tasks that are required by the deployment, such as creating the necessary table schemas in the underlying Cassandra database, creating the default 'ScopeID' based scope correlation policy in the Cloud Native Analytics Policy Registry service, and creating the initial training schedule for the available algorithms in the trainer.

Cloud Native Analytics Trainer service
Manages the training schedules for all of the algorithms, and starts retraining jobs in the Spark cluster.

Cloud Native Analytics UI server
This service hosts and serves the key UI components of the Cloud Native Analytics.

1.6.0.1 **Cloud Native Analytics Agile Service Manager Normaliser Mirror maker**
Mirrors the Agile Service Manager Kafka status topic.

1.6.0.1 **Cloud Native Analytics Agile Service Manager Normaliser Streams**
Combines ASM status with event data from Operations Management's Kafka topic.

Products and components on premises

Review the products and components included in Netcool Operations Insight.

IBM Netcool Operations Insight includes the product and component versions listed in the following topic: “V1.6.0.1 Product and component version matrix” on page 13. This topic also includes information on the eAssemblies and fix packs required to download and install.

Product and component details

Tivoli Netcool/OMNIbus core components V8.1.0
This product includes the following components. It is installed by Installation Manager. It is part of the base Netcool Operations Insight solution, so it must be installed, configured, and running before you can start the Networks for Operations Insight feature setup.

- Server components (includes ObjectServers)
- Probe and gateway feature
- Accelerated Event Notification (AEN) client

For systems requirements, see http://ibm.biz/Bd2LHA.

Important: The ObjectServer that manages the event data must be at V8.1.0.

Tivoli Netcool/OMNIbus Web GUI V8.1.0
This component includes the following subcomponents and add-ons. It is installed by Installation Manager. It is part of the base Netcool Operations Insight solution. The following extensions to the Web GUI are supplied in Netcool Operations Insight:

- Tools and menus for integration with Operations Analytics - Log Analysis.
- Extensions for Netcool Operations Insight: This supports the Event Analytics capability.

  Important: Both the Impact Server Extensions and the Web GUI extensions must be installed for the Event Analytics capability to work.

The Web GUI is installed into Dashboard Application Services Hub, which is part of Jazz® for Service Management. Jazz for Service Management is distributed as separate installation features in Installation Manager. For systems requirements, see http://ibm.biz/Bd2LHt .

Db2 Enterprise Server Edition database
Db2 is the default database used for the Netcool Operations Insight solution. Other types of databases are also possible.

- Db2 Enterprise Server Edition V**11.1** is for use with Operations Management components. For systems requirements, see http://ibm.biz/Bd2L4E.

Gateway for JDBC
This product is required for the base Netcool Operations Insight solution. It is installed by Installation Manager. The system requirements are the same as for Tivoli Netcool/OMNIbus V8.1. It is required for the transfer of event data from the ObjectServer to the IBM Db2 database.

Netcool/Impact V7.1.0
This product includes the following components. It is part of the base Netcool Operations Insight solution. It is installed by Installation Manager.

- Impact server
- GUI server
- Impact Server extensions: Includes the policies that are used to create the event analytics algorithms and the integration to IBM Connections.

  Important: Both the Impact Server Extensions and the Web GUI extensions must be installed for the Event Analytics capability to work.

For system requirements, see http://ibm.biz/Bd2L4Y.

IBM Operations Analytics - Log Analysis V1.3.5 and V1.3.6
Netcool Operations Insight works with IBM Operations Analytics - Log Analysis V1.3.3, 1.3.5 and 1.3.6. IBM Operations Analytics - Log Analysis is part of the base Netcool Operations Insight solution. It is installed by Installation Manager. For system requirements, search for "Hardware and software requirements" within the relevant version of IBM Operations Analytics - Log Analysis at https://www.ibm.com/support/knowledgecenter/SSPFMY.

Note: Operations Analytics - Log Analysis V1.3.6 is available with Netcool Operations Insight V1.6.0.1. Operations Analytics - Log Analysis 1.3.5 is available with earlier versions of Netcool Operations Insight.

Note: Operations Analytics - Log Analysis Service Desk Extension V1.1 is available with IBM Operations Analytics - Log Analysis V1.3.5.

Note: Operations Analytics - Log Analysis Standard Edition is included in Netcool Operations Insight. For more information about Operations Analytics - Log Analysis editions, search for "Editions" at the Operations Analytics - Log Analysis Knowledge Center, at https://www.ibm.com/support/knowledgecenter/SSPFMY.

OMNIbusInsightPack_v1.3.1 for IBM Operations Analytics - Log Analysis
This product is part of the base Netcool Operations Insight solution. It is required to enable the event search capability in Operations Analytics - Log Analysis. The Insight Pack is installed into Operations Analytics - Log Analysis.

Gateway for Message Bus V8.0
This product is part of the base Netcool Operations Insight solution. It is installed by Installation Manager. The system requirements are the same as for Tivoli Netcool/OMNIbus V8.1.0. It is used for the following purposes:

- Transferring event data to the IBM Operations Analytics - Log Analysis product.
- Supports the transfer of event data to Agile Service Manager by integrating with the Agile Service Manager Event Observer.

Jazz for Service Management V1.1.3.0 and V1.1.3.5

Note: Jazz for Service Management 1.1.3.5 is available with Netcool Operations Insight V1.6.0.1. Jazz for Service Management V1.1.3.0 is available with earlier versions of Netcool Operations Insight.

This component provides the GUI framework for the Netcool Operations Insight solution. It is installed by Installation Manager, and it includes the following subcomponents.

- Dashboard Application Services Hub V3.1.3.1
- Reporting Services (previously called Tivoli Common Reporting)

Note: For the cumulative patch to use for this version of Jazz for Service Management, see the web page for the relevant version of Netcool Operations Insight at this location: https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Netcool%20OMNIbus/page/Release%20details

For the system requirements for Dashboard Application Services Hub, see http://ibm.biz/BdiVYN. The instance of Dashboard Application Services Hub hosts the V8.1 Web GUI and the Seasonal Event Reports portlet. Jazz for Service Management is included in the Web GUI installation package but is installed as separate features.

You can set up Network Manager and Netcool Configuration Manager to work with Reporting Services by installing their respective reports when installing the products. Netcool/OMNIbus V8.1.0 and later can be integrated with Reporting Services V3.1 to support reporting on events. To configure this integration, connect Reporting Services to a relational database through a gateway. Then, import the report package that is supplied with Netcool/OMNIbus into Reporting Services. For more information about event reporting, see http://www-01.ibm.com/support/knowledgecenter/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/omnibus/wip/install/task/omn_con_ext_deploytcrreports.html.

Network Manager IP Edition V4.2.0

This product includes the core and GUI components for the optional Networks for Operations Insight feature.

For system requirements, see the following links:

- http://www-01.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/itnm/ip/wip/install/task/nmip_pln_planninginst.html
- http://ibm.biz/Bd2L4h

Network Manager Insight Pack V1.3.0.0 for IBM Operations Analytics - Log Analysis

This product is part of the Networks for Operations Insight feature. It is required to enable the topology search capability in Operations Analytics - Log Analysis. The Insight Pack is installed into Operations Analytics - Log Analysis. It requires that the OMNIbusInsightPack_v1.3.1 is installed.

Note: The Network Manager Insight Pack V1.3.0.0 can share a data source with the OMNIbusInsightPack_v1.3.1 only. It cannot share a data source with previous versions of the Tivoli Netcool/OMNIbus Insight Pack.

Probe for SNMP

This product is optional for the base Netcool Operations Insight solution. It is used in environments that have SNMP traps. It is required for the Networks for Operations Insight feature. For installations of the probe on the Tivoli Netcool/OMNIbus V8.1 server, use the instance of the probe that installs with IBM Installation Manager.

Syslog Probe

This product is optional for the base Netcool Operations Insight solution. It is required for the Networks for Operations Insight feature. For installations of the probe on the Tivoli Netcool/OMNIbus V8.1 server, use the instance of the probe that installs with IBM Installation Manager.

Netcool Configuration Manager V6.4.2

This product has the following components. It is part of the optional Networks for Operations Insight feature.

- Core components
- Drivers
- OOBC component

For system requirements, see http://ibm.biz/Bd2L4J.

IBM Network Performance Insight V1.3.1

Network Performance Insight is a network traffic performance monitoring system. It provides comprehensive and scalable visibility on network traffic with visualization and reporting of network performance data for complex, multivendor, multi-technology networks. The end user is able to perform the following tasks: visualize flow across selected interfaces, display performance anomaly

events in the Tivoli Netcool/OMNIbus **Event Viewer**, and view performance anomaly and performance timeline data in the **Device Dashboard**. For more information, see http://www-01.ibm.com/support/knowledgecenter/SSCVHB/welcome.

IBM Alert Notification

IBM Alert Notification provides instant notification of alerts for any critical IT issues across multiple monitoring tools. It gives IT staff instant notification of alerts for any issues in your IT operations environment. For more information, see http://www-01.ibm.com/support/knowledgecenter/SSY487/com.ibm.netcool_OMNIbusaas.doc_1.2.0/landingpage/product_welcome_alertnotification.html.

IBM Runbook Automation

IBM Runbook Automation empowers IT operations teams to be more efficient and effective. Operators can focus their attention where it is really needed and receive guidance to the best resolution with recommended actions and pre-filled context. With Runbook Automation you can:

- Investigate and delegate problems faster and more efficiently.
- Diagnose and fix problems faster and build operational knowledge.
- Easily create, publish, and manage runbooks and automations.
- Keep score to track achievements and find opportunities for improvement.

For more information, see http://www-01.ibm.com/support/knowledgecenter/SSZQDR/com.ibm.rba.doc/RBA_welcome.html.

More information

For more information about the component products of Netcool Operations Insight, see the websites listed in the following table.

Table 2. Product information

Product	Website
IBM Netcool Operations Insight	http://www.ibm.com/support/knowledgecenter/SSTPTP/welcome
IBM Tivoli Netcool/OMNIbus and Web GUI	http://www-01.ibm.com/support/knowledgecenter/SSSHTQ/landingpage/NetcoolOMNIbus.html
IBM Tivoli Netcool/Impact	http://www-01.ibm.com/support/knowledgecenter/SSSHYH/welcome
IBM Operations Analytics - Log Analysis	http://www-01.ibm.com/support/knowledgecenter/SSPFMY/welcome
Jazz for Service Management	http://www.ibm.com/support/knowledgecenter/SSEKCU/welcome
IBM Tivoli Network Manager	Network Manager Knowledge Center
IBM Tivoli Netcool Configuration Manager	http://www-01.ibm.com/support/knowledgecenter/SS7UH9/welcome
Network Performance Insight	http://www-01.ibm.com/support/knowledgecenter/SSCVHB/welcome
Agile Service Manager	https://www-01.ibm.com/support/knowledgecenter/SS9LQB/welcome
Runbook Automation	http://www-01.ibm.com/support/knowledgecenter/SSZQDR/com.ibm.rba.doc/RBA_welcome.html
Alert Notification	http://www.ibm.com/support/knowledgecenter/SSY487/com.ibm.netcool_OMNIbusaas.doc_1.2.0/landingpage/product_welcome_alertnotification.html

V1.6.0.1 Product and component version matrix

This page lists the products and components that are supported in the current on premises release of Netcool Operations Insight, and the packages for Netcool Operations Insight on IBM Cloud Private. The current version is 1.6.0.1. Only this combination of product and component releases is supported.

Installing on IBM Cloud Private

This section lists the Netcool Operations Insight on IBM Cloud Private V1.6.0.1 eAssemblies.

Table 3. Netcool Operations Insight on IBM Cloud Private V1.6.0.1 eAssemblies

eAssembly name	**Download from Passport Advantage**
IBM Cloud Private Foundation 3.2 Installation Packages eAssembly for Netcool Operations Insight	CJ5N0EN
IBM Netcool Operations Insight certified containers	CJ5MZEN

Parts

Download the parts from the CJ5N0EN IBM Cloud Private Foundation 3.2 Installation Packages eAssembly depending on your requirements.

- **CC1W1EN:** IBM Cloud Private 3.2 for Linux (x86_64) Docker
- **CC1W6EN:** IBM Cloud Private 3.2 Docker for Linux (x86_64)
- **CC1W3EN:** IBM Cloud Private 3.2 for Red Hat Enterprise Linux OpenShift (64-bit) Docker

Download the parts from the CJ5MZEN Netcool Operations Insight certified containers eAssembly depending on your requirements.

- **CC3V5EN:** Always download this part.
- **CC3V4EN:** Download this part only if you require the optional Service Management extension (IBM Agile Service Manager).
- **CC3PMML:** Download this part only if you require the optional Cloud Event Management extension (IBM Cloud Event Management for IBM Cloud Private).
- **CC3PNML:** Download this part only if you require the optional Cloud Event Management extension on OpenShift (IBM Cloud Event Management for IBM Cloud Private on OpenShift).

You can also download the following parts from the CJ5MZEN Netcool Operations Insight certified containers eAssembly:

- **CC3V6EN:** Netcool/OMNIBus Probe for Cloud Monitoring Integration 3.10.4.1 for IBM Cloud Private Linux x86-64
- **CC43PEN:** Netcool/OMNIBus Probe for Kafka Integration 3.10.2.1 for IBM Cloud Private Linux x86-64
- **CC43NEN:** Netcool/OMNIBus Gateway for Cloud Event Management Integration 3.2.1 for IBM Cloud Private Linux x86-64

Installing on premises

The following sections list the Netcool Operations Insight on premises V1.6.0.1 product and component versions.

Supported versions

In general, only where we provide the details of the supported versions should they be considered as being supported as part of the Netcool Operations Insight solution for a particular release. The reasons for this are as follows:

- The features of a given release of Netcool Operations Insight are delivered using the versions of the component software in that release. If you do not have the latest versions of the component software, then you will not get all of the features in the Netcool Operations Insight release.
- Testing is done only on the component versions that make up a Netcool Operations Insight release, and this combination of products is what is supported. If you install a different set of component versions as part of your Netcool Operations Insight solution, then this is not a tested combination and we do not guarantee to support it.

Version information

Click a link below to display the relevant version information.

- Product and component versions
- Integrations

Product and component versions

The following table lists the Netcool Operations Insight on premises V1.6.0.1 product and component versions.

Note: If a table cell in either the **Download from Passport Advantage** column or the **Download from Fix Central** column is empty then there is nothing to download from that location for that particular product or component.

Install Netcool Operations Insight 1.6.0.1 using IBM Installation Manager V1.8.x.

Table 4. Netcool Operations Insight on premises V1.6.0.1 product and component versions

Product or component	Version	Change from previous release?	Download from Passport Advantage	Download from Fix Central
IBM Cloud Event Management		Yes	Subscribe to IBM Cloud Event Management on Marketplace: https://www.ibm.com/cloud/event-management?mhq=Cloud%20Event%20Management&mhsrc=ibmsearch_p	
IBM Cloud Event Management for IBM Cloud Private	IBM Cloud Event Management runs on IBM Cloud Private V3.2.0. **Note:** Download the IBM Cloud Private Foundation 3.2 Installation Packages eAssembly for Netcool Operations Insight from Passport Advantage: CJ5N0EN	Yes	CC3PMML	

Table 4. Netcool Operations Insight on premises V1.6.0.1 product and component versions (continued)				
Product or component	**Version**	**Change from previous release?**	**Download from Passport Advantage**	**Download from Fix Central**
IBM Cloud Event Management for IBM Cloud Private on OpenShift		Yes	CC3PNML	
IBM Tivoli Netcool/ OMNIbus core components	V8.1.0.21	Yes	CJ5N1EN	V8.1.0 Fix Pack 21
Tivoli Netcool/ OMNIbus Web GUI	V8.1.0.17	Yes		V8.1.0 Fix Pack 17
IBM Tivoli Netcool/ OMNIbus 8 Plus Gateway for Message Bus	V8.0	No	CNVL8EN	
IBM Tivoli Netcool/ OMNIbus 8 Plus Gateway for JDBC	V8.0	No	CN4FUEN	
IBM Tivoli Netcool/ OMNIbus 8 Plus JDBC Gateway Configuration Scripts	V8.0	No	CN1FLEN	
IBM Tivoli Netcool/ Impact	V7.1.0.17	Yes	CJ5N2EN	V7.1.0 Fix Pack 17
Db2	V11.1 For use with Operations Management components	No	CJ3INML	
Operations Analytics - Log Analysis	V1.3.6	Yes	CJ5N3EN	
Operations Analytics - Log Analysis Service Desk Extension	V1.1.0	No	CJ1NHEN **Note:** Only available with CJ5N3EN	
IBM Operations Analytics Advanced Insights Multiplatform English eAssembly	V1.3.6	Yes	CJ47JEN	

Table 4. Netcool Operations Insight on premises V1.6.0.1 product and component versions (continued)

Product or component	Version	Change from previous release?	Download from Passport Advantage	Download from Fix Central
Event Analytics	IBM Tivoli Netcool/ Impact Server Extensions for IBM Netcool Operations Insight_7.1.0.17	Yes		Included in Netcool/ Impact V7.1.0 Fix Pack 17
	IBM Netcool Operations Insight Extension for IBM Tivoli Netcool/ OMNIbus Web GUI_8.1.0.17	Yes		Included in Web GUI V8.1.0 Fix Pack 17
Event Search	Tivoli Netcool/ OMNIbus Insight Pack V1.3.1	No	CNS6GEN Included in Operations Analytics - Log Analysis **V1.3.5** eAssembly.	
	Tivoli Netcool/ OMNIbus Insight Pack V1.3.0.2	No	CN8IPEN Included in Operations Analytics - Log Analysis **V1.3.3** eAssembly.	
Topology Search	Network Manager Insight Pack V1.3.0.0	No	CNZ43EN Included in Operations Analytics - Log Analysis eAssembly.	
IBM Tivoli Network Manager IP Edition	V4.2.0.7	Yes	CJ5N4EN	V4.2.0 Fix Pack 7
Device Dashboard	V1.1.0.2	No	CJ1U2EN	V1.1 Fix Pack 2
Network Health Dashboard	V4.2.0.7	No	CJ0S2EN	V4.2.0 Fix Pack 7
IBM Tivoli Netcool Configuration Manager	V6.4.2.8	Yes	CJ5N5EN	V6.4.2 Fix Pack 8
IBM Network Performance Insight	V1.3.1	Yes	CJ5N6EN	
IBM Agile Service Manager	V1.1.6	Yes	CJ5N7EN	

Table 4. Netcool Operations Insight on premises V1.6.0.1 product and component versions (continued)				
Product or component	**Version**	**Change from previous release?**	**Download from Passport Advantage**	**Download from Fix Central**
IBM Agile Service Manager Observers	V1.1.6	Yes	CJ5N8EN	
Jazz for Service Management	V1.1.3.5	Yes	CJ49VML	V1.1.3.5
WebSphere® Application Server	V8.5.5.15	Yes	Included in Jazz for Service Management.	V8.5.5 Fix Pack 15
Java™ SDK for WebSphere Application Server	V8.0.5.6			
IBM Cognos® Analytic Server	V11	Yes	CJ5LIML	

Integrations

The following table lists the products that can be integrated with NOI, together with the versions of these products that are compatible with this release NOI.

Note: If a table cell in either the **Download from Passport Advantage** column or the **Download from Fix Central** column is empty then there is nothing to download from that location for that particular product or component.

Table 5. Products that can be integrated with Netcool Operations Insight			
Product	**Version**	**Download from Passport Advantage**	**Download from Fix Central**
IBM Tivoli Monitoring	V6.3.0	CRS7JML	
IBM Tivoli Monitoring Agents		CRS7KML	
IBM Tivoli Monitoring Agents for Tivoli Network Manager IP Edition V4.2	V6.3.0.2	CRYY6ML	
Tivoli Business Service Manager	V6.1.1.5	CRL8FML	Fix Pack 5

V1.6.0 Product and component version matrix

This page lists the products and components that are supported in the on premises release of Netcool Operations Insight V1.6.0, and the packages for Netcool Operations Insight on IBM Cloud Private. Only this combination of product and component releases is supported in V1.6.0.

Installing on IBM Cloud Private

This section lists the Netcool Operations Insight on IBM Cloud Private V1.6.0 eAssemblies.

Table 6. Netcool Operations Insight on IBM Cloud Private V1.6.0 eAssemblies	
eAssembly name	**Download from Passport Advantage**
IBM Cloud Private Foundation 3.2 Installation Packages eAssembly for Netcool Operations Insight	CJ5N0EN
IBM Netcool Operations Insight certified containers	CJ5MZEN

Parts

Download the parts from the CJ5MZEN Netcool Operations Insight certified containers eAssembly depending on your requirements.

- **CC28XEN:** Always download this part.
- **CC28WEN:** Download this part only if you require the optional Service Management extension (IBM Agile Service Manager).

Installing on premises

The following sections list the Netcool Operations Insight on premises V1.6.0 product and component versions.

Supported versions

In general, only where we provide the details of the supported versions should they be considered as being supported as part of the Netcool Operations Insight solution for a particular release. The reasons for this are as follows:

- The features of a given release of Netcool Operations Insight are delivered using the versions of the component software in that release. If you do not have the latest versions of the component software, then you will not get all of the features in the Netcool Operations Insight release.
- Testing is done only on the component versions that make up a Netcool Operations Insight release, and this combination of products is what is supported. If you install a different set of component versions as part of your Netcool Operations Insight solution, then this is not a tested combination and we do not guarantee to support it.

Version information

Click a link below to display the relevant version information.

- Product and component versions
- Integrations

Product and component versions

The following table lists the Netcool Operations Insight on premises V1.6.0 product and component versions.

Note: If a table cell in either the **Download from Passport Advantage** column or the **Download from Fix Central** column is empty then there is nothing to download from that location for that particular product or component.

Install Netcool Operations Insight V1.6.0 using IBM Installation Manager V1.8.x.

Table 7. Netcool Operations Insight on premises V1.6.0 product and component versions

Product or component	Version	Change from previous release?	Download from Passport Advantage	Download from Fix Central
IBM Tivoli Netcool/ OMNIbus core components	V8.1.0.19	Yes	CJ5N1EN	V8.1.0 Fix Pack 19
Tivoli Netcool/ OMNIbus Web GUI	V8.1.0.16	Yes		V8.1.0 Fix Pack 16
IBM Tivoli Netcool/ OMNIbus 8 Plus Gateway for Message Bus	V8.0	No	CNVL8EN	
IBM Tivoli Netcool/ OMNIbus 8 Plus Gateway for JDBC	V8.0	No	CN1FMEN	
IBM Tivoli Netcool/ OMNIbus 8 Plus JDBC Gateway Configuration Scripts	V8.0	No	CN1FLEN	
IBM Tivoli Netcool/ Impact	V7.1.0.16	Yes	CJ5N2EN	V7.1.0 Fix Pack 16
Db2	V11.1 For use with Operations Management components	No	CJ3INML	
Operations Analytics - Log Analysis	V1.3.5.3	Yes	CJ5N3EN	V1.3.5 fix pack 3
Operations Analytics - Log Analysis Service Desk Extension	V1.1.0	No	CJ1NHEN **Note:** Only available with CJ5N3EN	
IBM Operations Analytics Advanced Insights Multiplatform English eAssembly	V1.3.6	Yes	CJ47JEN	

Table 7. Netcool Operations Insight on premises V1.6.0 product and component versions (continued)

Product or component	Version	Change from previous release?	Download from Passport Advantage	Download from Fix Central
Event Analytics	IBM Tivoli Netcool/ Impact Server Extensions for IBM Netcool Operations Insight_7.1.0.16	Yes		Included in Netcool/ Impact V7.1.0 Fix Pack 16
	IBM Netcool Operations Insight Extension for IBM Tivoli Netcool/ OMNIbus Web GUI_8.1.0.16	Yes		Included in Web GUI V8.1.0 Fix Pack 16
Event Search	Tivoli Netcool/ OMNIbus Insight Pack V1.3.1	No	CNS6GEN Included in Operations Analytics - Log Analysis **V1.3.5** eAssembly.	
	Tivoli Netcool/ OMNIbus Insight Pack V1.3.0.2	No	CN8IPEN Included in Operations Analytics - Log Analysis **V1.3.3** eAssembly.	
Topology Search	Network Manager Insight Pack V1.3.0.0	No	CNZ43EN Included in Operations Analytics - Log Analysis eAssembly.	
IBM Tivoli Network Manager IP Edition	V4.2.0.7	Yes	CJ5N4EN	V4.2.0 Fix Pack 7
Device Dashboard	V1.1.0.2	No	CJ1U2EN	V1.1 Fix Pack 2
Network Health Dashboard	V4.2.0.7	No	CJ0S2EN	V4.2.0 Fix Pack 7
IBM Tivoli Netcool Configuration Manager	V6.4.2.8	Yes	CJ5N5EN	V6.4.2 Fix Pack 8
IBM Network Performance Insight	V1.3.1	Yes	CJ5N6EN	
IBM Agile Service Manager	V1.1.5	Yes	CJ5N7EN	

Table 7. Netcool Operations Insight on premises V1.6.0 product and component versions (continued)				
Product or component	**Version**	**Change from previous release?**	**Download from Passport Advantage**	**Download from Fix Central**
IBM Agile Service Manager Observers	V1.1.5	Yes	CJ5N8EN	
Jazz for Service Management	V1.1.3.3	Yes	CJ49VML	V1.1.3.3
WebSphere Application Server	V8.5.5.15	Yes	Included in Jazz for Service Management.	V8.5.5 Fix Pack 15
Java SDK for WebSphere Application Server	V8.0.5.6			
IBM Cognos Analytic Server	V11	Yes	CJ5LIML	

Integrations

The following table lists the products that can be integrated with NOI, together with the versions of these products that are compatible with this release NOI.

Note: If a table cell in either the **Download from Passport Advantage** column or the **Download from Fix Central** column is empty then there is nothing to download from that location for that particular product or component.

Table 8. Products that can be integrated with NOI			
Product	**Version**	**Download from Passport Advantage**	**Download from Fix Central**
IBM Tivoli Monitoring	V6.3.0	CRS7JML	
IBM Tivoli Monitoring Agents		CRS7KML	
IBM Tivoli Monitoring Agents for Tivoli Network Manager IP Edition V4.2	V6.3.0.2	CRYY6ML	
Tivoli Business Service Manager	V6.1.1.5	CRL8FML	Fix Pack 5

About Netcool Operations Insight

IBM Netcool Operations Insight consists of a base operations management solution. It can be optionally extended by integrating Network Management, Performance Management, and Service Management solution extensions.

The full name of the Netcool Operations Insight base solution is Operations Management for Operations Insight. This base solution provides the capability of monitoring the health and performance of IT and network infrastructure across local, cloud and hybrid environments. It also incorporates strong event

management capabilities, and leverages real-time alarm and alert analytics, combined with broader historic data analytics.

You can optionally extend this base solution by adding the following solution extensions:

Network Management for Operations Insight
Network Management adds network discovery, visualization, event correlation, topology-based root-cause analysis, and configuration and compliance management capabilities. It also adds network dashboarding and topology search capabilities. The extension is provided by integrating the Network Manager and Netcool Configuration Manager products.

Performance Management for Operations Insight
Performance Management adds performance management capability, including a wide range of dashboarding and flow capabilities . The extension is provided by integrating the Network Performance Insight product.

Service Management for Operations Insight
Service Management adds service management capability, including up-to-date visibility and control over dynamic infrastructure and services. For example, you can query a specific networked resource, and then presents a configurable topology view of it within its ecosystem of relationships and states, both in real time and within a definable time window. The extension is provided by integrating the Agile Service Manager product.

About Operations Management

Use this information to understand more about the Netcool Operations Insight base solution.

Operations Management capabilities

Use this information to understand the capabilities of Operations Management.

Operations Management is made up of the following products and components:

- IBM Tivoli Netcool/OMNIbus
- Tivoli Netcool/OMNIbus Web GUI
- IBM Tivoli Netcool/Impact
- IBM Operations Analytics - Log Analysis
- Event Analytics (on premises) and Cloud Native Analytics (on ICP)
- Event Search
- IBM Connections Integration

Operations Management leverages real-time alarm and alert analytics, combined with broader historic data analytics. Netcool Operations Insight is powered by the fault management capabilities of IBM Tivoli Netcool/OMNIbus and IBM's leading big data technologies within IBM Operations Analytics - Log Analysis, providing powerful event search and historical analysis in a single solution. Operations Management integrates infrastructure and operations management into a single coherent structure across business applications, virtualized servers, network devices and protocols, internet protocols, and security and storage devices, and includes the following capabilities:

The components and capabilities of Operations Management are described below:

1.6.0.1 **“Topology Analytics on IBM Cloud Private” on page 356**
Provides topological context for events. (Requires IBM Agile Service Manager extension).

Event Analytics (on premises) and “Cloud Native Analytics on IBM Cloud Private” on page 342
These provide analysis of event data to correlate and group events, and reduce the number of events that are presented to the operator.

Event Search
Event search applies the search and analysis capabilities of Operations Analytics - Log Analysis to events that are monitored and managed by Tivoli Netcool/OMNIbus.

IBM Connections Integration
Netcool/Impact enables social collaboration through IBM Connections by automatically providing updates to key stake holders.

Operations Management tasks

Use this information to understand the tasks that users can perform using Operations Management.

Operations Management tasks fall into the following categories:

- Event Search tasks
- Event Analytics tasks

Event Search tasks

Using Event Search, Operations staff can use the analytics available in Operations Analytics - Log Analysis to determine how the monitoring environment is performing over time.

Using Event Search
Network operators can diagnose and triage events in the **Event Viewer** by using the search and analysis capabilities within Event Search. An example of this is the use of Event Search to narrow down the cause of an event storm by running the Event Search dashboard and timeline tools against selected events in the **Event Viewer**.

Configuring Event Search
Administrators can make extra event data available within Event Search to provide Operations with a more semantically rich set of data to use in Event Search dashboard and timeline tools. This, in turn, helps operators to more effectively diagnose and triage events using Event Search.

Administrators can also customize Event Search dashboards to enable Operations to more effectively analyze event data.

Event Analytics (NOI on premises) tasks

Using Event Analytics, Operations staff can determine event patterns, groups, and seasonality, and use this knowledge to build rules that create parent and synthetic events, thereby reducing event count and presenting operators with events that are closer to the underlying incidents.

Using Event Analytics
Operations staff can review generated analytics reports, and drill into the reports to see seasonality graphs, related event groups, and event patterns. Based on an analysis of the report data, they can set up rules to act on live events and thereby reduce event count and improve the quality of the events in the **Event Viewer**.

Configuring Event Analytics
Administrators can customize Event Analytics in a variety of ways:

- Making custom data available within seasonal and related event reports to provide Operations with a richer set of analytics data.
- Changing the mechanism used by seasonality to suppress events.
- Configuring how event pattern processing is performed.

In addition, administrators can set up configuration scans to run against historical data over a specified time range. They can specify which type of analytic to run and can set up a schedule so that analytics reports are automatically generated for Operations.

Cloud Native Analytics (NOI on ICP) tasks

Cloud Native Analytics analyzes historic and live event data to create event grouping policies that reduce the number of events displayed in the **Event Viewer**, and present operators with events that are closer to the underlying incident.

Using Cloud Native Analytics

Cloud Native Analytics analyzes historic and live event data and suggests policies to group events together in a more coherent way for the operator than ungrouped single events. Policies group events in temporal groups where events usually occur together, and in seasonal groups where events consistently occur at the same times or dates. Operators can also create scope-based policies that group events together based on a common attribute such as a resource (scope-based grouping). Users can choose which policies to deploy, and these policies then group incoming events under a parent event, from which the user can drill down to individual events if required. Scheduled training runs ensure that the grouping policies maintain their relevance to the stream of incoming events.

Configuring Cloud Native Analytics

Administrators can customize Cloud Native Analytics in various ways:

- Determine the event data to be analyzed to generate policies.
- Compare and rank the policy groupings.
- Reject or approve policies for deployment.

Topology Analytics on IBM Cloud Private tasks

1.6.0.1

(Requires IBM Agile Service Manager extension.) With Topology Analytics, events that have an associated resource in the topology are enriched with topological information. Users can see topological context when they are investigating an incident, and drill down into the topology that relates to an event, enabling faster identification and resolution of problems. Users can use the following features:

- Use the Event Viewer to see when an event is associated with a status in Agile Service Manager.
- Use the Incident Viewer with an embedded topology hop view to see the relationship between an incident's events and any associated topology.
- Launch to a full Topology Viewer for all topologically enriched events.

Related tasks

Using Event Search

Related reference

Configuring Event Search

You can customize Event Search to your specific needs by customizing the Netcool/OMNIbus Insight Pack . For example, you might want to send an extended set of Netcool/OMNIbus event fields to Operations Analytics - Log Analysis and chart results based on those fields.

Operations Management on premises data flow

Use this information to understand how event data is retrieved from a monitored application environment and transferred between the products and components of the base Netcool Operations Insight in order to provide Event Analytics and Event Search capabilities.

The following figure shows a simplified data flow between the products of the base Netcool Operations Insight solution.

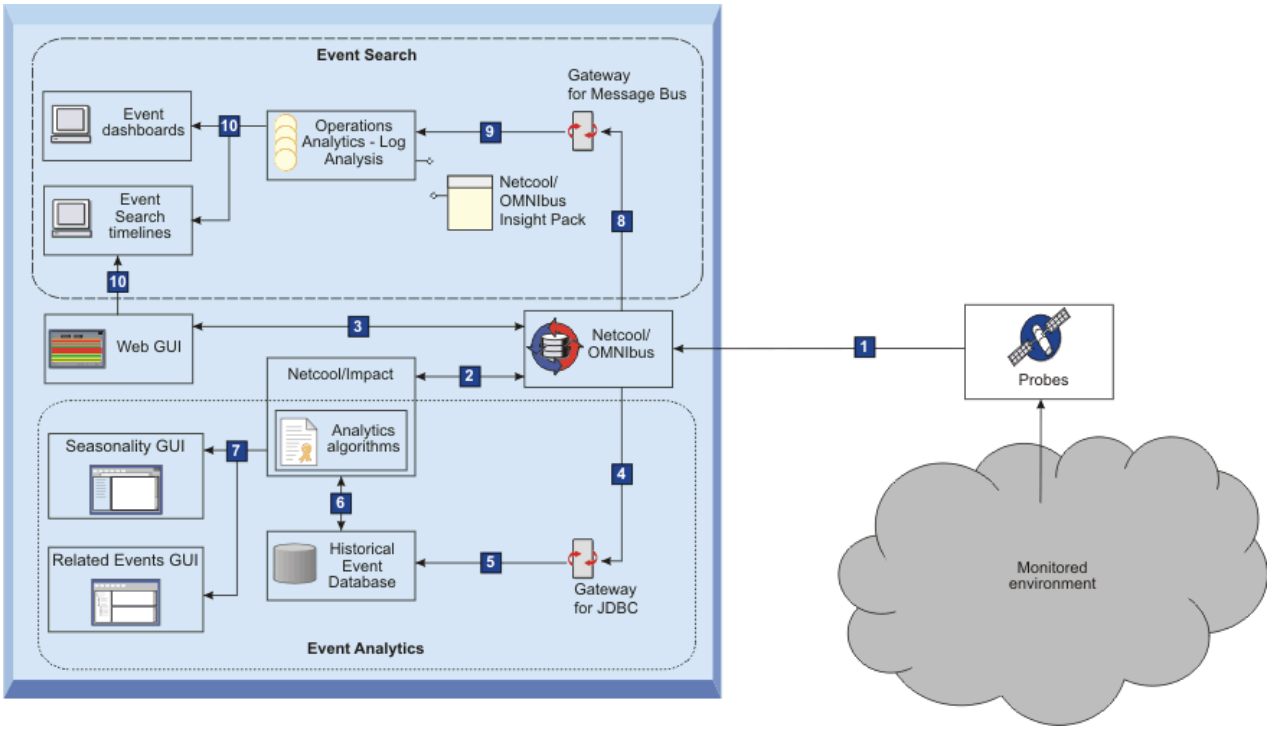

Figure 2. Data flow for the Netcool Operations Insight on premises base solution.

The stages of this data flow are as follows, indicated by the call-out graphics (for example, 1).

Capture of alert data

Probes monitor the devices and applications in the environment.

1: Alerts are received from applications and devices

Alert data is captured by the probes and forwarded to the Netcool/OMNIbus ObjectServer. Event data is then manipulated in various data flows.

Web GUI data flow

Event data is enriched and visualized in Web GUI.

2: Event data is read from the ObjectServer and enriched

Netcool/Impact reads the event data from the ObjectServer. In Netcool/Impact, the event data is enriched by information retrieved by Impact policies.

3: Event data is visualized and managed in the Web GUI

The Web GUI displays the application events that are in the ObjectServer. From the event lists, you can run tools that changes the event data; these changes are synchronized with the data in the ObjectServer.

Event Analytics data flow

Event data is archived and historical event data is used to generate analytics data.

4: Events are read from the ObjectServer by the Gateway for JDBC

The Gateway for JDBC reads events from the ObjectServer.

5: Event data is archived

The Gateway for JDBC sends the event data via an HTTP interface to the Historical Event Database. The figure shows an IBM Db2 database but any supported database can be used. The gateway must be configured in reporting mode. This data flow is a prerequisite for the event analytics capability.

6: Event analytics algorithms run on archived event data

After a set of historical alerts is archived, the seasonality algorithms of the Netcool/Impact policies can generate seasonal reports. The related events function analyzes Netcool/OMNIbus historical event data to determine which events have a statistical tendency to occur together and can therefore be grouped into related event groups. Pattern functions analyze the statistically

related event groups to determine if the groups have any generic patterns that can be applied to events on other network resources.

7: Analytics data is visualized and managed
The seasonality function helps you identify and examine seasonal trends while monitoring and managing events. This capability is delivered in a Seasonal Events Report portlet in Dashboard Application Services Hub. The portlet contains existing seasonal reports, which can be used to identify the seasonal pattern of the events in the Event Viewer. You can create new seasonal reports and edit existing ones. Statistically related groups can be analyzed in the Related Events GUI. Validated event groups can be deployed as Netcool/Impact correlation rules. Patterns in the statistically related event groups can also be analyzed in the Related Events GUI. These patterns can be extracted and deployed as Netcool/Impact generalized patterns.

Event Search data flow
Event data is indexed in Operations Analytics - Log Analysis and used to display event dashboard and timelines.

8: Events are read from the ObjectServer by Gateway for Message Bus
The Gateway for Message Bus reads events from the ObjectServer.

9: Event data is transferred for indexing to Operations Analytics - Log Analysis
The Gateway for Message Bus sends the event data via an HTTP interface to the Operations Analytics - Log Analysis product where the event data is indexed. The Tivoli Netcool/OMNIbus Insight Pack V1.3.0.0 parses the event data into a format suitable for use by Operations Analytics - Log Analysis. The diagram shows the default IDUC connection, which sends only event inserts. For event inserts and reinserts, the Accelerated Event Notification client can be deployed, which can handle greater event volumes. See "On-premises scenarios for Operations Management" on page 34.

10: Event search data is visualized
Event search results are visualized in Operations Analytics - Log Analysis event dashboards and timelines by performing right-click tools from event lists in Web GUI.

Related information
Tivoli Netcool/OMNIbus architecture
IBM Operations Analytics - Log Analysis architecture
Overview of Netcool/Impact deployments

About Network Management

Use this information to understand more about the Network Management for Operations Insight solution extension.

Network Management capabilities

Use this information to understand the capabilities of Network Management.

Network Management is made up of the following products and components:

- Network Manager
- Netcool Configuration Manager
- Topology Search

Networks for Operations Insight is an optional solution extension that can be added to a deployment of the base Netcool Operations Insight solution to provide service assurance in dynamic network infrastructures. The capabilities of Networks for Operations Insight include network discovery, visualization, event correlation and root-cause analysis, and configuration and compliance management. It contributes to overall operational insight into application and network performance management. The Networks for Operations Insight capability is provided through the Network Manager and Netcool Configuration Manager products.

The components and capabilities of Network Management are described below:

Network Health Dashboard
: This dashboard leverages the capabilities of Network Manager, Netcool/OMNIbus, and Netcool Configuration Manager products to display availability, performance, event, and configuration data for selected network views.

  The Network Health Dashboard displays device and interface availability within that network view. It also reports on performance by presenting graphs, tables, and traces of KPI data for monitored devices and interfaces. A dashboard timeline reports on device configuration changes and event counts, enabling correlation of events with configuration changes. The dashboard includes the **Event Viewer** for more detailed event information.

Topology Search
: This capability provides insight into network performance by determining lowest cost routes between two endpoints on the network over time. The topology search capability is an extension of the Networks for Operations Insight feature. It applies the search and analysis capabilities of Operations Analytics - Log Analysis to give insight into network performance. Events that have been enriched with network data are analyzed by the Network Manager Insight Pack and are used to calculate the lowest-cost routes between two endpoints on the network topology over time. The events that occurred along the routes over the specified time period are identified and shown by severity. The topology search requires the Networks for Operations Insight feature to be installed and configured.

Network Management tasks

Use this information to understand the tasks that users can perform using Network Management.

Network Management tasks fall into the following categories:

- Custom dashboard development tasks
- **Network Health Dashboard** tasks
- Topology Search tasks

Custom dashboard development tasks

Administrators can create pages that act as "dashboards" for displaying information on the status of parts of your network, and they can edit existing dashboards, such as the **Network Health Dashboard**. They can select from the widgets that are provided with Network Manager, Tivoli Netcool/OMNIbus Web GUI, and also from other products that are deployed in your Dashboard Application Services Hub environment.

Network Health Dashboard tasks

The Network Health Dashboard displays device and interface availability within that network view. It also reports on performance by presenting graphs, tables, and traces of KPI data for monitored devices and interfaces. A dashboard timeline reports on device configuration changes and event counts, enabling correlation of events with configuration changes. The dashboard includes the **Event Viewer** for more detailed event information.

Monitoring the network
: Operations staff can use **Network Health Dashboard** to monitor the network by selecting a network view within an area of responsibility, such as a geographical area, or a specific network service such as BGP or VPN, and reviewing the data that appears in the other widgets on the dashboard.

Administering the dashboard
: Administrators can configure how data is displayed, and which data is displayed in the **Network Health Dashboard**.

Topology Search tasks

This capability provides insight into network performance by determining lowest cost routes between two endpoints on the network over time.

Using Topology Search

Operations staff can use the analytics available in the Topology Search capability to obtain insight into network performance. For example, they can visualize the lowest-cost routes between two endpoints on the network topology over time.

Configuring Topology Search

Administrations staff can configure and customize these tools to match the network and alerting ecosystem.

Related concepts

About the Network Health Dashboard

Use the **Network Health Dashboard** to monitor a selected network view, and display availability, performance, and event data, as well as configuration and event history for all devices in that network view.

Related tasks

Developing custom dashboards

You can create pages that act as "dashboards" for displaying information on the status of parts of your network or edit existing dashboards, such as the **Network Health Dashboard**. You can select from the widgets that are provided with Network Manager, Tivoli Netcool/OMNIbus Web GUI, and also from other products that are deployed in your Dashboard Application Services Hub environment.

Using Topology Search

Configuring Topology Search

You can configure and customize the Network Manager Insight Pack to meet your requirements.

Network Management data flow

Use this information to understand how event data is retrieved from a monitored application environment and transferred between the products and components of Network Management in order to provide Topology Search, **Network Health Dashboard** and **Device Dashboard** capabilities.

The following figure shows a simplified data flow between the products of Network Management and, where appropriate, Operations Management.

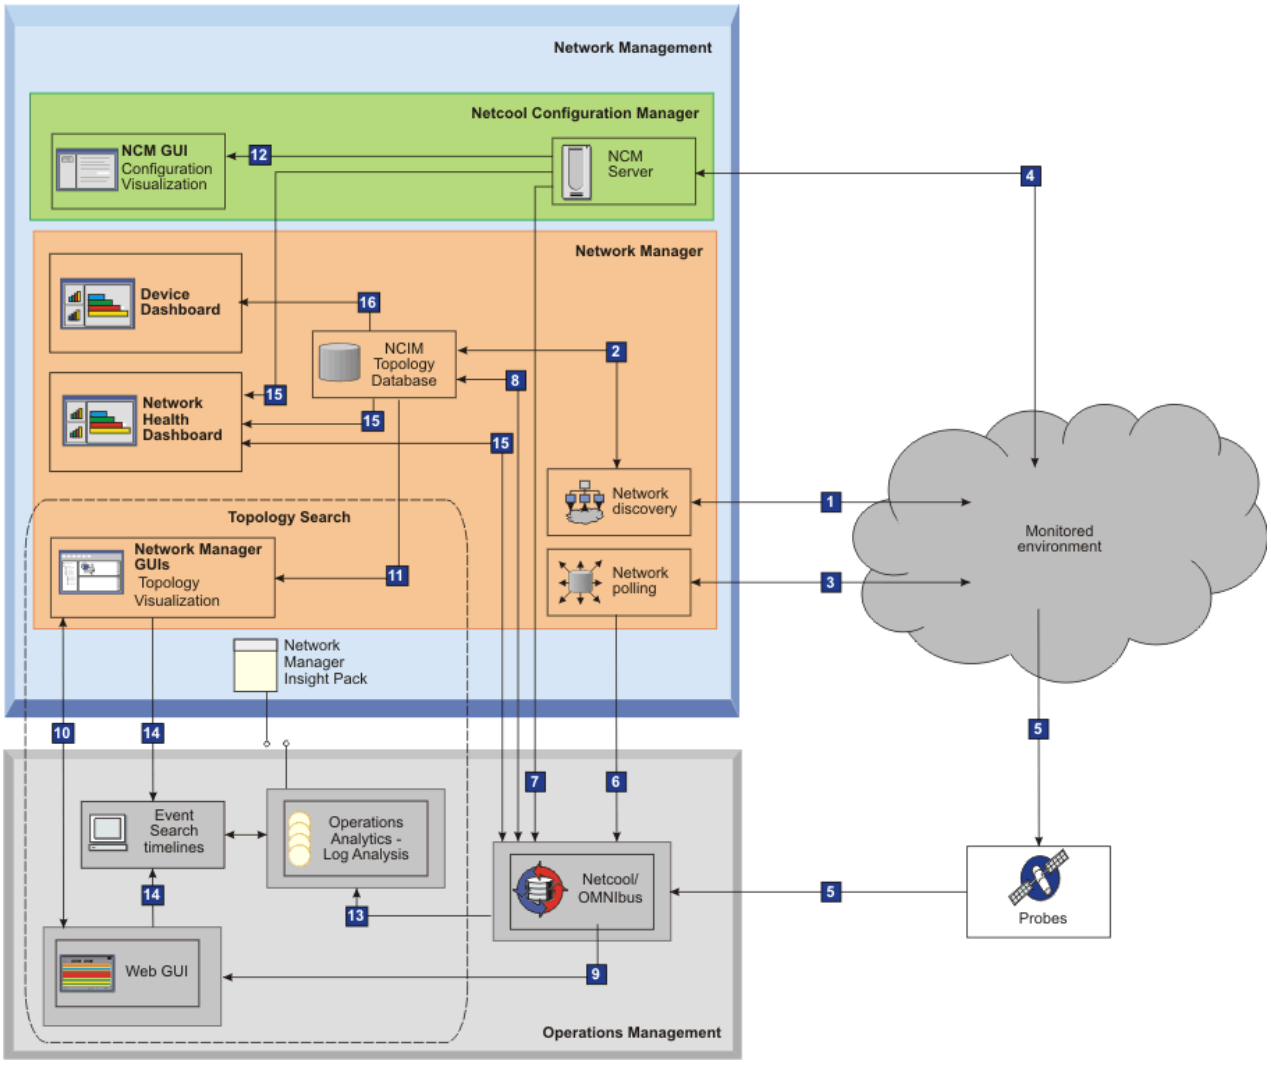

Figure 3. Simplified data flow

Collection of network topology, polling, and configuration data

1: Network discovery is run
Based on configurations set up by network administrators, Network Manager gathers data about the network. The discovery function identifies what entities, for example routers and switches, are on the network and interrogates them, for example, for connectivity information.

2 Network topology is stored
Network Manager classifies and stores the network topology that was discovered in step 1 in the NCIM topology database.

3: Network devices and interfaces are polled
Based on configurations set up by network administrators, Network Manager polling policies are run to determine whether a network device is up or down, whether it exceeds key performance parameters, and identifies inter-device link faults.

4: Changes to device configuration and policy changes are detected
Netcool Configuration Managerdiscovers whether there are any changes to device configuration or policy violations.

Collection and enrichment of alert data

5: Alerts are received from applications and devices
Alert data is captured by probes and forwarded to the ObjectServer.

6: Network events are generated if polls fail
Network Manager generates fault alerts if device and interface polls (step 2) fail. Network Manager converts the results of the relevant polls into events, and sends these network events to the ObjectServer.

7: Network configuration events are generated if device configurations change
Netcool Configuration Manager generates events for the configuration changes and policy violations (referred to hereafter as *network configuration events*) that were detected in step 3. Configuration change and policy violation events are sent via the Probe for SNMP to the ObjectServer.

8: Events are enriched using topology data
Network events (generated in step 6) and network configuration events (generated in step 7) are passed to the Event Gateway, where they are enriched with network topology data. For example, the system location, contact information, and product serial number can be added to the events. The events are returned to the ObjectServer.

Once steps 5 to 8 are complete the Netcool/OMNIbus ObjectServer contains the application events from the probes, network events from Network Manager, and the network configuration events from Netcool Configuration Manager.

Visualization of events and topology

9 Event are visualized and monitored
The Tivoli Netcool/OMNIbus Web GUI displays the application events, network events, and network configuration events that are in the ObjectServer.

10 Event information is shared
The event information is shared between the Web GUI and the Network Manager GUIs, for example, the Network Views and Hop View.

11 Network topology is visualized
The Network Manager GUIs display the network topology data that is in the NCIM database. This data is enriched by the configuration change and policy event information from the ObjectServer.

12 Network configuration events are analyzed
Configuration changes and policy violations are displayed for further analysis in the following GUIs:

- Network Manager GUIs
- Web GUI **Event Viewer**
- Netcool Configuration Manager Activity Viewer, wizards, and other Netcool Configuration Manager user interfaces

Using the right-click menus, operators can optionally launch-in-context across into Reporting Services, if it is installed. Reporting Services is not shown on this figure.

Topology search data flow

13 Event data is transferred for indexing to Operations Analytics - Log Analysis
The Gateway for Message Bus sends the event data via an HTTP interface to the Operations Analytics - Log Analysis product where the event data is indexed. The Network Manager Insight Pack parses the event data into a format suitable for use by Operations Analytics - Log Analysis.

14 Topology search data is visualized
Topology search results are visualized in Operations Analytics - Log Analysis event dashboards and timelines by performing right-click actions on two nodes in the network between which the analysis is required. This is done in one of the following ways: either select two network nodes in a network map within one of the Network Manager GUIs, or two events in the Web GUI Event Viewer.

Dashboard data flow

15 Network health information is visualized
In the **Network Health Dashboard**, selection of a network view enables you to visualize availability summary data, top 10 performance data, and configuration timeline data for the

devices in that network view. Data used to populate the **Network Health Dashboard** is retrieved from the ObjectServer, Network Manager polling databases, and Netcool Configuration Manager.

16 Device and interface health is visualized
You can right click to the **Device Dashboard** from any topology view or event list. In the **Device Dashboard**, selection of a device enables you to visualize top 10 performance data for the device or any of its interfaces. You can also visualize timeline data for any of the performance metrics associated with the device or any of its interfaces. Data used to populate the **Device Dashboard** is retrieved from the ObjectServer, Network Manager polling databases, and from Network Performance Insight.

Related concepts
Netcool Configuration Manager events

About Performance Management

Use this information to understand more about the Performance Management for Operations Insight solution.

Capabilities of Performance Management for Operations Insight

Use this information to understand the capabilities of Performance Management.

The extension is made up of the following product: Network Performance Insight

Network Performance Insight is a network traffic performance monitoring system. It provides comprehensive and scalable visibility on network traffic with visualization and reporting of network performance data for complex, multivendor, multi-technology networks. The end user is able to perform the following tasks: visualize flow across selected interfaces, display performance anomaly events in the Tivoli Netcool/OMNIbus **Event Viewer**, and view performance anomaly and performance timeline data in the **Device Dashboard**. For more information, see http://www-01.ibm.com/support/knowledgecenter/SSCVHB/welcome.

The components and capabilities of Performance Management are described below:

Device Dashboard (networks and performance)
This dashboard leverages the capabilities of Network Manager, and Netcool/OMNIbus to display event data, and top 10 performance metric data for a selected network device and its interfaces.

Note: This capability requires that both Network Management and Performance Management are integrated into your Netcool Operations Insight solution.

Traffic Details Dashboard
Use the Traffic Details dashboard to monitor network performance and flow details for a particular interface. Network Performance Insight provides built-in and interactive dashboards that cover the entire traffic data representation.

Network Performance Insight Dashboards
If you are a network planner or engineer, then use Network Performance Insight Dashboardsto view top 10 information on interfaces across your network, including the following:

- Congestion
- Traffic utilization
- Quality of service

About Service Management

Use this information to understand more about the Service Management for Operations Insight solution extension.

Capabilities of Service Management for Operations Insight

Use this information to understand the capabilities of Service Management.

The extension is made up of the following product: Agile Service Manager

Agile Service Manager provides operations teams with complete up-to-date visibility and control over dynamic infrastructure and services. Agile Service Manager lets you query a specific networked resource, and then presents a configurable topology view of it within its ecosystem of relationships and states, both in real time and within a definable time window. For more information, see https://www-01.ibm.com/support/knowledgecenter/SS9LQB.

The components and capabilities of Service Management are described below:

Topology Viewer

Using the Topology Viewer you can use elastic search features to easily locate and visualize near real-time and historical configurable views of multidomain topologies, including IT, network, storage, and application data. This enables you to reduce the complexity of managing modern and hybrid services, across vendors, data centers and traditional management silos.

Observer and API integration

Ease and rapidity of integration with any topology source is provided by means of observers and APIs. Observers are provided for a wide range of data, including event data, Network Manager topology data, TADDM topology data, OpenStack data, multiple file formats, REST, Docker, and VMware. This ensures rapid time-to-value, by providing up-to-date visibility and control over dynamic infrastructure and services.

Chapter 2. Deployment

Plan your deployment of Netcool Operations Insight

Deploying on premises

Use this information to understand the architecture of the Netcool Operations Insight deployment on premises. In this deployment mode, all of the Netcool Operations Insight products and components are installed onto servers and use the native computing resources of those servers.

Deployment scenarios for Operations Management

When you plan a deployment, it is important to consider the relationship between the event volumes that are supported by Netcool/OMNIbus and the capacity of Operations Analytics - Log Analysis to analyze events. The scenarios available depend on whether you are installing Operations Management on premises or in a private cloud using IBM Cloud Private.

Deployment considerations for on premises Operations Management
The desired volume of events determines whether a basic, failover, or desktop architecture or a multitier architecture is deployed. The Gateway for Message Bus can be configured to support event inserts only or both inserts and reinserts. Note that if you are deploying Operations Management on ICP then only a failover architecture is available and this means that the system will only support a low capacity event volume.

The following explains the architecture and event volume, and the event analysis capacity of Operations Analytics - Log Analysis in more detail.

Note: Operations Analytics - Log Analysis Standard Edition is included in Netcool Operations Insight. For more information about Operations Analytics - Log Analysis editions, search for "Editions" at the Operations Analytics - Log Analysis Knowledge Center, at https://www.ibm.com/support/knowledgecenter/SSPFMY.

Event volume

Event inserts are the first occurrence of each event and reinserts are every occurrence of each event. By default, the Gateway for Message Bus is configured to accept only event inserts from ObjectServers through an IDUC channel. To support event inserts and reinserts, you can configure event forwarding through the Accelerated Event Notification (AEN) client.

Note: Event Search functionality varies as follows depending on the choice of channel:

- IDUC channel: Event Search functionality is limited. Chart display functionality is fully available, but you will not be able to perform a deep dive into events or search for event modifications.
- AEN channel: All Event Search functionality is available. However, as part of your Netcool/OMNIbus configuration you will also have to install triggers in the ObjectServer.

For more information, search for *Integrating with Operations Analytics - Log Analysis* in the Gateway for Message Bus documentation.

Architecture of Netcool/OMNIbus

Basic, failover, and desktop architectures support low and medium capacity for analyzing events. Multitiered architectures support higher Operations Analytics - Log Analysis capacities. In a multitier architecture, the connection to the Gateway for Message Bus supports higher capacity at the collection layer than at the aggregation layer.

For more information about these architectures, see the Netcool/OMNIbus documentation and also the *Netcool/OMNIbus Best Practices Guide*.

Capacity of Operations Analytics - Log Analysis
The volume of events that Operations Analytics - Log Analysis is able to handle. For the hardware levels that are required for expected event volumes, see the Operations Analytics - Log Analysis documentation at http://www-01.ibm.com/support/knowledgecenter/SSPFMY/welcome. If capacity is limited, you can use the deletion tool to remove old data.

Connection layer
The connection layer is the layer of the multitier architecture to which the Gateway for Message Bus is connected. This consideration applies only when the Netcool/OMNIbus product is deployed in a multitier architecture. The connection layer depends on the capacity of Operations Analytics - Log Analysis. For more information about multitier architectures, see the Netcool/OMNIbus documentation and also the *Netcool/OMNIbus Best Practices Guide*.

Related information
Netcool/OMNIbus Best Practices GuideFor provisioning and sizing advice, refer to the planning chapter in the Netcool/OMNIbus Best Practices Guide.

On-premises scenarios for Operations Management

This topic presents the scenarios available in a deployment of Operations Management on premises together with the associated architectures.

The deployment scenarios and associated architecture diagrams are shown below.

- "Deployment scenarios" on page 34
- "Illustrations of architectures" on page 36

Deployment scenarios

This section describes possible deployment scenarios.

- "Deployment scenario 1: low capacity with IDUC channel" on page 34
- "Deployment scenario 2: medium capacity with AEN channel" on page 35
- "Deployment scenario 3: medium capacity with IDUC channel" on page 35
- "Deployment scenario 4: high capacity with IDUC channel" on page 35
- "Deployment scenario 5: high capacity with AEN channel" on page 35
- "Deployment scenario 6: very high capacity with AEN channel" on page 36

Deployment scenario 1: low capacity with IDUC channel

Table 9. Inserts only, standard architecture, low capacity

Event volume	Architecture of Netcool/ OMNIbus	Capacity of Operations Analytics - Log Analysis	Connection layer	IDUC or AEN	Illustration of this architecture
Inserts only	Basic, failover, and desktop architecture	Low	Not applicable	IDUC	See Figure 4 on page 37. Disregard the reference to reinserts in item 1.

Deployment scenario 2: medium capacity with AEN channel

Table 10. Inserts and reinserts, standard architecture, medium capacity

Event volume	Architecture of Netcool/ OMNIbus	Capacity of Operations Analytics - Log Analysis	Connection layer	IDUC or AEN	Illustration of this architecture
Inserts and reinserts	Basic, failover, and desktop architecture	Medium	Not applicable	AEN	See Figure 4 on page 37.

Deployment scenario 3: medium capacity with IDUC channel

Table 11. Inserts only, multitier architecture, medium capacity

Event volume	Architecture of Netcool/ OMNIbus	Capacity of Operations Analytics - Log Analysis	Connection layer	IDUC or AEN	Illustration of this architecture
Inserts only	Multitier	Medium	Aggregation layer	IDUC	See Figure 5 on page 38. Disregard the reference to reinserts in item 1.

Deployment scenario 4: high capacity with IDUC channel

Table 12. Inserts only, multitier architecture, high capacity

Event volume	Architecture of Netcool/ OMNIbus	Capacity of Operations Analytics - Log Analysis	Connection layer	IDUC or AEN	
Inserts only	Multitier	High	Collection layer	IDUC	See Figure 6 on page 39. Disregard the reference to reinserts in item 1.

Deployment scenario 5: high capacity with AEN channel

Table 13. Inserts and reinserts, multitier architecture, high capacity

Event volume	Architecture of Netcool/ OMNIbus	Capacity of Operations Analytics - Log Analysis	Connection layer	IDUC or AEN	
Inserts and reinserts	Multitier	High	Aggregation layer	AEN	See Figure 5 on page 38.

Deployment scenario 6: very high capacity with AEN channel

Table 14. Inserts and reinserts, multitier architecture, very high capacity

Event volume	Architecture of Netcool/ OMNIbus	Capacity of Operations Analytics - Log Analysis	Connection layer	IDUC or AEN	
Inserts and reinserts	Multitier	Very high	Collection layer	AEN	See Figure 6 on page 39.

Illustrations of architectures

The following sections show the architecture of Operations Analytics - Log Analysis deployments and how they fit into the various architectures of Netcool/OMNIbus deployments with the Gateway for Message Bus.

The data source that is described in the figures is the raw data that is ingested by the Operations Analytics - Log Analysis product. You define it when you configure the integration between the Operations Analytics - Log Analysis and Netcool/OMNIbus products.

- "Basic, failover, and desktop architectures" on page 36
- "Multitier architecture, events are sent from the Aggregation layer" on page 37
- "Multitier architecture, events are sent from the Collection layer" on page 38

Basic, failover, and desktop architectures

The following figure shows how the integration works in a basic, failover, or desktop Netcool/OMNIbus architecture. This figure is an illustration of the architectures that are described in Table 9 on page 34 and Table 10 on page 35. In the case of the architecture in Table 9 on page 34, disregard item 1 in this figure.

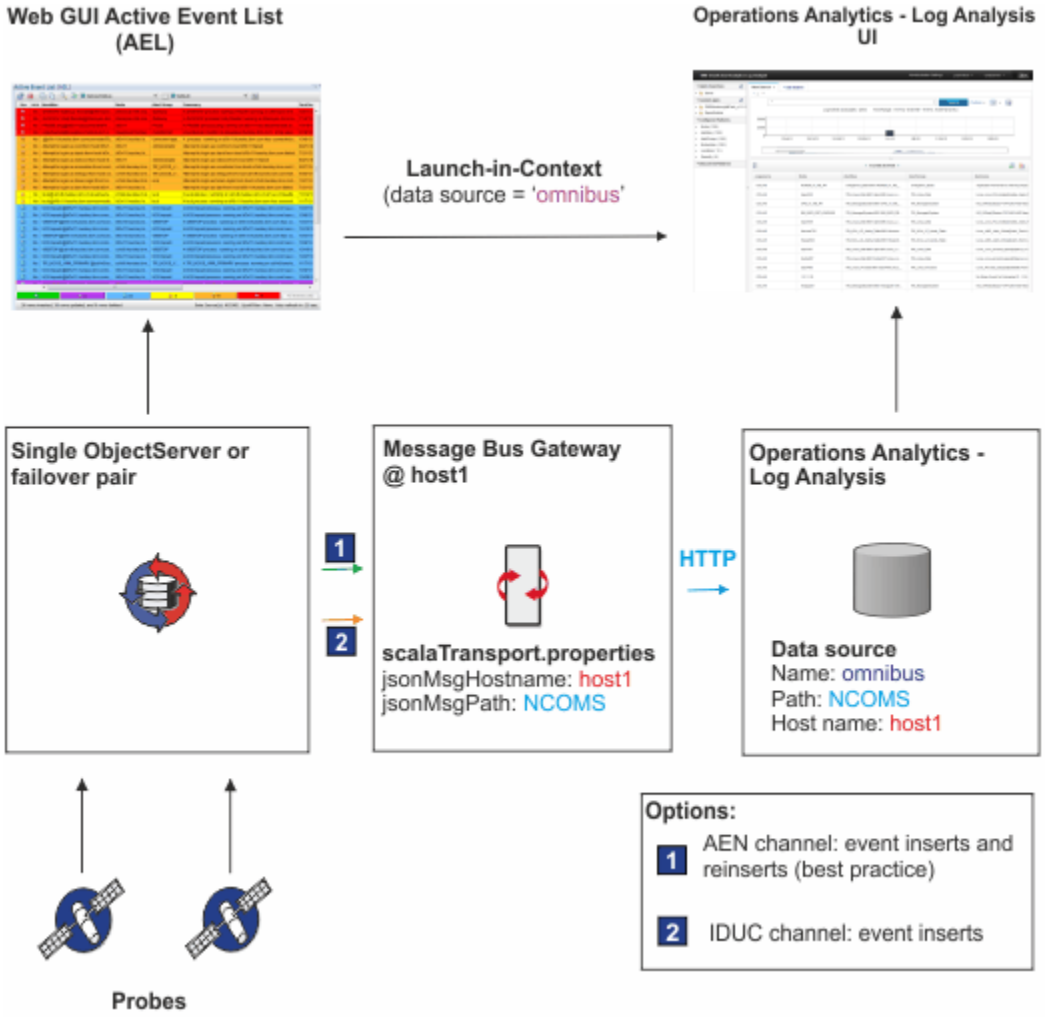

Figure 4. Basic, failover, and desktop deployment architecture

Multitier architecture, events are sent from the Aggregation layer

The following figure shows how the integration works in a multitier Netcool/OMNIbus architecture, with events sent from the Aggregation layer. This figure is an illustration of the architectures that are described in Table 11 on page 35 and Table 13 on page 35. In the case of the architecture in Table 11 on page 35, disregard item 1 in this figure.

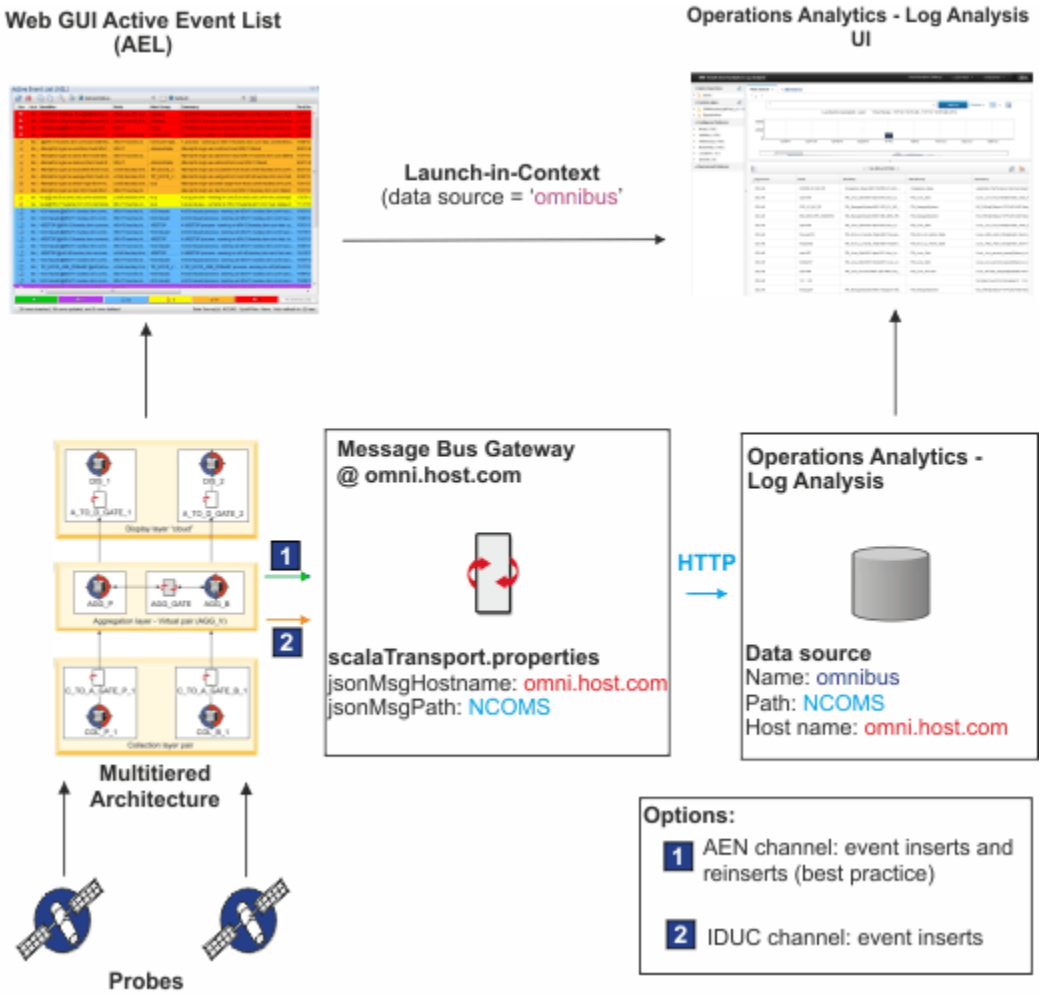

Figure 5. Multitier architecture deployment - Aggregation layer

Multitier architecture, events are sent from the Collection layer

The following figure shows how the integration works in a multitier Netcool/OMNIbus architecture, with events sent from the Collection layer. This is a best practice for integrating the components. This figure is an illustration of the architectures that are described in Table 12 on page 35 and Table 14 on page 36. In the case of the architecture in Table 12 on page 35, disregard item 1 in this figure.

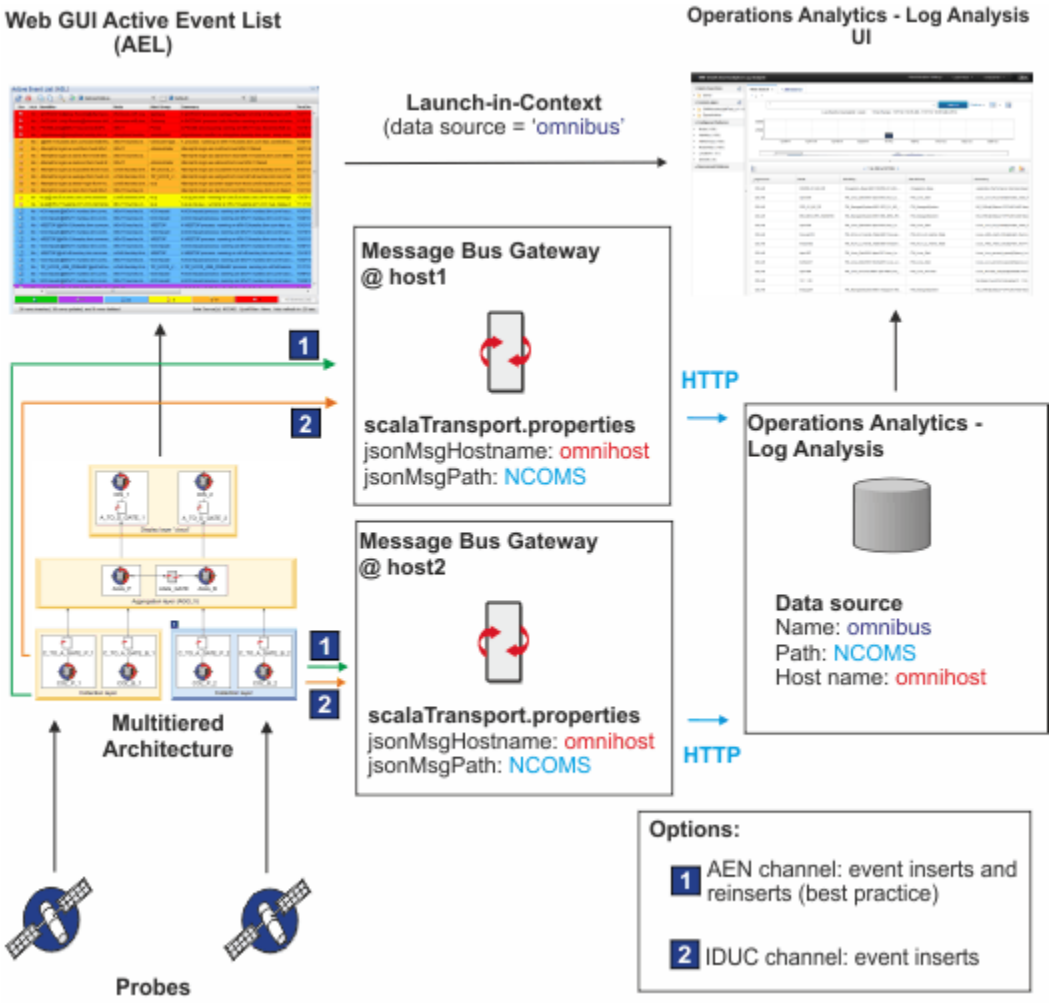

Figure 6. Multitier architecture deployment - Collection layer (best practice)

Related concepts

Overview of the standard multitiered architecture

Overview of the AEN client

Related tasks

Sizing your Tivoli Netcool/OMNIbus deployment

Configuring and deploying a multitiered architecture

Installing Netcool/OMNIbus and Netcool/Impact

Related reference

Failover configuration

Example Tivoli Netcool/OMNIbus installation scenarios (basic, failover, and desktop architectures)

Related information

Message Bus Gateway documentation

IBM Operations Analytics - Log Analysis documentation

IBM developerWorks: Tivoli Netcool OMNIbus Best PracticesClick here to access best practice documentation for Netcool/OMNIbus.

Deployment considerations for Network Management

Networks for Operations Insight is an optional feature that integrates network management products with the products of the base Netcool Operations Insight solution.

The Networks for Operations Insight feature includes Network Manager IP Edition and Netcool Configuration Manager. Deploying these products depends on your environment and the size and

complexity of your network. For guidance on deployment options, see guidance provided in the respective product documentation:

- Network Manager IP Edition: https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/install/concept/ovr_deploymentofitnm.html
- Netcool Configuration Manager: https://www.ibm.com/support/knowledgecenter/SS7UH9_6.4.2/ncm/wip/planning/concept/ncm_plan_planninginstallation.html

Example of an on-premises physical deployment

Use this example to familiarize yourself with the architecture of an on-premises physical deployment of Netcool Operations Insight. The architecture described in this example can be scaled up and extended for failover, a multitiered architecture, load balancing, and clustering.

This scenario assumes that there are no existing Netcool Operations Insight products in your environment, so no backup, restore, or upgrade information is given. The information supplied in this scenario is high-level and covers the most salient points and possible issues you might encounter that are specific to Netcool Operations Insight. The steps to install the Networks for Operations Insight feature are included, but skip these steps if you want to install only the base solution. This scenario is end-to-end and you must perform the tasks in the specified order.

For more information about each task in this scenario, see the Related concept, task, and information links at the bottom of each page.

The following figure shows the simplified installation architecture that this scenario adheres to.

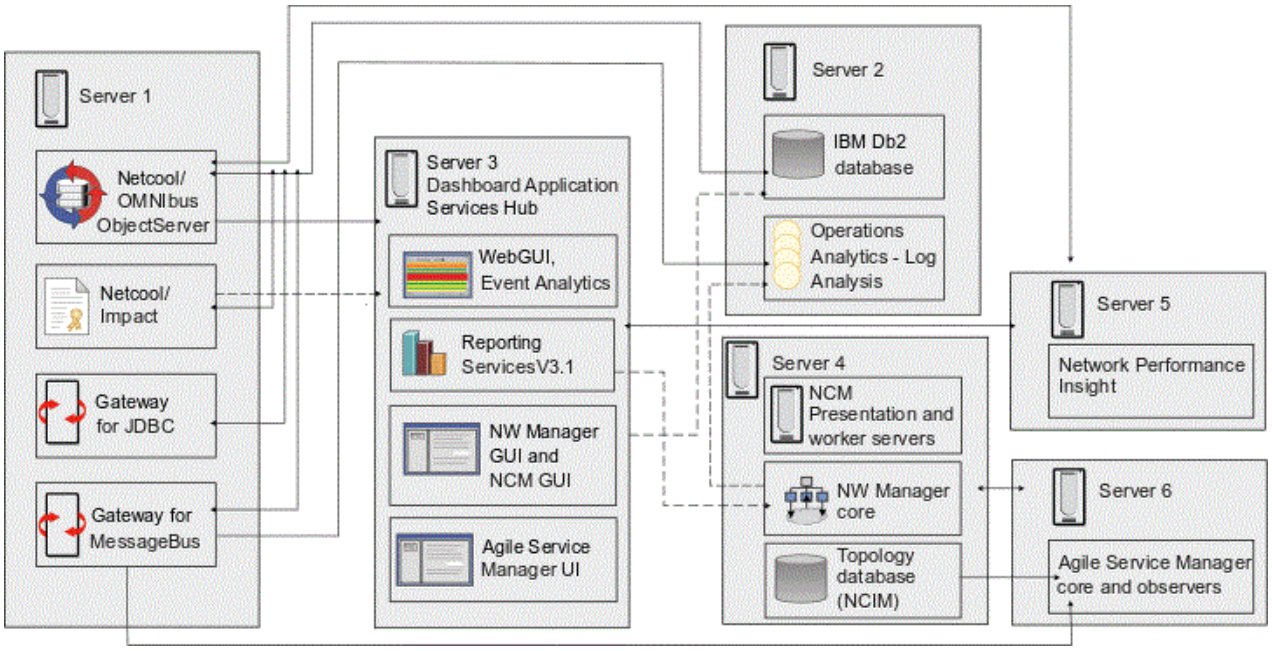

Figure 7. Simplified installation architecture for the installation scenario

For information on the product and component versions supported in the current version of Netcool Operations Insight including supported fix packs, see “Products and components on premises” on page 9.

Server 1

Hosts the Netcool/OMNIbus core components, the Gateway for JDBC, Gateway for Message Bus, and Netcool/Impact. Configurations are applied to the ObjectServer to support the event analytics and topology search capabilities. Event analytics is part of the base Netcool Operations Insight solution. Topology search is part of the Networks for Operations Insight feature. The default configuration of the Gateway for Message Bus is to transfer event inserts to Operations Analytics - Log Analysis through an IDUC channel. This connection can be changed to forward events reinserts and inserts through the Accelerated Event Notification client.

Server 2

Hosts an IBM Db2 database and Operations Analytics - Log Analysis. The Tivoli Netcool/OMNIbus Insight Pack and the Network Manager Insight Pack are installed into Operations Analytics - Log Analysis. The Tivoli Netcool/OMNIbus Insight Pack is part of the base Netcool Operations Insight solution. The Network Manager Insight Pack is part of the Networks for Operations Insight feature. The REPORTER schema is applied to the Db2 database so that events can be transferred from the Gateway for JDBC. Various installation methods are possible for Db2. For more information, see https://ibm.biz/BdEWtm.

Server 3

Hosts Dashboard Application Services Hub, which is a component of Jazz for Service Management. Jazz for Service Management provides the GUI framework and the Reporting Services component. The Netcool/OMNIbus Web GUI and the Event Analytics component are installed into Dashboard Application Services Hub. In this setup Reporting Services is also installed on this server, together with parts of the Networks for Operations Insight feature: the Network Manager IP Edition GUI components, Netcool Configuration Manager, and the Agile Service Manager UI. This simplifies the configuration of the GUI server, and provides the reporting engine and the report templates provided by the products on one host.

Note: You can set up Network Manager and Netcool Configuration Manager to work with Reporting Services by installing their respective reports when installing the products. Netcool/OMNIbus V8.1.0 and later can be integrated with Reporting Services V3.1 to support reporting on events. To configure this integration, connect Reporting Services to a relational database through a gateway. Then, import the report package that is supplied with Netcool/OMNIbus into Reporting Services. For more information about event reporting, see http://www-01.ibm.com/support/knowledgecenter/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/omnibus/wip/install/task/omn_con_ext_deploytcrreports.html.

Server 4

Hosts the Netcool Configuration Manager presentation and worker server, the Network Manager IP Edition core components, and the NCIM topology database, which are all components of the Networks for Operations Insight feature. This setup assumes large networks where discovering the network and creating and maintaining the network topology can require significant system resources.

Server 5

Hosts the Network Performance Insight components that support the performance management feature. For information on installation and configuration of Network Performance Insight, see “Installing Performance Management” on page 92.

Server 6

Hosts the Agile Service Manager components that support the service management feature, including the Agile Service Manager core and the Agile Service Manager observers. For information on installation and configuration of Agile Service Manager, see the Agile Service Manager documentation at https://www.ibm.com/support/knowledgecenter/SS9LQB_1.1.0/welcome_page/kc_welcome-444.html.

Deploying on IBM Cloud Private

Use this information to understand the architecture of the Netcool Operations Insight deployment on IBM Cloud Private.

Note: The current version of Netcool Operations Insight is compatible with IBM Cloud Private version 3.2 and all of the links in this documentation point to that version of the IBM Cloud Private documentation. If you want to view a different version of the IBM Cloud Private documentation, then navigate to IBM Cloud Private documentation: Welcome page and select the desired version.

Deployment considerations for Operations Management in a private cloud

The architecture for a deployment of Operations Management on IBM Cloud Private is preconfigured. For the Event Search feature, only a failover ObjectServer architecture is available and this means that the system will currently only support a medium capacity event volume.

For an overview of the architecture and event volume, and the event analysis capacity of Operations Analytics - Log Analysis see "Deployment considerations for on premises Operations Management" on page 33. The information below provides information specific to a deployment of Operations Management on IBM Cloud Private.

Event volume

By default, the Gateway for Message Bus that is implemented in Operations Management on IBM Cloud Private is configured to accept both event inserts and reinserts from ObjectServers through the Accelerated Event Notification (AEN) client.

Architecture of Netcool/OMNIbus

An ObjectServer failover architecture is implemented in Operations Management on IBM Cloud Private. Basic, failover, and desktop architectures support low and medium capacity for analyzing events. Multitiered ObjectServer architectures are currently not available.

For more information about these architectures, see the Netcool/OMNIbus documentation and also the *Netcool/OMNIbus Best Practices Guide*.

Hardware levels for event volumes

In a deployment of Operations Management on IBM Cloud Private, all of the components of the Netcool Operations Insight solution are deployed in containers within a single IBM Cloud Private cluster. Resources, including hardware, are allocated to the entire cluster. This is unlike an on-premises deployment, where hardware is assigned to individual components. For the Event Search feature, only a failover ObjectServer architecture is available. The resources required to support this architecture are detailed in "Requirements for an installation on IBM Cloud Private" on page 98.

Deployment scenarios for Operations Management on IBM Cloud Private

Learn about the deployment scenarios for Operations Management on IBM Cloud Private.

What to do next

- Read about how to prepare for the installation of Operations Management on IBM Cloud Private at "Preparing for installation on IBM Cloud Private" on page 98.

Medium capacity event volume scenario for Operations Management on IBM Cloud Private

In a deployment of Operations Management on IBM Cloud Private there is only one medium capacity event volume scenario available; this means that the transfer of events from Netcool/OMNIbus to Operations Analytics - Log Analysis is available at medium capacity only.

In a deployment of Operations Management on IBM Cloud Private, the deployment scenario available is a single ObjectServer pair with an Accelerated Event Notification (AEN) channel for communication with Operations Analytics - Log Analysis. Operations Management on IBM Cloud Private does not currently support multi-tier architectures, and the Accelerated Event Notification channel is the preconfigured channel of choice in this architecture.

The deployment scenario is described in the following table:

Table 15. Inserts and reinserts, standard architecture, medium capacity

Event volume	Architecture of Netcool/ OMNIbus	Capacity of Operations Analytics - Log Analysis	Connection layer	IDUC or AEN	Illustration of this architecture
Inserts and reinserts	Basic, failover, and desktop architecture	Medium	Not applicable	AEN	See Figure 4 on page 37.

Chapter 3. Installing Netcool Operations Insight

Plan the installation and complete any pre-installation tasks before installing Netcool Operations Insight.

Installing on premises

Follow these instructions to prepare for and install Netcool Operations Insight on premises.

For a step-by-step guide to installing your on-premises deployment, see the *Netcool Operations Insight Installation and Set-up Guide* on the IT Operations Management Developer Center: http://developer.ibm.com/itom/wp-content/uploads/sites/39/2018/05/NOI1.5.0.1-Installation-and-Set-up.pdf

Planning for an on-premises installation

Prepare of an on-premises installation of base Netcool Operations Insight and of Netcool Operations Insight solution extensions.

About this task

Procedure

Ports used by products and components

Use this information to understand which ports are used by the different products and components that make up the Netcool Operations Insight solution.

The following table lists sample ports that you might need to configure, and provides links to Netcool Operations Insight product and component documentation where you can access detailed information.

Table 16. Default port information

Product	Example default ports	Links
Netcool/OMNIbus	Aggregation ObjectServer primary port Process agent ports Gateway server port IBM Eclipse Help System server port Port numbers for individual Netcool/OMNIbus probes	ObjectServer ports can be configured using the Netcool Configuration wizard. See http://ibm.biz/BdskVc. Default ports used by Netcool/OMNIbus. See http://ibm.biz/BdskVr. Ports for a Netcool/OMNIbus basic architecture. See http://ibm.biz/BdsWi9. Ports for a Netcool/OMNIbus basic failover architecture. See http://ibm.biz/BdsWqw. Ports for a Netcool/OMNIbus desktop server architecture. See http://ibm.biz/BdsWqt.

Table 16. Default port information (continued)

Product	Example default ports	Links
Netcool/OMNIbus Web GUI	Jazz for Service Management WAS profile • HTTP port • HTTPS port	Jazz for Service Management port availability requirements. See http://ibm.biz/BdsWzf. Firewall Ports to open for DASH Services. See http://www-01.ibm.com/support/docview.wss?uid=swg21687730 for a technote.
Netcool/Impact	Netcool/Impact server • HTTP port • HTTPS port Netcool/Impact GUI • HTTP port • HTTPS port	Assigning Netcool/Impact ports. See http://ibm.biz/BdsWyK. Assigning Netcool/Impact data source and service ports. See http://ibm.biz/BdsWyG. **Note:** It is not possible to install Netcool/Impact GUI and Jazz for Service Management using the same default port numbers (16310/16311) on the same server. In this case you must modify the port numbers during installation.
Operations Analytics - Log Analysis	Application WebConsole Port Application WebConsole Secure Port Database Server Port Data Collection Server Port	Default ports used by Operations Analytics - Log Analysis: • V1.3.5: see http://ibm.biz/BdsWyn. • V1.3.3: see http://ibm.biz/BdiyPy
Db2 Enterprise Server Edition database	Port 50000. **Note:** This port is also configurable following installation.	
Network Performance Insight	Ports must be assigned for the following Network Performance Insight components: • Ambari Metrics • Hadoop Distributed File System (HDFS) • Apache Kafka	Default ports used by Network Performance Insight: • V1.3.1: see http://ibm.biz/npi_131_ports • V1.2.3: see http://ibm.biz/npi_123_ports

Related concepts

Installing Db2 and configuring the REPORTER schema

Netcool Operations Insight requires a Db2 database with the REPORTER schema for historical event archiving.

Related tasks

Installing Netcool/OMNIbus and Netcool/Impact

Installing IBM Operations Analytics - Log Analysis

Operations Analytics - Log Analysis supports GUI, console, and silent installations. The installation process differs for 64-bit and z/OS operating systems.

Installing Network Performance Insight
Install Network Performance Insight by performing the steps in the Installation section of the Network Performance Insight documentation.

Checking prerequisites
Before you install each product, run the IBM Prerequisite Scanner (PRS) to ensure that the target host is suitable, and no installation problems are foreseeable. Also check the maxproc and ulimit settings on the servers you are configuring to ensure they are set to the appropriate minimum values.

Before you begin

- For information about hardware and software compatibility of each component, and detailed system requirements, see the IBM Software Product Compatibility Reports website: http://www-969.ibm.com/software/reports/compatibility/clarity/index.html

  Tip: When you create a report, search for Netcool Operations Insight and select your version (for example, V1.4). In the report, additional useful information is available through hover help and additional links.
  For example, to check the compatibility with an operating system for each component, go to the **Operating Systems** tab, find the row for your operating system, and hover over the icon in the **Components** column. For more detailed information about restrictions, click the **View** link in the **Details** column.
- Download IBM Prerequisite Scanner from IBM Fix Central at http://www.ibm.com/support/fixcentral/. Search for "IBM Prerequisite Scanner".
- After you download the latest available version, decompress the `.tar` archive into the target directory on all hosts.
- On the IBM Tivoli Netcool/Impact host, set the environment variable `IMPACT_PREREQ_BOTH=True` so that the host is scanned for both the Impact Server and the GUI Server.

For a list of all product codes, see http://www.ibm.com/support/docview.wss?uid=swg27041454

About this task
Operations Analytics - Log Analysis and IBM Db2 are not supported by IBM Prerequisite Scanner. For the installation and system requirements for these products, refer to the documentation.

Procedure

Using the IBM Prerequisite Scanner

- On the IBM Tivoli Netcool/OMNIbus and IBM Tivoli Netcool/Impact host, run IBM Prerequisite Scanner as follows:

Product	Command
IBM Tivoli Netcool/OMNIbus	`prereq_checker.sh NOC detail`
IBM Tivoli Netcool/Impact	`prereq_checker.sh NCI detail`

- On the host for the GUI components:

Product	Command
Jazz for Service Management	`prereq_checker.sh ODP detail`
Dashboard Application Services Hub	`prereq_checker.sh DSH detail`
IBM Tivoli Netcool/OMNIbus Web GUI	`prereq_checker.sh NOW detail`

- On the Networks for Operations Insight host

Product	Command
IBM Tivoli Network Manager	`prereq_checker.sh TNM detail`
IBM Tivoli Netcool Configuration Manager	`prereq_checker.sh NCM detail`
Tivoli Common Reporting	`prereq_checker.sh TCR detail`

Check the maxproc settings.

- Open the following file: `/etc/security/limits.d/90-nproc.conf`
- Set nproc to a value of `131073`

Check the ulimit settings.

- Open the following file: `/etc/security/limits.conf`
- Set nofile to a value of `131073`

Related tasks

Installation prerequisites for Operations Analytics - Log Analysis V1.3.5
Installation prerequisites for Operations Analytics - Log Analysis V1.3.3
System requirements for Db2 products
Installation requirements for Db2 products

Obtaining IBM Installation Manager

Perform this task only if you are installing directly from an IBM repository or a local repository. IBM Installation Manager is required on the computers that host Netcool/OMNIbus, Netcool/Impact, Operations Analytics - Log Analysis, and the products and components that are based on Dashboard Application Services Hub. In this scenario, that is servers 1, 2, and 3. The installation packages of the products include Installation Manager.

Before you begin

Create an IBM ID at http://www.ibm.com. You need an IBM ID to download software from IBM Fix Central.

Note:

On Red Hat Enterprise Linux, the GUI mode of the Installation Manager uses the libcairo UI libraries. The latest updates for RHEL 6 contain a known issue that causes the Installation Manager to crash. Before installingInstallation Manager on Red Hat Enterprise Linux 6, follow the instructions in the following technote to configure libcairo UI libraries to a supported version: http://www.ibm.com/support/docview.wss?uid=swg21690056

Remember: The installation image of Netcool/OMNIbus V8.1.0 available from IBM Passport Advantage and on DVD includes Installation Manager. You only need to download Installation Manager separately if you are installing Netcool/OMNIbus directly from an IBM repository or from a local repository.

You can install Installation Manager in one of three user modes: Administrator mode, Nonadministrator mode, or Group mode. The user modes determine who can run Installation Manager and where product data is stored. The following table shows the supported Installation Manager user modes for products in IBM Netcool Operations Insight.

Table 17. Supported Installation Manager user modes

Product	Administrator mode	Nonadministrator mode	Group mode
IBM Tivoli Netcool/OMNIbus [1]	X	X	X
IBM Tivoli Netcool/Impact	X	X	X

[1] Includes OMNIbus Core, Web GUI, and the Gateways.

Table 17. Supported Installation Manager user modes (continued)			
Product	**Administrator mode**	**Nonadministrator mode**	**Group mode**
IBM Operations Analytics - Log Analysis [2]		X	

Procedure

The IBM Fix Central website offers two approaches to finding product files: **Select product** and **Find product**. The following instructions apply to the **Find product** option.

1. Go to IBM Fix Central at http://www.ibm.com/support/fixcentral/ and search for **IBM Installation Manager**.
   a) On the **Find product** tab, enter IBM Installation Manager in the **Product selector** field.
   b) Select V1.8.x from the **Installed Version** list.
   c) Select your intended host operating system from the **Platform** list and click **Continue**.
2. On the **Identity Fixes** page, choose **Browse for fixes** and **Show fixes that apply to this version (1.X.X.X)**. Click **Continue**.
3. On the **Select Fixes** page, select the installation file appropriate to your intended host operating system and click **Continue**.
4. When prompted, enter your IBM ID and password.
5. If your browser has Java enabled, choose the Download Director option. Otherwise, select the HTTP download option.
6. Start the installation file download. Make a note of the download location.

What to do next
Install Installation Manager. See http://www.ibm.com/support/docview.wss?uid=swg24034941.
Related information
IBM Installation Manager overview

Installing Installation Manager (GUI or console example)
You can install Installation Manager with a wizard-style GUI or an interactive console, as depicted in this example.

Before you begin

Take the following actions:

- Extract the contents of the Installation Manager installation file to a suitable temporary directory.
- Ensure that the necessary user permissions are in place for your intended installation, data, and shared directories.
- The console installer does not report required disk space. Ensure that you have enough free space before you start a console installation.

Before you run the Installation Manager installer, create the following target directories and set the file permissions for the designated user and group that Installation Manager to run as, and any subsequent product installations:

Main installation directory
Location to install the product binary files.

Data directory
Location where Installation Manager stores information about installed products.

[2] Insight Packs are installed by the Operations Analytics - Log Analysis pkg_mgmt command.

Shared directory
Location where Installation Manager stores downloaded packages that are used for rollback.

Ensure that these directories are separate. For example, run the following commands:

```
mkdir /opt/IBM/NetcoolIM
mkdir /opt/IBM/NetcoolIM/IBMIM
mkdir /opt/IBM/NetcoolIM/IBMIMData
mkdir /opt/IBM/NetcoolIM/IBMIMShared
chown -R netcool:ncoadmin /opt/IBM/NetcoolIM
```

About this task

The initial installation steps are different depending on which user mode you use. The steps for completing the installation are common to all user modes and operating systems.

Installation Manager takes account of your current umask settings when it sets the permissions mode of the files and directories that it installs. Using Group mode, Installation Manager ignores any group bits that are set and uses a umask of 2 if the resulting value is 0.

Procedure

1. Install in Group mode:
   a) Use the id utility to verify that your current effective user group is suitable for the installation. If necessary, use the following command to start a new shell with the correct effective group: `newgrp group_name`
   b) Use the umask utility to check your umask value. If necessary, change the umask value.
   c) Change to the temporary directory that contains the Installation Manager installation files.
   d) Use the following command to start the installation:

      GUI installation
      `./groupinst -dL data_location`

      Console installation
      `./groupinstc -c -dL data_location`

      In this command, *data_location* specifies the data directory. You must specify a data directory that all members of the group can access.

      Remember: Each instance of Installation Manager requires a different data directory.
2. Follow the installer instructions to complete the installation.
   The installer requires the following input at different stages of the installation:

   GUI installation

   - In the first page, select the `Installation Manager` package.
   - Read and accept the license agreement.
   - When prompted, enter an installation directory or accept the default directory.
   - Verify that the total installation size does not exceed the available disk space.
   - When prompted, restart Installation Manager.

   Console installation

   - Read and accept the license agreement.
   - When prompted, enter an installation directory or accept the default directory.
   - If required, generate a response file. Enter the directory path and a file name with a `.xml` extension. The response file is generated before installation completes.
   - When prompted, restart Installation Manager.

Results
Installation Manager is installed and can now be used to install IBM Netcool Operations Insight.

Note: If it is not possible for you to install Netcool Operations Insight components in GUI mode (for example, security policies at your site might limit the display of GUI pages) then you can use the Installation Manager web application to install the Netcool Operations Insight base solution components, which are as follows:

- IBM Tivoli Netcool/OMNIbus core components
- IBM Tivoli Netcool/OMNIbus 8 Plus Gateway for Message Bus
- Tivoli Netcool/OMNIbus Web GUI and the Web GUI extensions for Event Analytics
- IBM Tivoli Netcool/Impact and the Netcool/Impact extensions for Event Analytics

However, note that the following Netcool Operations Insight base solution components cannot be installed by using the Installation Manager web application:

- Dashboard Application Services Hub
- Operations Analytics - Log Analysis

Dashboard Application Services Hub also cannot be installed in console mode.

What to do next
If required, add the Installation Manager installation directory path to your PATH environment variable.
Related information
IBM Installation Manager V1.8.5 documentation: Working from a web browserClick here for information on how to use the Installation Manager web server to manage your installations.

Downloading for on-premises installation

You need to download products and components from Passport Advantage and Fix Central.

About this task

Refer to the following topic for information on where to obtain downloads for each product and component: “V1.6.0.1 Product and component version matrix” on page 13. For more information, see the IBM Support Download Document for Netcool Operations Insight, at http://www.ibm.com/support/docview.wss?uid=ibm10886869.

Installing on premises

Follow these instructions to install Operations Management and optionally install the solution extensions Network Management, Performance Management, and Service Management.

Quick reference to installing
Use this information as a quick reference if you are new to Netcool Operations Insight and want to perform an installation from scratch. This overview assumes detailed knowledge of the products in Netcool Operations Insight. It does not provide all the details. Links are given to more information, either within the Netcool Operations Insight documentation, or in the product documentation of the constituent products of Netcool Operations Insight.

This topic lists the high-level steps for installing Netcool Operations Insight.

“Installing Operations Management” on page 49
“Installing Network Management” on page 52
“Installing Performance Management” on page 54
“Installing Service Management” on page 54

Installing Operations Management

You can install Operations Management on premises, or within a private cloud using IBM Cloud Private.

Installing Operations Management on premises

The following table lists the high-level steps for installing Operations Management on premises.

For information on the product and component versions to install, including which fix packs to apply, see "Products and components on premises" on page 9.

Tip: To verify the versions of installed packages, select **View Installed Packages** from the File menu on the main IBM Installation Manager screen.

Table 18. Quick reference for installing Operations Management on premises

Item	Action	More information
1	Prepare for the installation by checking the prerequisites.	For information about hardware and software compatibility of each component, and detailed system requirements, see the IBM Software Product Compatibility Reports website: http://www-969.ibm.com/software/reports/compatibility/clarity/index.html "Checking prerequisites" on page 45
2	Install IBM Installation Manager on each host where components of the Netcool Operations Insight are to be installed. Installation Manager is included in the compressed file distribution of IBM Tivoli Netcool/OMNIbus and Operations Analytics - Log Analysis. Download Installation Manager separately if you are installing directly from an IBM repository or from a local repository. If you need to install Installation Manager separately, you can download it from IBM Fix Central.	http://www.ibm.com/support/fixcentral/ "Obtaining IBM Installation Manager" on page 46 "Installing Installation Manager (GUI or console example)" on page 47
3	Install the Netcool/OMNIbus core components, and apply the latest supported fix pack. Associated tasks include creating and starting ObjectServers, and setting up failover or a multitier architecture.	"Installing Netcool/OMNIbus and Netcool/Impact" on page 55 https://ibm.biz/BdE6tr https://ibm.biz/BdE6t4 https://ibm.biz/BdE6tF See "Products and components on premises" on page 9 for latest supported fix packs.
4	Install the Db2 database. Apply the REPORTER schema.	"Installing Db2 and configuring the REPORTER schema" on page 56
5	Install the Gateway for JDBC and the Gateway for Message Bus.	Gateway for JDBC documentation: https://ibm.biz/BdE9Db Gateway for Message Bus documentation: https://ibm.biz/BdEQaD

Table 18. Quick reference for installing Operations Management on premises (continued)

Item	Action	More information
6	Install Netcool/Impact, and apply the latest supported fix pack.	"Installing Netcool/OMNIbus and Netcool/Impact" on page 55 http://www-01.ibm.com/support/knowledgecenter/SSSHYH/welcome See "Products and components on premises" on page 9 for latest supported fix packs.
7	Configure the ObjectServer to support the related events function of the Event Analytics capability. Run the **nco_sql** utility against the `relatedevents_objectserver.sql` file, which is delivered with Netcool/Impact.	"Configuring the ObjectServer " on page 277
8	Install a supported version of IBM Operations Analytics - Log Analysis. Create a data source called "omnibus".	http://www-01.ibm.com/support/knowledgecenter/SSPFMY/welcome See step "3" on page 225 in "Configuring event search" on page 224. See "Products and components on premises" on page 9 for supported versions.
9	Configure the Gateway for Message Bus as the interface between the ObjectServer and Operations Analytics - Log Analysis. Optionally configure the Accelerated Event Notification Client if you do not want to use the default IDUC channel.	"Configuring the Gateway for JDBC and Gateway for Message Bus" on page 58 https://ibm.biz/BdEQaD
10	To support Event Search, install the latest supported version of Tivoli Netcool/OMNIbus Insight Pack.	For more information, see "Installing the Tivoli Netcool/OMNIbus Insight Pack" on page 61. See "Products and components on premises" on page 9 for latest supported versions.

Table 18. Quick reference for installing Operations Management on premises (continued)		
Item	**Action**	**More information**
11	Install the Netcool/OMNIbus Web GUI, and apply the latest supported fix pack. During the installation, ensure that the latest supported versions are selected for the following components, based on the information in "Products and components on premises" on page 9: • IBM WebSphere Application Server. • Jazz for Service Management. • Netcool/OMNIbus Web GUI In addition, make the following selections during the installation: • When installing Jazz for Service Management, Installation Manager discovers two required packages in the Jazz repository. Select the Jazz for Service Management extension for IBM WebSphere V8.5 and Dashboard Application Services Hub V3.1.3.0 packages for installation. • IBM WebSphere SDK Java Technology Edition V7.0.*x*. • Install the **Netcool Operations Insight Extensions for IBM Tivoli Netcool/OMNIbus Web GUI** feature. To install Event Analytics with seasonality reporting, ensure that **install Event Analytics** is selected.	"Installing Dashboard Application Services Hub and the UI components" on page 59 https://ibm.biz/BdE6kW See "Products and components on premises" on page 9 for latest supported fix packs.
12	Configure the Web GUI for integration with Operations Analytics - Log Analysis. In the `server.init` file, set the **`scala.*`** properties appropriately.	"Configuring event search" on page 224

Back to top

Installing Network Management

The following table lists the high-level steps for installing Network Management.

For information on the product and component versions to install, including which fix packs to apply, see "Products and components on premises" on page 9.

Tip: To verify the versions of installed packages, select **View Installed Packages** from the File menu on the main IBM Installation Manager screen.

Table 19. Quick reference for installing Network Management

Item	Action	More information
1	Installing the Probe for SNMP and Syslog Probe for Network Manager	"Installing the Probe for SNMP and Syslog Probe" on page 64
2	Optional: Configure the ObjectServer for integration with Network Manager by obtaining the **ConfigOMNI** script from the Network Manager installation package and running it against the ObjectServer.	"Optional: Preparing the ObjectServer for integration with Network Manager" on page 65 **Important:** If you have already installed Tivoli Netcool/OMNIbus, the Netcool/OMNIbus Knowledge Library, and the Probe for SNMP, you can now install Network Manager, and do not need to follow the steps in this task. The Network Manager installer configures Tivoli Netcool/OMNIbus for you during the installation process. If the ObjectServer setup changes after you have already installed and configured Network Manager and Tivoli Netcool/OMNIbus, then you must reintegrate the ObjectServer with Network Manager as described in this topic.
3	Prepare the topology database for use by Network Manager	"Preparing the database for Network Manager" on page 66
4	Install Network Manager core and GUI components, and apply the latest supported fix pack.	"Installing Network Manager IP Edition and Netcool Configuration Manager" on page 67 More information: https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/install/task/ins_installing.html See "Products and components on premises" on page 9 for latest supported fix packs.
5	Install and configure Netcool Configuration Manager and apply the latest supported fix pack. This involves configuring the integration with Network Manager.	"Installing Network Manager IP Edition and Netcool Configuration Manager" on page 67 "Configuring integration with Netcool Configuration Manager" on page 70 More information: http://www-01.ibm.com/support/knowledgecenter/SS7UH9_6.4.2/ncm/wip/install/concept/ncm_ins_installingncm.dita For more information about fix packs for Netcool Configuration Manager 6.4.2 , see https://www.ibm.com/support/knowledgecenter/SS7UH9_6.4.2/ncm/wip/relnotes/ncm_rn_top.html. See "Products and components on premises" on page 9 for latest supported fix packs.
6	To support Topology Search, install the latest supported version of Network Manager Insight Pack and configure the connection to the NCIM topology database.	For more information, see "Installing the Network Manager Insight Pack" on page 90. See "Products and components on premises" on page 9 for latest supported versions.

Table 19. Quick reference for installing Network Management (continued)		
Item	**Action**	**More information**
7	Configure the topology search capability. Run **nco_sql** against the `scala_itnm_configuration.sql` file, which is delivered in the Netcool/OMNIbus fix pack. Install the tools and menus to launch the custom apps of the Network Manager Insight Pack in the Operations Analytics - Log Analysis UI from the Web GUI.	"Configuring topology search" on page 362
8	Configure the Web GUI to launch the custom apps of the Network Manager Insight Pack from the event lists.	See step "3" on page 363 of "Configuring topology search" on page 362.

Back to top

Installing Performance Management

The following table lists the high-level steps for installing Performance Management.

For information on the product and component versions to install, including which fix packs to apply, see "Products and components on premises" on page 9.

Tip: To verify the versions of installed packages, select **View Installed Packages** from the File menu on the main IBM Installation Manager screen.

Table 20. Quick reference for installing Performance Management		
Item	**Action**	**More information**
1	To add the performance management feature, set up Network Performance Insight and follow the steps for integrating it with Netcool/OMNIbus and Network Manager.	"Installing Performance Management" on page 92
2	Install and configure the **Device Dashboard** to view performance information.	"Installing the Device Dashboard" on page 94

Back to top

Installing Service Management

The following table lists the high-level steps for installing Service Management.

Table 21. Quick reference for installing Service Management		
Item	**Action**	**More information**
1	To add the service management feature, install the Agile Service Manager core, observers, and UI, and then follow the steps for integrating the observers: • Integrate the Event Observer with the Netcool/OMNIbus gateway. • Integrate the ITNM Observer with the Network Manager ncp_model Topology manager process.	"Installing Agile Service Manager" on page 96

Back to top

Installing Operations Management on premises
Follow these instructions to install the Netcool Operations Insight base solution, also known as Operations Management for Operations Insight on premises.

Installing Netcool/OMNIbus and Netcool/Impact
Obtain and install the Netcool/OMNIbus core components, Netcool/Impact, and the Gateway for JDBC and Gateway for Message Bus. All these products are installed by IBM Installation Manager. You can use IBM Installation Manager to download the installation packages for these products and install them in a single flow. Extra configuration of each product is required after installation.

Procedure

- Install the Netcool/OMNIbus V8.1.0 core components.

  After the installation, you can use the Initial Configuration Wizard (**nco_icw**) to configure the product, for example, create and start ObjectServers, and configure automated failover or a multitier architecture. See related links later for instructions on installing Netcool/OMNIbus.
- Install the Netcool/Impact GUI server and Impact server. See related links later for instructions on installing Netcool/Impact.
- Apply the latest supported Netcool/OMNIbus core and Netcool/Impact fix packs. Also ensure that you apply the appropriate IBM Tivoli Netcool/Impact Server Extensions for Netcool Operations Insight feature. This is delivered in the Netcool/Impact fix pack.

  For information on the product and component versions supported in the current version of Netcool Operations Insight including supported fix packs, see "Products and components on premises" on page 9

  The IBM Tivoli Netcool/Impact Server Extensions for Netcool Operations Insight feature is required for the event analytics capability. Fix packs are available from IBM Fix Central, see http://www.ibm.com/support/fixcentral/.
- Create the connection from Netcool/Impact to the Db2 database as described in "Configuring Db2 database connection within Netcool/Impact" on page 259.
- Configure the ObjectServer to support the related events function of the event analytics capability.

  This requires a ParentIdentifier column in the alerts.status table.

  Add the column using the SQL utility as described in "Configuring the ObjectServer " on page 277.
- Configure the ObjectServer to support the topology search capability. In `$NCHOME/omnibus/extensions`, run the **nco_sql** utility against the `scala_itnm_configuration.sql` file.

```
./nco_sql -user root -password myp4ss -server NCOMS
< /opt/IBM/tivoli/netcool/omnibus/extensions/scala/scala_itnm_configuration.sql
```

Triggers are applied to the ObjectServer that delay the storage of events until the events are enriched by Network Manager IP Edition data from the NCIM database.

- Install the Gateway for JDBC and Gateway for Message Bus.

  After installation, create the connection between the ObjectServer and the gateways in the Server Editor (**nco_xigen**). See related links later for instructions on creating connections in the Server Editor.

- Configure the integration between Netcool/Impact and IBM Connections.

  This involves importing the $*IMPACT_HOME*/add-ons/IBMConnections/importData project into Netcool/Impact and adding IBM Connections properties to the $*IMPACT_HOME*/etc/NCI_server.props file. After you edit this file, restart the Netcool/Impact server. See related links later for instructions on configuring integration to IBM Connections and restarting the Impact server.

What to do next

Search on IBM Fix Central for available interim fixes and apply them. See http://www.ibm.com/support/fixcentral/.

Related concepts

Connections in the Server Editor

Related tasks

Installing Tivoli Netcool/OMNIbus

Creating and running ObjectServers

Configuring Db2 database connection within Netcool/Impact

You can configure a connection to a valid Db2 database from within IBM Tivoli Netcool/Impact.

Configuring the ObjectServer

Prior to deploying rules based on related event events or patterns you must run SQL to update the ObjectServer. This SQL introduces relevant triggers into the ObjectServer to enable to rules to be fully functional.

Configuring integration to IBM Connections

IBM Tivoli Netcool/Impact provides integration to IBM Connections by using a Netcool/Impact IBMConnections action function. The IBMConnections action function allows users to query forums and topics lists, create a new forum, create a new topic, and update existing topics. The IBMConnections action function package is available in the directory $IMPACT_HOME/integrations/IBMConnections.

Restarting the Impact server

Related reference

On-premises scenarios for Operations Management

This topic presents the scenarios available in a deployment of Operations Management on premises together with the associated architectures.

Initial configuration wizard

Ports used by products and components

Use this information to understand which ports are used by the different products and components that make up the Netcool Operations Insight solution.

Related information

Installing Netcool/Impact

Installing Db2 and configuring the REPORTER schema

Netcool Operations Insight requires a Db2 database with the REPORTER schema for historical event archiving.

Tip: For information on the housekeeping of historical Db2 event data, as well as sample SQL scripts, see the 'Historical event archive sizing guidance' section in the Netcool/OMNIbus*Best Practices Guide*, which can be found on the Netcool/OMNIbus best-practice Wiki: http://ibm.biz/nco_bps

Procedure

- Obtain and download the package for the Db2 database and the Gateway configuration scripts.
- Decompress the packages. Then, as the root system user, run the **db2setup** command to install the Db2 database on the host. The **db2setup** command starts the Db2 Setup wizard. Install as the root system user because the setup wizard needs to create a number of users in the process.
- Run IBM Installation Manager on the Netcool/OMNIbus host and install the Gateway configuration scripts. The SQL file that is needed to create the REPORTER schema is installed to `$OMNIHOME/gates/reporting/db2/db2.reporting.sql`.
- In the `db2.reporting.sql` file, make the following changes.
  - Uncomment the `CREATE DATABASE` line.
  - Set the default user name and password to match the Db2 installation:

    ```
    CREATE DATABASE reporter @
    CONNECT TO reporter USER db2inst1 USING db2inst1 @
    ```

  - Uncomment the following lines, so that any associated journal and details rows are deleted from the database when the corresponding alerts are deleted:

    ```
    -- Uncomment the line below to enable foreign keys
    -- This helps pruning by only requiring the alert to be
    -- deleted from the status table
    , CONSTRAINT eventref FOREIGN KEY (SERVERNAME, SERVERSERIAL) REFERENCES
    REPORTER_STATUS(SERVERNAME, SERVERSERIAL) ON DELETE CASCADE)
    ```

    This SQL appears twice in the SQL file: once in the details table definition and once in the journal table definition. Uncomment both instances.
- Run the SQL file against the Db2 database by running the following command as the db2inst1 system user:

  ```
  $ db2 -td@ -vf db2.reporting.sql
  ```

Result
The Db2 installer creates a number of users including db2inst1.
Related reference
Ports used by products and components
Use this information to understand which ports are used by the different products and components that make up the Netcool Operations Insight solution.

Related information
Installing Db2 servers using the Db2 Setup wizard (Linux and UNIX)
Gateway for JDBC configuration scripts for Reporting Mode

Installing IBM Operations Analytics - Log Analysis
Operations Analytics - Log Analysis supports GUI, console, and silent installations. The installation process differs for 64-bit and z/OS operating systems.

Procedure

Operations Analytics - Log Analysis can be installed by IBM Installation Manager or you can run the `install.sh` wrapper script.

Tip: The best practice is to install the Web GUI and Operations Analytics - Log Analysis on separate hosts.

Restriction: Operations Analytics - Log Analysis does not support installation in Group mode of IBM Installation Manager.

What to do next

- If the host locale is not set to English United States, set the locale of the command shell to `export LANG=en_US.UTF-8` before you run any Operations Analytics - Log Analysis scripts.
- Install the Tivoli Netcool/OMNIbus Insight Pack into Operations Analytics - Log Analysis to enable ingestion of event data into Operations Analytics. For more information, see "Installing the Tivoli Netcool/OMNIbus Insight Pack" on page 61.
- (Optional) Install the Network Manager Insight Pack into Operations Analytics - Log Analysis to use the topology search capability. For more information, see "Installing the Network Manager Insight Pack" on page 90.
- Search on IBM Fix Central for available interim fixes and apply them. See http://www.ibm.com/support/fixcentral/.

Related tasks

Installing Operations Analytics - Log Analysis V1.3.5

Installing Operations Analytics - Log Analysis V1.3.3

Related reference

Ports used by products and components

Use this information to understand which ports are used by the different products and components that make up the Netcool Operations Insight solution.

Configuring the Gateway for JDBC and Gateway for Message Bus

Configure the Gateway for JDBC to run in reporting mode, so it can forward event data to the Db2 database for archiving. Configure the Gateway for Message Bus gateway to forward event data to Operations Analytics - Log Analysis and run it in Operations Analytics - Log Analysis mode.

Before you begin

- Install the Db2 database and configure the REPORTER schema so that the Gateway for JDBC can connect.
- Install the gateways on the same host as Tivoli Netcool/OMNIbus core components (that is, server 1).
- Install Operations Analytics - Log Analysis and obtain the URL for the connection to the Gateway for Message Bus.

Procedure

- Configure the Gateway for JDBC.

  This involves the following steps:

  – Obtain the JDBC driver for the target database from the database vendor and install it according to the vendor's instructions. The drivers are usually provided as `.jar` files.

  – To enable the gateway to communicate with the target database, you must specify values for the **`Gate.Jdbc.*`** properties in the `$OMNIHOME/etc/G_JDBC.props` file. This is the default properties file, which is configured for reporting mode, that is supplied with the gateway.

  Here is a sample properties file for the Gateway for JDBC.

```
# Reporting mode properties
Gate.Jdbc.Mode: 'REPORTING'
# Table properties
Gate.Jdbc.StatusTableName: 'REPORTER_STATUS'
Gate.Jdbc.JournalTableName: 'REPORTER_JOURNAL'
Gate.Jdbc.DetailsTableName: 'REPORTER_DETAILS'
# JDBC Connection properties
Gate.Jdbc.Driver: 'com.ibm.db2.jcc.Db2Driver'
Gate.Jdbc.Url: 'jdbc:db2://server3:50000/REPORTER'
Gate.Jdbc.Username: 'db2inst1'
Gate.Jdbc.Password: 'db2inst1'
Gate.Jdbc.ReconnectTimeout: 30
Gate.Jdbc.InitializationString: ''
# ObjectServer Connection properties
Gate.RdrWtr.Username: 'root'
```

```
Gate.RdrWtr.Password: 'netcool'
Gate.RdrWtr.Server: 'AGG_V'
```

- Configure the Gateway for Message Bus to forward event data to Operations Analytics - Log Analysis.
  This involves the following steps:
  - Creating a gateway server in the Netcool/OMNIbus interfaces file
  - Configuring the `G_SCALA.props` properties file, including specifying the `.map` mapping file.
  - Configuring the endpoint in the `scalaTransformers.xml` file
  - Configuring the SSL connection, if required
  - Configuring the transport properties in the `scalaTransport.properties` file
- If you do not want to use the default configuration of the Gateway for Message Bus (an IDUC channel between the ObjectServer and Operations Analytics - Log Analysis and supports event inserts only), configure event forwarding through the AEN client.
  This support event inserts and reinserts and involves the following steps:
  - Configuring AEN event forwarding in the Gateway for Message Bus
  - Configuring the AEN channel and triggers in each ObjectServer by enabling the postinsert triggers and trigger group
- Start the Gateway for Message Bus in Operations Analytics - Log Analysis mode.
  For example:

  ```
  $OMNIHOME/bin/nco_g_xml -propsfile $OMNIHOME/etc/G_SCALA.props
  ```

  The gateway begins sending events from Tivoli Netcool/OMNIbus to Operations Analytics - Log Analysis.
- Start the Gateway for JDBC in reporter mode.
  For example:

  ```
  $OMNIHOME/bin/nco_g_jdbc -jdbcreporter
  ```

- As an alternative to starting the gateways from the command-line interface, put them under process control.

Related concepts

Tivoli Netcool/OMNIbus process control

Related information

Gateway for JDBC documentation

Gateway for Message Bus documentation

Installing Dashboard Application Services Hub and the UI components

Install the Dashboard Application Services Hub and all the UI components. This applies to the Netcool/OMNIbus Web GUI, the Event Analytics component, fix packs, and optionally Reporting Services.

The UI components are installed in two stages. First, IBM WebSphere Application Server and Jazz for Service Management are installed, which provide the underlying UI technology. Then, the Web GUI and the extension packages that support the Event Analytics component and the event search capability are installed. After installation, configure the Web GUI to integrate with Operations Analytics - Log Analysis and support the topology search capability.

You can optionally install Reporting Services V3.1 into Dashboard Application Services Hub. You can set up Network Manager and Netcool Configuration Manager to work with Reporting Services by installing their respective reports when installing the products. Netcool/OMNIbus V8.1.0 and later can be integrated with Reporting Services V3.1 to support reporting on events. To configure this integration, connect Reporting Services to a relational database through a gateway. Then, import the report package that is supplied with Netcool/OMNIbus into Reporting Services. For more information about event reporting, see http://www-01.ibm.com/support/knowledgecenter/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/omnibus/wip/install/task/omn_con_ext_deploytcrreports.html.

Before you begin

- Obtain the packages from IBM Passport Advantage. For information about the eAssembly numbers you need for the packages, see http://www-01.ibm.com/support/docview.wss?uid=swg24043698.
- To install Reporting Services V3.1, ensure that the host meets the extra requirements at http://www.ibm.com/support/knowledgecenter/SSEKCU_1.1.3.0/com.ibm.psc.doc/install/tcr_c_install_prereqs.html.

Procedure

1. Start Installation Manager and install Jazz for Service Management.

   The packages that you need to install are as follows.

Package	Description
IBM WebSphere Application Server V8.5.5.15 for Jazz for Service Management	Select V8.5.5.15. If V8.0.5 is also identified, clear it.
IBM WebSphere SDK Java Technology Edition V7.0.*x*.	
Jazz for Service Management V1.1.3.5	Select the following items for installation. • Jazz for Service Management extension for IBM WebSphere V8.5. • Dashboard Application Services Hub V3.1.3.0.
Reporting Services V3.1	This package is optional. Select it if you want to run reports for events and network management.

2. Install the packages that constitute the Web GUI and extensions.

Package	Description
Netcool/OMNIbus Web GUI	This is the base component that installs the Web GUI.
Install tools and menus for event search with IBM SmartCloud Analytics - Log Analysis	This package installs the tools that launch the custom apps of the Tivoli Netcool/OMNIbus Insight Pack from the event lists.
Netcool Operations Insight Extensions for IBM Tivoli Netcool/OMNIbus Web GUI	This package installs the Event Analytics GUIs.
Netcool/OMNIbus **Web GUI V8.1.0 fix pack, as specified in "Products and components on premises" on page 9.**	This is the fix pack that contains the extensions for the topology search capability.

3. Configure the Web GUI.
   For example, the connection to a data source (ObjectServer), users, groups, and so on.

   You can use the Web GUI configuration tool to do this. For more information, see https://ibm.biz/BdXqcP.
4. Configure the integration with Operations Analytics - Log Analysis.

   Ensure that the `server.init` file has the following properties set:

```
scala.app.keyword=OMNIbus_Keyword_Search
scala.app.static.dashboard=OMNIbus_Static_Dashboard
scala.datasource=omnibus
scala.url=protocol://host:port
scala.version=1.2.0.3
```

If you need to change any of these values, restart the Web GUI server.

5. Set up the Web GUI Administration API client.
6. Install the tools and menus to launch the custom apps of the Network Manager Insight Pack in the Operations Analytics - Log Analysis UI from the Web GUI.

   In $*WEBGUI_HOME*/extensions/LogAnalytics, run the **runwaapi** command against the scalaEventTopology.xml file.

   ```
   $WEBGUI_HOME/waapi/bin/runwaapi -user username -password password -file
   scalaEventTopology.xml
   ```

   Where *username* and *password* are the credentials of the administrator user that are defined in the $*WEBGUI_HOME*/waapi/etc/waapi.init properties file that controls the WAAPI client.

What to do next

- Search on IBM Fix Central for available interim fixes and apply them. See http://www.ibm.com/support/fixcentral/.
- Reconfigure your views in the Web GUI event lists to display the **NmosObjInst** column. The tools that launch the custom apps of the Network Manager Insight Pack work only against events that have a value in this column. For more information, see http://www-01.ibm.com/support/knowledgecenter/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/webtop/wip/task/web_cust_settingupviews.html.

Related concepts

Installing and uninstalling Event Analytics
Read the following topics before you install or uninstall Event Analytics.

Related tasks

Installing the Web GUI

Restarting the server

Related reference

server.init properties

Ports used by products and components
Use this information to understand which ports are used by the different products and components that make up the Netcool Operations Insight solution.

Installing the Tivoli Netcool/OMNIbus Insight Pack

This topic explains how to install the Netcool/OMNIbus Insight Pack into the Operations Analytics - Log Analysis product. Operations Analytics - Log Analysis can be running while you install the Insight Pack. This Insight Pack ingests event data into Operations Analytics - Log Analysis and installs custom apps.

Before you begin

- Install the Operations Analytics - Log Analysis product. The Insight Pack cannot be installed without the Operations Analytics - Log Analysis product.
- Download the relevant Insight Pack installation package from IBM Passport Advantage, ensuring that the downloaded version is compatible with the installed versions of Netcool Operations Insight and Operations Analytics - Log Analysis.
- Install Python 2.6 or later with the simplejson library, which is required by the Custom Apps that are included in Insight Packs.

Procedure

1. Create a new OMNIbus directory under $UNITY_HOME/unity_content.
2. Copy the Netcool/OMNIbus Insight Pack installation package to $UNITY_HOME/unity_content/OMNIbus.
3. Unpack and install the Insight Pack, using the following command as an example:

```
$UNITY_HOME/utilities/pkg_mgmt.sh -install $UNITY_HOME/unity_content/OMNIbus/
OMNIbusInsightPack_v1.3.0.2.zip
```

4. On the Operations Analytics - Log Analysis UI, use the Data Source Wizard to create a data source into which the event data is ingested.

   The "omnibus1100" data source can ingest data for both the Tivoli Netcool/OMNIbus Insight Pack and the Network Manager Insight Pack.

   a) In the **Select Location** panel, select **Custom** and type the Netcool/OMNIbus server host name.

      Enter the same host name that was used for the **JsonMsgHostname** transport property of the Gateway for Message Bus.

   b) In the **Select Data** panel, enter the following field values:

Field	Value
File path	NCOMS. This is the default value of the **jsonMsgPath** transport property of the Gateway for Message Bus. If you changed this value from the default, change the value of the **File path** field accordingly.
Type	This is the name of the data source type on which this data source is based. • To use the default data source type, specify OMNIbus1100. • To use a customized data source type, specify the name of the customized data source type; for example: customOMNIbus
Collection	OMNIbus1100-Collection

   c) In the **Set Attributes** panel, enter the following field values:

Field	Value
Name	omnibus. Ensure that the value that you type is the same as the value of the **scala.datasource** property in the Web GUI server.init file. If the **Name** field has a value other than omnibus, use the same value for the **scala.datasource** property.
Group	Leave this field blank.
Description	Type a description of your choice.

Results

The Insight Pack is installed to the directory specified in step 3. After the installation is completed, the Rule Set, Source Type, and Collection required for working with Netcool/OMNIbus events is in place. You can view these resources in the **Administrative Settings** page of the Operations Analytics - Log Analysis UI.

What to do next

- Use the **pkg_mgmt** command to verify the installations of the Insight Pack. See Verifying the Tivoli Netcool/OMNIbus Insight Pack.
- If you have several ObjectServers, use separate instances of the Gateway for Message Bus to connect to each ObjectServer. The best practice is for each gateway to send events to a single data source. For more information about configuring the gateway to send events to Operations Analytics - Log Analysis, see the Gateway for Message Bus documentation at https://ibm.biz/BdEQaD and search for *Integrating with Operations Analytics - Log Analysis*.

Related concepts

Data Source creation in Operations Analytics - Log Analysis V1.3.5

Data Source creation in Operations Analytics - Log Analysis V1.3.3

Related tasks

Installing the Tivoli Netcool/OMNIbus Insight Pack
This topic explains how to install the Netcool/OMNIbus Insight Pack into the Operations Analytics - Log Analysis product. Operations Analytics - Log Analysis can be running while you install the Insight Pack. This Insight Pack ingests event data into Operations Analytics - Log Analysis and installs custom apps.

Installing the Network Manager Insight Pack
This topic explains how to install the Network Manager Insight Pack into the Operations Analytics - Log Analysis product and make the necessary configurations. The Network Manager Insight Pack is required only if you deploy the Networks for Operations Insight feature and want to use the topology search capability. For more information, see "Network Manager Insight Pack" on page 360. Operations Analytics - Log Analysis can be running while you install the Insight Pack.

Related information

Gateway for Message Bus documentation

Installing Network Management

This installation scenario describes how to set up the Networks for Operations Insight feature in the Netcool Operations Insight solution. A sample system topology is given on which the installation tasks are based. It is assumed that the core products of the Netcool Operations Insight solution are already installed and running.

The information supplied in this scenario is high-level and covers the most salient points and possible issues you might encounter that are specific to the Networks for Operations Insight feature in the Netcool Operations Insight solution. This scenario is end-to-end and you should perform the tasks in the specified order.

For more information, see the Related concept, task, and information links at the bottom of this topic.

Before you begin

- Install the components of Netcool Operations Insight as described in "Products and components on premises" on page 9. The Networks for Operations Insight solution requires that the following products are installed, configured, and running as follows:
  - The Tivoli Netcool/OMNIbus V8.1 server components are installed and an ObjectServer is created and running. Ensure that the administrator user of the ObjectServer was changed from the default.
  - The Tivoli Netcool/OMNIbus V8.1 Web GUI is installed and running in an instance of Dashboard Application Services Hub. The ObjectServer is defined as the Web GUI data source.
  - An IBM Db2 database is installed and configured for event archiving, and the Gateway for JDBC is installed and configured to transfer and synchronize the events from the ObjectServer.
  - IBM Operations Analytics - Log Analysis is installed and running, and configured so that events are forwarded from Tivoli Netcool/OMNIbus to Operations Analytics - Log Analysis via the Gateway for Message Bus. See "Configuring event search" on page 224.
- Obtain the following information about the ObjectServer:
  - Host name and port number
  - Installation directory (that is, the value of the $OMNIHOME environment variable)
  - Name, for example, NCOMS
  - Administrator password
- Install and configure the event search and event seasonality features.

If any of the above products are not installed, or features not configured, they must be configured before you can set up the Networks for Operations Insight feature.

About this task

This task and the sub-tasks describe the scenario of a fresh deployment of the products in the Networks for Operations Insight feature. The system topology is a logical sample. It is not the only system topology

that can be used. It is intended for reference and to help you plan your deployment. The system topology is as follows:

- Tivoli Netcool/OMNIbus and Network Manager are installed on separate hosts (that is, a distributed installation). The version of Tivoli Netcool/OMNIbus is 8.1.
- The ObjectServer is configured to be the user repository for the products.

  Note: All the products of the Netcool Operations Insight solution also support the use of an LDAP directory as the user repository.
- Network Manager and Netcool Configuration Manager both use the V8.1 ObjectServer to store and manage events.
- In this topology, the default Db2 v10.5 Enterprise Server Edition database is used.

Related concepts

Network Management data flow

Use this information to understand how event data is retrieved from a monitored application environment and transferred between the products and components of Network Management in order to provide Topology Search, **Network Health Dashboard** and **Device Dashboard** capabilities.

Related tasks

Installing Performance Management

To add the Performance Management solution extension, install Network Performance Insight and then integrate it with Netcool Operations Insight.

Related reference

Supported products for Networks for Operations Insight

Release notes

IBM Netcool Operations Insight V1.6.0.1 is available. Compatibility, installation, and other getting-started issues are addressed in these release notes.

Installing the Probe for SNMP and Syslog Probe

The Networks for Operations Insight feature requires the Probe for SNMP and the Syslog Probe. It is important that you install the probes that are included in the entitlement for the Tivoli Netcool/OMNIbus V8.1 product. Although the probes are also available in the Network Manager IP Edition entitlement, do not install them from Network Manager IP Edition. The instances of the probes that are available with Tivoli Netcool/OMNIbus V8.1 are installed by IBM Installation Manager.

Procedure

1. Change to the `/eclipse/tools` subdirectory of the Installation Manager installation directory and run the following command to start Installation Manager:

   ```
   ./IBMIM
   ```

   To record the installation steps in a response file for use with silent installations on other computers, use the `-record` option. For example, to record to the `/tmp/install_1.xml` file:

   ```
   ./IBMIM -record /tmp/install_1.xml
   ```
2. Configure Installation Manager to download package repositories from IBM Passport Advantage.
3. In the main **Installation Manager** pane, click **Install** and follow the installation wizard instructions to complete the installation.

   The installer requires the following inputs at different stages of the installation:

   - If prompted, enter your IBM ID user name and password.
   - Read and accept the license agreement.
   - Specify an Installation Manager shared directory or accept the default directory.

   Select the `nco-p-syslog` feature for the Syslog Probe, and select the `nco-p-mttrapd` feature for the Probe for SNMP.

After the installation completes, click **Finish**.

Results
If the installation is successful, Installation Manager displays a success message and the installation history is updated to record the successful installation. If not, you can use Installation Manager to uninstall or modify the installation.

What to do next
Ensure that both probes are configured:

- For more information about configuring the Probe for SNMP, see http://www-01.ibm.com/support/knowledgecenter/SSSHTQ/omnibus/probes/snmp/wip/reference/snmp_config.html.
- For more information about configuring the Syslog Probe, see http://www-01.ibm.com/support/knowledgecenter/SSSHTQ/omnibus/probes/syslog/wip/concept/syslog_intro.html.

Related tasks
Netcool/OMNIbus V8.1 documentation: Obtaining IBM Installation Manager You can install IBM Installation Manager with a GUI or console, or do a silent installation. Before installation, you must determine which user mode you require.
Installing Tivoli Netcool/OMNIbus V8.1
Related reference
Tivoli Netcool/OMNIbus V8.1 installable features

Optional: Preparing the ObjectServer for integration with Network Manager
If you have already installed Tivoli Netcool/OMNIbus, the Netcool/OMNIbus Knowledge Library, and the Probe for SNMP, you can now install Network Manager, and do not need to follow the steps in this task. The Network Manager installer configures Tivoli Netcool/OMNIbus for you during the installation process. If the ObjectServer setup changes after you have already installed and configured Network Manager and Tivoli Netcool/OMNIbus, then you must reintegrate the ObjectServer with Network Manager as described in this topic.

To reintegrate the Network Manager product with the existing Tivoli Netcool/OMNIbus V8.1 ObjectServer, run the **ConfigOMNI** script against the ObjectServer.

Procedure

1. Use the **ConfigOMNI** script to configure an ObjectServer to run with Network Manager.

   The script creates the Network Manager triggers and GUI account information. If the ObjectServer is on a remote server, then copy the `$NCHOME/precision/install/scripts/ConfigOMNI` script and the support script `$NCHOME/precision/scripts/create_itnm_triggers.sql` and put them into the same directory on the remote ObjectServer. If the ObjectServer is local to Network Manager, then you can use both scripts as is.
2. On the ObjectServer host, change to the `scripts` directory and run the **ConfigOMNI** script.

   For example, the following configures the ObjectServer called `NCOMS2` using the administrative password `NC0M5password`, or creates the ObjectServer called `NCOMS2` if it does not exist, in the specified directory (OMNIHOME), and creates or modifies the `itnmadmin` and `itnmuser` users in the ObjectServer.

   ```
   ./ConfigOMNI -o NCOMS2 -p NC0M5password -h /opt/ibm/tivoli/netcool
   -u ITNMpassword
   ```
3. You might also need to update the Network Manager core settings and the Web GUI data source settings. For more information, see https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/install/task/ins_installingandconfiguringomnibus.html.

Related tasks
Installing and configuring Tivoli Netcool/OMNIbus

Preparing the database for Network Manager

After a supported database has been installed, you must install and run the database scripts to configure the topology database for use by Network Manager IP Edition. You must run the scripts before installing Network Manager IP Edition.

About this task

If you downloaded the compressed software package from Passport Advantage, the database creation scripts are included at the top level of the uncompressed software file. Copy the scripts to the database server and use them.

You can also install the Network Manager IP Edition topology database creation scripts using Installation Manager by selecting the **Network Manager topology database creation scripts** package. The database scripts are installed by default in the `precision/scripts/` directory in the installation directory (by default, `/opt/IBM/netcool/core/precision/scripts/`).

Procedure

1. Log in to the server where you installed Db2.
2. Change to the directory where your Db2 instance was installed and then change to the `sqllib` subdirectory.
3. Set up the environment by typing the following command:

Shell	Command
Bourne	`. ./db2profile`
C	`source db2cshrc`

   The Network Manager IP Edition application wrapper scripts automatically set up the Db2 environment.
4. Locate the compressed database creation file `db2_creation_scripts.tar.gz` and copy it to the server where Db2 is installed. Decompress the file.
5. Change to the `precision/scripts/` directory and run the `create_db2_database.sh` script as the Db2 administrative user for the instance (db2inst1):

   ```
   ./create_db2_database.sh database_name user_name -force
   ```

   Where *database_name* is the name of the database, *user_name* is the Db2 user to use to connect to the database, and `-force` an argument that forces any Db2 users off the instance before the database is created.

   Important: The *user_name* must not be the administrative user. This user must be an existing operating system and Db2 user.

   For example, to create a Db2 database that is called ITNM for the Db2 user ncim, type:

   ```
   ./create_db2_database.sh ITNM ncim
   ```
6. After you run `create_db2_database.sh`, restart the database as the Db2 administrative user as follows: **`run db2stop`** and then **`run db2start`**.
7. When running the Network Manager IP Edition installer later on, make sure you select the option to configure an existing Db2 database. The Network Manager IP Edition installer can then create the tables in the database either on the local or a remote host, depending on where your database is installed.

   The installer populates the connection properties in the following files, you can check these files for any problems with your connection to the database:

   - The `DbLogins.DOMAIN.cfg` and `MibDbLogin.cfg` files in `$NCHOME/etc/precision`. These files are used by the Network Manager IP Edition core processes.

- The `tnm.properties` file in `$NMGUI_HOME/profile/etc/tnm`. These files are used by the Network Manager IP Edition GUI.

Installing Network Manager IP Edition and Netcool Configuration Manager

Install Network Manager IP Edition and Netcool Configuration Manager to form the basis of the Networks for Operations Insight feature.

Before you begin

- Ensure you have installed and configured the base products and components of Netcool Operations Insight, including Tivoli Netcool/OMNIbus, Netcool/Impact, and Operations Analytics - Log Analysis, and the associated components and configurations. See Supported products for Networks for Operations Insight.
- Obtain the following information about the ObjectServer:
  - ObjectServer name, for example, NCOMS
  - Host name and port number
  - Administrator user ID
  - Administrator password
- Obtain the following information about your Db2 database:
  - Database name
  - Host name and port number
  - Administrator user ID with permissions to create tables
  - Administrator user password
- Obtain the packages from IBM Passport Advantage. For information about the eAssembly numbers you need for the packages, see http://www-01.ibm.com/support/docview.wss?uid=swg24043698.
- Obtain the latest supported fix packs for Network Manager IP Edition and Netcool Configuration Manager from IBM Fix Central, at http://www.ibm.com/support/fixcentral/. For information on the product and component versions supported in the current version of Netcool Operations Insight including supported fix packs, see “Products and components on premises” on page 9.
- Ensure that a compatible version of Python is installed on this server before you start. On Linux, Network Manager IP Edition core components require version 2.6 or 2.7 of Python to be installed on the server where the core components are installed. On AIX®, Network Manager IP Edition requires version 2.7.5 of Python.

About this task

These instructions describe the options that are presented in the Installation Manager in wizard mode. Other modes are also available with equivalent options.

Procedure

1. Start Installation Manager and install the following packages:

Package	Description
Network Manager Core Components Version V4.2.0.8	Installs the Network Manager IP Edition core components, sets up connection to the specified ObjectServer, sets up connection to the database to be used for the NCIM topology and creates the tables (needs to be selected), creates Network Manager IP Edition default users, sets up network domain, and configures the details for the poller aggregation. For more information, see https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/install/task/ins_installingcorecomponents.html.

Package	Description
	The Network Manager IP Edition core components can be installed on server 4 of the scenario described in Performing a fresh installation.
Network Manager GUI Components Version V4.2.0.8	Installs the Network Manager IP Edition GUI components, sets up connection to the specified ObjectServer, sets up connection to the NCIM topology database, and sets up the default users. For more information, see https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/install/task/ins_installingguicomponents.html. The Network Manager IP Edition GUI components can be installed on server 3 of the scenario described in Performing a fresh installation. The GUI components of other products in the solution, Netcool Configuration Manager, and Reporting Services would also be on this host.
Network Health Dashboard V4.2.0.8	Installs the **Network Health Dashboard**. Installing the Network Health Dashboard installs the following roles, which allow users to work with the Network Health Dashboard: • ncp_networkhealth_dashboard • ncp_networkhealth_dashboard_admin • ncp_event_analytics The new **Network Health Dashboard** is only available if you have Network Manager as part of Netcool Operations Insight. The **Network Health Dashboard** monitors a selected network view, and displays device and interface availability within that network view. It also reports on performance by presenting graphs, tables, and traces of KPI data for monitored devices and interfaces. A dashboard timeline reports on device configuration changes and event counts, enabling you to correlate events with configuration changes. The dashboard includes the **Event Viewer**, for more detailed event information. **Note:** The Network Health Dashboard must be installed on the same host as the Network Manager IP Edition GUI components.
Network Manager Reports V4.2.0.8	Installs the reports provided by Network Manager IP Edition that you can use as part of the Reporting Services feature. Reporting Services requires a Db2 database to store its data. This database must be running during installation. If the database is installed on the same server as Reporting Services, the installer configures the database during installation. If the database is on a different server, you must configure the database before you install Reporting Services. In the scenario described in Performing a fresh installation, where the Db2 database is on a different server, you must set up the remote Db2 database for Reporting Services as follows: a. From the Jazz for Service Management package, copy the `TCR_generate_content_store_db2_definition.sh` script to the server where Db2 is installed. b. Run the following command: `./TCR_generate_content_store_db2_definition.sh` *`database_name db2_username`*

Package	Description

	Where *database_name* is the name you want for the Reporting Services database, and *db2_username* is the user name to connect to the content store, that is, the database owner (db2inst1).
	c. Copy the generated SQL script to a temporary directory and run it against your Db2 instance as the Db2 user (db2inst1), for example:
	`$ cp tcr_create_db2_cs.sql /tmp/tcr_create_db2_cs.sql`
	`$ su – db2inst1 –c "db2 -vtf /tmp/tcr_create_db2_cs.sql"`

Netcool Configuration Manager V6.4.2.9	Installs the Netcool Configuration Manager components and loads the required database schema. For **Server Installation Type**, select **Presentation Server and Worker Server** to install both the GUI and worker servers. For more information, see http://www-01.ibm.com/support/knowledgecenter/SS7UH9_6.4.2/ncm/wip/install/task/ncm_ins_installingncm.html. The Netcool Configuration Manager components can be installed on server 3 of the scenario described in Performing a fresh installation. For more information about fix pack 2, see http://www.ibm.com/support/knowledgecenter/SS7UH9_6.4.2/ncm/wip/relnotes/ncm_rn_6422.html.
Reporting Services environment	Installs the reports provided by Netcool Configuration Manager (ITNCM-Reports) that you can use as part of the Reporting Services feature.

2. Apply the latest supported Network Manager IP Edition and Netcool Configuration Manager fix packs.

   For information on the product and component versions supported in the current version of Netcool Operations Insight including supported fix packs, see “Products and components on premises” on page 9
3. On the host where the Network Manager GUI components are installed, install the tools and menus to launch the custom apps of the Network Manager Insight Pack in the Operations Analytics - Log Analysis GUI from the Network Views.

   a) In `$NMGUI_HOME/profile/etc/tnm/topoviz.properties`, set the **`topoviz.unity.customappsui`** property, which defines the connection to Operations Analytics - Log Analysis.
   For example:

   ```
   # Defines the LogAnalytics custom App launcher URL
   topoviz.unity.customappsui=https://server3:9987/Unity/CustomAppsUI
   ```

   b) In the `$NMGUI_HOME/profile/etc/tnm/menus/ncp_topoviz_device_menu.xml` file, define the **Event Search** menu item.

   Add the item `<menu id="Event Search"/>` in the file as shown:

   ```
   <tool id="showConnectivityInformation"/>
                  <separator/>
                  <menu id="Event Search"/>
   ```
4. Optional: Follow the steps to configure the integration between Network Manager IP Edition and Netcool Configuration Manager as described in “Configuring integration with Netcool Configuration Manager” on page 70.

Results

The ports used by each installed product or component are displayed. The ports are also written to the `$NCHOME/log/install/Configuration.log` file.

What to do next

- To add the performance management feature, see “Installing Performance Management” on page 92.
- To set up the **Device Dashboard** for the network performance monitoring feature, see “Installing the Device Dashboard” on page 94.
- Search on IBM Fix Central for available interim fixes and apply them. See http://www.ibm.com/support/fixcentral/.

Related reference

Installing Network Manager

Related information

Installing Netcool Configuration Manager

V4.2 download document

Configuring integration with Netcool Configuration Manager

After installing the products, you can configure the integration between Network Manager and Netcool Configuration Manager.

Related tasks

Installing Netcool Configuration Manager

Related reference

Netcool Configuration Manager release notes

Installation information checklist

Related information

Preparing Db2 databases for Netcool Configuration Manager

User role requirements

Certain administrative user roles are required for the integration.

Note: For single sign-on information, see the related topic links.

DASH user roles

The following DASH roles are required for access to the Netcool Configuration Manager components that are launched from within DASH, such as the Netcool Configuration Manager Wizards and the Netcool Configuration Manager - Base and Netcool Configuration Manager - Compliance clients.

Either create a DASH user with the same name as an existing Netcool Configuration Manager user who already has the ‘IntellidenUser' role, or use an appropriate Network Manager user, such as itnmadmin, who is already set up as a DASH user. If you use the Network Manager user, create a corresponding new Netcool Configuration Manager user with the same name (password can differ), and assign the ‘IntellidenUser' role to this new user.

Important: If a DASH user is being created on Network Manager with the same name as an existing Netcool Configuration Manager user. then they also need to be added to an appropriate Network Manager user group, or alternatively be granted any required Network Manager roles manually.

Additionally, assign the following roles to your DASH user:

- ncp_rest_api
- ncmConfigChange
- ncmConfigSynch
- ncmIDTUser
- ncmPolicyCheck
- ncmActivityViewing
- ncmConfigViewing
- ncmConfigEdit
- ncmDashService

The following table cross-references security requirements between user interfaces, DASH roles, Netcool Configuration Manager functionality, and Netcool Configuration Manager realm content permissions. Use this information to assign DASH roles and define realm content permissions.

Table 22. UI security by DASH roles, Netcool Configuration Manager functionality, and realm content permissions

Access	**DASH role**	**Functionality**	**Realm content permissions**
Apply Modelled Command Set	ncmConfigChange	Execute Configuration Change	View, Execute
Apply Native Command Set	ncmConfigChange	Execute Configuration Change, Apply Native Command Sets	View, Execute
Synchronize (ITNCM to Device)	ncmConfigSynch	Execute Configuration Synchronization	View, Execute
Submit Configuration	ncmConfigChange	Execute Configuration Change	View, Execute
Apply Policy	ncmPolicyCheck	Execute Compliance Policy	View
View Configuration	ncmConfigViewing	n/a	View
Edit Configuration	ncmConfigEdit	n/a	View, Modify
Compare Configuration	ncmConfigViewing	n/a	View
IDT Automatic	ncmIDTUser	IDT Access, IDT Allow Auto Login	View
IDT Manual	ncmIDTUser	IDT Access, IDT Allow Manual Login	View
Find Device	n/a	n/a	View
View UOW Log	n/a	n/a	n/a
View IDT Log	n/a	n/a	View
Activity Viewer	ncmActivityViewing	n/a	n/a
Device Synchronization	ncp_rest_api	n/a	n/a
Access DASH services (through right-click menus)	ncmDashServices	n/a	n/a

Reporting Services user roles

Reporting Services and the Netcool Configuration Manager default reports are installed together with the DASH components.

Any user who needs to access reports requires the following permissions:

- The relevant Reporting Services roles for accessing the **Reporting** node in the DASH console. Assign these roles to enable users to run reports to which they are authorized from the Reporting Services GUI.
- The authorization to access the report set, and the relevant Reporting Services roles for working with the reports. Assign these permissions to enable users to run Netcool Configuration Manager reports from Network Manager topology displays, the Active Event List, and the Reporting Services GUI.

For information about authorizing access to a report set and assigning roles by user or group, go to the IBMTivoli Systems Management Information Center at http://www-01.ibm.com/support/knowledgecenter/SS3HLM/welcome, locate the Reporting Services documentation node, and search for *authorization* and *user roles*.

Other user roles

To configure the **Alerts** menu in the Web GUI, the ncw_admin role is required.

Installing the Dashboard Application Services Hub components
For integrated scenarios, Netcool Configuration Manager provides the following Dashboard Application Services Hub components: The Activity Viewer, the Dashboard Application Services Hub wizards, and the Netcool Configuration Manager thick-client launch portal.

Before you begin

From Version 6.4.2 onwards, Netcool Configuration Manager reporting is no longer installed as part of the Dashboard Application Services Hub components installation, but rather as part of the Netcool Configuration Manager main installation.

Important: Before installing the Dashboard Application Services Hub components, install Netcool Configuration Manager using the 'Integrated' option.

About this task

Restriction: The Netcool Configuration Manager Dashboard Application Services Hub components must be installed as the same user who installed Network Manager.

Procedure

1. Log onto the Dashboard Application Services Hub server.
2. Change to the `/eclipse` subdirectory of the Installation Manager Group installation directory and use the following command to start Installation Manager:
   `./IBMIM`

   To record the installation steps in a response file for use with silent installations on other computers, use the '-record response_file' option. For example:

   ```
   IBMIM -record C:\response_files\install_1.xml
   ```
3. Configure Installation Manager to download package repositories from IBM Passport Advantage:
   a) From the main menu, choose **File** > **Preferences**.
   You can set preferences for proxy servers in IBM Installation Manager. Proxy servers enable connections to remote servers from behind a firewall.
   b) In the **Preferences** window, expand the Internet node and select one of the following options:

   FTP Proxy
   Select this option to specify a SOCKS proxy host address and a SOCKS proxy port number.

   HTTP Proxy
   Select this option to enable an HTTP or SOCKS proxy.

   c) Select **Enable proxy server**.
   d) In the **Preferences** window, select the **Passport Advantage** panel.
   e) Select **Connect to Passport Advantage**.
   f) Click **Apply**, and then click **OK**.
4. In the main **Installation Manager** window, click **Install**, select **IBM Dashboard Applications for ITNCM**, and then follow the installation wizard instructions to complete the installation.
5. Accept the license agreement, select an installation directory, and supply the following details:

Netcool Configuration Manager database details

Sid/service name/database name(db2)

Database hostname

Port

Username

Password

Dashboard Application Services Hub administrative credentials

Dashboard Application Services Hub administrator username (default is `smadmin`)

Dashboard Application Services Hub administrator password

Network Manager administrative credentials

Default is `itnmadmin` (or the Dashboard Application Services Hub superuser, who must have the 'ncw_admin' role in Dashboard Application Services Hub

Password

Netcool Configuration Manager presentation server

Connection details to the Netcool Configuration Manager Presentation server

A skip validation option should the Presentation server be unavailable

Reporting Services server

Connection details to the Reporting Services server

A skip validation option should the Reporting Services server be unavailable

6. Complete the installation.

Example

Tip: Best practice recommendation: You can generate a response file through Installation Manager, as in the following example:

```
<?xml version='1.0' encoding='UTF-8'?>
<agent-input>
  <variables>
    <variable name='sharedLocation' value='/opt/IBM/IMShared'/>
  </variables>
  <server>
    <repository location='/opt/IBM/IM/output'/>
  </server>
  <profile id='IBM Netcool GUI Components' installLocation='/opt/IBM/netcool/gui'>
    <data key='eclipseLocation' value='/opt/IBM/netcool/gui'/>
    <data key='user.import.profile' value='false'/>
        <!--Update OS to aix for AIX-->
    <data key='cic.selector.os' value='linux'/>
        <!--Update architecture to ppc64 for AIX-->
    <data key='cic.selector.arch' value='x86_64'/>
    <data key='cic.selector.ws' value='gtk'/>
    <data key='user.org.apache.ant.classpath' value='/root/IBM/InstallationManager_Group/eclipse/plugins/

org.apache.ant_1.8.3.v201301120609/lib/ant.jar'/>
    <data key='user.org.apache.ant.launcher.classpath' value='/root/IBM/InstallationManager_Group/eclipse/

plugins/org.apache.ant_1.8.3.v201301120609/lib/ant-launcher.jar'/>
    <data key='cic.selector.nl' value='en'/>
    <data key='user.DashHomeDir' value='/opt/IBM/JazzSM/ui'/>
    <data key='user.WasHomeDir' value='/opt/IBM/WebSphere/AppServer'/>
    <data key='user.DashHomeUserID' value='smadmin'/>
    <data key='user.DashHomeContextRoot' value='/ibm/console'/>
    <data key='user.DashHomeWasCell' value='JazzSMNode01Cell'/>
    <data key='user.DashHomeWasNode' value='JazzSMNode01'/>
    <data key='user.DashHomeWasServerName' value='server1'/>
    <data key='user.SaasEnabled' value=''/>
    <data key='user.JAZZSM_HOME,com.ibm.tivoli.netcool.itnm.gui' value='/opt/IBM/JazzSM'/>
    <data key='user.WAS_SERVER_NAME,com.ibm.tivoli.netcool.itnm.gui' value='server1'/>
    <data key='user.WAS_PROFILE_PATH,com.ibm.tivoli.netcool.itnm.gui' value='/opt/IBM/JazzSM/profile'/>
    <data key='user.WAS_USER_NAME,com.ibm.tivoli.netcool.itnm.gui' value='smadmin'/>
    <data key='user.itnm.ObjectServerUsername,com.ibm.tivoli.netcool.itnm.gui' value='root'/>
    <data key='user.itnm.ObjectServer.skip.validation,com.ibm.tivoli.netcool.itnm.gui' value='false'/>
    <data key='user.itnm.ObjectServerHostname,com.ibm.tivoli.netcool.itnm.gui' value='NMGUIServerLocation'/>
    <data key='user.itnm.ObjectServerName,com.ibm.tivoli.netcool.itnm.gui' value='NCOMS'/>
    <data key='user.itnm.ObjectServer.create.instance,com.ibm.tivoli.netcool.itnm.gui' value='false'/>
    <data key='user.itnm.ObjectServerMainPort,com.ibm.tivoli.netcool.itnm.gui' value='4105'/>
    <data key='user.itnm.database.server.type,com.ibm.tivoli.netcool.itnm.gui' value='db2'/>
    <data key='user.itnm.database.skip.validation,com.ibm.tivoli.netcool.itnm.gui' value='false'/>
    <data key='user.itnm.database.name,com.ibm.tivoli.netcool.itnm.gui' value='NCIM'/>
    <data key='user.itnm.database.hostname,com.ibm.tivoli.netcool.itnm.gui' value='DatabaseServerLocation'/>
```

```
    <data key='user.itnm.database.username,com.ibm.tivoli.netcool.itnm.gui' value='db2inst1'/>
    <data key='user.itnm.database.create.tables,com.ibm.tivoli.netcool.itnm.gui' value='false'/>
    <data key='user.itnm.database.tables.prefix,com.ibm.tivoli.netcool.itnm.gui' value=''/>
    <data key='user.itnm.database.port,com.ibm.tivoli.netcool.itnm.gui' value='50001'/>
    <data key='user.WAS_USER_NAME' value='smadmin'/>
    <data key='user.itnm.ObjectServerItnmAdminUsername,com.ibm.tivoli.netcool.itnm.gui' value='itnmadmin'/>
    <data key='user.itnm.ObjectServerItnmAdminUsername' value='itnmadmin'/>
    <data key='user.itncm.database.port' value='1521'/>
    <data key='user.itncm.database.schema' value='itncm'/>
    <data key='user.itncm.database.type' value='ORACLE_12'/>
    <data key='user.itncm.database.username' value='DBUSER'/>
    <data key='user.itncm.database.hostname' value='DatabaseServerLocation'/>
    <data key='user.itncm.pres.server.port' value='16311'/>
    <data key='user.itncm.pres.server.hostname' value='PresentationServerLocation'/>
    <data key='user.itncm.pres.server.skip.conn.check' value='true'/>
    <data key='user.itncm.pres.server.scheme' value='https'/>
    <data key='user.itncm.reports.path' value='/tarf/servlet/dispatch'/>
    <data key='user.itncm.reports.skip.conn.check' value='true'/>
    <data key='user.itncm.reports.port' value='16311'/>
    <data key='user.itncm.reports.hostname' value='TCRServerLocation'/>
    <data key='user.itncm.reports.scheme' value='https'/>
    <data key='user.WAS_PASSWORD,com.ibm.tivoli.netcool.itnm.gui' value=''/>
    <data key='user.itnm.ObjectServerItnmUserPassword,com.ibm.tivoli.netcool.itnm.gui' value=''/>
    <data key='user.WAS_PASSWORD' value=''/>
    <data key='user.itnm.ObjectServerItnmUserPassword' value=''/>
    <data key='user.itncm.database.password' value=''/>
  </profile>
  <install modify='false'>
    <!-- IBM Dashboard Applications for ITNCM 6.4.2 -->
    <offering profile='IBM Netcool GUI Components' id='com.ibm.tivoli.netcool.itncm.ui.dash' version=
'6.4.2.20160202_1049'
features='main.feature.activityviewer,main.feature.wizard' installFixes='none'/>
  </install>
  <preference name='com.ibm.cic.common.core.preferences.eclipseCache' value='${sharedLocation}'/>
  <preference name='com.ibm.cic.common.core.preferences.connectTimeout' value='30'/>
  <preference name='com.ibm.cic.common.core.preferences.readTimeout' value='45'/>
  <preference name='com.ibm.cic.common.core.preferences.downloadAutoRetryCount' value='0'/>
  <preference name='offering.service.repositories.areUsed' value='false'/>
  <preference name='com.ibm.cic.common.core.preferences.ssl.nonsecureMode' value='false'/>
  <preference name='com.ibm.cic.common.core.preferences.http.disablePreemptiveAuthentication' value='false'/>
  <preference name='http.ntlm.auth.kind' value='NTLM'/>
  <preference name='http.ntlm.auth.enableIntegrated.win32' value='true'/>
  <preference name='com.ibm.cic.common.core.preferences.preserveDownloadedArtifacts' value='true'/>
  <preference name='com.ibm.cic.common.core.preferences.keepFetchedFiles' value='false'/>
  <preference name='PassportAdvantageIsEnabled' value='false'/>
  <preference name='com.ibm.cic.common.core.preferences.searchForUpdates' value='false'/>
  <preference name='com.ibm.cic.agent.ui.displayInternalVersion' value='false'/>
  <preference name='com.ibm.cic.common.sharedUI.showErrorLog' value='true'/>
  <preference name='com.ibm.cic.common.sharedUI.showWarningLog' value='true'/>
  <preference name='com.ibm.cic.common.sharedUI.showNoteLog' value='true'/>
</agent-input>
```

What to do next

Before you can access the Netcool Configuration Manager Dashboard Application Services Hub components, you must set up the Netcool Configuration Manager Dashboard Application Services Hub users and provide them with appropriate access permission.

Once users have been set up, you access the Netcool Configuration Manager Dashboard Application Services Hub components, that is, the Activity Viewer, the Dashboard Application Services Hub wizards, and the thick-client launch portal in the following ways:

- You launch the stand-alone Netcool Configuration Manager UIs (sometimes referred to as the thick-client UIs), from the Dashboard Application Services Hub thick-client launch portal.
- You access the Activity Viewer, the Dashboard Application Services Hub wizards, and a subset of reports **in context** from Network Manager and Tivoli Netcool/OMNIbus.
- You access the complete reports using the Dashboard Application Services Hub Reporting Services GUI.

Configuring separate database types

Under certain circumstances, such as when different or remote databases are used in an integrated environment, you must perform additional database configuration steps.

About this task

If you are installing Network Manager and ITNCM-Reports together, and if the Network Manager database is Db2 and on a different server, then its component databases must be cataloged.

If Network Manager uses an Informix® database in a distributed environment and Dashboard Application Services Hub is not installed on the same server as Network Manager, you ensure that the correct library jars are used.

Procedure

1. Required: **If Network Manager and ITNCM-Reports are installed together, and if the Network Manager database is Db2 and on a different server:**
   a) To connect to a Db2 database on a server remote from your TCR Installation, ensure that a Db2 client is installed and the remote database cataloged. When the database server is remote to the WebSphere Application Server node where configuration is taking place, enter the following command at the node to add a TCP/IP node entry to the node directory:

   ```
   db2 catalog tcpip node <NODENAME> remote <REMOTE> server <PORT>
   ```

   where

   NODENAME
   Specifies a local alias for the node to be cataloged.

   REMOTE
   Specifies the fully qualified domain name of the remote DB server.

   PORT
   Is the port on which the database is accessible, typically port 50000.

   ```
   db2 catalog database <database_name> at node <NODENAME>
   ```

   where

   database_name
   Specifies the Db2 database name.

   NODENAME
   Is the local alias specified in the previous step.

   b) Add 'source $HOME/sqllib/db2profile' to your *<install_user>*/.bash_profile.

   Where $HOME refers to the home directory of the user which was configured during the installation of the Db2 client to manage the client (usually db2inst1), and <install_user> is the user who installed Netcool Configuration Manager, usually 'icosuser'.

   Note: The .bash_profile is only used for bash shell, and it will be different for sh, csh or ksh.

   c) Restart your reporting server after this update. However, before restarting the Reporting Server, check that the amended login profile has been sourced.

   Tip: For installations which use a Db2 database, Cognos requires 32 bit Db2 client libraries, which will be installed by the 64 bit Db2 client. However, there maybe further dependencies on other 32 bit packages being present on the system; if such errors are reported, you can check this with **'ldd $library_name'**.

2. Required: **If Network Manager and ITNCM-Reports are installed together, and if the Network Manager database is Oracle:**
   a) To connect to an Oracle database from your TCR Installation, ensure that ITNCM-Reports have been installed, and then update the itncmEnv.sh file default location:

   ```
   /opt/IBM/tivoli/netcool/ncm/reports/itncmEnv.sh
   ```

   Export the following variables (where <install directory> is the Netcool Configuration Manager installation directory):

   ORACLE_HOME

   ```
   ORACLE_HOME=<install directory>/reports/oracle
   export ORACLE_HOME
   ```

TNS_ADMIN

```
TNS_ADMIN=<install directory>/reports/oracle/network/admin
export TNS_ADMIN
```

LD_LIBRARY_PATH

```
LD_LIBRARY_PATH=<install directory>/reports/oracle:$LD_LIBRARY_PATH
export LD_LIBRARY_PATH
```

b) Create a tnsnames.ora file located in `<install directory>/reports/oracle/network/admin/`

c) Add the NCIM database to the tnsnames.ora file.

For example: `NCIM = (DESCRIPTION = (ADDRESS_LIST = (ADDRESS = (PROTOCOL = TCP)(Host = <Database Server>)(Port = 1521)) ) (CONNECT_DATA = (SERVICE_NAME = NCIM) ))`

d) Add 'source <install directory>/reports/itncmEnv.sh' to your *<install_user>*`/.bash_profile`.

Note: The `.bash_profile` is only used for bash shell, and it will be different for sh, csh or ksh.

e) Restart your reporting server after this update. However, before restarting the Reporting Server, check that the amended login profile has been sourced.

Configuring integration with Tivoli Netcool/OMNIbus
Ensure that you have Netcool/OMNIbus Knowledge Library (NcKL) Enhanced Probe Rules for Netcool Configuration Manager installed on your Tivoli Netcool/OMNIbus server.

Before you begin

Deploy rules specific to Netcool Configuration Manager. These rules have been bundled with Netcool Configuration Manager and deployed on the Netcool Configuration Manager Presentation server during installation, and are located in the *<NCM-INSTALL-DIR>*`/nckl-rules` directory.

Note: This procedure is no longer required for device synchronization with Network Manager, and the mapping of devices between Netcool Configuration Manager and Network Manager.

The standard Netcool/OMNIbus Knowledge Library configuration must have been applied to the ObjectServer and to the Probe for SNMP as part of the prerequisite tasks for the integration. The $NC_RULES_HOME environment variable must also have been set on the computer where the probe is installed. This environment variable is set to `$NCHOME/etc/rules` on UNIX or Linux.

Tip: To source the Network Manager environment script, run the following script: `./opt/IBM/tivoli/netcool/env.sh` where `opt/IBM/tivoli/netcool` is the default Network Manager directory.

Note: If you have existing Probe for SNMP custom rules that you want to preserve, create backups as required before deploying the Netcool/OMNIbus Knowledge Library rules in step two.

About this task
The location denoted by $NC_RULES_HOME holds a set of Netcool/OMNIbus Knowledge Library lookup files and rules files within a number of sub-directories. In particular, the `$NC_RULES_HOME/include-snmptrap/ibm` subdirectory contains files that can be applied to the Probe for SNMP. To support the integration, you must add customized rules for Netcool Configuration Manager to this subdirectory.

Remember: If you have installed Netcool/OMNIbus Knowledge Library (NcKL) Enhanced Probe Rules Version 4.4 Multiplatform English (NcKL4.4) on your Tivoli Netcool/OMNIbus server, which is the recommended option, you do not need to install the ITNCM-specific Rules files, as documented here.

Procedure

Installing rules files specific to Netcool Configuration Manager (not the recommended option)

1. From the server where you have installed Netcool Configuration Manager, copy the following files:

- *ncm_install_dir*/nckl_rules/nckl_rules.zip

  where *ncm_install_dir* represents the installation location of Netcool Configuration Manager, for example */opt/IBM/tivoli/netcool/ncm*

  Copy these files to a temporary location on the computer where the Probe for SNMP is installed.

2. Extract the contents of the nckl_rules.zip file, and then copy the extracted files to the $NC_RULES_HOME/include-snmptrap/ibm subdirectory.
3. If object server failover has already been configured, proceed to step 4. Otherwise, perform the following steps:
   a) Go to the folder in which the mttrapd.props has been placed, for example $NCHOME/omnibus/probes/*AIX5*, where *AIX5* is specific to your operating system.
   b) Edit the mttrapd.props file by commenting out the backup object server reference: #ServerBackup : ''
4. To ensure that the probe can reference the enhanced lookup and rules files, edit the $NC_RULES_HOME/snmptrap.rules file by uncommenting the following include statements, as shown:

```
include "$NC_RULES_HOME/include-snmptrap/ibm/ibm.master.include.lookup
include "$NC_RULES_HOME/include-snmptrap/ibm/ibm.master.include.rules"
include "$NC_RULES_HOME/include-snmptrap/ibm/ibm-preclass.include.snmptrap.rules"
```

5. Run the probe. If the probe was already running, force the probe to re-read the rules file so that the changes can take effect, for example:

   Locate the PID of the probe by running the following command on the server running the probe. Look for a process named - nco_p_mttrapd

```
ps -eaf | grep mttrapd
kill -9 PID
```

   Note: If the probe is installed on a different computer from Network Manager or the DASH portal, you must restart the probe manually.

Configuring integration with Network Manager
Copy a number of jar files from the Network Manager GUI server into the Netcool Configuration Manager instance of WebSphere.

About this task

Note: The following default locations may differ depending on where WebSphere was installed on your Network Manager and Netcool Configuration Manager servers.

Procedure

Copy the following jars from the Network Manager GUI server into the corresponding folder in the Netcool Configuration Manager WebSphere instance.

- /opt/IBM/WebSphere/AppServer/etc/vmm4ncos.jks
- /opt/IBM/WebSphere/AppServer/lib/ext/com.ibm.tivoli.ncw.ncosvmm.jar
- /opt/IBM/WebSphere/AppServer/lib/ext/jconn3.jar

Configuring device synchronization
You configure device synchronization to enable Netcool Configuration Manager to use Network Manager for network device discovery.

Before you begin

During Netcool Configuration Manager 6.4.2 installation you are asked if the product is to be integrated or not. If you select **Yes** the installer will ask the necessary questions to set up the configuration of device synchronization between Netcool Configuration Manager and Network Manager.

A default value of 24 hours (1440mins) is defined in the Netcool Configuration Manager `rseries.properties` file for the periodic synchronization with Network Manager. For the initial synchronization, a large number of devices may already have been discovered by Network Manager, and it can take a considerable time before they are imported into Netcool Configuration Manager. (This also applies in a situation where the discovery scope is widened so that a significant number of new devices are added to Network Manager.) Consequently the devices may not yet appear in the NMENTITYMAPPING table in the Netcool Configuration Manager database, and therefore the context tools (right-click tools) from Network Manager will not be available for those devices.

Tip: You can reduce this time by editing the `rseries.properties` file, and changing the mapping period to 60 (for example). This will speed up the process by which devices are added to the autodiscovery queue on Netcool Configuration Manager, but will not change the actual time to import each device configuration.

Tip:

If the password for the itnmadmin user has changed on Network Manager, update the locally stored copy on Netcool Configuration Manager as follows:

Use the icosadmin script located in `/opt/IBM/tivoli/netcool/ncm/bin`.

For example:

```
icosadmin ChangeNmPassword -u itnmadmin -p <new_password>
```

About this task

The configuration is stored in the `rseries.properties` file located in the following directory: `<ncm-install-dir>/config/properties/`

Network Manager:

```
NMEntityMappingComponent/baseURL=https://nmguiservername:16311
NMEntityMappingComponent/uri=/ibm/console/nm_rest/topology/devices/domain/NCOMS
NMEntityMappingComponent/uriParam=
NMEntityMappingComponent/uriProps=
#####Note: Complete URL = baseURL+uri+uriProps&uriParam

NMEntityMappingComponent/delay=10 ## delay on startup before first run
NMEntityMappingComponent/importRealm=ITNCM/@DOMAINNAME
NMEntityMappingComponent/maxDevicesPerRealm=50
NMEntityMappingComponent/ncmUser=administrator
NMEntityMappingComponent/period=1440 ## Daily (in minutes)
NMEntityMappingComponent/user=itnmadmin
NMEntityMappingComponent/passwd=netcool ## Optional: Install stores securely
```

Note: You can edit this file and the component configuration properties after install if requirements change.

Before device synchronization runs for the first time ensure that the Network Manager Rest API user (in our example 'itnmadmin') has the ncp_rest_api role in DASH.

Device synchronization is now done by a new core component of Netcool Configuration Manager, and is therefore part of Netcool Configuration Manager Component configuration and started automatically when Netcool Configuration Manager starts. Component start up is configured in `<ncm-install-dir>/config/server/config.xml`

```
<component>
<name>NMEntityMappingComponent</name>
<class>com.intelliden.nmentitymapping.NMEntityMappingComponent</class>
</component>
```

Note: The NMEntityMappingComponent is configured by default so if you wish to stop it being started on Netcool Configuration Manager startup you can comment it out in the config.xml file.

Note: There is a limit of 50 imported devices per Realm in Netcool Configuration Manager. If there are more devices than this in a Network Manager domain, they will be added to sub-realms (labeled 001, 002, etc) in Netcool Configuration Manager.

Example

Troubleshooting NM Component

Verify that the component has started in file:

```
<NCM_INSTALL_DIR>/logs/Server.out
Fri Jul 31 13:30:06 GMT+00:00 2015 - Starting component : NMEntityMappingComponent
Fri Jul 31 13:30:06 GMT+00:00 2015 - All components started
```

Verify that the `config.xml` file has the component specified for startup

Verify that the NMEntityMapping table has the new columns required for the new component implementation:

```
"NMENTITYMAPPING" (
"UNIQUEKEY" BIGINT NOT NULL,
"ENTITYID" BIGINT NOT NULL DEFAULT 0,
"RESOURCEBROWSERID" BIGINT NOT NULL DEFAULT 0,
"DOMAINNAME" VARCHAR(64),
"JPAVERSION" BIGINT NOT NULL DEFAULT 1,
"ENTITYNAME" VARCHAR(255),
"ACCESSIPADDRESS" VARCHAR(64),
"SERIALNUMBER" VARCHAR(64),
"VENDORTYPE" VARCHAR(64),
"MODELNAME" VARCHAR(64),
"OSVERSION" VARCHAR(64),
"OSIMAGE" VARCHAR(255),
"OSTYPE" VARCHAR(64),
"HARDWAREVERSION" VARCHAR(64)
)
```

Ensure that the Network Manager Rest API user has the ncp_rest_api role in DASH.

Configuring the Alerts menu of the Active Event List

You must add access to the Activity Viewer from the Active Event List by configuring the Alerts menu.

Procedure

1. From the navigation pane, click **Administration** > **Event Management Tools** > **Menu Configuration**.
2. From the **Available menus** list on the right, select **alerts** and click **Modify**.
3. From the **Menus Editor** window, select **<separator>** from the drop-down list under **Available items**, and then click **Add selected item** to add the item to the **Current items** list.

   The **<separator>** item is added as the last item.
4. Under **Available items**, select **menu** from the drop-down list.

   The list of all menu items that can be added to the **Alerts** menu is shown.
5. Select the **Configuration Management** item and click **Add selected item**.

   The item is added below the **<separator>** item in the **Current items** list.
6. Click **Save** and then click **OK**.

Results

The **Configuration Management** submenu and tools are now available in the **Alerts** menu of the Active Event List, for use with Netcool Configuration Manager events.

Note: Reports Menu options will not be displayed if the selected event is not enriched.

What to do next

You can optionally create a global filter to restrict the events displayed in the Active Event List to Netcool Configuration Manager events only. You can add this filter to the Web GUI either by using the WAAPI client or by using the Filter Builder. When creating the filter, specify a meaningful name (for example, ITNCMEvents) and define the filter condition by specifying the following SQL WHERE clause:

```
where Class = 87724
```

Migrating reports
If you have custom Reporting Services reports in an existing Netcool Configuration Manager installation, and are integrating with Network Manager, which has its own Reporting Services solution, you migrate your custom reports from the stand-alone to the integrated version of Reporting Services.

Before you begin

If you are installing Network Manager on the same server as your existing Netcool Configuration Manager installation, you must export your custom reports before installing Network Manager.

About this task

The report migration procedure is different for single and multiple server integrations.

If you are installing Netcool Configuration Manager and Network Manager on the same server

1. Export the custom reports from the existing Netcool Configuration Manager version of Reporting Services and copy them to a safe location.

   Note: You export your custom reports before installing Network Manager to prevent the existing reports from being overwritten.
2. Disable and uninstall the existing Netcool Configuration Manager version of Reporting Services.
3. Install Network Manager and integrate it with the existing version of Netcool Configuration Manager as documented.
4. Import the custom reports into the Network Manager version of Reporting Services.

If you are installing Netcool Configuration Manager and Network Manager on different servers

1. Install Network Manager and integrate it with the existing version of Netcool Configuration Manager as documented.
2. Export the custom reports from the existing Netcool Configuration Manager version of Reporting Services and copy them to the Network Manager server.
3. Import the custom reports into the Network Manager version of Reporting Services.
4. Disable the existing Netcool Configuration Manager version of Reporting Services.

Exporting custom reports (distributed integration architecture)
After you have installed Network Manager on a server other than your existing Netcool Configuration Manager installation and performed all integration tasks, you export your custom Reporting Services reports. You also disable and uninstall the existing Netcool Configuration Manager version of Reporting Services.

Before you begin

You export reports **after** installing Network Manager when all of the following circumstances apply to your scenario:

- You are already running Reporting Services as part of an existing, non-integrated Netcool Configuration Manager installation.
- You are deploying a distributed integration architecture and have already installed Network Manager on a server other than your existing version of Netcool Configuration Manager.
- You have customized Netcool Configuration Manager reports that need to be migrated into your planned integrated solution.

About this task

When you install the Network Manager version of Reporting Services on a server other than your existing version of Netcool Configuration Manager, the previous reports as well as the previous version of Reporting Services remain on the Netcool Configuration Manager server. To migrate such reports into an integrated solution, you perform the following tasks:

If you are installing Netcool Configuration Manager and Network Manager on different servers

1. Install Network Manager and integrate it with the existing version of Netcool Configuration Manager as documented.
2. Export the custom reports from the existing Netcool Configuration Manager version of Reporting Services and copy them to the Network Manager server.
3. Import the custom reports into the Network Manager version of Reporting Services.
4. Disable the existing Netcool Configuration Manager version of Reporting Services.

Remember: You do not have to migrate the standard Netcool Configuration Manager reports, because these will be installed together with the Network Manager version of Reporting Services (in addition to a number of Network View reports). You only migrate reports you have customized since installing the standard reports, or new reports you have created.

Procedure

1. Log into the Netcool Configuration Manager version of Reporting Services using the following URL:
   `http://`*`hostname`*`:16310/ibm/console`
   where *`hostname`* is the name of your Netcool Configuration Manager server and *`16310`* is the default port number for Reporting Services.
2. Click **Reporting** > **Common Reporting**.
3. Click **Launch** on the toolbar, and then select **Administration** from the drop-down menu.
4. Select the **Configuration** tab, then click **Content Administration**.
5. Click **New Export** to launch the **New Export** wizard.
6. Enter a name and description for the report export, then click **Next**.
7. Accept the default deployment method and click **Next**.
8. Click the **Add** link and select the **ITNCM Reports** checkbox, then move ITNCM Reports to the **Selected Entries** list.
9. Click **OK**, then **Next** > **Next** > **Next**, accepting the default values.
10. Select **New archive**, then **Next** > **Next**, accepting the default values..
11. Click **Finish** > **Run** > **OK**.

    The reports are exported and the new export archive is displayed.
12. Navigate to the following directory:
    `/opt/IBM/tivoli/netcool/ncm/tipv2Components/TCRComponent/cognos/deployment`,
    where you can view the report archive, for example:
    `-rw-r--r-- 1 icosuser staff 262637 23 Feb 10:27 ncm_export.zip`
    where `ncm_export.zip` is the report archive.
13. Copy the file to the following directory on the Network Manager server:
    `$TIP_HOME/../TCRComponent/cognos/deployment`

Results
You have exported the custom reports and copied them to the Network Manager server.

What to do next
Next, you import the archived reports into the Network Manager version of Reporting Services, and then disable the Netcool Configuration Manager version of Reporting Services.

Importing reports (distributed integration architecture)
After exporting the custom reports and copying them to the Network Manager server, you import the archived reports into the Network Manager version of Reporting Services, and then disable the Netcool Configuration Manager version of Reporting Services.

Before you begin
You must have exported the custom reports and copied them to the Network Manager server.

About this task

Procedure

1. Log into the Network Manager Dashboard Application Services Hub.
2. Click **Reporting** > **Common Reporting**.
3. Click **Launch** on the toolbar, and then select **Administration** from the drop-down menu.
4. Select the **Configuration** tab, then click **Content Administration**.
5. Click **New Import** to launch the **New Import** wizard.
   A list of available report archives will be displayed.
6. Select the archive that you exported earlier and click **Next**.
7. Select **ITNCM Reports**, then **Next** and **Next** again, accepting the default values.
8. Click **Finish** > **Run** > **OK**.
   The reports are imported and the new archive is displayed in the list of archives.
9. Close the Common Reporting tab and click the Common Reporting link in the navigation pane.
   The custom reports will now be available in the Netcool Configuration Manager version of Reporting Services.
10. Navigate to the following directory:
    `/opt/IBM/tivoli/netcool/ncm/bin/utils/support`
11. Run the `setPlatform.sh` script:`bash-3.2$ ./setPlatform.sh` and disable Reporting, then exit.
    When the Netcool Configuration Manager server is restarted, Reporting Services will no longer be running.

Results
You have now completed the migration of custom reports in your distributed custom environment.

Importing reports (single server integration architecture)
After exporting the custom reports, disabling and uninstalling the Netcool Configuration Manager version of Reporting Services, and completing all other integration steps, you import the report archive into the Network Manager version of Reporting Services,

Before you begin

You must have exported the custom reports before installing Network Manager on the same server as your existing Netcool Configuration Manager installation.

About this task

Procedure

1. Log into the Network Manager Dashboard Application Services Hub.
2. Click **Reporting** > **Common Reporting**.
3. Click **Launch** on the toolbar, and then select **Administration** from the drop-down menu.
4. Select the **Configuration** tab, then click **Content Administration**.
5. Click **New Import** to launch the **New Import** wizard.
   A list of available report archives will be displayed.
6. Select the archive that you exported earlier and click **Next**.
7. Select **ITNCM Reports**, then **Next** and **Next** again, accepting the default values.
8. Click **Finish** > **Run** > **OK**.
   The reports are imported and the new archive is displayed in the list of archives.
9. Close the **Common Reporting** tab and click the **Common Reporting** link in the navigation pane.

Results
The custom reports will now be available in the Netcool Configuration Manager version of Reporting Services.

Configuring Single Sign-On for Netcool Configuration Manager
The single sign-on (SSO) capability in Tivoli® products means that you can log on to one Tivoli application and then launch to other Tivoli web-based or web-enabled applications without having to re-enter your user credentials.

The repository for the user IDs is the Tivoli Netcool/OMNIbus ObjectServer. A user logs on to one of the participating applications, at which time their credentials are authenticated at a central repository. With the credentials authenticated to a central location, the user can then launch from one application to another to view related data or perform actions. Single sign-on can be achieved between applications deployed to DASH servers on multiple machines.

Single sign-on capabilities require that the participating products use Lightweight Third Party Authentication (LTPA) as the authentication mechanism. When SSO is enabled, a cookie is created containing the LTPA token and inserted into the HTTP response. When the user accesses other Web resources in any other application server process in the same Domain Name Service (DNS) domain, the cookie is sent with the request. The LTPA token is then extracted from the cookie and validated. If the request is between different cells of application servers, you must share the LTPA keys and the user registry between the cells for SSO to work. The realm names on each system in the SSO domain are case sensitive and must match exactly. See Managing LTPA keys from multiple WebSphere® Application Server cells on the WebSphere Application Server Information Center.

When configuring ITNCM-Reports for an integrated installation, ensure you configure single sign-on (SSO) on the Reporting Services server. Specifically, you must configure SSO between the instance of WebSphere that is hosting the Network Manager GUI, and the instance of WebSphere that is hosting ITNCM Reports. This will prevent unwanted login prompts when launching reports from within Network Manager. For more information, see the related topic links.

Creating user groups for DASH
To configure single sign-on (SSO) between DASH and Netcool Configuration Manager, you must create Netcool Configuration Manager groups and roles for DASH.

Before you begin

Note: For SSO between DASH and Netcool Configuration Manager to work, the user groups specified in this procedure must exist in both DASH and Netcool Configuration Manager.

Network Manager and Netcool Configuration Manager users in DASH should use the same authentication type, for example, ObjectServer.

Note: The **IntellidenUser** role needs to be assigned to the **IntellidenUser** group. Similarly, the **IntellidenAdminUser** role needs to be given to the **IntellidenAdminUser** group.

About this task

Procedure

1. Log onto the WebSphere Administrative console of the Network Manager GUI server as the profile owner (for example, `smadmin`).
2. Create a group by selecting **Users and Groups** > **Manage Groups** > **Create**.
3. Enter `IntellidenUser` in the Group name field.
4. Click **Create**, then click **Create Like**.
5. Enter `IntellidenAdminUser` in the Group name field.

   IntellidenAdminUser is required for access to Account Management in Netcool Configuration Manager.
6. Click **Create**, then click **Close**.

7. Log off from the WebSphere Administrative console, then log on to the DASH GUI.
8. Select **Console Settings** > **Roles** > **IntellidenUser**.
9. Click **Users and Groups** > **Add Groups** > **Search**, then select the **IntellidenUser** group, and then click **Add**.
10. Select **Console Settings** > **Roles** > **IntellidenAdminUser**.
11. Click **Users and Groups** > **Add Groups** > **Search**, then select the **IntellidenAdminUser** group, and then click **Add**.

What to do next
After creating Netcool Configuration Manager groups and roles for DASH, you create Netcool Configuration Manager users for DASH.

Creating users for DASH
This section explains how to create the Netcool Configuration Manager **Intelliden** super-user as well as the default users: **administrator**, **operator**, and **observer** for DASH.

Before you begin

For single sign-on (SSO) between DASH and Netcool Configuration Manager to work, a user must exist (that is, have an account) in both DASH and Netcool Configuration Manager.

At install time Netcool Configuration Manager automatically creates four users: **Intelliden**, **administrator**, **operator**, and **observer**. Of these users, only the **Intelliden** user must be created in DASH. However, it is advisable that the other users are also created.

Note: Only the username must match, it is not necessary that the passwords also match. After single-sign on configuration is complete, the user password entered in DASH will be used to authenticate a Netcool Configuration Manager login.

About this task
This task describes how to create the previously listed Netcool Configuration Manager users for DASH.

Procedure

1. Log onto the WebSphere console of the Network Manager GUI server as the profile owner (foe example, `smadmin`).
2. Click **Users and Groups** > **Manage Users**, then click **Create**.
3. Enter `Intelliden` in the User ID, First name, and Last Name fields.
4. Enter the Intelliden user's password in the Password and Confirm Password fields.
5. Click on **Group Membership** and select **Search**.
6. Highlight the **IntellidenAdminUser** and **IntellidenUser** groups in the matching groups list, and click **Add**, then click **Close**.
7. Click **Create**, then click **Create Like**.
8. Enter `administrator` in the following fields:
   - User ID
   - First name
   - Last Name
   - Password
   - Confirm password
9. Click on **Group Membership** and select **Search**.
10. Highlight the **IntellidenAdminUser** and **IntellidenUser** groups in the matching groups list, and click **Add**, then click **Close**.
11. Click **Create** and then **Close**.

12. Click **Create**.
13. Enter `operator` in the following fields:
    - User ID
    - First name
    - Last Name
    - Password
    - Confirm password
14. Click on **Group Membership** and select **Search**.
15. Highlight the **`IntellidenUser`** group in the matching groups list, and click **Add** and then **Close**.
16. Click **Create**, then click **Create Like**.
17. Enter `observer` in the following fields:
    - User ID
    - First name
    - Last Name
    - Password
    - Confirm password
18. Click on **Group Membership** and select **Search**.
19. Highlight the **`IntellidenUser`** group in the matching groups list, and click **Add**, then click **Close**.
20. Click **Create** and then **Close**.

What to do next
After you have created the Netcool Configuration Manager users for DASH, you export the LTPA keystore to the Netcool Configuration Manager server.

Exporting the DASH LTPA keystore
For added security the contents of the LTPA token are encrypted and decrypted using a keystore (referred to in the subsequent procedure as the LTPA keystore) maintained by WebSphere. In order for two instances of WebSphere to share authentication information via LTPA tokens they must both use the same LTPA keystore. The IBM Admin Console makes this a simple process of exporting the LTPA keystore on one instance of WebSphere and importing it into another.

About this task
This task describes how to export the LTPA keystore from the instance of WebSphere running on the Network Manager DASH server to the instance of WebSphere running on the Netcool Configuration Manager server for keystore synchronization.

Procedure

1. Launch the DASH Admin Console. For example: `http://www.nm_gui_server_ip.com:16310/ibm/console`.
2. Navigate to **Settings** > **WebSphere Administrative Console**.
3. Click **Security** > **Global security**.
4. Under the **Authentication mechanisms and expiration** tab, click **LTPA**.
5. Under the **Cross-cell single sign-on** tab, enter a password in the Password and Confirm password fields.

   The password will subsequently be used to import the LTPA keystore on the Netcool Configuration Manager server.
6. Enter the directory and filename you want the LTPA keystore to be exported to in the **Fully qualified key file name** field.
7. Complete by clicking **Export keys**.

8. Transfer the LTPA keystore to the Netcool Configuration Manager server.

Results
You will receive a message indicating that the LTPA keystore has been exported successfully.

What to do next
You now configure the SSO attributes for DASH.
Related tasks
Importing the DASH LTPA keystore to the Netcool Configuration Manager server
For added security the contents of the LTPA token are encrypted and decrypted using a keystore maintained by WebSphere. In order for two instances of WebSphere to share authentication information via LTPA tokens they must both use the same keystore. The IBM admin console makes this a simple process of exporting the keystore on one instance of WebSphere and importing it into another.

Configuring Single Sign-On for Netcool Configuration Manager
Configuring SSO is a prerequisite to integrating products that are deployed on multiple servers. All DASH server instances must point to the central user registry.

About this task
Use these instructions to configure single sign-on attributes for the DASH.

Procedure

1. Launch the DASH Admin Console. For example: `http://www.nm_gui_server_ip.com:16310/ibm/console` .
2. Navigate to **Settings** > **WebSphere Administrative Console**.
3. Select **Security**, then click **Global Security** > **Web and SIP Security** > **Single sign on (SSO)**.
4. In the **Authentication** area, expand Web security, then click **Global Security** > **Web and SIP Security (on the Authentication area)** > **Single sign on (SSO)**.
5. Select the **Enabled** option if SSO is disabled.
6. Deselect **Requires SSL**.
7. Enter the fully-qualified domain names in the Domain name field where SSO is effective. If the domain name is not fully qualified, the DASH server does not set a domain name value for the LTPAToken cookie and SSO is valid only for the server that created the cookie. For SSO to work across Tivoli® applications, their application servers must be installed in the same domain (use the same domain name). See below for an example.
8. Deselect the **Interoperability Mode** option.
9. Deselect the **Web inbound security attribute propagation** option.
10. Click **OK**, then save your changes.
11. Stop and restart all the DASH server instances. Log out of the WebSphere Administrative Console.

Example

If DASH is installed on **`server1.ibm.com`** and Netcool Configuration Manager is installed on **`server2.ibm.com`**, then enter a value of **`.ibm.com`**.

What to do next
You enable SSO on Netcool Configuration Manager next.
Related tasks
Importing the DASH LTPA keystore to the Netcool Configuration Manager server
For added security the contents of the LTPA token are encrypted and decrypted using a keystore maintained by WebSphere. In order for two instances of WebSphere to share authentication information via LTPA tokens they must both use the same keystore. The IBM admin console makes this a simple process of exporting the keystore on one instance of WebSphere and importing it into another.

Enabling SSO for Netcool Configuration Manager
Both Netcool Configuration Manager and Netcool Configuration Manager WebSphere must be configured to enable SSO.

About this task
This task describes how to enable SSO for Netcool Configuration Manager if it was not enabled during installation.

Procedure

1. Navigate to $*NCM_installation_dir*/utils .
2. Run the configSSO.sh script, for example:

```
cd /opt/IBM/tivoli/netcool/ncm/bin/utils ./configSSO.sh enable
```

What to do next
When SSO is enabled, the interface to Netcool Configuration Manager must accept an LTPA token as a means of authentication. This is achieved by importing the LTPA keystore to the Netcool Configuration Manager server.

Importing the DASH LTPA keystore to the Netcool Configuration Manager server
For added security the contents of the LTPA token are encrypted and decrypted using a keystore maintained by WebSphere. In order for two instances of WebSphere to share authentication information via LTPA tokens they must both use the same keystore. The IBM admin console makes this a simple process of exporting the keystore on one instance of WebSphere and importing it into another.

Before you begin
You must have exported the LTPA keystore from the instance of WebSphere running on the Network Manager DASH server and copied it to the Netcool Configuration Manager server in a previous task.

About this task
In this procedure you will import that LTPA keystore to the instance of WebSphere running on the Netcool Configuration Manager server.

Procedure

1. Logon to the WebSphere Administrative Console for the Netcool Configuration Manager Presentation Server using the superuser name and password specified at install time (typically Intelliden).

   For example: http://NCM_presentation_server:16316/ibm/console
2. Click **Security** > **Global security**.
3. Under **Authentication mechanisms and expiration**, click **LTPA**.
4. Under **Cross-cell single sign-on**, enter the password in the Password and Confirm password fields. This password is the one that was used when the LTPA keystore was exported from DASH.
5. Enter the LTPA keystore file name in the Fully qualified key file name field. This is the LTPA keystore that was exported from DASH.
6. Click **Import keys**.
7. Click **Save directly to the master configuration**.

What to do next
You should now configure single sign-on attributes for the WebSphere instance running on the Netcool Configuration Manager server.
Related tasks
Exporting the DASH LTPA keystore

For added security the contents of the LTPA token are encrypted and decrypted using a keystore (referred to in the subsequent procedure as the LTPA keystore) maintained by WebSphere. In order for two instances of WebSphere to share authentication information via LTPA tokens they must both use the same LTPA keystore. The IBM Admin Console makes this a simple process of exporting the LTPA keystore on one instance of WebSphere and importing it into another.

Configuring single sign-on attributes for DASH
Configuring SSO is a prerequisite to integrating products that are deployed on multiple servers. All DASH server instances must point to the central user registry.

Configuring single sign-on attributes for Netcool Configuration Manager WebSphere
Configuring SSO is a prerequisite to integrating products that are deployed on multiple servers. All eWAS server instances must point to the central user registry.

About this task
This procedure is performed on the Netcool Configuration Manager eWAS instance running on the Netcool Configuration Manager server.

Procedure

1. Logon to the WebSphere Administrative Console for the Netcool Configuration Manager Presentation Server using the superuser name and password specified at install time (typically `Intelliden`).

   For example http://NCM_presentation_server:16316/ibm/console
2. In the **Authentication** area, expand Web security then click Single sign-on.
3. Select the **Enabled** option if SSO is disabled.
4. Deselect **Requires SSL**.
5. Leave the domain name blank in the **Domain name** field.
6. Deselect the **Interoperability Mode** option.
7. Deselect the **Web inbound security attribute propagation** option.
8. Click **Apply** to save your changes.
9. Click **Save Directly to the Master Configuration**.

What to do next
You create a federated user repository for Netcool Configuration Manager next.

Creating and configuring a federated user repository for Netcool Configuration Manager
The first step for authenticating by using a Tivoli Netcool/OMNIbus ObjectServer is to create a federated user repository for Netcool Configuration Manager.

Before you begin

Important: Before attempting this procedure, complete the following task: “Configuring integration with Network Manager” on page 77

About this task

A federated user repository is built on Virtual Member Manager (VMM), which provides the ability to map entries from multiple individual user repositories into a single virtual repository. The federated user repository consists of a single named realm, which is a set of independent user repositories. Each user repository may be an entire external user repository.

This task describes how to create and configure a federated user repository for Netcool Configuration Manager.

Procedure

1. Launch the WebSphere Administrative Console from `http://<ncmserver-hostname-ip>:<16316>/ibm/console` and login using the Netcool Configuration Manager superuser name and password specified during installation.

   Note: The port number may be different for a non-standard installation.
2. Select **Security** > **Global security**.
3. Under the User account repository, select **Federated repositories** from the Available realm definitions field, and click **Configure**.
4. Under Repositories in the realm, select **Add repositories (LDAP, custom, etc)**.
5. Under General Properties, select **New Repository** > **Custom Repository**
6. Update the ObjectServer VMM properties as described here (or per your custom repository):

   Repository identifier
   NetcoolObjectServer

   Repository adapter class name
   com.ibm.tivoli.tip.vmm4ncos.ObjectServerAdapter

   Custom Properties
   Add the following four properties:

   Note: Find the exact details from the repository viewable on the Network Manager Gui Administrative Console.

   Table 23. Custom Properties

Name (case-sensitive	Value
username	ObjectServer administrator user name
password	ObjectServer encrypted administrator user password
port1	Object Server port number
host1	Object Server hostname/IP address

7. Click **Apply** and save your changes directly to the master configuration.
8. Under General properties of Repository Reference, update the Unique distinguished name to `o=netcoolObjectServerRepository`
9. Click **OK** and save your changes directly to the master configuration, then click **OK** again.
10. The local repository may not contain IDs that are also in Netcool Configuration Manager. To mitigate, perform one of the following steps:
    - Remove the local file repository from the federation of repositories.
    - Remove all the conflicting users from the local file repository.
11. If prompted, enter the WebSphere Administrator user password in the **Password** and **Confirm Password** fields, and click **OK**.
12. In Global security under the User account repository, select **Federated Repositories** from the Available realm definitions field, and click **Set as current**.
13. Click **Apply** and save your changes directly to the master configuration.
14. Log out of the Administrative Console.
15. Stop the Netcool Configuration Manager server using the `./itncm.sh stop` command. Then start the Netcool Configuration Manager server using the `./itncm.sh start` command.

What to do next

Netcool Configuration Manager will now authenticate with the ObjectServer VMM.

The Netcool Configuration Manager Superuser has been reverted to the user created during the Dash profile Installation (which is `smadmin` by default)

Installing the Network Manager Insight Pack

This topic explains how to install the Network Manager Insight Pack into the Operations Analytics - Log Analysis product and make the necessary configurations. The Network Manager Insight Pack is required only if you deploy the Networks for Operations Insight feature and want to use the topology search capability. For more information, see "Network Manager Insight Pack" on page 360. Operations Analytics - Log Analysis can be running while you install the Insight Pack.

Before you begin

You already completed some of these prerequisites when you installed the Tivoli Netcool/OMNIbus Insight Pack. See "Installing the Tivoli Netcool/OMNIbus Insight Pack" on page 61 for more details.

- Install the Operations Analytics - Log Analysis product. For upgrades, migrate the data from previous instances of the product.
- Ensure that the Tivoli Netcool/OMNIbus Insight Pack is installed before a data source is created. For more information, see "Netcool/OMNIbus Insight Pack" on page 218.
- Download the Network Manager Insight Pack from IBM Passport Advantage. The Insight Pack image is contained within the Operations Analytics - Log Analysis download, see information about *Event Search integration* and *Topology Search integration* in http://www-01.ibm.com/support/docview.wss?uid=swg24043698. The file name of the Insight Pack is `NetworkManagerInsightPack_V1.3.0.0.zip`.
- Install Python 2.6 or later with the `simplejson` library, which is required by the custom apps that are included in the Insight Pack.
- Over large network topologies, the topology search can be performance intensive. It is therefore important to determine which parts of your network you want to use the topology search on. You can define those parts of the network into a single domain. Alternatively, implement the cross-domain discovery function in Network Manager IP Edition to create a single aggregation domain of the domains that you want to search. You can restrict the scope of the topology search to that domain or aggregation domain. To do so, set the **`ncp.dla.ncim.domain`** property to the name of the domain. If you still anticipate a detrimental impact on performance, you can also set the **`ncp.spf.multipath.maxLinks`** property. This property sets a threshold on the number of links that are processed when the paths between the two end points are retrieved. If the threshold number is breached, only the first identified route between the two end points is retrieved. Make these settings in step "Installing the Network Manager Insight Pack" on page 90 of this task. For more information about deploying Network Manager IP Edition to monitor networks of small, medium, and larger networks, see https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/overview/concept/ovr_deploymentseg.html. For more information about the cross-domain discovery function, see https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/disco/task/dsc_configuringcrossdomaindiscoveries.html.
- Obtain the details of the NCIM database that is used to store the network topology for the Network Manager IP Edition product.

Procedure

1. Copy the `NetworkManagerInsightPack_V1.3.0.0.zip` installation package to `$UNITY_HOME/unity_content`.

   Tip: For better housekeeping, create a new `$UNITY_HOME/unity_content/NetworkManager` directory and copy the installation package there.
2. Use the `$UNITY_HOME/utilities/pkg_mgmt.sh` command to install the Insight Pack.
   For example, to install into `$UNITY_HOME/unity_content/NetworkManager`, run the command as follows:

   ```
   $UNITY_HOME/utilities/pkg_mgmt.sh -install $UNITY_HOME/unity_content/
   /NetworkManagerNetworkManagerInsightPack_V1.3.0.0.zip
   ```

3. In $UNITY_HOME/AppFramework/Apps/NetworkManagerInsightPack_V1.3.0.0/Network_Topology_Search/NM_EndToEndSearch.properties, specify the details of the NCIM database.

   Tip: You can obtain most of the information that is required from the $NCHOME/etc/precision/DbLogins.cfg or DbLogins.*DOMAIN*.cfg files (where *DOMAIN* is the name of the domain).

   ncp.dla.ncim.domain
   Limits the scope of the topology search capability to a single domain in your topology. For multiple domains, implement the cross-domain discovery function in Network Manager IP Edition and specify the name of the aggregation domain. For all domains in the topology, comment out this property. Do not leave it blank.

   ncp.spf.multipath.maxLinks
   Sets a limit on the number of links that are processed when the paths between the two end points are retrieved. If the number of links exceeds the limit, only the first identified path is returned. For example, you specify ncp.spf.multipath.maxLinks = 1000. If 999 links are processed, all paths between the two end points are retrieved. If 1001 links are processed, one path is calculated and then processing stops.

   ncp.dla.datasource.type
   The type of database used to store the Network Manager IP Edition topology. Possible values are db2 or oracle.

   ncp.dla.datasource.driver
   The database driver. For Db2, type com.ibm.db2.jcc.Db2Driver. For Oracle, type oracle.jdbc.driver.OracleDriver.

   ncp.dla.datasource.url
   The database URL. For Db2, the URL is as follows:

   ```
   jdbc:db2://host:port/name
   ```

   For Oracle the URL is as follows:

   ```
   jdbc:oracle:thin:@host:port:name
   ```

   In each case, *host* is the database host name, *port* is the port number, and *name* is the database name, for example, NCIM.

   ncp.dla.datasource.schema
   Type the NCIM database schema name. The default is ncim.

   ncp.dla.datasource.ncpgui.schema
   Type the NCPGUI database schema name. The default is ncpgui.

   ncp.dla.datasource.username
   Type the database user name.

   ncp.dla.datasource.password.
   Type the database password.

   ncp.dla.datasource.encrypted
   If the password is encrypted, type true. If not, type false.

   ncp.dla.datasource.keyfile
   Type the name of and path to the cryptographic key file, for example $UNITY_HOME/wlp/usr/servers/Unity/keystore/unity.ks.

   ncp.dla.datasource.loginTimeout
   Change the number of seconds until the login times out, if required.

   Optionally change the logging information, which is specified by the **java.util.logging.*** properties.

Results

- The `NetworkManagerInsightPack_V1.3.0.0` Insight Pack is installed into the directory that you selected in step 2.
- The Rule Set, Source Type, and Collection are in place. You can view these resources in the **Administrative Settings** page of Operations Analytics - Log Analysis.

What to do next

- Use the **pkg_mgmt** command to verify the installation of the Insight Pack. See Verifying the Network Manager Insight Pack.
- If you are using an Oracle database, make the extra configurations that are required to support Oracle. See "Configuring topology search apps for use with Oracle databases" on page 182. Configure the products to support the topology search capability. See Chapter 11, "Enabling topology search," on page 359.

Related concepts

Data Source creation in Operations Analytics - Log Analysis V1.3.5

Data Source creation in Operations Analytics - Log Analysis V1.3.3

Related tasks

Installing the Tivoli Netcool/OMNIbus Insight Pack

This topic explains how to install the Netcool/OMNIbus Insight Pack into the Operations Analytics - Log Analysis product. Operations Analytics - Log Analysis can be running while you install the Insight Pack. This Insight Pack ingests event data into Operations Analytics - Log Analysis and installs custom apps.

Installing the Network Manager Insight Pack

This topic explains how to install the Network Manager Insight Pack into the Operations Analytics - Log Analysis product and make the necessary configurations. The Network Manager Insight Pack is required only if you deploy the Networks for Operations Insight feature and want to use the topology search capability. For more information, see "Network Manager Insight Pack" on page 360. Operations Analytics - Log Analysis can be running while you install the Insight Pack.

Configuring topology search

Before you can use the topology search capability, configure the Tivoli Netcool/OMNIbus core and Web GUI components, the Gateway for Message Bus and Network Manager IP Edition.

Related information

Gateway for Message Bus documentation

Installing Performance Management

To add the Performance Management solution extension, install Network Performance Insight and then integrate it with Netcool Operations Insight.

Related tasks

Installing the Device Dashboard

Follow these instructions to install the **Device Dashboard**.

Installing Network Performance Insight

Install Network Performance Insight by performing the steps in the Installation section of the Network Performance Insight documentation.

About this task

In particular, you must ensure that you perform the following tasks as part of the installation of Network Performance Insight:

- Enable Network Performance Insight integration with Jazz for Service Management by running the `npiDashIntegration` script.
- Add the relevant singer certificate to your browser to enable single sign-on.
- Configure the Netcool/OMNIbus Standard Input probe to enable performance anomaly events to be displayed in the Netcool/OMNIbus Event Viewer, and verify that the probe starts automatically once

installation of Network Performance Insight is complete and the Network Performance Insight Event Service has started.

Navigate to the Network Performance Insight documentation at https://www.ibm.com/support/knowledgecenter/SSCVHB_1.3.1/npi_kc_welcome.html for more information.

Related reference

Ports used by products and components
Use this information to understand which ports are used by the different products and components that make up the Netcool Operations Insight solution.

Enabling the integration with Network Performance Insight

Enable the integration by creating a `kafka.properties` file and populating it with relevant properties.

Before you begin

The Network Performance Insight Kafka server must be available and running in order to be able to enable the integration.

Procedure

1. Copy the `kafka.properties` file from its default location `$NCHOME/precision/storm/conf/default/` to the following location:

   ```
   $NCHOME/precision/storm/conf/
   ```

2. Edit the `kafka.properties` file as follows:
   a) Set the kafka producer properties under the `kafka.producer` property according to the information at the following link: http://kafka.apache.org/documentation.html#producerconfigs.
   b) Set the kafka consumer properties under the `kafka.consumer` property according to the information at the following link: http://kafka.apache.org/documentation.html#newconsumerconfigs.

   Note: The only mandatory properties are the following:

   - `kafka.consumer.bootstrap.servers`
   - `kafka.producer.bootstrap.servers`

3. (Optional) If you anticipate the need to enable and disable the integration often then you can add the `kafka.enabled` property to facilitate this. To do this, add the `kafka.enabled` property to one of the following properties files, and set this property to the value `true`.

   - `$NCHOME/precision/storm/conf/NMStormTopology.properties`
   - `$NCHOME/precision/storm/conf/kafka.properties`

   If the property is not present in either file, then this means that `kafka.enabled=true`.

4. Restart Apache Storm, by running the following commands:

   ```
   itnm_stop storm
   ```

   ```
   itnm_start storm
   ```

Results

To test the output of the integration use the `ncp_storm_validate.sh` script.

Related information

Kafka producer properties

Kafka consumer properties

Network Manager V4.2 documentation: Starting and stopping ncp_stormClick here to access information on starting and stopping Network Managerprocesses, including the ncp_storm process, in the Network Manager documentation.

Network Manager V4.2 documentation: ncp_storm_validate.shClick here to access information on the ncp_storm_validate.sh script in the Network Manager documentation.

Configuring Network Performance Insight
Integrate Network Performance Insight with Netcool Operations Insight by performing the steps in the Configuring section of the Network Performance Insight documentation.

About this task

Navigate to the Network Performance Insight documentation at https://www.ibm.com/support/knowledgecenter/SSCVHB_1.3.1/npi_kc_welcome.html for more information.

Installing the Device Dashboard
Install the Device Dashboard to view event and performance data for a selected device and its interfaces on a single dashboard.
Related concepts
Device Dashboard
Use the **Device Dashboard** to troubleshoot network issues by navigating the network topology and seeing performance anomalies and trends on any device, link, or interface.

About the Device Dashboard
Read this information before installing the **Device Dashboard**.

The content of the **Device Dashboard** varies depending on which Netcool Operations Insight solution extensions you have installed.

As a minimum you must have the base Netcool Operations Insight solution installed together with the Networks for Operations Insight solution extension in order to be able to install the **Device Dashboard**. If you have this combination installed, then the **Device Dashboard** displays the following content:

- **Network Hop View** portlet, enabling network navigation to a selected device.
- **Event Viewer** portlet, displaying events for a selected device.
- **Top Performers** portlet, displaying the top Network Manager polling data for a selected device or interface.

If you have the base Netcool Operations Insight solution installed together with both the Networks for Operations Insight and Performance Management for Operations Insight solution extensions, then, in addition to the **Network Hop View** and the **Event Viewer** portlets, the **Device Dashboard** also displays the following content:

- **Performance Insights** portlet, displaying Network Performance Insight performance anomaly data for a selected device and its interfaces.
- **Performance Timeline** portlet displaying Network Performance Insight performance timeline data for a selected device and its interfaces.
- Capability to launch from a selected interface in the **Performance Insights** portlet to the Network Performance Insight Traffic Details dashboard to see flow data for that interface

Related tasks
Upgrading the Device Dashboard

Installing the Device Dashboard
Follow these instructions to install the **Device Dashboard**.

Before you begin

Before you install the **Device Dashboard**, the following prerequisites must be met:

- If your Netcool Operations Insight system includes the Networks for Operations Insight solution extension only, then Network Manager and the **Network Health Dashboard** must be installed and configured.

- If your Netcool Operations Insight system includes the Networks for Operations Insight and the Performance Management for Operations Insight solution extensions, then the following prerequisites must *also* be met:
  - Network Performance Insight must be installed and configured.
  - Network Manager must be integrated with Network Performance Insight.
  - All post-installation Network Performance Insight configuration tasks must be complete.

Note: For more information about the Networks for Operations Insight and the Performance Management for Operations Insight solution extensions, see the Solution overview.

Procedure

On the host where you installed the Network Manager GUI components, start Installation Manager and install the following package: **Netcool Operations Insight Widgets *version* -> Netcool Operations Insight Device Dashboard**.

Where *version* is the VRMF version number for the current version of the **Device Dashboard**; for example, 1.1.0.2. For information on all current version numbers, see the following web page: https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Netcool%20OMNIbus/page/Release%20details

This installs the **Device Dashboard**, which you can use to troubleshoot network issues by navigating the network topology and viewing performance anomalies and trends on any device or interface.

Note: If you installed Network Performance Insight and integrated this as part of your Netcool Operations Insight solution, then at the start of this process, Installation Manager asks for Ambari credentials. At this point you must specify your Ambari credentials. Ambari is configured during the Network Performance Insight installation process and you would have configured Ambari credentials at that time. Default credentials are as follows:

- userid: `admin`
- password: `admin`

Related tasks

Installing Performance Management
To add the Performance Management solution extension, install Network Performance Insight and then integrate it with Netcool Operations Insight.

Configuring the Device Dashboard
Following installation of the **Device Dashboard**, you must perform these post-installation tasks.

Procedure

1. The **Device Dashboard** installation process automatically creates the `noi_npi` and `noi_npi_admin` roles, which allow users to work with the dashboard. Assign these **Device Dashboard** roles to relevant users. Log in as the Dashboard Application Services Hub administrator to assign the roles `noi_npi` and `noi_npi_admin` as follows:
   a) Go to **Console settings** > **Roles**.
   b) In **Users and Groups**, assign roles `noi_npi` and `noi_npi_admin` to the `npiadmin` user, and assign the `noi_npi` role to the `npiuser` user. The `noi_npi` role provides access to view the **Device Dashboard**, while the `noi_npi_admin` role provides edit access to the **Device Dashboard**.

      Note: You can assign the roles to other individual users or add the role to a group to control access.

      

      Warning: If you do not assign the roles and the user selects **Show Device Dashboard** in the right-click tools, the GUI will hang and the user will receive an encoding error message when logging out, requiring the restart of the browser.
   c) Click **Save**.

If Network Performance Insight is integrated this as part of your Netcool Operations Insight solution, then you must also complete the following configuration tasks.

2. On the same host save the Network Performance Insight profile by performing the following steps:
   a) Go to **Console Settings** > **Console Integrations**.
   b) Select **NPI**.
   c) Review the information, and click **Save**.

   The Network Performance Insight  icon appears on the left in the Navigation bar.
3. Configure access to traffic flow data by performing the following steps:
   a) Log into the Network Manager GUI server and navigate to the following file:

   ```
   $NMGUI_HOME/profile/etc/tnm/tnm.properties
   ```

   Where $NMGUI_HOME location where the Network Manager GUI components are installed. By default, this location is `/opt/IBM/netcool/gui/precision_gui`.
   b) Open the `tnm.properties` file for editing.
   c) Add the following property:

   ```
   tnm.npi.host.name=https://NPI_Server_Name:9443
   ```

   Where *NPI_Server_Name* is the hostname of the Network Performance Insight server.
   d) Save the `tnm.properties` file.
4. Specify the Network Performance Insight version by performing the following steps:
   a) Log into the Network Manager GUI server and navigate to the following file:

   ```
   $NMGUI_HOME/profile/etc/tnm/npi.properties
   ```

   Where $NMGUI_HOME location where the Network Manager GUI components are installed. By default, this location is `/opt/IBM/netcool/gui/precision_gui`.
   b) Open the `npi.properties` file for editing.
   c) Add the following property:

   ```
   npi.server.version=1.3.1
   ```

   d) Save the `npi.properties` file.
   e) Restart the Dashboard Application Services Hub server to enable these properties to take effect.

Related information

Network Manager V4.2 documentation: User rolesClick here to access information on user roles in the Network Manager documentation.

Tivoli Netcool/OMNIbus V8.1 documentation: Restarting the DASH serverClick here to access information how to restart the Dashboard Application Services Hub server.

Installing Agile Service Manager

Install the latest version of Agile Service Manager.

Procedure

1. To install the Agile Service Manager Core services, and Observers, proceed as follows:
   a) On the server where Agile Service Manager Core services are installed (in our example, server 6), extract the Agile Service Manager Base and Observer eAssembly archives.
   b) Follow the instructions in the Agile Service Manager to complete the installation.
2. To install the Agile Service Manager UI, proceed as follows:

a) On the server where Dashboard Application Services Hub is installed (in our example, server 3), extract the Agile Service Manager Base eAssembly archive.

b) Start Installation Manager and configure it to point to the following repository files: `repository.config` file for Agile Service Manager

c) See the Agile Service Manager documentation for more information on how to perform the installation, and perform post-installation configuration tasks, such as configuring the Gateway for Message Bus to support the Agile Service Manager Event Observer.

Related information

Netcool Agile Service Manager Knowledge Center

Configuring Single Sign-On

Single Sign-On (SSO) can be configured to support the launch of tools between the products and components of Netcool Operations Insight. Different SSO handshakes are supported; which handshake to configure for which capability is described here. Each handshake must be configured separately.

Procedure

Set up the SSO handshake as described in the following table.

The table lists which products and components are connected by SSO, which capabilities require which SSO handshake and additional useful information. See the related tasks at the end of the page for links to more information.

Table 24. SSO handshakes for Netcool Operations Insight

SSO handshake can be configured between these products or components		**Handshake is configured to support this capability**	**Additional notes**
Operations Analytics - Log Analysis	Dashboard Application Services Hub	Event search	Supports the launch of right-click tools from the event lists[3] of the Netcool/OMNIbus Web GUI to the custom apps of the Tivoli Netcool/OMNIbus Insight Pack.
Operations Analytics - Log Analysis	Dashboard Application Services Hub	Topology search	Supports the launch of right-click tools from the Web GUI event lists to the custom apps of the Network Manager Insight Pack. Supports the launch of right-click tools from the Network Views in the Network Manager product for the custom apps in the Network Manager Insight Pack.
Netcool Configuration Manager	Dashboard Application Services Hub	Networks for Operations Insight	Supports the launch of right-click tools from the Network Views to the Netcool Configuration Manager GUIs.

Related tasks

Configuring single sign-on for the event search capability

[3] That is, the Event Viewer and Active Event List.

Configure single sign-on (SSO) between Web GUI and Operations Analytics - Log Analysis so that users can switch between the two products without having to log in each time.

Configuring single sign-on for the topology search capability
Configuring SSO between Operations Analytics - Log Analysis V1.3.5 and Dashboard Application Services Hub
Configuring SSO between Operations Analytics - Log Analysis V1.3.3 and Dashboard Application Services Hub
Configuring Single Sign-On for Netcool Configuration Manager
Configuring SSO is a prerequisite to integrating products that are deployed on multiple servers. All DASH server instances must point to the central user registry.

Related information
Configuring Jazz for Service Management for SSO

Installing on IBM Cloud Private

Follow these instructions to prepare for and install Operations Management on IBM Cloud Private.

About this task

In the Private cloud mode, distinct features within Netcool Operations Insight products and components, such as probes, the ObjectServer, Web GUI, and Operations Analytics - Log Analysis, run within containers, and communication between these containers is managed and orchestrated using IBM Cloud Private running on a Kubernetes cluster. For more information about deploying on IBM Cloud Private, see “Deploying on IBM Cloud Private” on page 41.

When you install Operations Management on IBM Cloud Private, all of the components of Operations Management are automatically deployed as pods running within the cluster, as described in the architecture diagram in “Products and components on IBM Cloud Private” on page 3. These components include the following:

- Netcool/OMNIbus Core
- Netcool/OMNIbus Web GUI
- Netcool/Impact
- Operations Analytics - Log Analysis
- Message Bus gateway to support interaction between the components
- Cloud Native Analytics
- LDAP proxy if you are connecting to an existing LDAP, or OpenLDAP

Due to the nature of the cluster, you cannot install a subset of these components.

Note: The current version of Netcool Operations Insight is compatible with IBM Cloud Private version 3.2 and all of the links in this documentation point to that version of the IBM Cloud Private documentation. If you want to view a different version of the IBM Cloud Private documentation, then navigate to IBM Cloud Private documentation: Welcome page and select the desired version.

Preparing for installation on IBM Cloud Private

Before you can install Operations Management on IBM Cloud Private you must set up your Operations Management cluster.

Requirements for an installation on IBM Cloud Private
If you are installing Operations Management on IBM Cloud Private, then ensure that you have the following requirements in place.

Note: For operating system and other detailed system requirements, search for the latest version of the Netcool Operations Insight product in the Software Product Compatibility Reports website: https://www.ibm.com/software/reports/compatibility/clarity/softwareReqsForProduct.html

Requirements for the installation are listed under the following categories:

Note: The current version of Netcool Operations Insight is compatible with IBM Cloud Private version 3.2 and all of the links in this documentation point to that version of the IBM Cloud Private documentation. If you want to view a different version of the IBM Cloud Private documentation, then navigate to IBM Cloud Private documentation: Welcome page and select the desired version.

IBM Cloud Private

The following IBM Cloud Private prerequisites must be met:

- IBM Cloud Private is installed.
- Storage is configured.
- A set of default ports is available for the installation of IBM Cloud Private. For more information, see IBM Cloud Private documentation: Default ports.

IBM Cloud Private with OpenShift

The following IBM Cloud Private prerequisites must be met:

- IBM Cloud Private with OpenShift must be installed.
- Storage is configured.
- Your target namespace must be bound to the predefined SecurityContextConstraints object. For more information, see step "6" on page 110 in the *Configuring pod access control* topic.
- A set of default ports must be available for the installation of IBM Cloud Private with OpenShift. You must have different ports for the nginx ingress controller if you deploy **`nginx ingress`** to the OpenShift master node. Ports 80 and 443 are used by OpenShift services. For more information, see IBM Cloud Private with OpenShift documentation: Preparing to install IBM Cloud Private with OpenShift.

Architecture

The hardware architecture on which IBM Cloud Private is installed must be AMD64.

Kubernetes can have a mixture of worker nodes with different architectures, like AMD64, s390x (Linux on System z®), and ARM8.

Pod Deployment and High Availability (HA)

The deployment of pods across worker nodes is managed by IBM Cloud Private. Pods for a service are deployed on nodes that meet the specification and affinity rules that are defined in the service's yaml file. HA is achieved across the environment by the deployment of a primary ObjectServer pod and a backup ObjectServer pod in a failover configuration. Kubernetes automatically restarts pods if they fail.

Hardware requirements

The bare minimum resource requirements for only the Netcool Operations Insight applications are given in Table 1. These requirements do not give spare capacity for redundancy, resiliency or increased workloads, and give no allowance for the resources required to support IBM Cloud Private.

Table 25. NOI application requirements			
	Number of CPUs	**Memory (Gigabytes)**	**Storage (Gigabytes)**
Trial deployment	18	33	160
Production deployment	49	74	380

For trial environment sizing, see "Trial Sizing" on page 100.

For production environment sizing, see "Production Sizing" on page 101.

Note: If you are also planning to install Agile Service Manager within the IBM Cloud Private cluster, then extra resources are required. These extra resources are described in the relevant topic in the Agile Service Manager Knowledge Center.

Storage

The recommended storage classes for Operations Management on IBM Cloud Private are vSphere and local storage. For more information, see "Storage" on page 104. Hardware requirements are unaffected by the choice of storage class.

Utilities

The following utilities must be installed on the node that is used to connect to the Kubernetes cluster:

- Kubernetes client
- Helm client

Related information

IBM Cloud Private documentation: Preparing your cluster for installation.Click here to access information on how to set up and configure an IBM Cloud Private cluster.

Trial Sizing

A trial deployment has a reduced hardware requirement. It does not provide high availability and is only suitable for a trial, demonstration or proof of concept. It is only suitable for small data sets such as the sample data set, see "Loading sample data: scenario for Operations Management on IBM Cloud Private" on page 347.

Hardware sizing for a trial environment

Set the parameter *Environment size* to `size0` when you configure your installation. For more information, see the *Environment size* parameter in "Configuring Installation Parameters for Operations Management on IBM Cloud Private" on page 125.

Table 1 and Table 2 give an indication of the minimum recommendations for a trial single node and a trial mulit-node deployment of Operations Management on IBM Cloud Private. The minimum required CPU core speed is 2.4Ghz. The hardware requirements for a deployment of Operations Management on IBM Cloud Private vary with workload, and additional nodes or higher specification nodes might be required. For more information on sizing, see https://www.ibm.com/support/knowledgecenter/SSBS6K_3.2.0/installing/plan_capacity.html.

Table 26. Minimum recommendations for a single node trial deployment				
Node	**Number of CPUs**	**Memory (Gigabytes)**	**Storage (Gigabytes)**	**Notes®**
Master node with NOI components	24	48	500	Local storage. Minimum ICP installation with all management services disabled, including logging. Less performant than a multinode size0 deployment.

Table 27. Minimum recommendations for a multi-node cluster trial deployment				
Node	**Number of CPUs**	**Memory (Gigabytes)**	**Storage (Gigabytes)**	**Notes**
Master node	8	16	300	Management and proxy installed on same node.
Worker node 1	8	16	350	150 GB for IBM Cloud Private. 200 GB for local storage.
Worker node 2	8	16	350	150 GB for IBM Cloud Private. 200 GB for local storage.
Worker node 3	8	16	350	150 GB for IBM Cloud Private. 200 GB for local storage.
WORKER TOTAL	24	48	1050	
CLUSTER TOTAL	32	64	1350	

Note: Helm upgrade is not supported on a trial install due to the additional resource footprint requirements of an upgrade. Upgrade from a trial to a production deployment is not supported. Scaling up of StatefulSet pods is not supported for trial deployments.

Production Sizing

A production environment is required for production deployments. The multiple master, proxy, and worker nodes provide redundancy and promote high availability.

Hardware sizing for a production environment

Set the parameter *Environment size* to `size1` when you configure your installation. For more information, see the *Environment size* parameter in "Configuring Installation Parameters for Operations Management on IBM Cloud Private" on page 125.

Table 1 gives an indication of the minimum recommendations for a production deployment of Operations Management on IBM Cloud Private. The minimum required CPU core speed is 2.4Ghz. The hardware requirements for a deployment of Operations Management on IBM Cloud Private vary with workload, and

additional nodes or higher specification nodes might be required. For more information on sizing, see https://www.ibm.com/support/knowledgecenter/SSBS6K_3.2.0/installing/plan_capacity.html.

Table 28. Minimum recommendations for a production environment

Node	Number of CPUs	Memory (Gigabytes)	Storage (Gigabytes)	Notes
Master node 1	8	16	800	
Master node 2	8	16	800	
Master node 3	8	16	800	
Management node 1	8	16	300	
Proxy node 1	2	4	350	
Proxy node 2	2	4	350	
Proxy node 3	2	4	350	
Worker node 1	16	32	550	150 GB for IBM Cloud Private. 400® GB for local storage.
Worker node 2	16	32	550	150 GB for IBM Cloud Private. 400 GB for local storage.
Worker node 3	16	32	550	150 GB for IBM Cloud Private. 400 GB for local storage.
Worker node 4	16	32	550	150 GB for IBM Cloud Private. 400 GB for local storage.
Worker node 5	16	32	550	150 GB for IBM Cloud Private. 400 GB for local storage.
WORKER TOTAL	80	160	2750	
CLUSTER TOTAL	118	236	6500	

Note: Additional worker nodes are required during a helm upgrade.

Preparing your cluster

Perform these steps to prepare a fully functioning Operations Management cluster on IBM Cloud Private.

About this task

Follow the steps in the table to prepare a fully functioning Operations Management cluster on IBM Cloud Private. This is a prerequisite to installing Operations Management on the cluster.

Note: The current version of Netcool Operations Insight is compatible with IBM Cloud Private version 3.2 and all of the links in this documentation point to that version of the IBM Cloud Private documentation. If

you want to view a different version of the IBM Cloud Private documentation, then navigate to IBM Cloud Private documentation: Welcome page and select the desired version.

Table 29. Preparing an Operations Management cluster on IBM Cloud Private

Item	Action	More information
1	Provision the required virtual machines. For further information, see “Requirements for an installation on IBM Cloud Private” on page 98.	For generic IBM Cloud Private requirements, see IBM Cloud Private requirements at IBM Cloud Private documentation: Hardware requirements and recommendations.
2	Create and configure persistent volumes for storage.	For more information, see “Storage” on page 104
3	Download and install IBM Cloud Private.	For the download, see the Operations Management Download Document, at http://www.ibm.com/support/docview.wss?uid=ibm10886865. The IBM Cloud Private e Assembly package is listed in the **Prerequisite eAssembly** section. For the installation, see IBM Cloud Private documentation: Installing IBM Cloud Private. This link provides information on: • Differences between the IBM Cloud Private editions • Preparing to install IBM Cloud Private • Setting up Docker for IBM Cloud Private • Installing IBM Cloud Private
4	Install and configure the Kubernetes command line interface `kubectl` to enable command-line access to the cluster.	See IBM Cloud Private documentation: Accessing your IBM Cloud Private cluster by using the kubectl command-line interface.
5	Install the IBM Cloud Private command line interface (**`cloudctl`** commands) to enable command-line management of the cluster.	See IBM Cloud Private documentation: Installing the IBM Cloud Private command-line interface.
6	Familiarize yourself with the command-line interfaces that you will need to perform the installation and communicate with the cluster.	You can find more information at the following locations: • Helm CLI commands: Helm documentation: Commands • IBM Cloud Private CLI commands: IBM Cloud Private documentation: CLI command reference • Kubernetes CLI commands: Kubernetes documentation: Overview

Table 29. Preparing an Operations Management cluster on IBM Cloud Private (continued)		
Item	**Action**	**More information**
7	Create a custom namespace to deploy your Operations Management installation into, and set the existing `ibm-privileged-psp` as the pod Security Policy. If you are installing into the default namespace (not recommended), then your user profile automatically has the correct privileges. **Note:** There are no pod security policies when installing on IBM Cloud Private with OpenShift **Optional**: If you want multiple independent installations of Operations Management within the cluster, then create multiple namespaces within your cluster. Run each installation in a separate namespace. **Note:** Additional diskspace and worker nodes are required to support multiple installations.	See • IBM Cloud Private documentation: Creating namespaces within your cluster • Creating a custom pod security policy. • Binding a pod security policy to a namespace.

Storage
You must create storage prior to your installation of Operations Management on IBM Cloud Private.

Due to the high I/O bandwidth and low network latency that is required by Operations Management on IBM Cloud Private services, network-based storage options such as Network File System (NFS) and GlusterFS are not supported. vSphere or local storage are the currently supported storage classes.

Configuring persistent volumes with local storage
You can use local storage for Operations Management on IBM Cloud Private.

Local storage is not recommended for production environments, as the loss of a worker node may result in the loss of locally stored data. You could overcome some of these limitations with a file system with the required performance, scale, and redundancy requirements.

For trial or development systems, you can download the `createStorageAllNodes.sh` script from the IT Operations Management Developer Center:

- http://ibm.biz/local_storage_script

The script facilitates the creation of persistent storage volumes (PVs) using local storage. The PVs are mapped volumes, which are mapped to directories off the root file system on the parent node. The script also generates example SSH scripts that create the directories on the local file system of the node. The SSH scripts create directories on the local hard disk associated with the virtual machine and are only suitable for proof of concept or development work.

Configuring vSphere storage
Learn how to configure vSphere storage for your Operations Management on IBM Cloud Private deployment.

Procedure

1. Configure a storage class for vSphere volume with a name such as *vsphere*. For more information, see the IBM Cloud Private documentation: Creating a storage class for vSphere volume
2. To determine the storage class name, run the **`kubectl get storageclass`** command, as in the following example:

```
kubectl get storageclass
NAME                    PROVISIONER                    AGE
image-manager-storage   kubernetes.io/no-provisioner   9d
```

```
mongodb-storage         kubernetes.io/no-provisioner   9d
vsphere                 kubernetes.io/vsphere-volume   9d
```

3. Optional: If you want to configure vSphere storage in the IBM Cloud Private UI, set the following fields to vSphere in the configuration options:
   - Cassandra Data Storage Class
   - Cassandra Backup Storage Class
   - Zookeeper Data Storage Class
   - Kafka Data Storage Class
   - CouchDB Data Storage Class
4. Optional: If you want to configure vSphere storage in the command line, complete the following steps:
   a) Override the settings in the Helm chart by setting the following parameters in the `values.yaml` file to *vsphere*:

```
persistence:
storageClassOption:
cassandrabak: vsphere
cassandradata: vsphere
couchdbdata: vsphere
kafkadata: vsphere
zookeeperdata: vsphere

impactgui:
pvc:
storageClassName: vsphere
nciserver:
pvc:
storageClassName: vsphere
ncobackup:
pvc:
storageClassName: vsphere
ncoprimary:
pvc:
storageClassName: vsphere
openldap:
pvc:
storageClassName: vsphere
scala:
pvc:
storageClassName: vsphere
```

   b) Save and close the `values.yaml` file.

What to do next
Install Operations Management on IBM Cloud Private. For more information, see “Installing on IBM Cloud Private” on page 119.

Preparing secrets for TLS encryption
Operations Management on IBM Cloud Private provides the ability to automatically generate TLS certificates for customers who do not have their own Certificate Authority (CA). However, if you have your own CA and you want to deploy on your own site then you can manually create a certificate and use your own CA to sign the certificate.

Before you begin
A Kubernetes ingress is a collection of rules that can be configured to give services externally reachable URLs. Each Operations Management component requires its own ingress. Each ingress must have an associated secret which has the TLS encrypted certificate. If you are using your own CA then you will need to prepare these secrets using the procedure below.

Procedure

1. Create self signed certificate key pairs using a command similar to the following:

```
openssl req -newkey rsa:2048 -nodes -keyout key.pem -x509 -days 365 -out
```

```
certificate.pem -subj "/CN=ingress.release_name.fqdn/"
```

Where:

- `key.pem` and `certificate.pem` are secure key files and must have already been created by your CA when you created your self signed certificate key pair.
- Common name (CN) is made up of the following parameters:

  Important: Ensure that the common name (CN) value that you specify in this command for each Operations Management component matches *exactly* the parameters specified below.

  – *ingress* is the name of the ingress for the relevant Operations Management component. Run this command for each of the following ingress values:

    - `impact`
    - `nci-x` where *x* is for each of the impact instances from 0 - (number of Impact server instances-1). For example, for 4 impact servers, run for each of these values: nci-0,nci-1,nci-2,nci-3.
    - `netcool`
    - `proxy`
    - `scala`
    - `was`
    - `ibm-hdm-common-ui`

  – *release_name* is the name of the Helm release. Ensure that this is the same release name that you will use when you install your Operations Management on IBM Cloud Private deployment, as described in "Configuring Installation Parameters for Operations Management on IBM Cloud Private" on page 125.

  – *fqdn* is the certificate authority (CA) domain to be set in the **Master node FQDN (Fully Qualified Domain Name)** field when you install your Operations Management on IBM Cloud Private deployment, as described in "Configuring Installation Parameters for Operations Management on IBM Cloud Private" on page 125.

2. Create secrets with your TLS certificate by running the following command for each ingress:

```
kubectl create secret tls release_name-ingress-tls-secret
--cert=./certificate.pem --key=./key.pem [--namespace namespace]
```

Where

- *release_name* is the name of the Helm release. Ensure that this is the same release name that you will use when you install your Operations Management on IBM Cloud Private deployment, as described in "Configuring Installation Parameters for Operations Management on IBM Cloud Private" on page 125.
- *ingress* is the ingress which you are running the command for, and is one of:

  – `impact`

  – `nci-x` where *x* is for each of the impact instances from 0 - (number of Impact server instances-1). For example, for 4 impact servers, run for each of these values: nci-0,nci-1,nci-2,nci-3.

  – `netcool`

  – `proxy`

  – `scala`

  – `was`

  – `ibm-hdm-common-ui`

- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.

3. When installing, you must set the `Use existing TLS certificate secrets` flag to true. For more information, see "Configuring Installation Parameters for Operations Management on IBM Cloud Private" on page 125.

Example

To make this process easier, you could create a script similar to the following:

```
#!/bin/bash
if [ -z "$1" ]; then
    RELNAME=noi
fi

for INGRESS in impact nci-0 nci-1 netcool proxy scala was ibm-hdm-common-ui
do
    kubectl delete secret $RELNAME-$INGRESS-tls-secret
    openssl req -newkey rsa:2048 -nodes -keyout key.pem -x509 -days 365
-out certificate.pem -subj "/CN=$INGRESS.$RELNAME.{{global.cluster.fqdn}}/"
    kubectl create secret tls $RELNAME-$INGRESS-tls-secret --cert=./certificate.pem
--key=./key.pem
    rm ./key.pem ./certificate.pem
done
```

Related tasks

Configuring TLS encryption with a custom certificate

The proxy requires a public certificate and private key pair to be supplied through a Kubernetes secret called `{{ .Release.Name }}-proxy-tls-secret`. If you want to use a custom certificate, for example, one signed by your own public key infrastructure Certificate Authority (CA), create your own proxy secret, containing the public certificate and private key pair, before deployment. To enable a successful Transport Layer Security (TLS) handshake, import the CA signer certificate into the keystore of any client application as a trusted source.

Creating a registry secret to support namespace access

You must create a Docker registry secret to enable IBM Cloud Private to pull the Operations Management on IBM Cloud Private archive from your local Docker registry into your cluster's namespace.

Procedure

1. Run the following command to create the secret:

   ```
   kubectl create secret docker-registry noi-registry-secret \
     --docker-server=cluster-CA-domain:8500 \
     --docker-username=account-username \
     --docker-password=account-password \
     -n namespace
   ```

   Where:

   - *noi-registry-secret* is the name of the secret that you are creating. Suggested value is `noi-registry-secret`
   - *cluster-CA-domain* is the name of the certificate authority domain that was used in the `config.yaml` file during IBM Cloud Private installation as configured in "Preparing your cluster" on page 102
   - *account-username* is the username for the master node.
   - *account-password* is the password for the master node.
   - *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.

   Note: If you are installing on IBM Cloud Private with OpenShift, the command uses different parameters, as in the following example:

   ```
   kubectl create secret docker-registry noi-registry-secret
   --docker-username=ocadmin
   --docker-password=$(oc whoami -t)
   ```

```
--docker-server=docker-registry.default.svc:5000
-n default
```

2. Make a note of the name of the secret that you specified for *noi-registry-secret* as you will need to specify it during the installation of Operations Management on IBM Cloud Private.

Configuring pod access control

Pod access to the services and config maps of Operations Management on IBM Cloud Private is configured through service accounts, roles, and role-based access control (RBAC) settings.

Before you begin

- If you are installing Operations Management on IBM Cloud Private as the cluster administrator, then you can select the *Create required RBAC binding* check box in the IBM Cloud Private GUI. This action causes the required service account, roles, and role bindings to be created in the namespace that is being deployed to. For more information, see "Configuring Installation Parameters for Operations Management on IBM Cloud Private" on page 125.
- If you are not installing as the cluster administrator, or you want to configure the RBAC settings yourself, then you must clear the *Create required RBAC binding* check box in the IBM Cloud Private GUI. For more information, see "Configuring Installation Parameters for Operations Management on IBM Cloud Private" on page 125. You must then manually create the service account, roles, and role bindings that are required, by using the following procedure.

Procedure

1. Create a service account.

   ```
   kubectl create serviceaccount noi-service-account -n namespace
   ```

   Where:

   - *noi-service-account* is the name of the service account name. It is recommended that the service account is named `noi-service-account`.

     Note: Make a note of the name of the service account name as you need to specify it in the `global.rbac.serviceAccountName` parameter during the installation of Operations Management on IBM Cloud Private.

   - *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.

2. Give the service account access to the registry.

   Patch your service account with the NOI registry secret that you created in "Creating a registry secret to support namespace access" on page 107.

   ```
   kubectl patch serviceaccount noi-service-account -p '{"imagePullSecrets":
   [{"name": "noi-registry-secret"}]}' -n namespace
   ```

   Where

   - *noi-service-account* is the name of the service account name.
   - *noi-registry-secret* is the name of the NOI docker registry secret.
   - *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.

3. Run the following command to check whether the new service account `noi-service-account` has access to Kubernetes services on this namespace.

   ```
   kubectl auth can-i patch services --namespace namespace
    --as system:serviceaccount:namespace:noi-service-account
   ```

   Where:

- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.
- *noi-service-account* is the name of the service account name.

If this command returns the response no, then as an administrator user run the following commands to create a role and a role binding, both called noiservice-reader.

```
kubectl create role noiservice-reader --verb=update --verb=patch --verb=get
--verb=list --verb=watch --resource=services -n namespace
```

```
kubectl create rolebinding noiservice-reader --role noiservice-reader
--serviceaccount namespace:noi-service-account -n namespace
```

Where:

- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.
- *noi-service-account* is the name of the service account name.

4. Run the following command to check whether the new service account noi-service-account has access to Kubernetes config maps on this namespace.

```
kubectl auth can-i watch configmaps --namespace namespace
 --as system:serviceaccount:namespace:noi-service-account
```

Where:

- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.
- *noi-service-account* is the name of the service account name.

If this command returns the response no, then as an administrator user run the following commands to create a role and a role binding, both called noiconfigmap-reader.

```
kubectl create role noiconfigmap-reader --verb=update --verb=patch --verb=get
--verb=list --verb=watch --resource=configmaps -n namespace
```

```
kubectl create rolebinding noiconfigmap-reader --role noiconfigmap-reader
--serviceaccount namespace:noi-service-account -n namespace
```

Where:

- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.
- *noi-service-account* is the name of the service account name.

5. **Note:** If you are installing on IBM Cloud Private with OpenShift, skip this step.

Run the following command to check whether the new service account noi-service-account can use the privileged podsecuritypolicy in this namespace.

```
kubectl auth can-i use podsecuritypolicy/ibm-privileged-psp --namespace
namespace --as system:serviceaccount:namespace:noi-service-account
```

Where:

- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.
- *noi-service-account* is the name of the service account name.

If this command returns the response no, then as an administrator user run the following command to create the *noi-privileged* role binding.

```
kubectl create rolebinding noi-privileged --clusterrole ibm-privileged-clusterrole
 --serviceaccount namespace:noi-service-account -n namespace
```

Where:

- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.
- *noi-service-account* is the name of the service account name.

6. If you are installing on IBM Cloud Private with OpenShift, then you must bind your target namespace to the predefined SecurityContextConstraints object, *ibm-privileged-scc*, by using the following command:

```
oc adm policy add-scc-to-group ibm-privileged-scc system:serviceaccounts:namespace
```

Where:

- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.

7. Secrets can be created manually or by the installer. For more information, see "Configuring passwords and secrets" on page 111. If you are creating the required passwords and secrets manually, then miss out this step. If passwords and secrets are being created by the installer, then run the following command to verify whether the new service account `noi-service-account` can be used by the installer to create secrets.

```
kubectl auth can-i create secrets --namespace namespace --as
system:serviceaccount:namespace:noi-service-account
```

Where:

- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.
- *noi-service-account* is the name of the service account name.

If this command returns the response `no`, then as an administrator user you must run the following command to create a role and a role binding, both called `noisecret-creater`.

```
kubectl create role noisecret-creater --verb=list --verb=create --verb=get --verb=list
 --resource=secrets -n namespace
```

```
kubectl create rolebinding noisecret-creater --serviceaccount namespace:noi-service-account
 --role noisecret-creater -n namespace
```

Where:

- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.
- *noi-service-account* is the name of the service account name.

8. Run the following command to check whether the redis pod can view the status of other pods.

```
kubectl auth can-i watch pods --namespace namespace --as
system:serviceaccount:namespace:noi-service-account
```

Where:

- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.
- *noi-service-account* is the name of the service account name.

If this command returns the response `no`, then as an administrator user you must run the following command to create a role and a role binding, both called *releasename*`-redis`.

```
kubectl create role releasename-redis --verb=get --verb=list --verb=update
--verb=patch --verb=watch --resource=pods -n namespace
```

```
kubectl create rolebinding releasename-redis --serviceaccount
```

```
namespace:noi-service-account --role releasename-redis -n namespace
```

Where:

- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.
- *releasename* is the name of the helm release .
- *noi-service-account* is the name of the service account name.

Configuring passwords and secrets

Passwords are stored in secrets. You can either manually create the passwords and secrets that are required by Operations Management on IBM Cloud Private, or the required passwords and secrets can be generated for you by the installer.

Procedure

1. The following user passwords and secrets are required.

Users requiring password	Corresponding secret	Data key(s) in secret
smadmin	*helm-release-name*-was-secret	WAS_PASSWORD
impactadmin	*helm-release-name*-impact-secret	IMPACT_ADMIN_PASSWORD
unityadmin	*helm-release-name*-la-secret	UNITY_ADMIN_PASSWORD
icpadmin	*helm-release-name*-icpadmin-secret	ICP_ADMIN_PASSWORD
OMNIbus root	*helm-release-name*-omni-secret	OMNIBUS_ROOT_PASSWORD
LDAP admin	*helm-release-name*-ldap-secret	LDAP_BIND_PASSWORD
couchdb	*helm-release-name*-couchdb-secret	password username=root secret=couchdb
internal user	*helm-release-name*-ibm-hdm-common-ui-session-secret	session
internal user	*helm-release-name*-systemauth-secret	password username=system
hdm	*helm-release-name*-cassandra-auth-secret	username password
redis	*helm-release-name*-ibm-redis-authsecret	username password
kafka	*helm-release-name*-kafka-admin-secret	username password
admin	*helm-release-name*-kafka-client-secret	username password

   Create these passwords and secrets manually, or leave the installer to create the passwords and secrets automatically and then retrieve the passwords post-install.

2. Automatic creation of passwords and secrets.

   The Operations Management on IBM Cloud Private installer uses existing passwords and secrets. If any of the required passwords and secrets do not exist, then the installer automatically creates random passwords for the required passwords and then creates the required secrets from these passwords.

   For automatic creation of passwords and secrets, use the following procedure.

a) Ensure that the service account has permissions to create secrets. For more information, see "Configuring pod access control" on page 108.

b) Proceed with the installation, using "Installing Operations Management on IBM Cloud Private" on page 119. If you set the *LDAP mode* to `proxy`, then you MUST manually configure the passwords and secrets for LDAP `admin`, `impactadmin`, and `unityadmin`. For information on how to do this, refer to step "3" on page 112, `Manual creation of Secrets.`

c) Proceed with the installation, using "Installing Operations Management on IBM Cloud Private" on page 119, and ensure that the check box *Indicates that all password secrets have been created prior to install* is cleared.

d) After installation has successfully completed, you can extract the passwords from the secrets. See "Retrieving passwords from secrets" on page 132

3. Manual creation of passwords and secrets

All passwords must be less than 16 characters long and contain only alphanumeric characters. To create all the required passwords and secrets manually, use the following procedure.

a) Create passwords for the users requiring passwords in step 1, if these do not already exist.

b) Create *helm-release-name*-icpadmin-secret with the following command:

```
kubectl create secret generic helm-release-name-icpadmin-secret
--from-literal=ICP_ADMIN_PASSWORD=password --namespace namespace
```

Where

- *password* is the password for icpadmin.
- *helm_release_name* is the name that you are planning to use for your Operations Management on IBM Cloud Private Helm release name in "Installing Operations Management on IBM Cloud Private" on page 119
- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.

c) Create *helm-release-name*-impactadmin-secret with the following command:

```
kubectl create secret generic helm-release-name-impactadmin-secret
--from-literal=IMPACT_ADMIN_PASSWORD=password --namespace namespace
```

Where

- *password* is the password for impactadmin.
- *helm_release_name* is the name that you are planning to use for your Operations Management on IBM Cloud Private Helm release name in .
- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.

d) Create *helm-release-name*-la-secret with the following command:

```
kubectl create secret generic helm-release-name-la-secret
--from-literal=UNITY_ADMIN_PASSWORD=password --namespace namespace
```

Where

- *password* is the password for unityadmin.
- *helm_release_name* is the name that you are planning to use for your Operations Management on IBM Cloud Private Helm release name in "Installing Operations Management on IBM Cloud Private" on page 119
- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.

e) Create *helm-release-name*-ldap-secret with the following command:

```
kubectl create secret generic helm-release-name-ldap-secret
```

```
--from-literal=LDAP_BIND_PASSWORD=password --namespace namespace
```

Where

- *password* is the password of your organization's LDAP server.
- *helm_release_name* is the name that you are planning to use for your Operations Management on IBM Cloud Private Helm release name in "Installing Operations Management on IBM Cloud Private" on page 119
- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.

f) Create *helm-release-name*-omni-secret with the following command:

```
kubectl create secret generic helm-release-name-omni-secret
--from-literal=OMNIBUS_ROOT_PASSWORD=password --namespace namespace
```

Where

- *password* is the root password to set for the Netcool/OMNIbus ObjectServer.
- *helm_release_name* is the name that you are planning to use for your Operations Management on IBM Cloud Private Helm release name in "Installing Operations Management on IBM Cloud Private" on page 119.
- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.

g) Create *helm-release-name*-was-secret (smadmin user) with the following command:

```
kubectl create secret generic helm-release-name-was-secret
--from-literal=WAS_PASSWORD=password --namespace namespace
```

Where

- *password* is the password for OMNIbus admin user.
- *helm_release_name* is the name that you are planning to use for your Operations Management on IBM Cloud Private Helm release name in "Installing Operations Management on IBM Cloud Private" on page 119.
- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.

h) Create *helm-release-name*-couchdb-secret with the following command:

```
kubectl create secret generic helm-release-name-couchdb-secret
--from-literal=password=password --from-literal=secret=couchdb
--from-literal=username=root --namespace namespace
```

Where

- *password* is the password for the internal couch.
- *helm_release_name* is the name that you are planning to use for your Operations Management on IBM Cloud Private Helm release name in "Installing Operations Management on IBM Cloud Private" on page 119.
- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.

i) Create secret for communication between pods with the following command:

```
kubectl create secret generic helm_release_name-systemauth-secret
--from-literal=password=password --from-literal=username=system
--namespace namespace
```

Where

- *password* is the password for communication between pods.

- *helm_release_name* is the name that you are planning to use for your Operations Management on IBM Cloud Private Helm release name in "Installing Operations Management on IBM Cloud Private" on page 119.
- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.

j) Create secret for user interface communication between pods with the following command:

```
kubectl create secret generic helm_release_name-ibm-hdm-common-ui-session-secret
--from-literal=session=password --namespace namespace
```

Where

- *password* is the password for user interface communication between pods.
- *helm_release_name* is the name that you are planning to use for your Operations Management on IBM Cloud Private Helm release name in "Installing Operations Management on IBM Cloud Private" on page 119.
- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.

k) Create *helm_release_name*-cassandra-auth-secret with the following command:

```
kubectl create secret generic helm_release_name-cassandra-auth-secret
--from-literal=username=username --from-literal=password=passwordpassword
--namespace namespace
```

Where

- *username* default is hdm. Do not use cassandra.
- *password* is the password for user interface communication between pods.
- *helm_release_name* is the name that you are planning to use for your Operations Management on IBM Cloud Private Helm release name in "Installing Operations Management on IBM Cloud Private" on page 119.
- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.

l) Create *helm_release_name*-ibm-redis-authsecret with the following command:

```
kubectl create secret generic helm_release_name-ibm-redis-authsecret
--from-literal=username=username --from-literal=password=password --namespace namespace
```

Where

- *username* default is redis.
- *password* is the password for user interface communication between pods.
- *helm_release_name* is the name that you are planning to use for your Operations Management on IBM Cloud Private Helm release name in "Installing Operations Management on IBM Cloud Private" on page 119.
- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.

m) Create *helm_release_name*-kafka-admin-secret with the following command:

```
kubectl create secret generic helm_release_name-kafka-admin-secret
--from-literal=username=username --from-literal=password=password
--namespace namespace
```

Where

- *username* default is kafka.
- *password* is the password for user interface communication between pods.

- *helm_release_name* is the name that you are planning to use for your Operations Management on IBM Cloud Private Helm release name in "Installing Operations Management on IBM Cloud Private" on page 119.
- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.

n) Create *helm_release_name*-kafka-client-secret with the following command:

```
kubectl create secret generic helm_release_name-kafka-client-secret
--from-literal=username=username --from-literal=password=password
--namespace namespace
```

Where

- *username* default is `admin`.
- *password* is the password for user interface communication between pods.
- *helm_release_name* is the name that you are planning to use for your Operations Management on IBM Cloud Private Helm release name in "Installing Operations Management on IBM Cloud Private" on page 119.
- *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private.

o) Proceed with the installation, using "Installing Operations Management on IBM Cloud Private" on page 119, and ensure that the check box 'Indicates that all password secrets have been created prior to install' is checked.

If you wish to change a password after installation, see "Changing passwords and recreating secrets" on page 133

Loading the archive into IBM Cloud Private

Run these commands to load the archive into IBM Cloud Private.

Before you begin

Before you perform this task make sure you have met the following prerequisites:

- You must have downloaded the eAssembly CJ5MZEN from IBM Passport Advantage. For more information, see the IBM Support Download Document for Operations Management on IBM Cloud Private, at http://www.ibm.com/support/docview.wss?uid=ibm10886865.
- If you installed a previous version of Operations Management on IBM Cloud Private, and you are not upgrading, then you must uninstall that version, including all images and Helm charts before loading the current archive into IBM Cloud Private. To do this, follow the instructions at "Uninstalling on IBM Cloud Private" on page 143.

About this task

Note: The current version of Netcool Operations Insight is compatible with IBM Cloud Private version 3.2 and all of the links in this documentation point to that version of the IBM Cloud Private documentation. If you want to view a different version of the IBM Cloud Private documentation, then navigate to IBM Cloud Private documentation: Welcome page and select the desired version.

Procedure

1. Log into the cluster master node on IBM Cloud Private.

   Issue the following command:

   ```
   cloudctl login -a https://master-node-address:8443 -u username -p password
   -c account -n namespace
   ```

   Where:

   - *master-node-address* is the hostname or IP address of the master ICP node.

- *username* is the username for the master node. By default this is admin.
- *password* is the password for the master node. By default this is admin.
- *account* is the IBM Cloud Private account.
- *namespace* is the namespace.

  Note: Only specify this parameter if you are installing in a custom namespace.

Refer to the IBM Cloud Private general CLI commands (Cloudant®) topic link at the bottom of this topic for full syntax details.

2. Log into the Docker repository.

```
docker login cluster-CA-domain:8500 -u username -p password
```

Where:

- *cluster-CA-domain* is the name of the certificate authority domain that was used in the config.yaml during IBM Cloud Private installation. See "Preparing your cluster" on page 102
- *username* is the username for the master node. By default this is admin.
- *password* is the password for the master node. By default this is admin.

3. Unset the proxy connection.

The proxy connection enables you to connect to the Internet using a proxy server. You must unset this for ICP to be able to work.

```
unset http_proxy
unset HTTP_PROXY
```

4. Load the Operations Management on IBM Cloud Private archive into IBM Cloud Private.

```
cloudctl catalog load-archive --archive ibm-netcool-prod.2.1.1-x86_64.tar.gz
--repo local-charts --registry cluster-CA-domain:8500/namespace
```

Note: Ensure your connection doesn't time out while loading the archive.

Where:

- *cluster-CA-domain* is the name of the certificate authority domain that was used in the config.yaml file during IBM Cloud Private installation. See "Preparing your cluster" on page 102
- *namespace* is the name of the namespace which you wish to deploy to.

This command takes a considerable amount of time to complete. The exact amount of time depends on your system but you should allow from at least 25 minutes up to an hour or more.

5. Optional: Load the Agile Service Manager on IBM Cloud Private archive into IBM Cloud Private.

```
cloudctl catalog load-archive --archive NOI1.6_ASM_V1.1.5_FOR_ICP.tar.gz
--repo local-charts --registry cluster-CA-domain:8500/namespace
```

Where:

- *cluster-CA-domain* is the name of the certificate authority domain that was used in the config.yaml file during IBM Cloud Private installation. See "Preparing your cluster" on page 102
- *namespace* is the name of the namespace which you wish to deploy to.

Related information

IBM Cloud Private documentation: cloudctl commandsLearn about the general cloudctl commands that you can run to access your IBM® Cloud Private cluster.

IBM Cloud Private documentation: cloudctl catalog commandsLearn about the cloudctl catalog commands that you can run to manage your Helm charts.

Loading the archive into IBM Cloud Private with OpenShift

Run these commands to load the archive into IBM Cloud Private with OpenShift.

Before you begin

Before you perform this task, make sure that you have met the following prerequisites:

- You must have downloaded the eAssembly CJ5MZEN from IBM Passport Advantage. For more information, see the IBM Support Download Document for Operations Management on IBM Cloud Private, at http://www.ibm.com/support/docview.wss?uid=ibm10886865.
- If you installed a previous version of Operations Management on IBM Cloud Private, and you are not upgrading, then you must uninstall that version, including all images and Helm charts before loading the current archive into IBM Cloud Private. To do this, follow the instructions at "Uninstalling on IBM Cloud Private" on page 143.
- OpenShift CLI (oc). For details about installing the CLI, see Installing the CLI.
- You need to configure Docker to authenticate with IBM Cloud Private with OpenShift. Configure the service for the docker registry on OpenShift to make it accessible to clients outside the cluster. For more information, see Accessing the Registry. Create a regular user account with the following roles:
  - `system:registry` role
  - `admin` role for the project
  - `system:image-builder` role

  For more information about creating and configuring the user account to access the Docker registry on OpenShift, see User Prerequisites.

Procedure

1. Log in to your IBM Cloud Private with OpenShift cluster with your credentials by using the following command (replace <MasterNode_IP> with the IP address of the master node and select the account and namespace from the list):

```
cloudctl login -a https://<MasterNode_IP>:5443

Username> admin

Password>
Authenticating...
OK

Targeted account mycluster Account (id-mycluster-account)

Select a namespace:
1. cert-manager
2. default
3. glusterfs
4. istio-system
5. kube-public
6. kube-service-catalog
7. kube-system
8. management-infra
9. openshift
10. openshift-ansible-service-broker
11. openshift-infra
12. openshift-logging
13. openshift-node
14. openshift-sdn
15. openshift-template-service-broker
16. openshift-web-console
17. test-automation
Enter a number> 2
Targeted namespace default

Configuring kubectl ...
Property "clusters.mycluster" unset.
Property "users.mycluster-user" unset.
Property "contexts.mycluster-context" unset.
Cluster "mycluster" set.
```

```
User "mycluster-user" set.
Context "mycluster-context" created.
Switched to context "mycluster-context".
OK

Configuring helm: /Users/myUser/.helm
OK
```

Note: You can use the `--skip-ssl-validation` option to skip Transport Layer Security (TLS) certificate validation, but this option is not recommended for production environments.

2. Log in to the OpenShift Container Platform server using the following command (replace <MasterNode_IP> with the IP address of the master node and select the account and namespace from the list):

```
$ oc login
Authentication required for https://<MasterNode_IP>:7443 (openshift)
Username: admin
Password:
Login successful.

You have access to the following projects and can switch between them
with 'oc project <projectname>':

  cert-manager
* default
  glusterfs
  istio-system
  kube-public
  kube-service-catalog
  kube-system
  management-infra
  openshift
  openshift-ansible-service-broker
  openshift-infra
  openshift-logging
  openshift-node
  openshift-sdn
  openshift-template-service-broker
  openshift-web-console
  test-automation

Using project "default".
```

After logging in successfully, switch to the correct project if needed.

3. Identify the name of the Docker registry on the OpenShift server. The following command returns the <DockerRegistry> name:

```
$ oc get route docker-registry --template '{{.spec.host}}'
<DockerRegistry>
```

4. Log in to the Docker registry on the OpenShift server by entering the following command (replace <DockerRegistry> with the name of the Docker registry):

```
$ docker login -u $(oc whoami) -p $(oc whoami -t) <DockerRegistry>

Login Succeeded
```

Note: The **`oc whoami`** and **`oc whoami -t`** sub-commands enter the user name and token required for the login.

5. Load the PPA archive using the following command (replace <Archive> with the PPA archive file name, for example `ibm-netcool-prod.2.1.1-x86_64.tar.gz`, <DockerRegistry> with the Docker registry, and <NameSpace> with the target namespace):

```
$ cloudctl catalog load-archive --archive <Archive> --registry
<DockerRegistry>/<Namespace>
Expanding archive
OK

Importing docker image(s)
    Loading Image
```

```
Processing image: netcool-probe-messagebus:9.0.9.2-amd64
    Tagging Image
    Pushing image as: <DockerRegistry>/
<Namespace>/netcool-probe-messagebus:9.0.9.2-amd64
    Loading Image
Processing image: netcool-integration-util:1.0.0-amd64
    Tagging Image
    Pushing image as: <DockerRegistry>/
<Namespace>/netcool-integration-util:1.0.0-amd64
OK

Uploading helm chart(s)
Processing chart: charts/ibm-netcool-probe-4.0.0.tgz
Updating chart values.yaml
Uploading chart
Loaded helm chart
OK

Synch charts
Synch started
OK

Archive finished processing
```

Installing on IBM Cloud Private

Follow these instructions to install Operations Management and optionally the solution extension for Service Management, IBM Agile Service Manager, on a private cloud with IBM Cloud Private.

Note: If you plan to install the optional Agile Service Manager extension, then this must be installed before Operations Management is installed.

Related tasks

Viewing Kubernetes logs
To view information on the success of the installation process for Netcool Operations Insight on IBM Cloud Private, run the `kubectl logs` command from the command line.

Installing Operations Management on IBM Cloud Private

Follow these instructions to install the IBM Netcool Operations Insight base solution, also known as Operations Management for Operations Insight on IBM Cloud Private.

Before you begin

Before you install, check that the following prerequisites are met:

- Your cluster is prepared with the required tools, and a custom namespace. For more information, see “Preparing your cluster” on page 102.
- Your cluster has the required storage. For more information, see “Storage” on page 104
- You have created your own TLS certificate secrets, or it is your intention that they will be created by the installer. For more information, see “Preparing secrets for TLS encryption” on page 105.
- You have created a docker registry secret. For more information, see “Creating a registry secret to support namespace access” on page 107.
- You have created a service account that has the required access to the registry, Kubernetes services, configmaps, and other pods. For more information, see “Configuring pod access control” on page 108.
- You have created the required passwords and secrets, or it is your intention that they will be created by the installer. For more information, see “Configuring passwords and secrets” on page 111.
- The Operations Management on IBM Cloud Private archive has been loaded. For more information, see “Loading the archive into IBM Cloud Private” on page 115 or “Loading the archive into IBM Cloud Private with OpenShift” on page 117.

Note: If you plan to install the optional Agile Service Manager extension, then it must be installed before Operations Management on IBM Cloud Private is installed.

Operations Management for Operations Insight can be installed through the IBM Cloud Private UI, or by using the command line.

To install through the IBM Cloud Private UI, see "Installing Operations Management with the IBM Cloud Private UI" on page 120.

To install by using the command line, see "Installing Operations Management on IBM Cloud Private with the command line" on page 122.

Related tasks

Viewing Kubernetes logs
To view information on the success of the installation process for Netcool Operations Insight on IBM Cloud Private, run the `kubectl logs` command from the command line.

Installing Operations Management with the IBM Cloud Private UI

Follow these instructions to install Operations Management on IBM Cloud Private using the IBM Cloud Private UI.

Before you begin

If you plan to install the optional Agile Service Manager extension, then this must be installed before Operations Management is installed.

Procedure

1. Login to the IBM Cloud Private GUI by pointing your internet browser to a URL similar to:

   ```
   https://master_node_ip_or_hostname:8443/oidc/login.jsp
   ```

   Where *master_node_ip_or_hostname* is the IP address or host name of your cluster's master node.

   Note: For IBM Cloud Private with OpenShift, access your cluster by using a different URL.
2. In the IBM Cloud Private banner at the top of the page, click **Catalog** on the right side of the banner.
3. Enter `netcool` in the **Filter** and select the **ibm-netcool-prod** Helm chart entry.
4. Click **Configure**.

   Configure the installation parameters. See "Configuring Installation Parameters for Operations Management on IBM Cloud Private" on page 125.

   Note: If you want the installer to create and configure storage for your deployment, then leave the storage parameters with their default values. See https://www.ibm.com/support/knowledgecenter/SSBS6K_3.2.0/manage_cluster/cluster_storage.html if you want to manually set up storage, and then set the storage parameters for the persistent volumes (PVs) and persistent volume claims (PVCs) that are configured.
5. Click **Install**.

   Note: 1.6.0.1 If you have enabled the optional Agile Service Manager extension for Topology Analytics, there is a known issue that must be rectified after the install has completed. See step "7" on page 122.
6. Monitor the overall progress of the installation by using the following commands:

   `kubectl describe job`
   This command monitors jobs, as in the following example:

   ```
   kubectl describe job releasename-verifysecrets
   ```

   where `releasename` is the name that you chose in step 2 for your Operations Management on IBM Cloud Private deployment.

   `kubectl get pods`
   This command retrieves the status of all the Operations Management on IBM Cloud Private pods. Use this command for an overview of installation progress, and to verify that all the pods are running. Here is an example of the output from this command, where `releasename` is the name chosen for your Operations Management on IBM Cloud Private deployment in step 2.

```
NAME                                                              READY
STATUS            RESTARTS   AGE
releasename-cassandra-0                                           1/1
Running           0          10m
releasename-cassandra-change-superuser-post-install               1/1
Completed         0          10m
releasename-couchdb-0                                             1/1
Running           0          10m
releasename-db2ese-0                                              2/2
Running           0          10m
releasename-ea-noi-layer-eanoiactionservice                       1/1
Running           0          10m
releasename-ea-noi-layer-eanoigateway                             1/1
Running           0          10m
releasename-ea-ui-api-graphql                                     1/1
Running           0          10m
releasename-ibm-hdm-analytics-dev-collater-aggregationservice     1/1
Running           0          10m
releasename-ibm-hdm-analytics-dev-dedup-aggregationservice        1/1
Running           0          10m
releasename-ibm-hdm-analytics-dev-normalizer-aggregationservice   1/1
Running           0          10m
releasename-ibm-hdm-analytics-dev-archivingservice                1/1
Running           0          10m
releasename-ibm-hdm-analytics-dev-eventsqueryservice              1/1
Running           0          10m
releasename-ibm-hdm-analytics-dev-inferenceservice                1/1
Running           0          10m
releasename-ibm-hdm-analytics-dev-ingestionservice                1/1
Running           0          10m
releasename-ibm-hdm-analytics-dev-policyregistryservice           1/1
Running           0          10m
releasename-ibm-hdm-analytics-dev-servicemonitorservice           1/1
Running           0          10m
releasename-ibm-hdm-analytics-dev-setup                           2/2
Completed         0          10m
releasename-ibm-hdm-analytics-dev-trainer                         1/1
Running           0          10m
releasename-ibm-hdm-common-ui-uiserver                            1/1
Running           0          10m
releasename-impactgui-0                                           2/2
Running           0          10m
releasename-kafka-0                                               2/2
Running           0          10m
releasename-nciserver-0                                           2/2
Running           0          10m
releasename-nciserver-1                                           2/2
Running           0          10m
releasename-ncobackup-0                                           2/2
Running           0          10m
releasename-ncoprimary-0                                          1/1
Running           0          10m
releasename-openldap-0                                            1/1
Running           0          10m
releasename-proxy-0                                               1/1
Running           0          10m
releasename-redis-sentinel                                        1/1
Running           0          10m
releasename-redis-server                                          1/1
Running           0          10m
releasename-scala-0                                               3/3
Running           0          10m
releasename-spark-master                                          1/1
Running           0          10m
releasename-spark-slave                                           1/1
Running           0          10m
releasename-webgui-0                                              2/2
Running           0          10m
releasename-zookeeper-0                                           1/1
Running           0          10m
1.6.0.1releasename-ibm-ea-asm-normalizer-normalizerstreams        1/1
```

```
Running           0          10m
releasename-ibm-ea-asm-normalizer-mirrormaker            1/1
Running           0          10m
```

1.6.0.1 The 1.6.0.1 pods are only seen when the topology analytics integration with Agile Service Manager is enabled.

Here is a brief explanation of the output columns.

NAME
Name of the pod. This name is made up of the deployment name, the name of the primary image in the pod, and a random sequence of alphanumeric characters if the deployment is not a stateful set.

READY
The number of containers within the pod that are ready/the number of containers in the pod. For example, in the code snippet above, the backup ObjectServer pod, named `ncobackup` is made up of two containers, and both of them are ready.

STATUS
Status of the pod.

RESTARTS
How many times the pod has been restarted.

AGE
How long the pod has been in existence.

Note: The `kubectl get pods` command only provides overall status of the installation. To view more detailed status and to monitor the status of connections between components, you must view the Kubernetes logs.

When all of the pods have the status `Running` or `Completed`, then the deployment is ready.

7. 1.6.0.1
If you have enabled the optional Agile Service Manager extension for Topology Analytics, then you must update your `values.yaml` configuration file.

See this known issue "Incorrect path in ibm-ea-asm-normaliser section in values.yaml" on page 562, and follow the instructions in it. Your deployment must then be upgraded to use the updated path. Run the following command:

```
helm upgrade helm_release_name ./ibm-netcool-prod-2.1.1.tgz -f custom-values-upgrade.yaml --
tls
```

where *helm_release_name* is the name of your Operations Management on IBM Cloud Private release.

Installing Operations Management on IBM Cloud Private with the command line
Follow these instructions to install Operations Management on IBM Cloud Private with the command line interface (CLI).

Before you begin
If you plan to install the optional Agile Service Manager extension, then this must be installed before Operations Management is installed.

Procedure

1. Extract the charts and `values.yaml` file from the PPA archive.

```
tar -xvf ibm-netcool-prod.2.1.1-x86_64.tar.gz charts/ibm-netcool-prod-2.1.1.tgz
cd charts
tar -xvf ibm-netcool-prod-2.1.1.tgz ibm-netcool-prod/values.yaml
```

2. Configure the installation parameters for your Operations Management on IBM Cloud Private deployment by editing the supplied configuration file `values.yaml`, or by creating your own yaml

configuration file. For more information, see “Configuring Installation Parameters for Operations Management on IBM Cloud Private” on page 125.

3. 1.6.0.1
   If you are enabling the optional Agile Service Manager extension for Topology Analytics, then you **must** update your *values.yaml* configuration file. See this known issue “Incorrect path in ibm-ea-asm-normaliser section in values.yaml” on page 562 and follow the instructions to avoid installation errors.
4. Run the following command to install Operations Management on IBM Cloud Private from the *ibm-netcool-prod* helm chart with the parameters specified in your `values.yaml` file.

   ```
   helm install ibm-netcool-prod -f values.yaml --tls --name releasename --namespace namespace
   ```

   Where

   - *values.yaml* is the name of the configuration file that contains the installation parameters. See “Configuring Installation Parameters for Operations Management on IBM Cloud Private” on page 125.
   - *releasename* is the name that you want to call your Operations Management on IBM Cloud Private installation.
   - *namespace* is the name of the namespace into which you want to install Operations Management on IBM Cloud Private. For more information, see “Preparing your cluster” on page 102.
5. Monitor the overall progress of the installation by using the following commands:

   `kubectl describe job`
   This command monitors jobs, as in the following example:

   ```
   kubectl describe job releasename-verifysecrets -n namespace
   ```

   Where

   - *releasename* is the name that you chose in step 3 for your Operations Management on IBM Cloud Private deployment.
   - *namespace* is the name of the namespace in which Operations Management on IBM Cloud Private is deployed.

   `kubectl get pods -n namespace`
   This command retrieves the status of all the Operations Management on IBM Cloud Private pods where *namespace* is the name of the namespace in which Operations Management on IBM Cloud Private is deployed. Use this command for an overview of installation progress, and to verify that all the pods are running. Here is an example of the output from this command, where *releasename* is the name that is chosen for your Operations Management on IBM Cloud Private deployment in step 3.

```
NAME                                                              READY
STATUS            RESTARTS   AGE
releasename-cassandra-0                                             1/1
Running           0          10m
releasename-cassandra-change-superuser-post-install               1/1
Completed         0          10m
releasename-couchdb-0                                               1/1
Running           0          10m
releasename-db2ese-0                                                2/2
Running           0          10m
releasename-ea-noi-layer-eanoiactionservice                       1/1
Running           0          10m
releasename-ea-noi-layer-eanoigateway                             1/1
Running           0          10m
releasename-ea-ui-api-graphql                                       1/1
Running           0          10m
releasename-ibm-hdm-analytics-dev-collater-aggregationservice   1/1
Running           0          10m
releasename-ibm-hdm-analytics-dev-dedup-aggregationservice      1/1
Running           0          10m
releasename-ibm-hdm-analytics-dev-normalizer-aggregationservice 1/1
```

```
Running            0          10m
releasename-ibm-hdm-analytics-dev-archivingservice              1/1
Running            0          10m
releasename-ibm-hdm-analytics-dev-eventsqueryservice            1/1
Running            0          10m
releasename-ibm-hdm-analytics-dev-inferenceservice              1/1
Running            0          10m
releasename-ibm-hdm-analytics-dev-ingestionservice              1/1
Running            0          10m
releasename-ibm-hdm-analytics-dev-policyregistryservice         1/1
Running            0          10m
releasename-ibm-hdm-analytics-dev-servicemonitorservice         1/1
Running            0          10m
releasename-ibm-hdm-analytics-dev-setup                         2/2
Completed          0          10m
releasename-ibm-hdm-analytics-dev-trainer                       1/1
Running            0          10m
releasename-ibm-hdm-common-ui-uiserver                          1/1
Running            0          10m
releasename-impactgui-0                                         2/2
Running            0          10m
releasename-kafka-0                                             2/2
Running            0          10m
releasename-nciserver-0                                         2/2
Running            0          10m
releasename-nciserver-1                                         2/2
Running            0          10m
releasename-ncobackup-0                                         2/2
Running            0          10m
releasename-ncoprimary-0                                        1/1
Running            0          10m
releasename-openldap-0                                          1/1
Running            0          10m
releasename-proxy-0                                             1/1
Running            0          10m
releasename-redis-sentinel                                      1/1
Running            0          10m
releasename-redis-server                                        1/1
Running            0          10m
releasename-scala-0                                             3/3
Running            0          10m
releasename-spark-master                                        1/1
Running            0          10m
releasename-spark-slave                                         1/1
Running            0          10m
releasename-webgui-0                                            2/2
Running            0          10m
releasename-zookeeper-0                                         1/1
Running            0          10m
1.6.0.1 releasename-ibm-ea-asm-normalizer-normalizerstreams      1/1
Running            0          10m
1.6.0.1 releasename-ibm-ea-asm-normalizer-mirrormaker            1/1
Running            0          10m
```

1.6.0.1 The 1.6.0.1 pods are only seen when the topology analytics integration with Agile Service Manager is enabled.

Here is a brief explanation of the output columns.

NAME
Name of the pod. This name comprises the deployment name, the name of the primary image in the pod, and a random sequence of alphanumeric characters if the deployment is not a stateful set.

READY
The number of containers within the pod that are ready/the number of containers in the pod. For example, in the code snippet above, the backup ObjectServer pod, named `ncobackup` is made up of two containers, and both of them are ready.

STATUS
Status of the pod.

RESTARTS
How many times the pod has restarted.

AGE
How long the pod has existed.

Note: The `kubectl get pods` command provides overall status of the installation. To view more detailed status and to monitor the status of connections between components, you must view the Kubernetes logs.

When all of the pods have the status `Running` or `Completed`, then the deployment is ready.

Configuring Installation Parameters for Operations Management on IBM Cloud Private
Find out how to configure the parameters for your Operations Management on IBM Cloud Private installation.

Parameter Name (UI)	Parameter Name in `values.yaml` (CLI)	Description
Helm release name	Non-applicable. Specify as parameter to the `helm install` command with `--name`.	Name your helm release. Make sure that the name starts with a lowercase letter, the body contains only alphanumeric characters and hyphens, and it ends with an alphanumeric character. Helm release name must not exceed 30 characters. It is recommended that the helm release name length is 8 characters or less.
Target namespace	Non-applicable. Specify as parameter to the `helm install` command with `--namespace`.	The namespace to use for your deployment. Use the namespace that you created earlier. See “Preparing your cluster” on page 102. If Agile Service Manager is also installed then Operations Management on IBM Cloud Private and Agile Service Manager must be installed in the same namespace.
Target cluster	Non-applicable. This is the cluster that you are logged in to with `cloudctl`.	The cluster to use for your deployment. Available clusters with the required Kubernetes version are shown for the selected namespace.
Target namespace policies	Read only field that is not used by the CLI installation.	Non-editable. Displays suitable pod security policies for your selected namespace and cluster.
I have read and agreed to the license agreement	global.license	• **CLI**: `accept` or `not accepted` • **UI:** checkbox selected (true) or checkbox unselected (false) Select the checkbox/set to 'accept' if you agree with the license terms. Defaults to 'not accepted'.
Master node	global.cluster.fqdn	Specify the fully qualified domain name (FQDN) of the master node in your cluster. This value is used to construct ingress URLs to access NOI services. The FQDN must match the value of your certificate authority domain, as specified by *cluster_CA_domain* in the cluster's *config.yaml* file. This value must also be mapped to the master node IP address in the /etc/hosts file.

Parameter Name (UI)	Parameter Name in `values.yaml` (CLI)	Description
HTTPS Port	global.ingress.port	The cluster's HTTPS ingress port. Defaults to 443. **Note:** If you are installing on IBM Cloud Private with OpenShift, ports 80 and 443 are used by OpenShift services. For more information, see IBM Cloud Private with OpenShift documentation: Preparing to install IBM Cloud Private with OpenShift.
Image repository	global.image.repository	Docker repository for all the component charts and images. This will be `cluster-CA-domain:8500` if the charts and images are loaded in the IBM Cloud Private local-charts repository, where `cluster-CA-domain` is the name of the certificate authority domain that was used in the `config.yaml` file during IBM Cloud Private installation.
Docker image repository secret	global.image.secret	Specify the name of the Kubernetes secret that contains the credentials to access the Docker registry. This secret must be created before the installation. By default this secret is called `noi-registry-secret`.
ASM release name	global.integrations.asm.releaseName	Must match the helm release name of Agile Service Manager, if it is installed. For more information, see "Installing Agile Service Manager on IBM Cloud Private" on page 131.
Enable ASM integration	global.integrations.asm.enabled	• **CLI**: `true` or `false` • **UI:** checkbox selected (true) or checkbox unselected (false) Set to `true` if you are planning to integrate with the optional Agile Service Manager extension.
Environment size	global.environmentSize	Controls the resource sizes. The value can be either 'size1' or 'size0'. 'size1' must be used for production environments. 'size0' is a minimal specification for evaluation or development purposes. Defaults to 'size0'.
ServiceAccount under which your pods run	global.rbac.serviceAccount	Specify the name of the service account that was created when pod access control was configured, as described here. By default this service account is called `noi-service-account`.

Parameter Name (UI)	Parameter Name in `values.yaml` (CLI)	Description
Create required RBAC RoleBindings	global.rbac.create	• **CLI**: `true` or `false` • **UI:** checkbox selected (true) or checkbox unselected (false) This setting must be set to true only if the installation is being done by a cluster administrator. When set to `true` this setting creates the required service account, roles and role bindings. For installations done as non-cluster administrators, these bindings must be created manually before installation. For more information, see here. Defaults to `false`.
Use existing TLS certificate secrets	global.tls.certificate.useExistingSecret	• **CLI**: `true` or `false` • **UI:** checkbox selected (true) or checkbox unselected (false) Set to `true` if you want to use your own TLS certificate secrets instead of automatically generated ones. For more information, see “Preparing secrets for TLS encryption” on page 105. Defaults to `false`.
Indicate that all password secrets have been created prior to install	global.users.secretsCreatedPreInstall	• **CLI**: `true` or `false` • **UI:** checkbox selected (true) or checkbox unselected (false) Set to `true` if all of the required passwords and secrets have been created manually, and you do not want the installer to create any of the passwords or secrets. If this is set to `false` then the required passwords are randomly created. For more information, see “Configuring passwords and secrets” on page 111. Defaults to `false`.
Enable subchart resource requests	global.resource.requests.enable	• **CLI**: `true` or `false` • **UI:** checkbox selected (true) or checkbox unselected (false) Set to `true` to enable subchart resource requests, or set to `false` to disable the requests so that there is no check on the resources that are required for the release. For example, to specify a small demonstration system, you must disable the subchart resource requests by setting this field to `false`.
Enable anti-affinity	global.antiAffinity.enabled	• **CLI**: `true` or `false` • **UI:** checkbox selected (true) or checkbox unselected (false) Set to `true` to prevent primary and backup server pods from being installed on the same worker node.

Parameter Name (UI)	Parameter Name in `values.yaml` (CLI)	Description
Enable data persistence (Recommended)	global.persistence.enabled	• **CLI**: `true` or `false` • **UI:** checkbox selected (true) or checkbox unselected (false) Set to `true` to ensure that data continues to be available if the pod needs to restart. If set to `false`, data is lost between pod restarts.
Use dynamic provisioning	global.persistence.useDynamicProvisioning	• **CLI**: `true` or `false` • **UI:** checkbox selected (true) or checkbox unselected (false) Set to `true` if you are using a distributed storage solution such as vSphere to ensure that storage volumes are created automatically in the cluster as and when required. Set to `false` if you are using local storage.
Cassandra Data Storage Class	global.persistence.storageClassOption.cassandradata	The persistent volume storage class for the cassandra service. Can be disabled by not specifying any option.
Cassandra Backup Storage Class	global.persistence.storageClassOption.cassandrabak	The persistent volume storage class for the cassandra backup service.
Zookeeper Data Storage Class	global.persistence.storageClassOption.zookeeperdata	The persistent volume storage class for the zookeeper service.
Kafka Data Storage Class	global.persistence.storageClassOption.kafkadata	The persistent volume storage class for the kafka service.
CouchDB Data Storage Class	global.persistence.storageClassOption.couchdbdata	The persistent volume storage class for the couchDB service.
Cassandra Data Storage Size	global.persistence.storageSize.cassandradata	Cassandra data storage size option.
Cassandra Backup Storage Size	global.persistence.storageSize.cassandrabak	Cassandra backup storage size option.
Zookeeper Data Storage Size	global.persistence.storageSize.zookeeperdata	Zookeeper data storage size option.
Kafka Data Storage Size	global.persistence.storageSize.kafkadata	Kafka data storage size option.
CouchDB Data Storage Size	global.persistence.storageSize.couchdbdata	CouchDB data storage size option.
Number of Impact server instances	global.nciservers.replicaCount	Set to the required number of nciserver pods.

Parameter Name (UI)	Parameter Name in `values.yaml` (CLI)	Description
LDAP mode	global.ldapservice.mode	Set to `standalone` to install the built-in standalone LDAP server (openLDAP server) that comes with Netcool Operations Insight. Do not edit any *global.ldapservice.internal.** values. Set to `proxy` to install a proxy LDAP server and connect to your organization's LDAP server. Define the following mandatory users in your LDAP repository: **`icpadmin`**, **`impactadmin`**, and **`unityadmin`**. For more information, see "Creating users on an external LDAP server" on page 210. You must define custom values for all of the other LDAP parameters to match your organization's LDAP server. Advanced: The LDAP proxy can be further configured with the "LDAP Proxy configmap" on page 550. **Note:** If you select `proxy` mode, then regardless of the value of **Indicates that all password secrets have been created prior to install**, you must still manually configure passwords and secrets for `LDAP admin`, `impactadmin` , and `unityadmin` to enable Netcool Operations Insight to connect to your organization's LDAP server. For more information, see "Configuring passwords and secrets" on page 111. Defaults to `standalone`.
LDAP Server port	global.ldapservice.internal.ldapPort	If you set *LDAP mode* to `proxy`, then you must configure the port of your organization's LDAP server. Defaults to 389. Do not edit if *LDAP mode* is `standalone`.
LDAP Server SSL port:	global.ldapservice.internal.ldapSSLPort	If you set *LDAP mode* to `proxy`, then you must configure the SSL port of your organization's LDAP server. Defaults to 636. Do not edit if *LDAP mode* is `standalone`.
LDAP Server URL	global.ldapservice.internal.url	If you set *LDAP mode* to `proxy`, then you must configure the URL of your organization's LDAP server. Use the format ldap://<IP address or hostname>:port. Do not edit if *LDAP mode* is `standalone`.
LDAP Directory Information Tree top entry	global.ldapservice.internal.suffix	If you set *LDAP mode* to `proxy`, then you must configure the top entry in the LDAP Directory Information Tree (DIT). Use the standard domain settings as configured in your organization. Defaults to dc=mycluster,dc=icp. Do not edit if *LDAP mode* is `standalone`.
LDAP base entry	global.ldapservice.internal.baseDN	If you set *LDAP mode* to `proxy`, then you must configure the LDAP base entry by specifying the base distinguished name (base DN). Defaults to dc=mycluster,dc=icp. Do not edit if *LDAP mode* is `standalone`.

Parameter Name (UI)	Parameter Name in `values.yaml` (CLI)	Description
LDAP bind userid	global.ldapservice.internal.bindDN	If you set *LDAP mode* to `proxy`, then you must configure the LDAP bind user identity by specifying the bind distinguished name (bind DN). You must also create a Kubernetes secret that contains password information for your organization's LDAP server, as described in “Configuring passwords and secrets” on page 111. Defaults to cn. Do not edit if *LDAP mode* is `standalone`.
Storage Class Name (In sections 'DB2®', 'Primary Object Server','Backup Object Server', 'Primary Impact Server', 'Impact GUI', Log Analysis', 'LDAP Authentication'.)	`service_name`.pvc.storageClassName where `service_name` is one of (db2, nco_primary, ncobackup, nciserver, impactgui, scala, openldap)	Specify the persistent volume storage class for the service. Defaults to local-storage-`service-name`.
Existing storage claim name (In sections 'DB2', 'Primary Object Server','Backup Object Server', 'Primary Impact Server', 'Impact GUI', Log Analysis', 'LDAP Authentication'.)	`service_name`.pvc.existingNameClaim where `service_name` is one of (db2, nco_primary, ncobackup, nciserver, impactgui, scala, openldap)	Leave empty or specify the persistent volume claim for the service.
PVC Selector label (In sections 'DB2', 'Primary Object Server','Backup Object Server', 'Primary Impact Server', 'Impact GUI', Log Analysis', 'LDAP Authentication'.)	`service_name`.pvc.selector.label where `service_name` is one of (db2, nco_primary, ncobackup, nciserver, impactgui, scala, openldap)	Leave empty if you are using dynamic provisioning. In any other case, selectors can be used refine the binding process.
PVC Selector name (In sections 'DB2', 'Primary Object Server','Backup Object Server', 'Primary Impact Server', 'Impact GUI', Log Analysis', 'LDAP Authentication'.)	`service_name`.pvc.selector.name where `service_name` is one of (db2, nco_primary, ncobackup, nciserver, impactgui, scala, openldap)	Leave empty when using dynamic provisioning. In any other case, selectors can be used refine the binding process.
Storage Size (In sections 'DB2', 'Primary Object Server','Backup Object Server', 'Primary Impact Server', 'Impact GUI', Log Analysis', 'LDAP Authentication'.)	`service_name`.pvc.size where `service_name` is one of (db2, nco_primary, ncobackup, nciserver, impactgui, scala, openldap)	Size allocated for persistent storage in bytes. Do not change. See https://ibm.biz/BdzrBa.

Parameter Name (UI)	Parameter Name in `values.yaml` (CLI)	Description
Temporal Group Policies Deploy First	ibm-hdm-analytics-dev.common.temporalGroupingDeployFirst	Set to true to deploy all Cloud Native Analytics temporal group policies after training. Set to false to manually control deployment of the temporal group policies.
1.6.0.1 ASM status to event enrichment join window	global.ibm-ea-asm-normalizer.joinWindowSize	Maximum waiting time, in minutes, to allow an Object server event to be matched to an ASM status.

Installing Agile Service Manager on IBM Cloud Private
If you want the Service Management extension, then install Agile Service Manager on IBM Cloud Private.

About this task

The steps for installing Agile Service Manager on IBM Cloud Private are described in the Agile Service Manager Knowledge Center.

Note: If Agile Service Manager is required, then it should be installed before Operations Management on IBM Cloud Private is installed. Operations Management on IBM Cloud Private and Agile Service Manager must be installed in the same namespace.

Related tasks

Viewing Kubernetes logs
To view information on the success of the installation process for Netcool Operations Insight on IBM Cloud Private, run the `kubectl logs` command from the command line.

Related information

IBM Netcool Agile Service Manager Version 1.1.3 documentation

Installing Agile Service Manager after Operations Management, or changing its secrets
If you install Agile Service Manager for the first time when Operations Management on IBM Cloud Private 1.6.0.1 is already installed, or the Agile Service Manager secrets are changed, then you must restart some services.

About this task
The usual installation path is for Agile Service Manager to be installed first, followed by Operations Management on IBM Cloud Private. If Operations Management on IBM Cloud Private 1.6.0.1 is already installed, then you must complete some extra steps after your deployment of Agile Service Manager.

Procedure

1. Set your Operations Management on IBM Cloud Private release name in the shell

   ```
   RELEASE_NAME=helm_release_name
   ```

   Where *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private 1.6.0.1 deployment.
2. Restart the UI server.

   ```
   kubectl delete pod -l app.kubernetes.io/component=uiserver,app.kubernetes.io/instance=
   $RELEASE_NAME
   ```
3. Restart Web GUI.

   ```
   kubectl delete pod -l app.kubernetes.io/name=webgui,app.kubernetes.io/instance=$RELEASE_NAME
   ```

Red Hat OpenShift support

The Operations Management on IBM Cloud Private solution can be deployed on IBM Cloud Private with OpenShift.

When you install IBM Cloud Private with OpenShift, IBM Cloud Private provides the IBM Cloud Private experience, management, and operations for applications and uses OpenShift's Kubernetes and Docker registry that is already installed by Red Hat. For more information, see IBM Cloud Private documentation: IBM Cloud Private with OpenShift.

If you are installing on IBM Cloud Private with OpenShift, check that you meet all the prerequisites that are listed in “IBM Cloud Private with OpenShift” on page 99.

Post-installation tasks

Perform the following tasks to configure your Operations Management on IBM Cloud Private release.

Retrieving passwords from secrets

After a successful installation of Operations Management on IBM Cloud Private, passwords can be retrieved from the secrets that contain them.

About this task

To retrieve a password from Operations Management on IBM Cloud Private, use the following procedure.

icpadmin password

```
kubectl get secret helm_release_name-icpadmin-secret -o json -n namespace | grep
ICP_ADMIN_PASSWORD  | cut -d : -f2 | cut -d '"' -f2 | base64 -d;echo
```

sysadmin password

```
kubectl get secret helm_release_name-was-secret -o json -n namespace | grep WAS_PASSWORD | cut -
d : -f2 | cut -d '"' -f2 | base64 -d;echo
```

unity admin password

```
kubectl get secret helm_release_name-la-secret -o json -n namespace | grep UNITY_ADMIN_PASSWORD
| cut -d : -f2 | cut -d '"' -f2 | base64 -d;echo
```

impact admin password

```
 kubectl get secret helm_release_name-impact-secret -o json -n namespace | grep
IMPACT_ADMIN_PASSWORD | cut -d : -f2 | cut -d '"' -f2 | base64 -d;echo
```

omnibus password

```
kubectl get secret helm_release_name-omni-secret -o json -n namespace | grep
OMNIBUS_ROOT_PASSWORD | cut -d : -f2 | cut -d '"' -f2 | base64 -d;echo
```

couchdb password

```
kubectl get secret helm_release_name-couchdb-secret -o json -n namespace | grep password | cut -
d : -f2 | cut -d '"' -f2 | base64 -d;echo
```

LDAP admin password

```
kubectl get secret helm_release_name-ldap-secret -o json -n namespace | grep password | cut -
d : -f2 | cut -d '"' -f2 | base64 -d;echo
LDAP_BIND_PASSWORD
```

Where

- *helm_release_name* is the Operations Management on IBM Cloud Private Helm release name.

- *namespace* is the name of the namespace in which Operations Management on IBM Cloud Private is installed.

Creating users
In Operations Management on IBM Cloud Private, all users must be created either in the openLDAP user interface if you are running on the standalone LDAP pod that comes with Operations Management on IBM Cloud Private, or on your enterprise LDAP server if you are running the LDAP proxy option. Users must not be created in Netcool/OMNIbus, Operations Analytics - Log Analysis, or other Netcool Operations Insight components as these users will not be recognized in Operations Management on IBM Cloud Private.

About this task

For more information on how to create users using the openLDAP user interface, see "Administering users on IBM Cloud Private" on page 207.

Changing passwords and recreating secrets
Changes to any of the passwords used by Operations Management on IBM Cloud Private will require the secrets that use those passwords to be recreated, and the pods that use those secrets to be restarted. Use the following procedure if you need to change any of these passwords.

Procedure

Use this table to help you identify the secret that uses a password, and the pods that use a secret.

Password	Corresponding secret	Dependent pods
smadmin	*helm-release-name*-was-secret	*helm-release-name*-webgui-0 *helm-release-name*- ea-noi-layer-eanoiactionservice *helm-release-name*-ea-noi-layer-eanoigateway *helm-release-name*-ibm-hdm-common-ui-uiserver
impactadmin	*helm-release-name*-impact-secret	*helm-release-name*-impactgui-0 *helm-release-name*- nciserver-0 *helm-release-name*- nciserver-1 *helm-release-name*-webgui-0
unityadmin	*helm-release-name*-la-secret	*helm-release-name*-scala *helm-release-name*-webgui-0
icpadmin	*helm-release-name*-icpadmin-secret	none
OMNIbus root	*helm-release-name*-omni-secret	*helm-release-name*-webgui-0 *helm-release-name*-ea-noi-layer-eanoiactionservice *helm-release-name*-ea-noi-layer-eanoigateway *helm-release-name*-ibm-hdm-analytics-dev-aggregationnormalizerservice

Password	Corresponding secret	Dependent pods
		helm-release-name-ncobackup *helm-release-name*-ncoprimary *helm-release-name*-nciserver-0 *helm-release-name*-nciserver-1
LDAP admin	*helm-release-name*-ldap-secret	*helm-release-name*-openldap-0 *helm-release-name*-impactgui-0 *helm-release-name*-nciserver-0 *helm-release-name*-nciserver-1 *helm-release-name*-ncobackup *helm-release-name*-ncoprimary *helm-release-name*-scala *helm-release-name*-webgui-0
couchdb	*helm-release-name*-couchdb-secret	*helm-release-name*-couchdb *helm-release-name*-ibm-hdm-analytics-dev-aggregationcollaterservice *helm-release-name*-ibm-hdm-analytics-dev-trainer
internal password for inter pod communication	*helm-release-name*-ibm-hdm-common-ui-session-secret	*helm-release-name*-ibm-hdm-common-ui-uiserver
internal password	*helm-release-name*-systemauth-secret	*helm-release-name*-couchdb *helm-release-name*-ibm-hdm-analytics-dev-aggregationcollaterservice *helm-release-name*-ibm-hdm-analytics-dev-trainer
hdm	*helm-release-name*-cassandra-auth-secret	*helm-release-name*-cassandra
redis	*helm-release-name*-ibm-redis-authsecret	*helm-release-name*-ibm-hdm-analytics-dev-collater-aggregationservice *helm-release-name*-ibm-hdm-analytics-dev-dedup-aggregationservice
kafka	*helm-release-name*-kafka-admin-secret	*helm-release-name*-ibm-hdm-analytics-dev-archivingservice *helm-release-name*-ibm-hdm-analytics-dev-collater-aggregationservice

Password	Corresponding secret	Dependent pods
		helm-release-name-ibm-hdm-analytics-dev-dedup-aggregationservice *helm-release-name*-ibm-hdm-analytics-dev-inferenceservice *helm-release-name*-ibm-hdm-analytics-dev-ingestionservice *helm-release-name*-ibm-hdm-analytics-dev-normalizer-aggregationservice
admin	*helm-release-name*-kafka-client-secret	*helm-release-name*-ibm-hdm-analytics-dev-archivingservice *helm-release-name*-ibm-hdm-analytics-dev-collater-aggregationservice *helm-release-name*-ibm-hdm-analytics-dev-dedup-aggregationservice *helm-release-name*-ibm-hdm-analytics-dev-inferenceservice *helm-release-name*-ibm-hdm-analytics-dev-ingestionservice *helm-release-name*-ibm-hdm-analytics-dev-normalizer-aggregationservice

Where *helm_release_name* is the helm release name of the Operations Management on IBM Cloud Private installation.

To change a password use the following procedure.

1. Change the password that you wish to change.
2. Use the table at the start of this topic to find the secret that corresponds to the password that has been changed, and delete this secret.

   ```
   kubectl delete secret secretname --namespace namespace
   ```

   Where

   - *secretname* is the name of the secret to be recreated.
   - *namespace* is the name of the namespace in which the secret to be recreated exists.
3. Recreate the secret with the desired new password. See “Configuring passwords and secrets” on page 111 for instructions on how to create the required secret.
4. Use the table at the start of this topic to find which pods depend on the secret that you have recreated and will require restarting.
5. Restart the required pods using

   ```
   kubectl delete pod podname
   ```

   Where *podname* is the name of the pod that requires restarting.

Customizing applications using config maps

Customize the one or more applications by editing the relevant config map and restarting the associated pod.

Procedure

1. Log into the IBM Cloud Private GUI.
2. Click the Menu icon at the top left of the screen and select **Configuration** > **ConfigMaps**.

   The **ConfigMaps** list displays config maps for all releases across all namespaces.
3. Filter the list in one or more of the following ways:
   - Filter by namespace: at the top right click **All namespaces** and select the namespace where your release is installed.
   - Type the initial letters of your release name in the filter field above the list.

   The config maps are listed using a convention similar to the following, where *helm_release_name* is the name of the release of interest.

Table 30. Config maps in an Operations Management on IBM Cloud Private release

Component or Capability	Application	Pod name	Name of config map	More details
Netcool/ OMNIbus	Primary ObjectServer	`ncoprimary`	*helm_release_name*`-objserv-agg-primary-config`	"Primary Netcool/ OMNIbus ObjectServer configmap" on page 543
	Backup ObjectServer	`ncobackup`	*helm_release_name*`-objserv-agg-backup-config`	"Backup Netcool/ OMNIbus ObjectServer configmap" on page 544
Netcool/Impact	Primary Netcool/ Impact core server	`nciserver`	*helm_release_name*`-nciserver-config`	"Netcool/Impact core server configmap" on page 546
	Netcool/Impact GUI server	`impactgui`	*helm_release_name*`-impactgui-config`	"Netcool/Impact GUI server configmap" on page 548
Event Search	Gateway for Message Bus	`scala`	*helm_release_name*`-scala-config`	"Gateway for Message Bus configmap " on page 551

Table 30. Config maps in an Operations Management on IBM Cloud Private release (continued)

Component or Capability	Application	Pod name	Name of config map	More details
Dashboard Application Services Hub GUIs	Dashboard Application Services Hub GUI: • Event Viewer • Event Analytics GUIs • Event Search GUIs	webgui	*helm_release_name*-webgui-init-config	"Dashboard Application Services Hub configmap " on page 550
Proxy	Proxy	proxy	*helm_release_name*-proxy-config	"Proxy configmap" on page 549
LDAP Proxy	LDAP Proxy	openldap	*helm_release_name*-ldap-proxy-config	"LDAP Proxy configmap" on page 550
Agile Service Manager	Configures integration with ASM	asm-ui	*helm_release_name*-asm-ui-config	"ASM-UI configmap" on page 553
Configuration Share	Configures file sharing between pods. *Must not be edited*	configuration-share	*helm_release_name*-configuration-share	"Configuration share configmap" on page 553
Cloud Native Analytics	Created during helm install. *Must not be edited*	cassandra	*helm_release_name*-cassandra-bootstrap-config	"Cassandra configmap" on page 553
	Created during helm install. *Must not be edited*	couchdb	*helm_release_name*-couchdb	"CouchDB configmap" on page 553
	Created during helm install. *Must not be edited*	ea-noi-layer-eanoigateway	*helm_release_name*-ea-noi-layer-eanoigateway	"Cloud Native Analytics gateway configmap" on page 553
	Created during helm install. *Must not be edited*	kafka	*helm_release_name*-kafka	"Kafka configmap" on page 553
	Created during helm install. *Must not be edited*	zookeeper	*helm_release_name*-zookeeper	"Zookeeper configmap" on page 553

4. Edit the config map of interest.
   a) Identify the row containing the config map to be edited.

b) Within that row click the **Open and close a list of options** icon under the **Action** column on the right-hand side of the screen.

c) Click **Edit**.

d) Make desired changes to the config map.

For information on the structure and content of each config map, refer to the "Configmap reference" on page 543.

When editing the config map, be careful not to introduce any extraneous whitespace such as tabs or spaces, as this can cause the container to fail.

Note: It is also possible to edit the config map from the command line using the following command:

```
kubectl edit configmap name_of_configmap -n namespace
```

Where:

- *name_of_configmap* is the name of the config map.
- *namespace* is the namespace in which the container and its associated config map is located.

For example:

```
kubectl edit configmap myrelease-objserv-agg-primary-config -n noi
```

5. Delete the relevant pod or pods.

The pod automatically restarts following pod deletion.

Related reference

Configmap reference

This section lists the pods that have configmaps and explains which parameters you can configure in each configmap.

Loading a TLS certificate into Kubernetes

In a Netcool Operations Insight Kubernetes cluster environment, Operations Management components such as Web GUI, WebSphere Application Server, and Netcool/Impact are known as services. In order for your users to be able to use these Netcool Operations Insight services, the URL for each service requires its own Transport Layer Security (TLS) certificate. Operations Management on IBM Cloud Private ships with automatically generated certificates signed by the cluster Certificate Authority (CA).

About this task

For more information about cert-manager certificates, see Creating IBM Cloud Private Certificate manager (cert-manager) certificates.

A Kubernetes ingress is a collection of rules that can be configured to give services externally reachable URLs. Each Operations Management component requires its own ingress. Each of these ingresses has an associated secret for TLS encryption. Here is an example of the default TLS secrets:

```
kubectl get secret
NAME                                                    TYPE                  DATA   AGE
release_name-impact-tls-secret                          kubernetes.io/tls      2      1h
release_name-nci-0-tls-secret                           kubernetes.io/tls      2      1h
release_name-nci-1-tls-secret                           kubernetes.io/tls      2      1h
release_name-netcool-tls-secret                         kubernetes.io/tls      2      1h
release_name-proxy-tls-secret                           kubernetes.io/tls      2      1h
release_name-scala-tls-secret                           kubernetes.io/tls      2      1h
release_name-was-tls-secret                             kubernetes.io/tls      2      1h
release_name-ibm-hdm-common-ui-tls-secret               kubernetes.io/tls      2      1h
```

You might want to use your own certificates for TLS encryption. If so, then you will have to replace each of the default secrets using the following procedure.

Procedure

1. Generate your own certificate with the fully qualified domain name set to *ingress.release_name.fqdn*

   ```
   openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout /tmp/tls.key
   -out /tmp/tls.crt -subj "/CN=ingress.release_name.fqdn"
   ```

   Where

   - *ingress* is the name of the ingress for the relevant Operations Management component. Run this command for each of the following ingress values:
     - `impact`
     - `nci-x` where *x* is for each of the impact instances from 0 - (number of Impact server instances-1). For example if you have 4 impact servers, run for each of these values: nci-0,nci-1,nci-2,nci-3.
     - `netcool`
     - `proxy`
     - `scala`
     - `was`
     - `ibm-hdm-common-ui`
   - *release_name* is name of the Helm release. Ensure that this is the same release name that you will use when you install your Operations Management on IBM Cloud Private deployment, as described in "Installing Operations Management on IBM Cloud Private" on page 119.
   - *fqdn* is the certificate authority (CA) domain to be set in the **Master node FQDN (Fully Qualified Domain Name)** field when you install your Operations Management on IBM Cloud Private deployment, as described in "Installing Operations Management on IBM Cloud Private" on page 119.
2. Delete the existing secret for the ingress that was used in the previous step.

   ```
   kubectl delete secret release_name-ingress-tls-secret
   ```

   Where

   - *ingress* is the name of the ingress for the relevant Operations Management component. Run this command for each of the following ingress values:
     - `impact`
     - `nci-` *x* where *x* is for each of the impact instances from 0 - (number of Impact server instances-1). For example, for 4 impact servers, run for each of these values: nci-0,nci-1,nci-2,nci-3.
     - `netcool`
     - `proxy`
     - `scala`
     - `was`
     - `ibm-hdm-common-ui`
   - *release_name* is name of the Helm release. Ensure that this is the same release name that you will use when you install your Operations Management on IBM Cloud Private deployment, as described in "Installing Operations Management on IBM Cloud Private" on page 119.
3. Check that the URL for the ingress used in the previous steps no longer works.

   Try to access the following URL. You should be unable to connect.

   ```
   https://ingress.release_name.fqdn/ibm/console
   ```

   Where

   - *ingress* is the name of the ingress for the relevant Operations Management component. Run this command for each of the following ingress values:

  - `impact`
  - `nci-x` where *x* is for each of the impact instances from 0 - (number of Impact server instances-1). For example, for 4 impact servers, run for each of these values: nci-0,nci-1,nci-2,nci-3.
  - `netcool`
  - `scala`
  - `was`
  - `ibm-hdm-common-ui`
- *release_name* is name of the Helm release. Ensure that this is the same release name that you will use when you install your Operations Management on IBM Cloud Private deployment, as described in "Installing Operations Management on IBM Cloud Private" on page 119.
- *fqdn* is the certificate authority (CA) domain to be set in the **Master node FQDN (Fully Qualified Domain Name)** field when you install your Operations Management on IBM Cloud Private deployment, as described in "Installing Operations Management on IBM Cloud Private" on page 119.

4. Recreate the secret for the ingress used in the previous steps.

```
kubectl create secret tls release_name-ingress-tls-secret
--cert=./certificate.pem --key=./key.pem [--namespace namespace]
```

   Where

   - *ingress* is the name of the ingress for the relevant Operations Management component. Run this command for each of the following ingress values:
     - `impact`
     - `nci-x` where *x* is for each of the impact instances from 0 - (number of Impact server instances-1). For example, for 4 impact servers, run for each of these values: nci-0,nci-1,nci-2,nci-3.
     - `netcool`
     - `proxy`
     - `scala`
     - `was`
     - `ibm-hdm-common-ui`
   - *release_name* is name of the Helm release. Ensure that this is the same release name that you will use when you install your Operations Management on IBM Cloud Private deployment, as described in "Installing Operations Management on IBM Cloud Private" on page 119.
   - *namespace* is the namespace in which Operations Management on IBM Cloud Private is installed.

5. Check that the URL for the ingress now works.

   Try to access the following URL. You should be able to connect.

```
https://ingress.release_name.fqdn/ibm/console
```

   Where

   - *ingress* is the name of the ingress for the relevant Operations Management component. Run this command for each of the following ingress values:
     - `impact`
     - `nci-x` where *x* is for each of the impact instances from 0 - (number of Impact server instances-1). For example, for 4 impact servers, run for each of these values: nci-0,nci-1,nci-2,nci-3. Use `nameserver/services` instead of `ibm/console`.
     - `netcool`
     - `scala` Use `Unity` instead of `ibm/console`.
     - `was`
     - `ibm-hdm-common-ui`

- *release_name* is name of the Helm release. Ensure that this is the same release name that you will use when you install your Operations Management on IBM Cloud Private deployment, as described in "Installing Operations Management on IBM Cloud Private" on page 119.
- *fqdn* is the certificate authority (CA) domain to be set in the **Master node FQDN (Fully Qualified Domain Name)** field when you install your Operations Management on IBM Cloud Private deployment, as described in "Installing Operations Management on IBM Cloud Private" on page 119. If you are installing on IBM Cloud Private with OpenShift, use the `ingress_https_port` 3443 for this value.

Related tasks

Getting started with Operations Management on IBM Cloud Private
Log in to the components of Operations Management, such as Web GUI and Operations Analytics - Log Analysis.

Scaling the Netcool/Impact service

The number of IBM Tivoli Netcool/Impact core server pods is configured at installation through the *Number of Impact server instances* parameter in the IBM Cloud Private GUI, and corresponds to the value of the **`global.nciservers.replicaCount`** parameter in the `values.yaml` file. The number of Netcool/Impact core servers can be scaled up or down as required.

For more information, see "Installing Operations Management on IBM Cloud Private" on page 119.

Scaling the Impact Service with the command line

The number of Netcool/Impact core servers can be scaled up or down as required by using the helm upgrade command.

About this task

Scale the number of Netcool/Impact pods up or down by using the following procedure.

Procedure

1. If *global.tls.certificate.useExistingSecret* is set to true, then when scaling up you will need to create secrets for each of the new Netcool/Impact core server pods. See "Preparing secrets for TLS encryption" on page 105
2. Edit the `values.yaml` file and set the value of *global.nciservers.replicaCount* to the required number of Netcool/Impact core server pods. This number must be at least 1 and be less than the number of worker nodes. For example, to scale up or down to 4 Netcool/Impact core server pods, add the following *replicaCount* for *nciservers* to the global section of the `values.yaml` file:

   ```
   global:
     nciservers:
       # Expose the number of nciservers we need.
       replicaCount: 4
   ```

3. Run helm upgrade with the updated `values.yaml` file to create or delete Netcool/Impact core server pods to meet the new *replicaCount* by using the following command:

   ```
   helm upgrade helm_release_name helm_chart_repo --values values.yaml --tls
   ```

   Where

   - *helm_release_name* is the Operations Management on IBM Cloud Private Helm release name.
   - *helm_chart_repo* is the name of the helm chart repository for Operations Management on IBM Cloud Private.
4. When the upgrade is completed, verify that the expected number of `nciserver` pods are running by using the following command:

   ```
   kubectl get pods -n namespace | grep nciserver
   ```

Where *namespace* is the name of the namespace in which Operations Management on IBM Cloud Private is installed.

5. Delete all of the `nciserver` pods, as old pods are unaware of `nciserver` pod additions or deletions. New `nciserver` pods are then created, in line with the updated `values.yaml` file. Delete each impact server pod in turn so that the impact service is maintained, by using the following command:

   ```
   kubectl delete pod pod_name -n namespace
   ```

   Where

   - *namespace* is the name of the namespace in which Operations Management on IBM Cloud Private is installed.
   - *pod_name* is the name of the `nciserver` pod to delete

6. Verify that the expected number of `nciserver` pods are running by using the following command:

   ```
   kubectl get pods -n namespace | grep nciserver
   ```

   Where *namespace* is the name of the namespace in which Operations Management on IBM Cloud Private is installed.

Scaling the Impact Service with the ICP console

The number of Netcool/Impact core servers can be scaled up or down as required by using the ICP console.

About this task

Scale the number of Netcool/Impact pods up or down by using the following procedure.

Procedure

1. If *Use existing TLS certificate secrets* is set to true, then when scaling up you will need to create secrets for each of the newNetcool/Impact core server pods. See “Preparing secrets for TLS encryption” on page 105
2. Open the IBM Cloud Private GUI by pointing your internet browser to a URL similar to

   ```
   https://master_node_ip_or_hostname:8443/oidc/login.jsp
   ```

   Where *master_node_ip_or_hostname* is the IP address or host name of your cluster's master node.
3. Log in to the IBM Cloud Private GUI.
4. On the left side of the page, go to **Workloads** > **Helm Releases** > **Select release name** > **Upgrade** > **Select upgrade version**. Select the current release version.
5. Go to **All Parameters** > **Number of Impact server instances** and set this field to the required number of impact servers.
6. Click **Upgrade.**
7. When the upgrade completes, go to **Workloads** > **StatefulSets** > ***<helm_release_name>-nciserver*** and verify that the expected number of `nciserver` pods are running.
8. Click the menu next to each of the `nciserver` pods and select **Remove**, as old pods are unaware of `nciserver` pod additions or deletions. Remove each impact server pod in turn so that the impact service is maintained.
9. New `nciserver` pods are then created, in line with the updated **Number of Impact server instances** parameter.

Exposing an ObjectServer port in an Operations Management on IBM Cloud Private deployment

About this task

Use this information to learn how to expose an ObjectServer port, and then verify that the port is exposed by sending events to it via a REST API.

Procedure

1. Find the configmap for the primary ObjectServer, as in the following example:

```
kubectl get configmap |grep objserv-agg-primary
m101-objserv-agg-primary-config   2      5h24m
```

2. Edit the primary ObjectServer's configmap.

```
kubectl edit configmap helm_release_name-objserv-agg-primary-config
```

   where *helm_release_name* is the Operations Management on IBM Cloud Private helm release name, and include the following:

```
data:
  agg-p-props-append: |
    NRestOS.Enable: TRUE
    NHttpd.EnableHTTP: TRUE
    NHttpd.ListeningPort: 8080
```

3. Restart the *ncoprimary* pod.

```
kubectl delete pod helm_release_name-ncoprimary-0
```

   Where *helm_release_name* is the Operations Management on IBM Cloud Private Helm release name.

4. Expose port 8080.

```
kubectl expose po helm_release_name-ncoprimary-0 --port=8080 --type=NodePort --name=objserve-
http-client-external-portforward
```

   Where *helm_release_name* is the Operations Management on IBM Cloud Private Helm release name

5. Check that port 8080 is exposed.

```
kubectl get svc | grep objserve-http-client-external-portforward
objserve-http-client-external-portforward                NodePort    10.0.43.199
<none>        8080:31729/TCP                                        169m
```

6. Send events to the ObjectServer with CURL.

```
$ curl -X POST -v -u root:object_server_password -H "Accept: application/json" -H "Content-
Type: application/json" -d @server1.json http://objserve-http-client-external-portforward.
{master server}:31729/objectserver/restapi/alerts/status;
```

   Where *object_server_password* is the password for the ObjectServer. See “Retrieving passwords from secrets” on page 132.

7. Check that the ObjectServer received the events.

```
http://objserve-http-client-external-portforward.{master server}:31729/objectserver/restapi/
alerts/status
```

Uninstalling on IBM Cloud Private

Uninstall Operations Management on IBM Cloud Private by performing the following steps.

About this task

This procedure uninstalls the deployed version of NOI, but does not remove the load images or charts.

Note: When you uninstall a release, any unused docker images that were used as part of the installation remain on the master node and on a local repository on one or more worker nodes. From time to time, Kubernetes removes unused images as part of its garbage collection process. For more information, see Configuring kubelet Garbage Collection.

Procedure

1. Run the following Helm command to determine which releases of Operations Management on IBM Cloud Private are installed.

   ```
   helm ls --tls
   ```

2. Delete the desired NOI releases identified in the previous step by running the following Helm command against each release that is to be uninstalled in turn.

   ```
   helm delete --purge --tls release_name
   ```

   Where *release_name* is the name of the release to be deleted.

3. Run the following command to verify that the desired NOI release has been removed and that its' pods are no longer running.

   ```
   kubectl get pods -n namespace | egrep release_name
   ```

   Where *namespace* is the name of the namespace where you installed Operations Management on IBM Cloud Private.

   If any pods are still running, manually delete them by running the following command:

   ```
   kubectl delete pod <pod_name>
   ```

4. Remove all unwanted secrets by running the following command:

   ```
   kubectl delete secret secret-name -n namespace
   ```

   Where *secret-name* is the name of the secret that you do not want to persist.

5. Clean up your local storage.

6. Persistent Volume Claims (PVCs) are not deleted from StatefulSet pods. These can be retained and will be successfully bound in subsequent deployments if desired. To see which PVCs are remaining for the release which is being deleted, use the following command:

   ```
   kubectl get pvc -n namespace | grep release_name
   ```

7. Use the following command to delete PVCs that are not required.

   ```
   kubectl delete pvc pvc_name -n namespace
   ```

   Where *pvc_name* is the name of the PVC to be deleted.

8. Repeat the previous steps until all of the desired releases are deleted.

What to do next

Reinstall Operations Management on IBM Cloud Private. For more information, see “Installing on IBM Cloud Private” on page 119.

Related information

Helm documentation: Helm commandsClick here to view information on all Helm commands.

Chapter 4. Upgrading Netcool Operations Insight

Plan the upgrade and complete any pre-upgrade tasks before upgrading Netcool Operations Insight.

Upgrading on premises

Follow these instructions to upgrade Netcool Operations Insight on premises.

Before you begin
Back up all products and components in the environment.
Related concepts
Connections in the Server Editor
Related tasks
Restarting the Web GUI server
Configuring Reporting Services for Network Manager
Related reference
Web GUI server.init file
Related information
Gateway for Message Bus documentation
Operations Analytics - Log Analysis Welcome pageWithin the Operations Analytics - Log Analysis Welcome page, proceed as follows: (1) Select the version of interest. (2) For information on backing up and restoring Operations Analytics - Log Analysis data, or on installing Operations Analytics - Log Analysis, perform a relevant search.

Updated versions in the V1.6.0.1 release

In order to perform the most recent upgrade of Netcool Operations Insight, from V1.5.0.1 to V1.6.0.1, you must download the software described in this table, from either Passport Advantage or from Fix Central.

Note: If a table cell in either the **Download from Passport Advantage** column or the **Download from Fix Central** column is empty then there is nothing to download from that location for that particular product or component.

Table 31. Software downloads for upgrade from Netcool Operations Insight from V1.5.0.1 to V1.6.0.1

Product or component	Target release	Download from Passport Advantage	Download from Fix Central
IBM Tivoli Netcool/ OMNIbus core components	V8.1.0.21	CJ5N1EN	V8.1 Fix Pack 21
Tivoli Netcool/OMNIbus Web GUI	V8.1.0.17		V8.1 Fix Pack 17
IBM Tivoli Netcool/Impact	V7.1.0.17	CJ5N2EN	V7.1 Fix Pack 17
Db2	V11.1	CJ3INML	
Operations Analytics - Log Analysis	V1.3.6	CJ5N3EN	

Table 31. Software downloads for upgrade from Netcool Operations Insight from V1.5.0.1 to V1.6.0.1 (continued)			
Product or component	**Target release**	**Download from Passport Advantage**	**Download from Fix Central**
IBM Tivoli Network Manager IP Edition	V4.2.0.7	CJ5N4EN	V4.2 Fix Pack 7
Network Health Dashboard	V4.2.0.6	CJ0S2EN	V4.2 Fix Pack 6
IBM Tivoli Netcool Configuration Manager	V6.4.2.8	CJ5N5EN	V6.4.2 Fix Pack 8
IBM Network Performance Insight	V1.3.1	CJ5N6EN	
IBM Agile Service Manager	V1.1.6	CJ5N7EN	
IBM Agile Service Manager Observers	V1.1.6	CJ5N8EN	
Jazz for Service Management	V1.1.3.5	CJ49VML	

Updated versions in the V1.6.0 release

In order to perform the most recent upgrade of Netcool Operations Insight, from V1.5.0.1 to V1.6.0, you must download the software described in this table, from either Passport Advantage or from Fix Central.

Note: If a table cell in either the **Download from Passport Advantage** column or the **Download from Fix Central** column is empty then there is nothing to download from that location for that particular product or component.

Table 32. Software downloads for upgrade from Netcool Operations Insight from V1.5.0.1 to V1.6.0.			
Product or component	**Target release**	**Download from Passport Advantage**	**Download from Fix Central**
IBM Tivoli Netcool/ OMNIbus core components	V8.1.0.19	CJ5N1EN	V8.1 Fix Pack 21
Tivoli Netcool/OMNIbus Web GUI	V8.1.0.16		V8.1 Fix Pack 17
IBM Tivoli Netcool/Impact	V7.1.0.16	CJ5N2EN	V7.1 Fix Pack 17
Db2	V11.1	CJ3INML	
Operations Analytics - Log Analysis	V1.3.5.3	CJ5N3EN	
IBM Tivoli Network Manager IP Edition	V4.2.0.7	CJ5N4EN	V4.2 Fix Pack 7

Table 32. Software downloads for upgrade from Netcool Operations Insight from V1.5.0.1 to V1.6.0. (continued)			
Product or component	**Target release**	**Download from Passport Advantage**	**Download from Fix Central**
Network Health Dashboard	V4.2.0.6	CJ0S2EN	V4.2 Fix Pack 6
IBM Tivoli Netcool Configuration Manager	V6.4.2.8	CJ5N5EN	V6.4.2 Fix Pack 8
IBM Network Performance Insight	V1.3.1	CJ5N6EN	
IBM Agile Service Manager	V1.1.5	CJ5N7EN	
IBM Agile Service Manager Observers	V1.1.5	CJ5N8EN	
Jazz for Service Management	V1.1.3.3	CJ49VML	

Downloading product and components

In order to upgrade from V1.5.0 to V1.6.0.1, you must download software from Passport Advantage and Fix Central.

About this task

This scenario describes how to upgrade Netcool Operations Insight from V1.5.0 to the current version, V1.6.0.1. The scenario assumes that Netcool Operations Insight is deployed as shown in the simplified architecture in the following figure. Depending on how your Netcool Operations Insight system is deployed, you will need to download the software and run the upgrade on different servers.

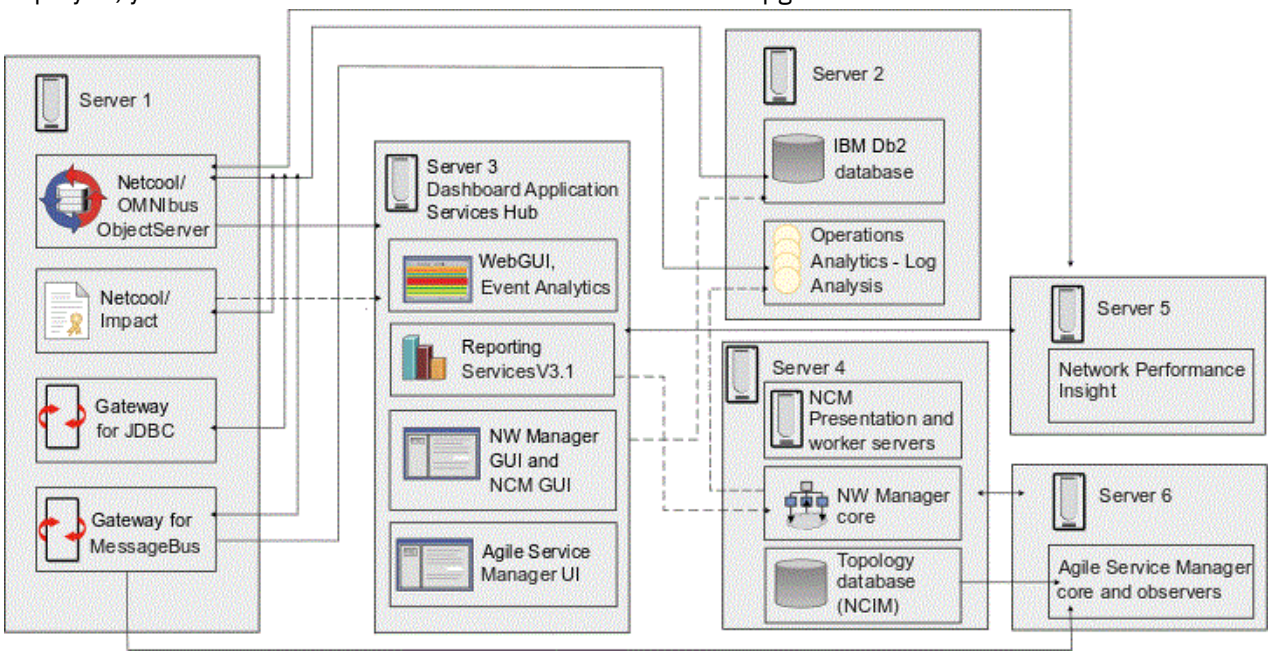

Figure 8. Simplified architecture for the upgrade scenario

Procedure

1. Refer to the following web page for information on where to obtain downloads for each product and component.

   https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Netcool%20OMNIbus/page/Upgrading

   Note: You will need to log into IBM Passport Advantage or Fix Central, as appropriate, to download the software.
2. Download the software to the servers listed in the table.

Table 33. Which server to download software to

If your current Netcool Operations Insight installation includes...	Then download any software related to the following products...	To the following server	For more details, see...
Netcool Operations Insight base solution only	Netcool/OMNIbus Netcool/Impact	Server 1	"Applying the latest fix packs" on page 148
	Operations Analytics - Log Analysis	Server 2	"Applying the latest fix packs" on page 148
	Jazz for Service Management WebSphere Application Server	Server 3	"Applying the latest fix packs" on page 148
Network Management for Operations Insight solution extension	Network Manager core components	Server 4	
	Netcool Configuration Manager core components	Server 4	
Performance Management for Operations Insight solution extension	Network Performance Insight	Server 5	"Upgrading Network Performance Insight" on page 149
Service Management for Operations Insight solution extension	Agile Service Manager Base	Server 3 Server 6	"Installing Agile Service Manager" on page 96
	Agile Service Manager Observers	Server 6	

Related information

Passport AdvantageClick here to go to the IBM Passport Advantage website.

Fix CentralClick here to go to the Fix Central website.

Applying the latest fix packs

Apply any latest available fix packs to upgrade to the latest version of Netcool Operations Insight.

About this task

Fix packs can be full image fix packs containing the full product image, or upgrade fix packs, containing just the code for fix updates from the last release. Full image fix packs are made available on Passport

Advantage. Upgrade fix packs are made available on Fix Central. For a list of any full image fix packs required for upgrade to the latest version of NOI, see the following link:

https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Netcool%20OMNIbus/page/Upgrading

Procedure

1. For each fix pack upgrade, start Installation Manager and configure it to point to the `repository.config` file for the fix pack.
2. In the main Installation Manager window, click **Update** and complete wizard instructions similar to the following:
   a) In the **Update Packages** tab, select the product group to find related update packages, and click **Next**. A list of the available update packages displays.
   b) From the list of available update packages, select the relevant version, and click **Next**.
   c) In the **Licenses** tab, review the licenses. Select **I accept the terms in the license agreements** and click **Next**.
   d) In the **Features** tab, select the features for your update package, and click **Next**.
   e) Complete the configuration details, and click **Next**.
   f) In the **Summary** tab, review summary details. If you need to change any detail click **Back**, but if you are happy with summary details click **Update** and wait for the installation of the update package to complete.
   g) When the installation of the update package completes, the window updates with details of the installation. Click **Finish**.

Related information

Fix CentralClick here to go to the Fix Central website.

Upgrading Network Performance Insight

Upgrade Network Performance Insight to the latest release.

About this task

Network Performance Insight has an built-in mechanism for performing product upgrades.

Procedure

1. On the server where Network Performance Insight is installed (in our example, server 5), extract the Network Performance Insight eAssembly archive.
2. See the Network Performance Insight documentation at https://www.ibm.com/support/knowledgecenter/SSCVHB_1.3.1/npi_kc_welcome.html for more information on how to perform the upgrade.

Installing Agile Service Manager

Install the latest version of Agile Service Manager.

Procedure

1. To install the Agile Service Manager Core services, and Observers, proceed as follows:
   a) On the server where Agile Service Manager Core services are installed (in our example, server 6), extract the Agile Service Manager Base and Observer eAssembly archives.
   b) Follow the instructions in the Agile Service Manager to complete the installation.
2. To install the Agile Service Manager UI, proceed as follows:
   a) On the server where Dashboard Application Services Hub is installed (in our example, server 3), extract the Agile Service Manager Base eAssembly archive.

b) Start Installation Manager and configure it to point to the following repository files: `repository.config` file for Agile Service Manager

c) See the Agile Service Manager documentation for more information on how to perform the installation, and perform post-installation configuration tasks, such as configuring the Gateway for Message Bus to support the Agile Service Manager Event Observer.

Related information

Agile Service Manager Knowledge CenterSearch for the Installing section within the Agile Service Manager Knowledge Center.

Migrating data for Operations Management on IBM Cloud Private

Upgrading is only supported for Operations Management on IBM Cloud Private version 1.6.0 to version 1.6.0.1. For other versions follow these instructions to migrate data from one release of Operations Management on IBM Cloud Private to another release.

Before you begin

Back up all products and components in the environment. Then install another release side-by-side and migrate the data.

Migrating data for Netcool/OMNIbus

Follow these instructions to migrate data from one release of IBM Tivoli Netcool/OMNIbus to another release.

Before you begin

Back up all products and components in the environment.

About this task

The following steps provide an example of how to migrate your data. For more information, see https://www.ibm.com/support/knowledgecenter/en/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/omnibus/wip/install/concept/omn_con_cnfp_importingexportobjserv.html.

Procedure

1. Export data from the previous release of Netcool/OMNIbus. Run the following commands:

```
$OMNIHOME/bin/nco_osreport -list -file confpack.list -server NCOMS -user root
$OMNIHOME/bin/nco_osreport -export -file confpack.list -package testpack.jar -user root
```

2. On the destination object server, import data from the previous release to the latest release of Netcool/OMNIbus. Run the following command:

```
$OMNIHOME/bin/nco_osreport -import -package testpack.jar -server NCOMS -user root
```

3. Verify that the data migration was successful:

   a. Log in to the administrator GUI (nco_config) of NCOMS (the default interface):

   b. Add multiple events to the object server.

   c. Generate an Osreport from NCOMS by running the following command:

```
/$OMNIHOME/bin/nco_osreport -server NCOMS -user root -password ""
```

   d. Copy the `alertsdata.sql` file generated from **`nco_osreport`** to IBM Cloud Private and load the file into the object server pod by running the following command:

```
kubectl cp alertsdata.sql default/my_object_server_pod:$OMNIHOME/alertsdata.sql
```

e. Run the following command on IBM Cloud Private:

```
/$OMNIHOME/bin/nco_sql -server AGG_P -user root -password '' < alertsdata.sql
```

f. Verify that the events are imported to AGG_P.

Migrating data for Web GUI

Follow these instructions to migrate data from one release of IBM Netcool/OMNIbus Web GUI to another release.

Before you begin

Back up all products and components in the environment.

About this task

The following steps provide an example of how to migrate your data. For more information, see https://www.ibm.com/support/knowledgecenter/en/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/webtop/wip/task/web_upg_upgrading.html.

Procedure

1. Export data from the previous release of Web GUI.
2. Copy the `data.zip` file to the Web GUI pod by running the following command:

   ```
   kubectl cp data.zip default/{webgui_pod}:$JazzSM_HOME/ui/input
   ```

3. Access the Web GUI pod by running the following command:

   ```
   kubectl exec -ti my_webgui_pod bash
   ```

   Where *my_webgui_pod* is the name of your pod.
4. Update the `/home/netcool/app/gui/omnibus_webgui/integration/plugins/OMNIbusWebGUI_DASH_clone.properties` properties file.
5. Run the following `consolecli.sh` script in the Web GUI pod to import data from the previous release:

   ```
   $JAZZSM_HOME/ui/bin/consolecli.sh Import --username my-username --password
   my-password
   --settingFile
   /home/netcool/app/gui/omnibus_webgui/integration/plugins/OMNIbusWebGUI_settings.properties
   ```

6. Restart Web GUI by deleting the Web GUI pod with the **`kubectl delete`** command:

   ```
   kubectl delete my_webgui_pod
   ```

   Where *my_webgui_pod* is the name of your pod.
7. Check that the pages are imported correctly.

Migrating data for Netcool/Impact

Follow these instructions to migrate data from one release of IBM Tivoli Netcool/Impact to another release.

Before you begin

Back up all products and components in the environment.

About this task

The following steps provide an example of how to migrate your data. For more information, see https://www.ibm.com/support/knowledgecenter/en/SSSHYH_7.1.0.16/com.ibm.netcoolimpact.doc/admin/imag_tools_nci_export_c.html.

Procedure

1. Export data from the previous release of Netcool/Impact. Run the following command:

```
$IMPACT_HOME/bin/nci_export NCI_1 --project Project_1 /tmp/Project_1
```

2. Compress the export as the `Project_1.tar` file. Copy the compressed file to the latest release of Netcool/Impact.
3. Bring down the secondary Impact Server, leaving only the primary server running. Run the following command:

```
kubectl scale --replicas=1 Statefulset/my_release_name-nciserver
```

   Where *my_release_name* is your release name.
4. Copy the `Project_1.tar` file into the Netcool/Impact pod by running the following command:

```
kubectl cp Project_1.tar default/my_nciserver_pod :/tmp/Project_1.tar
```

   Where *my_nciserver_pod* is your Impact Server pod.
5. Log in to the Netcool/Impact pod by running the following command:

```
kubectl exec -ti ll03-master42-nciserver-0 bash
```

6. Extract the compressed file:

```
tar -xvf /tmp/Project_1.tar
```

7. Import the data:

```
$IMPACT_HOME/bin/nci_import NCI_0 /tmp/Project_1
```

   Note: If there are errors complaining about lock files, check the `$IMPACT_HOME/etc/*lock*` directory to see what user is locked. Log in to the Netcool/Impact UI and click **Global** > **Clear my locks**. Clear locks for each user.
8. Bring up the secondary Impact Server by running the following command:

```
 kubectl scale --replicas=2 Statefulset/my_release_name-nciserver
```

9. Verify that the data migration was successful. Log in to the Netcool/Impact UI and check that the polices and services are imported correctly.

Migrating data for Event Analytics

Follow these instructions to migrate from Event Analytics to Cloud Native Analytics.

About this task

To migrate from Event Analytics to Cloud Native Analytics, you must migrate the data from your Db2 reporter database to Cassandra and then run training on the newly imported data to generate Cloud Native Analytics policies.

Procedure

See “Migrating historical data from a reporter database: scenario for Operations Management on IBM Cloud Private” on page 352

Upgrading Operations Management on IBM Cloud Private from 1.6.0 to 1.6.0.1

Use these instructions to upgrade your Operations Management on IBM Cloud Private deployment.

If your current version of Operations Management on IBM Cloud Private was installed from the CLI, then you must upgrade from the CLI.

If your current version of Operations Management on IBM Cloud Private was installed from the IBM Cloud Private UI, then you must upgrade from the IBM Cloud Private UI.

To upgrade from the IBM Cloud Private UI, see "Upgrading from the IBM Cloud Private UI" on page 153.

To upgrade from the CLI, see "Upgrading from the command line" on page 156.

Note: If the optional IBM Agile Service Manager extension is installed, then it must be upgraded to the latest version before Operations Management on IBM Cloud Private is installed. For more information, see https://www.ibm.com/support/knowledgecenter/SS9LQB_1.1.6/ReleaseNotes/asm_rn_11_6_upgradeicp.html#asm_rn_11_6_upgradeicp.

Upgrading from the IBM Cloud Private UI

Learn how to upgrade with the UI.

Before you begin

Upgrading from V1.6.0 to V1.6.0.1 involves deleting pre-existing IBM Operations Analytics - Log Analysisdata. For more information, see step "4" on page 153.

About this task

Upgrade Operations Management on IBM Cloud Private from V1.6.0 to V1.6.0.1 using the IBM Cloud Private UI.

Procedure

1. Upgrade IBM Agile Service Manager.

   If the optional Agile Service Manager extension is installed, then it must be upgraded to the latest version before Operations Management on IBM Cloud Private is upgraded. For more information about upgrading Agile Service Manager, see https://www.ibm.com/support/knowledgecenter/SS9LQB_1.1.6/ReleaseNotes/asm_rn_11_6_upgradeicp.html#asm_rn_11_6_upgradeicp.
2. Download the Operations Management on IBM Cloud Private version 1.6.0.1 PPA package and load it. For more information, see "Loading the archive into IBM Cloud Private" on page 115 or "Loading the archive into IBM Cloud Private with OpenShift" on page 117.
3. The replication factor for existing Kafka topics `ea-events`, `ea-collatedactions`, and `ea-actions` must be corrected before upgrade is run.

   There is a script on the IT Operations Management Developer Center that you can use:

   http://developer.ibm.com/itom/2019/10/31/updating-kafka-topic-partitions/

   To run the script, login to the cluster from the command line and run

   ```
   ./kafka_topic_reassigner.sh helm_release_name
   ```

   Where *helm_release_name* is the name of your Operations Management on IBM Cloud Private helm release.

   Run with a '-D' flag for additional debug information.
4. Delete pre-existing IBM Operations Analytics - Log Analysisdata, by completing the following steps:

   a. Scale down your scala pods by running the following command:

```
kubectl scale --replicas=0 Statefulsets/<helm_release_name>-scala
```

b. Find where the persistent volume claims (PVCs) are mounted, by running the following commands:

```
kubectl get pvc | grep scala
Bound    data-local-storage-scala-1-<worker node IP>             20Gi
RWO            local-storage-scala           9d

kubectl get pv -o json data-local-storage-scala-1-172.16.40.89 | jq '.spec.local.path'
"/root/tmp/1/m76/scala-1"
```

c. Go to the path of the worker node and delete the contents of the scala-1 directory, as in the following example:

```
ssh root@
rm -rf /root/tmp/1/m76/scala-1/*
```

5. From the Cassandra container, use the Cassandra CLI `nodetool` to verify that the Cassandra nodes are up, as in the following example:

```
kubectl exec -ti helm_release_name-cassandra-0 bash
[cassandra@m76-cassandra-0 /]$ nodetool status
Datacenter: datacenter1
=======================
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
--  Address      Load       Tokens       Owns (effective)  Host
ID                              Rack
UN  x.x.x.x     636.99 KiB  256          100.0%            d439ea16-7b55-4920-
a9a3-22e878feb844  rack1
DN  x.x.x.x     ?           256          100.0%            e121ab2e-fb9f-4aaf-9a45-
aef13d4d2317  rack1
DN  x.x.x.x     ?           256          100.0%            5464bd28-
e984-40dc-898a-012c80fbebf2  rack1
```

Where *helm_release_name* is the name of your Operations Management on IBM Cloud Private Helm release.

If a Cassandra node is up, then it shows a status of `'UN'`, if it is down then it shows a status of `'DN'`. If any Cassandra node shows a status of 'DN', then follow the resolution steps in “Restart of all Cassandra pods causes errors for connecting services” on page 535.

6. Delete any existing Operations Management on IBM Cloud Private jobs. To find these jobs, run the following command:

```
kubectl get jobs -n namespace
```

Where *namespace* is the namespace that your Operations Management on IBM Cloud Private deployment is in.

Delete each of the returned jobs whose name starts with the helm release name of your Operations Management on IBM Cloud Private deployment.

```
kubectl delete job jobname -n namespace
```

Where

- *jobname* is the name of the job to be deleted.
- *namespace* is the namespace that your Operations Management on IBM Cloud Private deployment is in.

7. Delete the *spark-master* service

```
kubectl delete service helm_release_name-spark-master -n namespace
```

Where

- *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment.

- *namespace* is the namespace that your Operations Management on IBM Cloud Private deployment is in.

8. Log in to the IBM Cloud Private UI by pointing your internet browser to a URL similar to:

   ```
   https://master_node_ip_or_hostname:8443/oidc/login.jsp
   ```

   Where *master_node_ip_or_hostname* is the IP address or hostname of your cluster's master node.

   Note: For IBM Cloud Private with OpenShift, access your cluster by using a different URL.

9. From the navigation menu, click **Workloads->Helm Releases**, then find and select the current Helm release, version 2.1.0 (App version 1.6.0), for *ibm-netcool-prod*.
10. Select **Upgrade** from the menu on the right.
11. Select version 2.1.1 (App version 1.6.0.1) from the drop-down list.
12. If IBM Agile Service Manager is not installed, then select **Reuse Values**. If IBM Agile Service Manager is installed, then verify that parameters that are new or changed in V1.6.0.1 to support the Cloud Native Analytics topology integration to IBM Agile Service Manager are correct, by selecting **All parameters** and then editing the following fields:

    - *(Optional) ASM status to event enrichment join window* is set to the required interval.
    - *ASM release name* is set to the name of the upgraded IBM Agile Service Manager deployment.
    - The *Enable ASM Integration* checkbox is selected.

    For more information, see "Configuring Installation Parameters for Operations Management on IBM Cloud Private" on page 125.

13. Select **Upgrade**. The message `Chart upgrade is in progress and may take a few minutes` is displayed until the upgrade is completed. **Workloads->Helm Releases** then shows the Operations Management on IBM Cloud Private deployment with a **Current version** of 2.1.1 (App version 1.6.0.1), and a **Status** of Deployed.

    Note: If you have enabled the optional Agile Service Manager extension for Topology Analytics, there is a known issue that must be rectified after the upgrade to V1.6.0.1 has completed. See step "18" on page 156.

14. Monitor the progress of the upgrade with the following commands:

    kubectl get jobs -n *namespace*
    Where *namespace* is the namespace that your Operations Management on IBM Cloud Private deployment is in.

    Use this command to see your deployment's jobs.

    kubectl describe job *jobname*
    Use this command to describe a job, where *jobname* is one of your deployment's jobs.

    kubectl get pods -n *namespace*
    Where *namespace* is the namespace that your Operations Management on IBM Cloud Private deployment is in.

    Use this command to retrieve the status of all the Operations Management on IBM Cloud Private pods.

15. Scale up your scala pods by running the following command:

    ```
    kubectl scale replicas=1 StatefulSets/<helm_release_name>-scala
    ```

16. If IBM Agile Service Manager is installed, you must run the tenant migration script.

    The static tenant ID for Operations Management on IBM Cloud Private V1.6.0.1 is updated to the tenant ID used by IBM Agile Service Manager so that topological enrichments from IBM Agile Service Manager can be seen. Historic data and policies must also be moved to the new tenant ID. A script to migrate the historic data and policies to the new tenant ID is provided within the event tools. For more information, see "Migrating data and policies to a new tenant ID" on page 159

17. Assign role *inasm_operator*, *inasm_editor* , or *inasm_admin* in Dashboard Application Services Hub to each user that needs to view topological data. For more information, see https://www.ibm.com/support/knowledgecenter/SS9LQB_1.1.6/Installing/t_asm_configuring.html.
18. 1.6.0.1

    If you have enabled the optional Agile Service Manager extension for Topology Analytics, then you must update your `values.yaml` configuration file.

    See this known issue "Incorrect path in ibm-ea-asm-normaliser section in values.yaml" on page 562. You **must** follow the steps in here to ensure that `ibm-ea-asm-normaliser.kafkaImage.name` is set correctly. Your deployment must then be upgraded to use the updated path. Run the following command:

    ```
    helm upgrade helm_release_name ./ibm-netcool-prod-2.1.1.tgz -f custom-values-upgrade.yaml --
    tls
    ```

    where *helm_release_name* is the name of your Operations Management on IBM Cloud Private release.

What to do next

If upgrade fails, you can roll back to the previous version of Operations Management on IBM Cloud Private. See "Rolling back Operations Management on IBM Cloud Private from V1.6.0.1 to V1.6.0" on page 161.

Upgrading from the command line

Learn how to upgrade with the command line.

Before you begin

Upgrading from V1.6.0 to V1.6.0.1 involves deleting pre-existing IBM Operations Analytics - Log Analysisdata. For more information, see step "5" on page 157.

About this task

Upgrade Operations Management on IBM Cloud Private from V1.6.0 to V1.6.0.1 using the command line.

Procedure

1. Upgrade IBM Agile Service Manager.

   If the optional Agile Service Manager extension is installed, then it must be upgraded to the latest version before Operations Management on IBM Cloud Private is upgraded. For more information about upgrading Agile Service Manager, see https://www.ibm.com/support/knowledgecenter/SS9LQB_1.1.6/ReleaseNotes/asm_rn_11_6_upgradeicp.html#asm_rn_11_6_upgradeicp.
2. Download the Operations Management on IBM Cloud Private version 1.6.0.1 PPA package and load it. For more information, see "Loading the archive into IBM Cloud Private" on page 115 or "Loading the archive into IBM Cloud Private with OpenShift" on page 117.
3. Extract the charts and `values.yaml` file from the PPA archive.

   ```
   tar -xvf ibm-netcool-prod.2.1.1-x86_64.tar.gz charts/ibm-netcool-prod-2.1.1.tgz
   cd charts
   tar -xvf ibm-netcool-prod-2.1.1.tgz ibm-netcool-prod/values.yaml
   ```
4. The replication factor for existing Kafka topics `ea-events`, `ea-collatedactions`, and `ea-actions` must be corrected before upgrade is run.

   There is a script on the IT Operations Management Developer Center that you can use to do this:

   http://developer.ibm.com/itom/2019/10/31/updating-kafka-topic-partitions/

   To run the script, login to the cluster from the command line and run

```
./kafka_topic_reassigner.sh helm_release_name
```

Where *helm_release_name* is the name of your Operations Management on IBM Cloud Private helm release.

Run with a '-D' flag for additional debug information.

5. Delete pre-existing IBM Operations Analytics - Log Analysisdata, by completing the following steps:

   a. Scale down your scala pods by running the following command:

   ```
   kubectl scale replicas=0 Statefulsets/<helm_release_name>-scala
   ```

   b. Find where the persistent volume claims (PVCs) are mounted, by running the following commands:

   ```
   kubectl get pvc | grep scala
   Bound    data-local-storage-scala-1-<worker node IP>             20Gi
   RWO            local-storage-scala           9d

   kubectl get pv -o json data-local-storage-scala-1-172.16.40.89 | jq '.spec.local.path'
   "/root/tmp/1/m76/scala-1"
   ```

   c. Go to the path of the worker node and delete the contents of the scala-1 directory, as in the following example:

   ```
   ssh root@
   rm -rf /root/tmp/1/m76/scala-1/*
   ```

6. From the Cassandra container, use the Cassandra CLI `nodetool` to verify that the Cassandra nodes are up, as in the following example:

```
kubectl exec -ti helm_release_name-cassandra-0 bash
[cassandra@m76-cassandra-0 /]$ nodetool status
Datacenter: datacenter1
=======================
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
--  Address      Load       Tokens       Owns (effective)  Host
ID                              Rack
UN  x.x.x.x     636.99 KiB  256          100.0%            d439ea16-7b55-4920-
a9a3-22e878feb844  rack1
DN  x.x.x.x     ?           256          100.0%            e121ab2e-fb9f-4aaf-9a45-
aef13d4d2317  rack1
DN  x.x.x.x     ?           256          100.0%            5464bd28-
e984-40dc-898a-012c80fbebf2  rack1
```

Where *helm_release_name* is the name of your Operations Management on IBM Cloud Private Helm release.

If a Cassandra node is up, then it shows a status of `'UN'`, if it is down then it shows a status of `'DN'`. If any Cassandra node shows a status of 'DN', then follow the resolution steps in “Restart of all Cassandra pods causes errors for connecting services” on page 535.

7. Delete any existing Operations Management on IBM Cloud Private jobs. To find these jobs, run the following command:

```
kubectl get jobs -n namespace
```

Where *namespace* is the namespace that your Operations Management on IBM Cloud Private deployment is in.

Delete each of the returned jobs whose name starts with the helm release name of your Operations Management on IBM Cloud Private deployment.

```
kubectl delete job jobname -n namespace
```

Where

- *jobname* is the name of the job to be deleted.

- *namespace* is the namespace that your Operations Management on IBM Cloud Private deployment is in.

8. Delete the *spark-master* service

```
kubectl delete service helm_release_name-spark-master -n namespace
```

Where

- *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment.
- *namespace* is the namespace that your Operations Management on IBM Cloud Private deployment is in.

9. If IBM Agile Service Manager is not installed, then do not change your Operations Management on IBM Cloud Private version 1.6.0 configuration file, *values.yaml*. If IBM Agile Service Manager is installed, then edit your Operations Management on IBM Cloud Private version 1.6.0 configuration file, *values.yaml*, so that parameters that are new or changed in Operations Management on IBM Cloud Private version 1.6.0.1 to support the Cloud Native Analytics topology integration to IBM Agile Service Manager are correct. Make the following changes:

- (Optional) Add parameter *global.ibm-ea-asm-normalizer.joinWindowSize*, and set it to the required time interval.
- Add parameter *global.integrations.asm.releaseName*, and set it to the name of the upgraded IBM Agile Service Manager deployment.
- Add the parameter *global.integrations.asm.enabled* and set it to `true`.
- Remove the parameter *webgui.asm.releasename*.
- 1.6.0.1 See this known issue "Incorrect path in ibm-ea-asm-normaliser section in values.yaml" on page 562. You **must** follow the steps in here to ensure that `ibm-ea-asm-normaliser.kafkaImage.name` is set correctly.

For more information, see "Configuring Installation Parameters for Operations Management on IBM Cloud Private" on page 125

10. Run the upgrade with the following command:

```
helm upgrade --tls helm_release_name -f values.yaml ibm-netcool-prod -n namespace
```

Where

- *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment.
- `values.yaml` is the name of the configuration file that contains the installation parameters. See "Configuring Installation Parameters for Operations Management on IBM Cloud Private" on page 125.
- *namespace* is the namespace that your Operations Management on IBM Cloud Private deployment is in.

11. Monitor the progress of the upgrade with the following commands:

kubectl get jobs -n *namespace*
Where *namespace* is the namespace that your Operations Management on IBM Cloud Private deployment is in.

Use this command to see your deployment's jobs.

kubectl describe job *jobname*
Use this command to describe a job, where *jobname* is one of your deployment's jobs.

kubectl get pods -n *namespace*
Where *namespace* is the namespace of your current Operations Management on IBM Cloud Private deployment.

Use this command to retrieve the status of all the Operations Management on IBM Cloud Private pods.

12. Scale up your scala pods by running the following command:

```
kubectl scale replicas=1 StatefulSets/<helm_release_name>-scala
```

13. If IBM Agile Service Manager is installed, you must run the tenant migration script.

    The static tenant ID for Operations Management on IBM Cloud Private is updated to the tenant ID used by IBM Agile Service Manager so that topological enrichments from IBM Agile Service Manager can be seen. Historic data and policies must also be moved to the new tenant ID. A script to migrate the historic data and policies to the new tenant ID is provided within the event tools. For more information, see “Migrating data and policies to a new tenant ID” on page 159

14. Assign role *inasm_operator*, *inasm_editor* , or *inasm_admin* in Dashboard Application Services Hub to each user that needs to view topological data. For more information, see https://www.ibm.com/support/knowledgecenter/SS9LQB_1.1.6/Installing/t_asm_configuring.html.

What to do next

If upgrade fails, you can roll back to the previous version of Operations Management on IBM Cloud Private. See “Rolling back Operations Management on IBM Cloud Private from V1.6.0.1 to V1.6.0” on page 161

Migrating data and policies to a new tenant ID

When Operations Management on IBM Cloud Private 1.6.0 is upgraded to Operations Management on IBM Cloud Private 1.6.0.1, use this information to move historical data and policies to a new tenant ID.

About this task

On completion of a successful Operations Management on IBM Cloud Private upgrade from release 1.6.0 to 1.6.0.1, the default tenant ID changes from `ea-generic-tenant` to `cfd95b7e-3bc7-4006-a4a8-a73a79c71255` to match the tenant ID in ASM. The historic data and policies of the deployment must be migrated from the old tenant ID to the new tenant ID.

Procedure

1. Get the name of the policy registry image from its pod.

```
kubectl get po \
$(kubectl get po -l app.kubernetes.io/component=policyregistryservice \
| awk '{if(NR==2) print $0}' \
| awk {'print $1'} ) -o jsonpath="{..image}" \
| tr -s '[[:space:]]' '\n' \
| sort \
| uniq \
| grep policy-registry-service
```

2. Get the name of the archiving image from its pod.

```
kubectl get po \
$(kubectl get po -l app.kubernetes.io/component=policyregistryservice \
| awk '{if(NR==2) print $0}' \
| awk {'print $1'} ) -o jsonpath="{..image}" \
| tr -s '[[:space:]]' '\n' \
| sort \
| uniq \
| grep archiving-service
```

3. Find the location and version of your image repository.

```
helm status helm_release_name --tls | grep -e '--image=' | cut -d'=' -f2 | cut -d' ' -f1
```

   Where *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment.

4. If the migration script has previously been run, remove any existing jobs or pods for it.

```
kubectl delete po migrate-data --ignore-not-found && \
kubectl delete job migrate-noi-historic-tenant-data --ignore-not-found && \
```

5. The command to migrate or delete historic data and policies is:

```
kubectl run migrate-data -it --command=true --restart=Never --env=LICENSE=accept --
image=image_repo migrateHistoricData.sh -- \
  -r "helm_release_name" \
  -t "target_tenant_ID"
  -n "new_tenant_ID" \
  -d "delete_flag" \
  -m "migrate_flag" \
  -s "service_account_name>" \
  -a "archiving_image" \
  -p "policy_registry_image" \
  | kubectl apply -f -
```

Where

- *image_repo* is the docker repository containing your images, as found in step 3.
- *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment. Defaults to `noi`.
- *target_tenant_ID* is the tenant ID that data is copied to or deleted from. This must be the tenant ID that Operations Management on IBM Cloud Private 1.6.0 used, which is `ea-generic-tenant.`
- *new_tenant_ID* is the tenant ID that data is copied to. This must be the tenant ID that IBM Agile Service Manager uses, which is `cfd95b7e-3bc7-4006-a4a8-a73a79c71255`.
- *delete_flag* is set to true if you want to delete data from the *target_tenant_ID* (the old tenant ID). Set to false to retain data from the *target_tenant_ID.*
- *migrate_flag* is set to true if you want to migrate data from the *target_tenant_ID* to the *new_tenant_ID*. Set to false to not migrate the data.
- *service_account_name* is the service account to use for the migration. The default *noi-service-account* can be used for the migration as it has the access required to pull the policy registry service and archiving service images. For more information, see *global.rbac.serviceAccount* in "Configuring Installation Parameters for Operations Management on IBM Cloud Private" on page 125.
- *policy_registry_image* is the fully qualified path to the policy image, as found in step 1.
- *archiving_image* is the fully qualified path to the archiving image, as found in step 2.

a) Run a trial of the `migrate-data` command.

   Run the migrate-data command without `| kubectl apply -f.` This command does not migrate or delete the data, and only outputs the rendered template that the script will create, along with any error messages.

b) Migrate the data to the new tenant ID.

   If the output from the previous step is error-free, then you can run the migration. Run the `migrate-data` command with *migrate_flag* set to true.

c) Delete the data from the old tenant ID.

   Run the `migrate-data` command with the *migrate_flag* set to false and the *delete_flag* set to true to delete historic data and policies from the old tenant ID.

Here is an example of the `migrate-data` command where data and policies are migrated to a new tenant ID, but data and policies are not deleted from the old tenant ID. In this example, the event tooling uses the default service account which has an appropriate image pull secret provided to the default service account.

```
# Set your release name - this example is to noi
RELEASE_NAME=noi

# Delete pods and jobs from previous runs
```

```
kubectl delete po migrate-data --ignore-not-found && \
kubectl delete job migrate-noi-historic-tenant-data --ignore-not-found && \


# Next run the following code snippet in order to generate the job
kubectl run migrate-data \
-it --command=true --restart=Never --env=LICENSE=accept \
--image=$(helm status $RELEASE_NAME --tls | grep -e '--image=' | cut -d'=' -f2 | cut -d' ' -
f1) migrateHistoricData.sh -- \
  -r "$RELEASE_NAME" -t "ea-generic-tenant" -n "cfd95b7e-3bc7-4006-a4a8-a73a79c71255" \
  -d "false" -m "true" \
  -s "default" \
  -a "$(kubectl get po \
    $(kubectl get po -l app.kubernetes.io/component=policyregistryservice \
    | awk '{if(NR==2) print $0}' \
    | awk {'print $1'} ) -o jsonpath="{..image}" \
    | tr -s '[[:space:]]' '\n' \
    | sort \
    | uniq \
    | grep archiving-service)" \
  -p "$(kubectl get po \
    $(kubectl get po -l app.kubernetes.io/component=policyregistryservice \
    | awk '{if(NR==2) print $0}' \
    | awk {'print $1'} ) -o jsonpath="{..image}" \
    | tr -s '[[:space:]]' '\n' \
    | sort \
    | uniq \
    | grep policy-registry-service)" \
  | kubectl apply -f -
```

Rolling back Operations Management on IBM Cloud Private from V1.6.0.1 to V1.6.0

Use these instructions to roll back your Operations Management on IBM Cloud Private deployment.

If your current version of Operations Management on IBM Cloud Private was installed from the CLI, then you must roll back from the CLI.

If your current version of Operations Management on IBM Cloud Private was installed from the IBM Cloud Private UI, then you must roll back from the IBM Cloud Private UI.

To roll back from the IBM Cloud Private UI, see “Rollback from the IBM Cloud Private UI” on page 161

To rollback from the CLI, see “Rollback from the command line” on page 163.

Note: Rollback should be performed only if an upgrade is unsuccessful and you want to revert to your previous version. It must be performed as soon as possible after the failed upgrade to ensure that persisted data is in a state recognizable by the original version of Operations Management on IBM Cloud Private.

Rollback from the IBM Cloud Private UI

About this task

Roll back Operations Management on IBM Cloud Private from 1.6.0.1 to 1.6.0 using the IBM Cloud Private UI.

Procedure

1. From the Cassandra container, use the Cassandra CLI `nodetool` to verify that the Cassandra nodes are up, as in the following example:

```
kubectl exec -ti helm_release_name-cassandra-0 bash
[cassandra@m76-cassandra-0 /]$ nodetool status
Datacenter: datacenter1
=======================
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
--  Address      Load       Tokens       Owns (effective)  Host
```

```
ID                                   Rack
UN  10.1.106.37  636.99 KiB  256          100.0%            d439ea16-7b55-4920-
a9a3-22e878feb844  rack1
DN  10.1.2.21    ?          256          100.0%            e121ab2e-fb9f-4aaf-9a45-
aef13d4d2317  rack1
DN  10.1.118.15  ?          256          100.0%            5464bd28-
e984-40dc-898a-012c80fbebf2  rack1
```

Where *helm_release_name* is the name of your Operations Management on IBM Cloud Private Helm release.

If a Cassandra node is up, then it shows a status of `'UN'`, if it is down then it shows a status of `'DN'`. If any Cassandra node shows a status of 'DN', then follow the resolution steps in "Restart of all Cassandra pods causes errors for connecting services" on page 535.

2. Delete any existing Operations Management on IBM Cloud Private jobs. To find these jobs, run the following command:

```
kubectl get jobs -n namespace
```

Where *namespace* is the namespace that your Operations Management on IBM Cloud Private deployment is in.

Delete each of the returned jobs whose name starts with the helm release name of your Operations Management on IBM Cloud Private deployment.

```
kubectl delete job jobname -n namespace
```

Where

- *jobname* is the name of the job to be deleted.
- *namespace* is the namespace that your Operations Management on IBM Cloud Private deployment is in.

3. Delete the *spark-master* service

```
kubectl delete service helm_release_name-spark-master -n namespace
```

Where

- *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment.
- *namespace* is the namespace that your Operations Management on IBM Cloud Private deployment. is in.

4. Log in to the IBM Cloud Private UI by pointing your internet browser to a URL similar to:

```
https://master_node_ip_or_hostname:8443/oidc/login.jsp
```

Where *master_node_ip_or_hostname* is the IP address or hostname of your cluster's master node.

Note: For IBM Cloud Private with OpenShift, access your cluster by using a different URL.

5. From the navigation menu, click **Workloads->Helm Releases**
6. Find and select the currently deployed Helm release, version 2.1.1 (App version 1.6.0.1), for *ibm-netcool-prod*.
7. Right-click **Rollback**, or select **Rollback** from the menu on the right.
8. Select version 2.1.0 (App version 1.6.0) from the drop-down list, select **Reuse values** and do not edit any parameters.
9. Select **Rollback**. When rollback is completed **Workloads->Helm Releases** shows the Operations Management on IBM Cloud Private deployment with a **Current version** of version 2.1.0 (App version 1.6.0), and a **Status** of Deployed.

Rollback from the command line

About this task

Roll back Operations Management on IBM Cloud Private from 1.6.0.1 to 1.6.0 using the command line.

Procedure

1. From the Cassandra container, use the Cassandra CLI `nodetool` to verify that the Cassandra nodes are up, as in the following example:

```
kubectl exec -ti helm_release_name-cassandra-0 bash
[cassandra@m76-cassandra-0 /]$ nodetool status
Datacenter: datacenter1
=======================
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
--  Address      Load       Tokens       Owns (effective)  Host
ID                              Rack
UN  10.1.106.37  636.99 KiB  256          100.0%            d439ea16-7b55-4920-
a9a3-22e878feb844  rack1
DN  10.1.2.21    ?          256          100.0%            e121ab2e-fb9f-4aaf-9a45-
aef13d4d2317  rack1
DN  10.1.118.15  ?          256          100.0%            5464bd28-
e984-40dc-898a-012c80fbebf2  rack1
```

   Where *helm_release_name* is the name of your Operations Management on IBM Cloud Private Helm release.

   If a Cassandra node is up, then it shows a status of `'UN'`, if it is down then it shows a status of `'DN'`. If any Cassandra node shows a status of 'DN', then follow the resolution steps in “Restart of all Cassandra pods causes errors for connecting services” on page 535.

2. Delete any existing Operations Management on IBM Cloud Private jobs. To find these jobs, run the following command:

```
kubectl get jobs -n namespace
```

   Where *namespace* is the namespace of your current Operations Management on IBM Cloud Private deployment.

   Delete each of the returned jobs whose name starts with the helm release name of your Operations Management on IBM Cloud Private deployment.

```
kubectl delete job jobname -n namespace
```

   Where

   - *jobname* is the name of the job to be deleted.
   - *namespace* is the namespace of your current Operations Management on IBM Cloud Private deployment.

3. Delete the *spark-master* service

```
kubectl delete service helm_release_name-spark-master -n namespace
```

   Where

   - *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment.
   - *namespace* is the namespace of your current Operations Management on IBM Cloud Private deployment.

4. Do not edit the `values.yaml` configuration file.
5. Run the following command to roll back to the previous deployment of Operations Management on IBM Cloud Private.

```
helm rollback helm_release_name --tls
```

Where *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment.

6. Monitor the progress of the rollback with the following commands:

   kubectl get jobs -n *namespace*
   Where *namespace* is the namespace of your current Operations Management on IBM Cloud Private deployment.

   Use this command to see your deployment's jobs.

   kubectl describe job *jobname*
   Use this command to describe a job, where *jobname* is one of your deployment's jobs.

   kubectl get pods -n *namespace*
   Where *namespace* is the namespace of your current Operations Management on IBM Cloud Private deployment.

   Use this command to retrieve the status of all the Operations Management on IBM Cloud Private pods.

Chapter 5. Configuring

Perform the following tasks to configure the components of Netcool Operations Insight.

Configuring Operations Management

Perform the following tasks to configure the components of Operations Management.

Configuring Event Search

You can customize Event Search to your specific needs by customizing the Netcool/OMNIbus Insight Pack . For example, you might want to send an extended set of Netcool/OMNIbus event fields to Operations Analytics - Log Analysis and chart results based on those fields.

Note: Do not directly modify the Netcool/OMNIbus Insight Pack. Instead, use it as a base for creating customized Insight Packs and Custom apps.

Related concepts

Operations Management tasks
Use this information to understand the tasks that users can perform using Operations Management.

Checking the version of the Insight Pack

To ensure compatibility between the versions of the Tivoli Netcool/OMNIbus Insight Pack, the Web GUI and the Operations Analytics - Log Analysis product, run the **pkg_mgmt** command to check which version of the Insight Pack is installed.

Procedure

To check which version of the Insight Pack is installed, run the **pkg_mgmt** as follows:

```
$SCALA_HOME/utilities/pkg_mgmt.sh -list
```

Search the results for a line similar to the following example. In this example, V1.3.0.2 is installed.

```
[packagemanager] OMNIbusInsightPack_v1.3.0.2     /home/myhome/IBM/LogAnalysis/
unity_content/OMNIbus
```

Related reference

Troubleshooting event search
How to resolve problems with your event search configuration.

Event annotations

The event annotations defined by the Insight Pack index configuration are described here.

The following table lists the Netcool/OMNIbus event fields that are defined in the index configuration file. It also lists the index configuration attributes assigned to each field. These annotations are displayed in the Operations Analytics - Log Analysis Search workspace, and can be used to filter or search the events.

Tip: The fields are not necessarily listed in Table 34 on page 166 in the same order as the data source properties file. If you need to know which order the fields are given, see the `omnibus1100.properties` file, which is in the `docs` directory of the Tivoli Netcool/OMNIbus Insight Pack.

Best practice for filtering on annotations is as follows:

- To avoid a negative impact on performance, set `filterable` attributes to `true` only on required fields.
- Do not set the `filterable` attribute to `true` for fields with potentially long strings, for example, the `Summary` field.

Table 34. Event annotations

Field	Attributes
LastOccurrence	dataType: DATE retrievable: true retrieveByDefault: true sortable: true filterable: true searchable: true
logRecord This is a default field required by Operations Analytics - Log Analysis.	dataType: TEXT retrievable: true retrieveByDefault: true sortable: false filterable: false searchable: true
Class	dataType: TEXT retrievable: true retrieveByDefault: true sortable: false filterable: false searchable: true
AlertGroup	dataType: TEXT retrievable: true retrieveByDefault: true sortable: true filterable: true searchable: true
Severity	dataType: TEXT retrievable: true retrieveByDefault: true sortable: true filterable: true searchable: true
AlertKey	dataType: TEXT retrievable: true retrieveByDefault: true sortable: true filterable: true searchable: true
Tally	dataType: TEXT retrievable: true retrieveByDefault: true sortable: false filterable: false searchable: true
NmosObjInst	dataType: TEXT retrievable: true retrieveByDefault: true sortable: true filterable: true searchable: true

Table 34. Event annotations (continued)

Field	Attributes
NodeAlias	dataType: TEXT retrievable: true retrieveByDefault: true sortable: true filterable: true searchable: true
timestamp	dataType: DATE retrievable: true retrieveByDefault: true sortable: true filterable: true searchable: true
Type	dataType: TEXT retrievable: true retrieveByDefault: true sortable: true filterable: false searchable: true
Location	dataType: TEXT retrievable: true retrieveByDefault: true sortable: true filterable: true searchable: true
Identifier	dataType: TEXT retrievable: true retrieveByDefault: true sortable: true filterable: false searchable: true
Node	dataType: TEXT retrievable: true retrieveByDefault: true sortable: true filterable: true searchable: true
Summary	dataType: TEXT retrievable: true retrieveByDefault: true sortable: true filterable: false searchable: true
OmniText This field contains a concatenated string of the event fields. By default, its value is: `'@Manager' + ' ' + '@Agent' + ' ' + TO_STRING('@Grade')` In the default configuration, the names of the individual event fields contained in the `OmniText` string are not visible in search results.	dataType: TEXT retrievable: true retrieveByDefault: true sortable: false filterable: false searchable: true

Table 34. Event annotations (continued)	
Field	**Attributes**
PubType **I** when an event is inserted. **U** when an existing entry is re-inserted or updated.	dataType: TEXT retrievable: true retrieveByDefault: true sortable: false filterable: true searchable: true
ServerName Corresponds to ServerName in alerts.status.	dataType: TEXT retrievable: true retrieveByDefault: true sortable: false filterable: true searchable: true
ServerSerial Corresponds to ServerSerial in alerts.status.	dataType: LONG retrievable: true retrieveByDefault: true sortable: false filterable: false searchable: true

Example queries
The following examples show you how to issue search queries on events.

- "Track changes to a specific event" on page 168
- "Search on one instance of an event only" on page 168

Track changes to a specific event
Issue the following search to track changes to a specific event:

```
ServerSerial:NNNN AND ServerName:NCOMS
```

Where NNNN is the serial number and NCOMS is the name of the server.

Search on one instance of an event only
Issue the following search to search on one instance of an event only:

```
NOT PubType:U
```

Customizing the events used in Event Search
You can send an extended set of Netcool/OMNIbus event fields to Operations Analytics - Log Analysis and chart results based on those fields.

Customizing events used in Event Search using the addIndex.sh script
Use the addIndex.sh script if you are running Netcool Operations Insight v1.4.1.2 or later versions.

Setting up and activating an initial set of custom events
You must first create a datasource type containing the extended set of Netcool/OMNIbus event fields to send to Operations Analytics - Log Analysis, and then update the Netcool/OMNIbus datasource already defined in Operations Analytics - Log Analysis to use this type. You then modify and restart the Gateway for Message Bus and Event Search will then chart results based on the custom fields you specified.

Before you begin
Ensure that the following prerequisites are in place before performing this task:

- Operations Analytics - Log Analysis 1.3.3 or 1.3.5 is installed.

- On the Operations Analytics - Log Analysis server, the $UNITY_HOME environment variable is set to the directory where Operations Analytics - Log Analysis is installed.
- The DSV toolkit is installed in the $UNITY_HOME/unity_content/DSVToolkit_v1.1.0.4 directory and you have write access to the directory.
- A version of python, which is compatible with the Operations Analytics - Log Analysis release, is installed and is on the system path.
- Tivoli Netcool/OMNIbus Insight Pack V1.3.1 is installed on the Operations Analytics - Log Analysis server.

About this task

You have added a new custom field to the ObjectServer alerts.status table and you want to update the datasource to send this field to Operations Analytics - Log Analysis along with the other fields, so that Event Search can present charts and dashboards using this new custom field. For the purposes of this task, we will assume that the new custom field stores a trouble ticket number, and is called TicketNumber.

Creating a custom data source type

Create a custom data source type to include the new custom field in addition to the existing default fields.

About this task

In Operations Analytics - Log Analysis a data source is an entity that enables Operations Analytics - Log Analysis to ingest data from a specific source. In order for Operations Analytics - Log Analysis to ingest data from Netcool/OMNIbus, a data source is required.

A data source type is a template for a data source, and lists out the event fields to send to Operations Analytics - Log Analysis, together with relevant control parameters for each field. You can have multiple data source types, each set up with a different set of event fields; the advantage of this is that you can easily change the events in the data source using a predefined data source type.

The default datasource type is called OMNIbus1100, and this datasource type contains the default set of events that are sent to Operations Analytics - Log Analysis.

Procedure

1. Log into the Operations Analytics - Log Analysis server and open a terminal there.
2. Unzip the contents of Tivoli Netcool/OMNIbus Insight Pack V1.3.1 to a local directory.

   This procedure assumes that the archive has been unzipped to the following location:

   ```
   /home/user/OMNIbusInsightPack_v1.3.1
   ```
3. Go to the docs sub-directory within the location to which you unzipped the file.

   ```
   cd /home/user/OMNIbusINsightPack_v1.3.1/docs
   ```
4. Edit the omnibus1100_template.properties file using a text editor of your choice, for example, vi:

   ```
   vi omnibus1100_template.properties
   ```

   The omnibus1100_template.properties file contains index definitions corresponding to one or more fields to be sent using the data source. For the Netcool/OMNIbus data source all of the event fields must be indexed, so the omnibus1100_template.properties file contains an index entry for each event field to be sent to Operations Analytics - Log Analysis.
5. Add the new custom field to the end of the omnibus1100_template.properties file.

   The following code snippet shows the beginning of the file and the end of the file.

```
# Properties file controlling data source specification
# Add new fields at the end of the file
# 'moduleName' specifies the name of the data source.
```

```
# Update the version number if you have created a data source with the same name previously and want
# to upgrade it to add an additional field.

<existing index definitions, one for each of the default event fields>

# ------------------------------------------------------------------------------------------------
# Insert new fields after this point.
# Number each field sequentially, starting with 'field19'.
# See the IBM Smart Cloud Log Analytics documentation for the DSV Toolkit for an explanation
# of the field values.
# ------------------------------------------------------------------------------------------------
#[field19_indexConfig]
#name: <INDEX NAME>
#dataType: TEXT
#retrievable: true
#retrieveByDefault: true
#sortable: false
#filterable: true
```

The end of the file includes a commented out section which you can uncomment to add the new field. In that section, replace the `name:` attribute with the name of the field that you are adding.

Here is what the end of file looks like when an index has been added for new custom field `TicketNumber`:

```
# ------------------------------------------------------------------------------------------------
# Insert new fields after this point.
# Number each field sequentially, starting with 'field19'.
# See the IBM Smart Cloud Log Analytics documentation for the DSV Toolkit for an explanation
# of the field values.
# ------------------------------------------------------------------------------------------------
[field19_indexConfig]
name: Ticket_Number
dataType: TEXT
retrievable: true
retrieveByDefault: true
sortable: false
filterable: true
```

Note: The order of indexes is important; it must match the order of values specified in the Gateway for Message Bus mapping file. This mapping file will be modified later in the procedure.

For more information on the other attributes of the index, see the following Operations Analytics - Log Analysis topics:

- 1.3.5: Editing an index configuration
- 1.3.5: Example properties file with edited index configuration fields

6. Change the name of the new custom data source type that you are about to create. Within the `[DSV]` section of the `omnibus1100_template.properties` file find the attribute specification `moduleName` and change the value specified there.

   By default `moduleName` is set to `CloneOMNIbus`. You can change this to a more meaningful name; for example, `customOMNIbus`.
7. Save the `omnibus1100_template.properties` file and exit the file editor.
8. From within the `/home/user/OMNIbusINsightPack_v1.3.1/docs` directory, run the `addIndex.sh` script to create the new data source type.

   ```
   addIndex.sh -i
   ```
9. Check that the data source type was created and installed onto the Operations Analytics - Log Analysis server by running the following command:

   ```
   $UNITY_HOME/utilities/pkg_mgmt.sh -list
   ```

   Where `$UNITY_HOME` is the Operations Analytics - Log Analysis home directory; for example, `/home/scala/IBM/LogAnalysis/`.

Results
The following two artifacts are also created. Store them in a safe place and make a note of the directory where you stored them, as you might need them later:

Insight pack image archive
By default this archive is called `CloneOMNIbusInsightPack_v1.3.1.0.zip`. If you followed the suggested example in this procedure, then this archive will be called `customOMNIbusInsightPack_v1.3.1.0.zip`. This archive contains the new custom data source type; you need a copy of this image if you ever want to delete it from the system in the future. The archive is located in the following directory:

```
/home/user/OMNIbusINsightPack_v1.3.1/dist
```

Template properties file
This is the `omnibus1100_template.properties` file that you edited during this procedure. Keep a copy of this file in case you want to modify the data source type settings at a later time.

Creating a custom data source
Create a new data source based on the custom data source type you created earlier.

Before you begin

The Web GUI right click tools and Event Search dashboards and charts are coded to use a data source named `omnibus`. The prerequisite for this task varies depending on whether you have ever ingested data into the existing `omnibus` data source.

- If you *have* already ingested data into the existing `omnibus` data source then you must delete the data. For information on how to do this, see Operations Analytics - Log Analysis 1.3.5 documentation: Deleting data.
- If you have not yet ingested data into the existing `omnibus` data source then simply delete this data source.

Procedure

1. In Operations Analytics - Log Analysis, start the **Add Data Source** wizard and configure an "omnibus" data source for Netcool/OMNIbus events.

   Only a single data source is required. The event management tools in the Web GUI support a single data source only.

   a) In the **Select Location** panel, select **Custom** and type the Netcool/OMNIbus server host name.

   Enter the same host name that was used for the **`JsonMsgHostname`** transport property of the Gateway for Message Bus.

   b) In the **Select Data** panel, enter the following field values:

Field	Value
File path	`NCOMS`. This is the default value of the **`jsonMsgPath`** transport property of the Gateway for Message Bus. If you changed this value from the default, change the value of the **File path** field accordingly.
Type	This is the name of the data source type on which this data source is based. • To use the default data source type, specify `OMNIbus1100`. • To use a customized data source type, specify the name of the customized data source type; for example: `customOMNIbus`
Collection	`OMNIbus1100-Collection`

   c) In the **Set Attributes** panel, enter the following field values:

Field	Value
Name	omnibus. Ensure that the value that you type is the same as the value of the **scala.datasource** property in the Web GUI server.init file. If the **Name** field has a value other than omnibus, use the same value for the **scala.datasource** property.
Group	Leave this field blank.
Description	Type a description of your choice.

2. Configure access to the data source you set up in the previous step. This involves the following steps in the administrative settings for Operations Analytics - Log Analysis:
   a) Create a role using the **Roles** tab, for example, noirole, and ensure you assign the role permission to access the data source.
   b) Add a user, for example, noiuser, and assign the role you created that has permissions to access the data source (in this example, noirole).

   For information about creating and modifying users and roles in Operations Analytics - Log Analysis, see one of the following links:

   - V1.3.5: see https://www.ibm.com/support/knowledgecenter/SSPFMY_1.3.5/com.ibm.scala.doc/config/iwa_config_pinstall_userrole_ovw_c.html
   - V1.3.3: see https://www.ibm.com/support/knowledgecenter/SSPFMY_1.3.3/com.ibm.scala.doc/config/iwa_config_pinstall_userrole_ovw_c.html

   .

   Note: The contents of the Netcool/OMNIbus Insight Pack dashboards are empty unless you log in with a user that has a role assigned with permissions to access the data source.

Modifying the Gateway for Message Bus mapping
You must modify the Gateway for Message Bus mapping to include the new custom field or fields.

Before you begin

The mapping that you configure in this task must match the order of fields configured in your custom data source type, as specified in “Creating a custom data source type” on page 169.

You must collect the following information prior to performing this task:

- Location of a Gateway for Message Bus mapping file that is compliant with the current version of the Netcool/OMNIbus Insight Pack.
- Name and location of the Gateway for Message Bus properties file.

Procedure

1. Copy a Gateway for Message Bus mapping file that is compliant with the current version of the Netcool/OMNIbus Insight Pack.

   For example, for Netcool/OMNIbus Insight Pack V1.3.0.2 and above, you can copy the file $OMNIHOME/gates/xml/scala/xml1302.map. In the following example, the file is copied to a file called xmlCustom1302.map:

```
cp  $OMNIHOME/gates/xml/scala/xml1302.map  $OMNIHOME/gates/xml/scala/xmlCustom1302.map
```

2. Add a command after the last entry in the CREATE MAPPING StatusMap section of the file.

   For example, if the last entry is a line specifying the ServerSerial field, then add a comma at the end of that line, like this:

```
        'ServerSerial'     =    '@ServerSerial',
  );
```

3. For the purposes of this task, we assume that you are adding a new custom field that stores a trouble ticket number, and this field is called `TicketNumber`. Add this custom field after the last entry in the `CREATE MAPPING StatusMap` section of the Gateway for Message Bus mapping file, before the terminating parenthesis.

```
        'ServerSerial'     =    '@ServerSerial',
        'TicketNumber'     =    '@TicketNumber'
  );
```

4. Save the Gateway for Message Bus mapping file, `xmlCustom1302.map`.
5. Locate the Gateway for Message Bus properties file.

   By default this file is called `G_SCALA.properties` and it is located in the following directory:

```
$OMNIHOME/gates/xml/scala/G_SCALA.props
```

   Where `$OMNIHOME` is `/opt/IBM/tivoli/netcool/omnibus/`.
6. Edit the Gateway for Message Bus properties file `G_SCALA.properties` using a text editor of your choice; for example, vi.

```
vi G_SCALA.properties
```

7. Change the `Gate.MapFile` parameter to refer to the new mapping file.

   For example:

```
Gate.MapFile        :'$OMNIHOME/gates/xml/scala/xmlClone1302.map'
```

8. Save and close the Gateway for Message Bus properties file `G_SCALA.properties`.
9. Restart the Gateway for Message Bus.

Updating the set of custom events
You can update the custom events used by Event Search.

Updating the custom data source type
Update the custom data source type that you created earlier to include additional fields. All data sources based on this data source type will automatically update to include the fields that you added.

Before you begin
You must have already created a custom data source type, as described in "Creating a custom data source type" on page 169.

About this task

The following procedure describes how to add a field to an existing data source type.

Restriction: You can add additional fields to the data source type at any time; however, once data has been ingested using a data source based on this type, you cannot modify or delete any of the added fields.

Procedure

1. Go to the following location:
2. Edit the `omnibus1100_template.properties` file using a text editor of your choice, for example, vi:

```
vi omnibus1100_template.properties
```

   The `omnibus1100_template.properties` file contains index definitions corresponding to one or more fields to be sent using the data source. For the Netcool/OMNIbus data source all of the event fields must be indexed, so the `omnibus1100_template.properties` file contains an index entry for each event field to be sent to Operations Analytics - Log Analysis.

3. Add the new custom field to the end of the `omnibus1100_template.properties` file.

   The following code snippet shows the beginning of the file and the end of the file.

```
# Properties file controlling data source specification
# Add new fields at the end of the file
# 'moduleName' specifies the name of the data source.
# Update the version number if you have created a data source with the same name previously and want
# to upgrade it to add an additional field.

<existing index definitions, one for each of the default event fields>

# ------------------------------------------------------------------------------------------------
# Insert new fields after this point.
# Number each field sequentially, starting with 'field19'.
# See the IBM Smart Cloud Log Analytics documentation for the DSV Toolkit for an explanation
# of the field values.
# ------------------------------------------------------------------------------------------------
#[field19_indexConfig]
#name: <INDEX NAME>
#dataType: TEXT
#retrievable: true
#retrieveByDefault: true
#sortable: false
#filterable: true
```

   The end of the file includes a commented out section which you can uncomment to add the new field. In that section, replace the `name:` attribute with the name of the field that you are adding.

   Here is what the end of file looks like when an index has been added for new custom field `TicketNumber`:

```
# ------------------------------------------------------------------------------------------------
# Insert new fields after this point.
# Number each field sequentially, starting with 'field19'.
# See the IBM Smart Cloud Log Analytics documentation for the DSV Toolkit for an explanation
# of the field values.
# ------------------------------------------------------------------------------------------------
[field19_indexConfig]
name: Ticket_Number
dataType: TEXT
retrievable: true
retrieveByDefault: true
sortable: false
filterable: true
```

   Note: The order of indexes is important; it must match the order of values specified in the Gateway for Message Bus mapping file. This mapping file will be modified later in the procedure.

   For more information on the other attributes of the index, see the following Operations Analytics - Log Analysis topics:

   - 1.3.5: Editing an index configuration
   - 1.3.5: Example properties file with edited index configuration fields

4. In the [DSV] section of the file, increase the value of the version parameter.
5. Save the `omnibus1100_template.properties` file and exit the file editor.
6. From within the `/home/user/OMNIbusINsightPack_v1.3.1/docs` directory, run the `addIndex.sh` script to update the data source type.

```
addIndex.sh -u
```

   Restriction: You can add additional fields to the data source type at any time; however, once data has been ingested using a data source based on this type, you cannot modify or delete any of the added fields.

Modifying the Gateway for Message Bus mapping
You must modify the Gateway for Message Bus mapping to include the new custom field or fields.

Before you begin

The mapping that you configure in this task must match the order of fields configured in your custom data source type, as specified in "Creating a custom data source type" on page 169.

You must collect the following information prior to performing this task:

- Location of a Gateway for Message Bus mapping file that is compliant with the current version of the Netcool/OMNIbus Insight Pack.
- Name and location of the Gateway for Message Bus properties file.

Procedure

1. Copy a Gateway for Message Bus mapping file that is compliant with the current version of the Netcool/OMNIbus Insight Pack.

   For example, for Netcool/OMNIbus Insight Pack V1.3.0.2 and above, you can copy the file `$OMNIHOME/gates/xml/scala/xml1302.map`. In the following example, the file is copied to a file called `xmlCustom1302.map`:

```
cp  $OMNIHOME/gates/xml/scala/xml1302.map  $OMNIHOME/gates/xml/scala/xmlCustom1302.map
```

2. Add a command after the last entry in the `CREATE MAPPING StatusMap` section of the file.

   For example, if the last entry is a line specifying the ServerSerial field, then add a comma at the end of that line, like this:

```
        'ServerSerial'     =    '@ServerSerial',
  );
```

3. For the purposes of this task, we assume that you are adding a new custom field that stores a trouble ticket number, and this field is called `TicketNumber`. Add this custom field after the last entry in the `CREATE MAPPING StatusMap` section of the Gateway for Message Bus mapping file, before the terminating parenthesis.

```
        'ServerSerial'     =    '@ServerSerial',
        'TicketNumber'     =    '@TicketNumber'
  );
```

4. Save the Gateway for Message Bus mapping file, `xmlCustom1302.map`.
5. Locate the Gateway for Message Bus properties file.

   By default this file is called `G_SCALA.properties` and it is located in the following directory:

```
$OMNIHOME/gates/xml/scala/G_SCALA.props
```

   Where `$OMNIHOME` is `/opt/IBM/tivoli/netcool/omnibus/`.
6. Edit the Gateway for Message Bus properties file `G_SCALA.properties` using a text editor of your choice; for example, vi.

```
vi G_SCALA.properties
```

7. Change the `Gate.MapFile` parameter to refer to the new mapping file.

   For example:

```
Gate.MapFile        :'$OMNIHOME/gates/xml/scala/xmlClone1302.map'
```

8. Save and close the Gateway for Message Bus properties file `G_SCALA.properties`.
9. Restart the Gateway for Message Bus.

Customizing events used in Event Search using the DSV toolkit

You must use the DSV toolkit to add events to Event Search if you are running Netcool Operations Insight v1.4.1.1 or lower, and are therefore using Tivoli Netcool/OMNIbus Insight Pack V1.3.0.2.

Before you begin

You can use the DSV toolkit to generate a customized insight pack. The DSV toolkit is provided with the Operations Analytics - Log Analysis product, in `$UNITY_HOME/unity_content/DSVToolkit_v1.1.0.4`. In the properties file, you can change the index configurations to meet your requirements

A new source type with an updated index configuration is created when you install the insight pack. An insight pack contains the following elements:

- An index configuration: defines how the fields are indexed in Operations Analytics - Log Analysis.
- A splitter: splits the ingested data into individual log entries.
- An annotator: splits the log entries into fields to be indexed.

The Netcool/OMNIbus insight pack requires the `newlineSplitter.aql` custom splitter, and the insight pack can use an annotator built using the DSV toolkit. To modify the index configuration, and generate a data source to ingest the data with the new index, you need to create a new insight pack using the DSV toolkit and modify it to use the Netcool Operations Insight `newlineSplitter.aql`.

- See the DSV toolkit documentation in the `$UNITY_HOME/unity_content/DSVToolkit_v1.1.0.4/docs` directory for information about specifying field properties and generating insight packs.
- See the Gateway for Message Bus documentation for information about mapping event fields to insight pack properties. Testing insight packs requires a Gateway for Message Bus to transfer events from Netcool/OMNIbus to Operations Analytics - Log Analysis.

About this task

Use the DSV toolkit to generate an insight pack that contains a new rule set (annotator and splitter) for the Netcool/OMNIbus event fields that you want.

The procedure describes how to create an insight pack called ITEventsInsightPack_V1.1.0.1, based on the Tivoli Netcool/OMNIbus Insight Pack V1.3.0.2. Use your own naming as appropriate.

Procedure

1. Make a copy of the `omnibus1100.properties` file, which is in the `docs` directory of the Tivoli Netcool/OMNIbus Insight Pack installation directory (`$UNITY_HOME/unity_content/OMNIbusInsightPack_v1.3.0.2`), and rename it.
   For example, rename it to `ITEvents.properties`.
2. Copy the `ITEvents.properties` file that you created in step "1" on page 176 to the DSV toolkit directory `$UNITY_HOME/unity_content/DSVToolkit_v1.1.0.4`.
3. Edit the `ITEvents.properties` file.
   For example, change the default value of the **aqlModuleName** field to `ITEvents`, and add, modify, or remove event field properties as required. To obtain the version number V1.1.0.1, change the **version** property to `1.1.0.1`.
4. If you added or removed fields from the file, change the value of the **totalColumns** field so that it specifies the total number of fields in the file.
5. Use the following command to generate an insight pack:

   ```
   python dsvGen.py ITEvents.properties -o
   ```

   The insight pack is named `ITEventsInsightPack_V1.1.0.1`.
6. Add the customized splitter into the insight pack as follows:

a) Create the following directory for the splitter: $UNITY_HOME/unity_content/DSVToolkit_v1.1.0.4/build/ITEventsInsightPack_v1.1.0.1/extractors/ruleset/splitter

b) Extract the Netcool/OMNIbus insight pack and copy the following file:

Insight Pack Extract Directory/OMNIbusInsightPack_v1.3.0.2/extractors/ruleset/splitter/newlineSplitter.aql

to

$UNITY_HOME/unity_content/DSVToolkit_v1.1.0.4/build/ITEventsInsightPack_v1.1.0.1/extractors/ruleset/splitter/

c) Open the file $UNITY_HOME/unity_content/DSVToolkit_v1.1.0.4/build/ITEventsInsightPack_v1.1.0.1/metadata/filesets.json and remove the following text:

```
,{"name":"ITEvents-
Split","type":0,"fileType":0,"fileName":"Dsv.jar","className":
"com.ibm.tivoli.unity.content.insightpack.dsv.extractor.splitter.
DsvSplitter"}
```

d) Open the file $UNITY_HOME/unity_content/DSVToolkit_v1.1.0.4/build/ITEventsInsightPack_v1.1.0.1/metadata/ruleset.json and add the following text:

```
[{"name":"ITEvents-Split","type":0,"rulesFileDirectory":"extractors\/
ruleset\/splitter"}]
```

e) Open the file $UNITY_HOME/unity_content/DSVToolkit_v1.1.0.4/build/ITEventsInsightPack_v1.1.0.1/metadata/sourcetypes.json and change the following text:

```
"splitter":{"fileSet":"ITEvents-Split","ruleSet":null,"type":1}
```

to

```
"splitter":{"fileSet":null,"ruleSet":"ITEvents-Split","type":1}
```

f) Go to $UNITY_HOME/unity_content/DSVToolkit_v1.1.0.4/build and compress the contents of the insight pack directory using the zip command utility to create a new insight pack. Ensure you run the command from the /build directory to preserve the directory structure in the resulting .zip file (in this example, the directory is /ITEventsInsightPack_v1.1.0.1, so the file would be ITEventsInsightPack_v1.1.0.1.zip).

For example:

```
zip -r ITEventsInsightPack_v1.3.0.2.zip ITEventsInsightPack_v1.3.0.2
```

g) Install the insight pack using the $UNITY_HOME/utilities/pkg_mgmt.sh command as described in "Installing the Tivoli Netcool/OMNIbus Insight Pack" on page 61.

7. Test the insight pack:

a) Create a temporary data source in Operations Analytics - Log Analysis for the new Source Type and Collection created by the DSV toolkit.

b) Change the Gateway for Message Bus map file to match the fields that you defined in the ITEvents.properties file.

Important: The order of the column entries must match exactly the order of the alert field entries in the gateway map file.

See the Gateway for Message Bus documentation for information about configuring the map file.

c) In the gateway scalaTransport.properties file, modify the values of the **jsonMsgHostname** and **jsonMsgLogPath** properties to match the attributes of the new data source that you created in step "7.a" on page 177.

d) Test the new configuration.

8. Create a new data source called "omnibus" by using the new source type defined in the ITEvents Insight Pack.

   Important: You cannot rename an existing data source to the default name `omnibus` or use an existing data source that is named `omnibus`. You must delete the existing data source, then create the new data source and name it `omnibus`.

What to do next
Test the new insight pack in the Operations Analytics - Log Analysis UI.

Customizing the Apps
The following procedure describes how to generate customized versions of the Custom Apps provided with the Insight Pack.

Before you begin

Consult the Operations Analytics - Log Analysis documentation for information about creating Custom Apps and building Insight Packs with the Eclipse-based Insight Pack Tooling.

Testing Insight Packs requires a Gateway for Message Bus to transfer events from Netcool/OMNIbus to Operations Analytics - Log Analysis.

About this task

The procedure is divided into two main parts:

1. Create a new Custom App, based on the Custom App provided with the Insight Pack.
2. Use the Operations Analytics - Log Analysis Eclipse-based Insight Pack Tooling to build a new Insight Pack that contains the new Custom App.

Procedure

Create a new Custom App

1. Copy all the files in the `$UNITY_HOME/AppFramework/Apps/OMNIbusInsightPack_`*`version_number`* directory to a new directory under `$UNITY_HOME/AppFramework/Apps/`.

   For example, copy them to `$UNITY_HOME/AppFramework/Apps/ITEvents`.
2. In the `$UNITY_HOME/AppFramework/Apps/ITEvents` directory, rename all the App files from `OMNIbus_*.app` to *`NewName`*`_*.app`.

   For example, rename `OMNIbus_Event_Distribution.app` to `ITEvents_Distribution.app`.
3. Modify `ITEvents_Distribution.app` as required.

   For example customizations, see “Example customization: OMNIbus Static Dashboard” on page 179.
4. Test the new Custom App.

Build a new Insight Pack with the Eclipse Tooling

5. Import the Netcool/OMNIbus Insight Pack, `$UNITY_HOME/unity_content/OMNIbusInsightPack_`*`version_number`*`.zip`, into the Insight Pack Tooling.
6. Use the **Refactor** > **Rename** command to create a new name for the Insight Pack, for example, `ITEvents`.
7. Import your Custom App files from the `$UNITY_HOME/AppFramework/Apps/ITEvents` directory to the `src-files/unity_apps/apps` Tooling directory.
8. When the files are successfully imported, remove the `$UNITY_HOME/AppFramework/Apps/ITEvents` directory.
9. Build the Insight Pack.

   The new Insight Pack is named `ITEventsInsightPack_V1.1.0.1.zip` and is located in the `dist` Tooling directory.

10. Use the following command to install the new Insight Pack:

    ```
    pkg_mgtm.sh -install directory_path/
    ITEventsInsightPack_V1.1.0.1.zip
    ```

    Where *directory_path* is the directory path to the `dist` Tooling directory.

What to do next
Test the new Insight Pack and Custom App.

Example customization: OMNIbus Static Dashboard
The following examples show you how to customize the OMNIbus Static Dashboard app.

Note: Do not directly modify the custom apps supplied with the Insight Pack. Before you implement any of the following examples, create a custom app as described in "Customizing the Apps" on page 178.

- "Removing a chart" on page 179
- "Adding a new chart" on page 179

Removing a chart

To remove a chart from the event distribution dashboard, edit your custom `.app` file (for example, `ITEvents_Distribution.app`) and remove the corresponding JSON element.

For example, to remove the Hotspot by Node and AlertGroup chart from the dashboard, remove the following element from the `charts` array:

```
{
      "type": "Heat Map",
      "title": "Hotspot by Node and AlertGroup",
      "data": {
      "$ref": "AlertGroupVsNode"
       },
      "parameters": {
      "yaxis": "AlertGroup",
      "xaxis": "Node",
      "category": "count"
      }
},
```

When you run the App, the Hotspot by Node and AlertGroup chart is no longer available in the event dashboard.

Adding a new chart

Use the following steps to add a new heat map that shows Hotspot by Node and Location:

1. Open your custom copy of the `OMNIbus_Static_Dashboard.py` file for editing.

   In the **Create a new Custom App** procedure described in "Customizing the Apps" on page 178, this file is in the `$UNITY_HOME/AppFramework/Apps/ITEvents` example directory.
2. To generate the search data, add the line formatted in bold type to the file:

```
dashboardObj.getSingleFieldCount(chartdata, querystring,
 'Node', '5');
dashboardObj.getSingleFieldCount(chartdata, querystring,
'AlertGroup', '5');
dashboardObj.getSingleFieldCount(chartdata, querystring,
'Severity', '5');
dashboardObj.getSingleFieldVsTimestamp(chartdata, querystring,
'AlertGroup', '1');
dashboardObj.getSingleFieldVsTimestamp(chartdata, querystring,
'Severity', '10');
dashboardObj.getSingleFieldVsTimestamp(chartdata, querystring,
'Node', '1');
dashboardObj.getFieldaVsFieldb(chartdata, querystring,
'AlertGroup', 'Node', '5', '5');
dashboardObj.getFieldaVsFieldb(chartdata, querystring,
```

```
'Severity', 'AlertGroup', '5', '5');
dashboardObj.getFieldaVsFieldb(chartdata, querystring,
'Location', 'Node', '5');
```

3. Add the following JSON element to the end of the `charts` array in your custom `.app` file (for example, `ITEvents_Distribution.app`):

```
{
      "type": "Heat Map",
      "title": "Hotspot by Node and Location",
      "data": {
        "$ref": "NodeVsLocation"
        },
      "parameters": {
        "yaxis": "Node",
        "xaxis": "Location",
        "category": "count"
      }
    },
```

The following table lists the values allowed for each JSON object.

Object	Values
`type`	Specifies the chart type. Depending on the functions that you add to `OMNIbus_Static_Dashboard.py`, the following chart types are available: • The `getSingleFieldCount` function enables the following types: `Line Chart`, `Pie Chart`, `Simple Bar Chart`, `Point Chart`. • The `getSingleFieldVsTimeStamp` function enables the following types: `Bubble Chart`, `Bar Chart`, `Two Series Line Chart`, `Stacked Bar Chart`, `Heat Map`, `Cluster Bar`, `Stacked Line Chart`. • The `getFieldaVsFieldb` function enables the following types: `Bubble Chart`, `Bar Chart`, `Two Series Line Chart`, `Stacked Bar Chart`, `Heat Map`, `Cluster Bar`, `Stacked Line Chart`.
`title`	You can specify a title of your choice.
`$ref`	Specifies the data reference. Depending on the functions that you add to `OMNIbus_Static_Dashboard.py`, the following data reference values are available: • The `getSingleFieldCount` function enables the following value: *`FieldName`*`Count` • The `getSingleFieldVsTimeStamp` function enables the following value: *`FieldName`*`VsTime` • The `getFieldaVsFieldb` function enables the following value: *`FieldNameA`*`Vs`*`FieldNameB`* Where *`FieldName`* is case-sensitive and matches the event field name in the index configuration of your source type.

Object	Values
`parameters`	Specifies the `yaxis`, `xaxis`, and `categories` chart parameters. Depending on the functions that you add to `OMNIbus_Static_Dashboard.py`, the following chart parameter values are available: • The `getSingleFieldCount` function enables the following values: *`FieldName`*`, count`. • The `getSingleFieldVsTimeStamp` function enables the following values: *`FieldName`*`, date, count`. • The `getFieldaVsFieldb` function enables the following values: *`FieldNameA`*`, `*`FieldNameB`*`, count`. Where *`FieldName`* is case-sensitive and matches the event field name in the index configuration of your source type.

When you run the App, the new Node over AlertGroup chart is displayed in the event dashboard.

Customizing dynamic dashboards
You can create new dynamic dashboards, or customize the OMNIbus Dynamic Dashboard and OMNIbus_Operational_Efficiencyapps, and the apps in the Event_Analysis_And_Reduction wizard. You can change the data source and time range filters of these apps.

For more information about creating new dynamic dashboards and customizing the apps, see the following topics, depending on your version of Operations Analytics - Log Analysis:

- V1.3.5: https://www.ibm.com/support/knowledgecenter/SSPFMY_1.3.5/com.ibm.scala.doc/use/scla_extend_create_dboard_t.html.
- V1.3.3: https://www.ibm.com/support/knowledgecenter/SSPFMY_1.3.3/com.ibm.scala.doc/use/scla_extend_create_dboard_t.html

Important: Do not directly modify the custom apps that are supplied with the Insight Pack. Before you customize an app, create a custom app as described in "Customizing the Apps" on page 178.

To avoid inconsistent results, change the data source or the time range filter for each chart in your custom `.app` file.

Configuring Network Management

Perform the following tasks to configure the components of Network Management.

Configuring Topology Search

You can configure and customize the Network Manager Insight Pack to meet your requirements.
Related concepts
Network Management tasks
Use this information to understand the tasks that users can perform using Network Management.

Verifying that the Insight Pack was installed
Run the **pkg_mgmt** command to verify that the Insight Pack was installed, or check the Operations Analytics - Log Analysis.

Procedure

- Run the **pkg_mgmt** as follows:

```
$SCALA_HOME/utilities/pkg_mgmt.sh -list
```

Search the results for a line similar to the following example. In this example, V1.3.0.0 is installed.

```
[packagemanager] NetworkManagerInsightPack_V1.3.0.0        /home/myhome/IBM/
LogAnalysis/unity_content/NetworkManager
```

- On the Operations Analytics - Log Analysis UI, check for the following items in the **Search Dashboards** area: **NetworkManagerInsightPack** > **Find events between two nodes on layer 3 topology** or **NetworkManagerInsightPack** > **Find events between two nodes on layer 2 topology**

Event annotations
The event annotations that are defined by the Network Manager Insight Pack index configuration file. These fields are the same as those in the Tivoli Netcool/OMNIbus Insight Pack.

For more information, see "Netcool/OMNIbus Insight Pack" on page 218.

Related tasks
Customizing the index configuration

Configuring topology search apps for use with Oracle databases
If you are using an Oracle database for topology storage (that is, as the NCIM database), obtain the Oracle driver and point to the driver in the custom apps. You can skip this task if you are using a Db2 database. You can make this configuration while the products and Insight Pack are still running.

Before you begin
Install and configure the supported products, and install and configure the Network Manager Insight Pack

Procedure

1. Obtain and install the Oracle driver from the Oracle website.
2. On the Operations Analytics - Log Analysis host, save the driver to $*UNITY_HOME*/wlp/usr/servers/Unity/apps/Unity.war/WEB-INF/lib.
3. Point the custom apps to the driver:
   a) Open the following files for editing:
      - $*UNITY_HOME*/AppFramework/Apps/NetworkManagerInsightPack_v1.3.0.0/Network_Topology_Search/NM_Show_Alerts_Between_Two_Nodes_Layer2.sh
      - $*UNITY_HOME*/AppFramework/Apps/NetworkManagerInsightPack_v1.3.0.0/Network_Topology_Search/NM_Show_Alerts_Between_Two_Nodes_Layer3.sh

   b) Insert the path to the Oracle driver into the **classpath** parameter.
   For example, change it from this **classpath**:

   ```
   $UNITY_LIB_PATH/db2jcc.jar:$UNITY_LIB_PATH/log4j-1.2.16.jar
   ```

   To this **classpath**:

   ```
   To: $UNITY_LIB_PATH/db2jcc.jar:$UNITY_LIB_PATH/ojdbc14.jar:
   $UNITY_LIB_PATH/log4j-1.2.16.jar
   ```

4. In the $UNITY_HOME/AppFramework/Apps/NetworkManagerInsightPack_V1.3.0.0/Network_Topology_Search/NM_EndToEndSearch.properties file, which point the Insight Pack to the NCIM database, ensure that the following properties for Oracle are correctly set.

   ncp.dla.datasource.type
   Enter oracle.

   ncp.dla.datasource.driver
   Enter the driver name, for example oracle.jdbc.driver.OracleDriver.

   ncp.dla.datasource.url
   Enter the URL, in the format jdbc:oracle:thin:@*host*:*port*:*name*, where *name* is the name of the database. For example, jdbc:oracle:thin:@192.168.1.2:1521:itnm

Customizing the index configuration
Because the rule set, source type, and collection are provided in the OMNIbusInsightPack_v1.3.1, any customizations must be made to that Insight Pack, not the Network Manager Insight Pack. Refer to the *OMNIbusInsightPack_v1.3.1 README* for more information about customizing that Insight Pack.

For more information, see "Netcool/OMNIbus Insight Pack" on page 218.

Related reference

Event annotations

Related information

Operations Analytics Log Analysis V1.3.5 documentation: Configuring the DSV toolkit The DSV toolkit is used to create Insight Packs that allow you to load Delimiter Separated Value (DSV) data into Operations Analytics - Log Analysis.

Operations Analytics Log Analysis V1.3.3 documentation: Configuring the DSV toolkit The DSV toolkit is used to create Insight Packs that allow you to load Delimiter Separated Value (DSV) data into Operations Analytics - Log Analysis.

Configuring integration to IBM Connections

IBM Tivoli Netcool/Impact provides integration to IBM Connections by using a Netcool/Impact `IBMConnections` action function. The `IBMConnections` action function allows users to query forums and topics lists, create a new forum, create a new topic, and update existing topics. The `IBMConnections` action function package is available in the directory `$IMPACT_HOME/integrations/IBMConnections`.

About this task

Complete the following steps to configure Netcool/Impact to integrate with the IBM Connections Server.

Procedure

1. Go to the directory `$IMPACT_HOME/add-ons/IBMConnections`, this directory is the IBM Connections integration package. The package includes the following subdirectory:

   importData
   A project that includes policies, data source, data type, and a service. The project serves an example to show how to connect, create, update, and topics in IBM Connections
2. Import the project `$IMPACT_HOME/bin/nci_ import <ServerName> __Extraction_Directory__/importData`.
3. Within the `$IMPACT_HOME/etc/<NCI>_server.props` file, add the following parameters:

   `impact.ibmconnections.forum.title.maxsize=` *number*. Default value is 0. Any string size can be used.
   `impact.ibmconnections.forum.content.maxsize=` *number*. Default value is 0. Any string size can be used
   `impact.ibmconnections.topic.title.maxsize=` *number*. Default value is 255, for a topic and a reply.
   `impact.ibmconnections.topic.content.maxsize =` *number*. Default value is 0. Any string size can be used
4. Restart Netcool/Impact servers.

Related tasks

Installing Netcool/OMNIbus and Netcool/Impact

IBM Connections Overview

IBM Connections is a leading social software platform that can help your organization to engage the right people, accelerate innovation, and deliver results.

This integrated, security-rich platform helps people engage with networks of experts in the context of critical business processes. Now everyone can act with confidence and anticipate and respond to emerging opportunities

For information about IBM Connections, refer to

http://www-03.ibm.com/software/products/en/conn

Parameters for the `IBMConnections` function

The integration between Netcool/Impact and IBM Connections uses a new policy action function that is called the `IBMConnections` function. The `IBMConnections` function can be used within any policy to connect to the IBM Connections Server and to perform an action on the IBM Connections Server. The function accepts two input parameters, the **`Action Option`** parameter and the **`Impact Object`** parameter.

`Action Option` Parameter

The **`Action Option`** parameter accepts one of the following action entries, which are case insensitive. Some action entries require property information that is case-sensitive.

In the content of an IBM Connections forum, topic, or reply, you can use HTML formatting tags `br`, b, and a. For more information about the supported HTML tags, see http://www-03.ibm.com/software/products/en/conn.

CREATEFORUM

Creates a forum.

Enter the following property information that is case-sensitive. The tags must be created before they pass to a variable name.

`props.ForumTitle=`*`title`*`;`
`props.ForumContent=`*`full text of the body`*`;`
`props.ForumTags=List_Of_Tags;` Is optional, the object must be a Netcool/Impact object.
`Tags=NewObject();`
`Tags.Tag1=`*`some tag`*`;`
`Tags.Tag`*`2`*`=`*`some tag2`*`;` Is optional if want more than one tag.

CREATETOPIC

Creates a topic.

Enter the following property information that is case-sensitive:

`props.TopicTitle=`*`title`*`;`
`props.TopicContent=`*`full text of the body`*`;`
`props.ForumId=`*`forum id`*`:` Where the *forum id* is an ID and not a forum name.

DELETEFORUM

Deletes the forum name that was created by the logged in user and any topic or reply belonging to it.

Enter the following property information that is case-sensitive:

`props.ForumId=`*`forumId`*`;` Or `props.ForumId=`*`forumTitle`*`;`
`props.FirstMatchOnly=`*`true`*`;` Or `props.FirstMatchOnly=`*`false`*`;` The `props.FirstMatchOnly` property deletes the first matching forum or matching topic that it finds, or else it deletes any matching forum or matching topic and its default value is `true`.

DELETEPUBLICFORUM

Deletes the given public forum name and any topic or reply belonging to it.

Enter the following property information that is case-sensitive:

`props.ForumId=forumId`; Or `props.ForumId=forumTitle`;
`props.FirstMatchOnly=true`; Or `props.FirstMatchOnly=false`;The `props.FirstMatchOnly` property deletes the first matching forum or matching topic that it finds, or else it deletes any matching forum or matching topic and its default value is `true`.

DELETEREPLY

Deletes a topic reply in the topic in the forum id.

Enter the following property information that is case-sensitive:

`props.ForumId=forumId`; Where the *forumId* is an ID and not a title.
`props.TopicTitle=title`; Or `props.TopicTitle=id`;
`props.ReplyTitle=replytitle`; Or `props.ReplyTitle=replyid`;
`props.FirstMatchOnly=true`; Or `props.FirstMatchOnly=false`;The `props.FirstMatchOnly` property deletes the first matching forum or matching topic that it finds, or else it deletes any matching forum or matching topic and its default value is `true`.

DELETETOPIC

Deletes a topic in the forum id.

Enter the following property information that is case-sensitive:

`props.ForumId=forumId`; Where the *forumId* is an ID and not a title.
`props.TopicTitle=title`; Or `props.TopicTitle=id`;
`props.FirstMatchOnly=true`; Or `props.FirstMatchOnly=false`;The `props.FirstMatchOnly` property deletes the first matching forum or matching topic that it finds, or else it deletes any matching forum or matching topic and its default value is `true`.

GETCOMMUNITYFORUMID

Gets the ID of the community that is created by the logged in user

Enter the following property information that is case-sensitive:

`props.CommunityName=Community name`;
`props.ForumName=forum name`;

GETFORUMTOPICS

Gets list of topics for the forum id

Enter the following property information that is case-sensitive:

`props.ForumId=forum id`;

GETMYCOMMUNITYID

Gets the ID of the community for the logged in user

Enter the following property information that is case-sensitive:

`props.CommunityName=Community name`;

GETMYFORUMID

Gets the ID of the forum that is created by the logged in user.

Enter the following property information that is case-sensitive:

`props.ForumName=actual forum name that is created by the logged in user`

GETMYFORUMS

Gets all the forums for the logged in user

GETPUBLICCOMMUNITYID

Gets the ID of the given public community ID

Enter the following property information that is case-sensitive:

`props.CommunityName=Community name`;

GETPUBLICFORUMID

Gets the ID of the given public forum name

Enter the following property information that is case-sensitive:

```
props.ForumName=actual public forum name
```

GETPUBLICFORUMS
Gets list of all public forums

GETTOPICREPLIES
Gets list of replies for a topic

Enter the following property information that is case-sensitive:

`props.TopicId=topic id`; Where *topic id* must be the topic id not the topic name.

REPLYTOTOPIC
Creates a reply to an existing topic

Enter the following property information that is case-sensitive:

`props.TopicId=topic id`; Where *topic id* is a topic ID not a topic name
`props.ReplyTitle=title`;
`props.ReplyContent=full text of the body`;

`Impact Object` Parameter
The **`Impact Object`** parameter accepts the following property information. The authentication, and connection property information is mandatory.

```
props = NewObject();
props.Protocol=https;
props.Host=IBM Connections Server Host/IP;
props.Port=_PORT_;
props.Username=userName;
```

`props.Password=password`; The password can be encrypted by using either the Netcool/Impact `nci_crypt` tool or the policy function `Encrypt()`. If the password is encrypted, you must use the property `props.DecryptPassword=true;`

`IBMConnections` Project and artifacts

The imported `IBMConnections` project includes artifacts that are categorized into four categories.

Data sources

- `IBMConnectionsObjectServerDSA`
- `Internal`

Data types

- `TopicCreationTracker`
  - An internal data type that is used to track the topic creation to avoid duplicate names.
- `InternetOutageEvents`
  - ObjectServer data type that can be used to view the critical events in the UI Data Provider widgets. It populates the severity as status data type to show colorful images.

Policies

- `IBMConnectionsUtils`
  - Includes a utility function to extract value from a data item.
- `IBMConnectionsUtilsCaller`
  - Shows how to call the utility function in `IBMConnectionsUtils`.

- IBMConnectionsUtilsJS
  - For JavaScript policies.
- IBMConnectionsUtilsCallerJS
  - For JavaScript policies.
- NetworkMonitorExample
  - Example policy that is run by an event reader to create and update topics.
- NetworkMonitorForOpView
  - Example policy that is run by the operator view from the ObjectServer Event List tool.
- Opview_IBMConnectionsOpView
  - Is run by the AEL tool.

Services

- NetworkMonitorExample
  - Connects to the object server data source and uses a default filter of Node in ('US','France','UK')and Identifier Like 'Monitoring Network for'. You can change the default filter at any time. It runs the policy NetworkMonitorExample.

Automatic topic management

The IBM Connections integration package includes NetworkMonitorExample event reader service that connects to the Netcool/OMNIbus Object Server and filters for specific events. When there is a match, the service runs the policy that either updates the topic by sending a reply to the topic, or creates a topic if a topic does not exist.

The forum that is used in this example is the same name as the AlertKey in the Object Server event, forumName = @IBMConnections_Forum

You can use the sql scripts in the $IMPACT_HOME/integrations/IBMConnections/db directory to create extra fields in the Object Server or you can use existing fields.

The topic title is created as a combination of a hardcoded string and the node field from the event:

```
topicTitleVar="Network Monitor is down on node: @Node@" ;
IBMConnectionsUtils.extractParametersAndSubstitute
(topicTitleVar,EventContainer,result);
topicTitle = result;
```

The policy checks if the topic was created by querying the internal data type TopicCreationTracker. If the topic exists, the policy sends a reply instead of creating a new one.

Automatic topic management with event management tools

The operator view policy is updated to run the NetworkMonitorForOpView policy that automatically updates an existing topic or creates a new topic.

Procedure

1. Create the event management tool, refer to the Netcool/OMNIbus documentation.
2. In the event management tool, select the executable box tab and enter the following text: start "" "https://<impactgui_server>:<port>/opview/displays/NCICLUSTER-IBMConnectionsOpView.html?Node=@Node&Serial=@Serial&Severity=@Severity&Acknowledged=@Acknowledged&AlertKey=@AlertKey&AlertGroup=@AlertGroup&Summary=@Summary"

   NCICLUSTER Is the default cluster but if you are using a different cluster name then update the URL with your cluster name.

The above text is an example of text to enter for an Object Server on Windows.

3. Add the new event management tool to the `AlertsStatus` tools menu, refer to the Netcool/OMNIbus documentation.
4. When you right-click on an event, click the tool to start the URL and run the policy. The operator view is started and gives a notification that the topic is created or updated, along with a link to the topic URL to use.

Related information

Creating event management tools

Enabling historical events

Create a connection to the historical database in Impact to view historical events in the Event Viewer.

Procedure

1. Log in to Dashboard Application Services Hub and select the **Data Model** tab.
2. Click the **New Data Source** icon.
3. Point to **Database SQL** and select your database type. For example **Db2**.
4. In the **Data Source Name** field enter: **`historicalEventsDatasource`**.
5. Enter your Username and Password in the fields provided and click the **Save** icon.
6. In the left-hand navigation pane, right-click **ImpactHistoricalEventData** and select **New Data Type**.
7. In the **Data Type Name** field enter: **`historicalEventData`**.
8. Click **Refresh**.

Chapter 6. Connecting to event sources

Learn how to connect to event sources.

Connecting event sources to your IBM Netcool Operations Insight on premises deployment

Your IBM Netcool Operations Insight deployment provides rich capabilities to integrate event sources from virtually any event source across local, hybrid, and cloud environments.

Set up fast and simple integrations to connect private or public cloud event sources and view data in Web GUI.

For information about connecting local event sources, using probes and gateways, to your Operations Management on IBM Cloud Private deployment, see “Connecting event sources to your Operations Management on IBM Cloud Private deployment” on page 190.

For more information about connecting local event sources, using probes and gateways, to your IBM Netcool Operations Insight on premises deployment, see “Quick reference to installing” on page 49.

Connecting event sources from your private cloud environment

Learn how to integrate cloud events and view cloud event data in Web GUI.

Complete the following steps to integrate private cloud events with your IBM Netcool Operations Insight on premises deployment:

1. Download and install the IBM Tivoli Netcool/OMNIbus Probe for Message Bus. To download the probe, see Netcool/OMNIbus documentation: Generic integrations using the Message Bus Probe. To install the probe, see Netcool/OMNIbus documentation: Installing probes and gateways on Tivoli Netcool/OMNIbus V8.1.
2. Configure the Message Bus Probe with the required ObjectServer fields. For more information, see Netcool/OMNIbus documentation: Integrating IBM Cloud Event Management with Netcool Operations Insight.
3. Install Cloud Event Management in your IBM Cloud Private environment. For more information see Cloud Event Management documentation: Installing in IBM Cloud Private.
4. Create an outgoing integration. For more information, see Cloud Event Management documentation: Sending events to Netcool/OMNIbus.
5. Configure an event policy to forward events to Netcool Operations Insight. In the **Action** section, select **Forward events**. For more information, see Cloud Event Management documentation: Setting up event policies.

It may take a moment before Cloud Event Management starts to forward events to Netcool Operations Insight. Verify that Cloud Event Management events appear in the **Event List**.

Sending events to Cloud Event Management

You can also view Netcool Operations Insight events in Cloud Event Management, by creating an incoming integration and installing and configuring the Netcool/OMNIbus Gateway for Cloud Event Management. For more information, see Cloud Event Management documentation: Configuring Netcool/OMNIbus as an event source.

Connecting event sources from your public cloud environment

Learn how to integrate public cloud events and view cloud event data in Web GUI.

Complete the following steps to connect public cloud events with your IBM Netcool Operations Insight on premises deployment:

1. Download and install the IBM Tivoli Netcool/OMNIbus Probe for Message Bus. To download the probe, see Netcool/OMNIbus documentation: Generic integrations using the Message Bus Probe. To install the probe, see Netcool/OMNIbus documentation: Installing probes and gateways on Tivoli Netcool/OMNIbus V8.1.
2. Configure the Message Bus Probe with the required ObjectServer fields. For more information, see Netcool/OMNIbus documentation: Integrating IBM Cloud Event Management with Netcool Operations Insight.
3. Subscribe to IBM Cloud Event Management for your public cloud environment. For more information, see IBM Cloud Event Management on Marketplace.
4. Create an outgoing integration. For more information, see Cloud Event Management documentation: Sending events to Netcool/OMNIbus.
5. Configure an event policy to forward events to Netcool Operations Insight. In the **Action** section, select **Forward events**. For more information, see Cloud Event Management documentation: Setting up event policies.
6. Configure the secure gateway. For more information, see Cloud Event Management documentation: Sending events to Netcool/OMNIbus via the IBM Secure Gateway.

It may take a moment before Cloud Event Management starts to forward events to Netcool Operations Insight. Verify that Cloud Event Management events appear in the **Event List**.

Sending events to Cloud Event Management

You can also view Netcool Operations Insight events in Cloud Event Management, by creating an incoming integration and installing and configuring the Netcool/OMNIbus Gateway for Cloud Event Management. For more information, see Cloud Event Management documentation: Configuring Netcool/OMNIbus as an event source.

Connecting event sources to your Operations Management on IBM Cloud Private deployment

After you successfully deploy Netcool Operations Insight on IBM Cloud Private, connect your event source, which is usually an IBM Tivoli Netcool/OMNIbus on-premises Probe or Gateway. You can connect directly to the ObjectServer NodePort or you can connect with a proxy NodePort.

You can configure connections to Operations Management on IBM Cloud Private in two ways. The primary method to connect your event sources is to make a direct Transmission Control Protocol (TCP) connection to the ObjectServer NodePort. This method supports plain text connections. If Transport Layer Security (TLS) encryption is required, you can connect with a proxy NodePort. This method supports plain text and TLS encrypted connections.

Connecting with the proxy NodePort

Learn how to connect to the ObjectServer from outside the IBM Cloud Private deployment by using the secure connection proxy with Transport Layer Security (TLS) encryption.

The proxy provides a secure TLS encrypted connection for clients that require a direct Transmission Control Protocol (TCP) connection to the ObjectServer instance running on IBM Cloud Private. Typically, clients such as Netcool/OMNIbus Probes and Gateways require this type of connection. The clients can be installed in a traditional on-premise installation or deployed in another IBM Cloud Private cluster.

The proxy is deployed automatically as part of the ibm-netcool-prod deployment. By default, the proxy is deployed with a TLS certificate, which is automatically created and signed by the IBM Cloud Private cluster Certificate Authority (CA) during deployment. However, it is also possible to use a custom certificate that has been signed by an external CA.

For more information about configuring the proxy and the proxy config map, see “Proxy configmap” on page 549.

Identifying the proxy listening port

To connect to the IBM Tivoli Netcool/OMNIbus Object Server pair from outside the IBM Cloud Private cluster with Transport Layer Security (TLS) encryption, you must identify the externally accessible NodePort where the proxy listens for connections.

About this task

The proxy defines a Kubernetes service, called *helm_release_name*-proxy, where *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment. The *helm_release_name*-proxy service defines the NodePorts that clients must use when connecting to the Object Server pair.

Procedure

1. Describe the proxy service by running the following command:

   ```
   kubectl get service -o yaml helm_release_name-proxy -n namespace
   ```

   Where

   - *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment.
   - *namespace* is the name of the namespace in which Operations Management on IBM Cloud Private is installed.

2. Identify the NodePorts from the command output, for example:

   ```
   ports:
   - name: aggp-proxy-port
     nodePort: 30135
     port: 6001
     protocol: TCP
     targetPort: 6001
   - name: aggb-proxy-port
     nodePort: 30456
     port: 6002
     protocol: TCP
     targetPort: 6002
   ```

   In the example, the NodePort for the primary Object Server is 30135 and the NodePort for the backup Object Server is 30456.

Results

Make a note of the NodePorts that you identified. This information is required when configuring the client's Secure Sockets Layer (SSL) connection. For more information, see “3” on page 192 in the “Configuring TLS encryption with the default certificate” on page 191 topic.

Configuring TLS encryption with the default certificate

The proxy requires a public certificate and private key pair to be supplied through a Kubernetes secret called *{{ .Release.Name }}*-proxy-tls-secret, where *{{ .Release.Name }}* is the unique name that is assigned to the deployment. If the secret does not exist, the secret key pair is automatically created during deployment. Certficates that are created automatically are signed by the IBM Cloud Private Certificate Authority (CA). To enable a successful Transport Layer Security (TLS) handshake, import the default CA signer certificate into the keystore of any client application as a trusted source.

About this task

Follow this procedure when the proxy certificate has been automatically created and signed by the IBM Cloud Private cluster CA during deployment.

Procedure

1. Extract the cluster certificate by running the following command:

```
kubectl get secret -n kube-system cluster-ca-cert -o go-template='{{ index .data
"tls.crt" }}' | base64 --decode > cluster-ca-cert.pem
```

2. **Note:** For a successful TLS connection, only the IBM Cloud Private cluster CA signer certificate is required. You do not need to add the proxy certificate to the client keystore.

   To establish a successful TLS connection, import the default IBM Cloud Private cluster CA signer certificate into the client keystore as a trusted certificate. Complete the following steps:

   a. If necessary, create the keystore by using one of the following commands:

   ```
   $NCHOME/bin/nc_ikeyman
   ```

   or

   ```
   •    $NCHOME/bin/nc_gskcmd -keydb -create -db "$NCHOME/etc/security/keys/omni.kdb" -pw
   password -stash -expire 366
   ```

   For more information about creating a keystore, see https://www.ibm.com/support/knowledgecenter/en/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/omnibus/wip/install/task/omn_con_ssl_creatingkeydbase.html.

   b. Import a privacy enhanced mail (PEM) encoded signer certificate by running one of the following commands:

   ```
   $NCHOME/bin/nc_ikeyman
   ```

   or

   ```
   •    $NCHOME/bin/nc_gskcmd -cert -add -file mycert.pem -db $NCHOME/etc/security/keys/
   omni.kdb -stashed
   ```

   For more information about adding certificates from CAs, see https://www.ibm.com/support/knowledgecenter/en/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/omnibus/wip/install/task/omn_con_ssl_addingcerts.html.

3. **Note:** To successfully complete the TLS handshake and establish a secure TLS connection, the ObjectServer address, which is specified in the `omni.dat` file, must exactly match the certificate subject Common Name (CN) value. Certificates that are automatically generated have a subject CN value in the following format:

   ```
   proxy.{{ .Release.Name }}.{{global.cluster.fqdn}}
   ```

   Where *{{ .Release.Name }}* is the name of the Helm release and *{{global.cluster.fqdn}}* is the fully qualified domain name (FQDN) of the cluster's master node. For more information about the FQDN, see Master node.

   Edit the client's `omni.dat` file to configure a Secure Sockets Layer (SSL) connection. Specify the SSL for each ObjectServer entry and add the server address and port number in the `omni.dat` file, as displayed in the following example:

   ```
   [AGG_P]
   {
   Primary: proxy.{{ .Release.Name }}.{{global.cluster.fqdn}} ssl 3XXXX
   }
   [AGG_B]
   {
   Primary: proxy.{{ .Release.Name }}.{{global.cluster.fqdn}} ssl 3XXXX
   }
   ```

   For more information, see “Identifying the proxy listening port” on page 191.

4. Run the following command to generate the interfaces file:

   ```
   $NCHOME/bin/nco_igen
   ```

Configuring TLS encryption with a custom certificate

The proxy requires a public certificate and private key pair to be supplied through a Kubernetes secret called `{{ .Release.Name }}-proxy-tls-secret`. If you want to use a custom certificate, for example, one signed by your own public key infrastructure Certificate Authority (CA), create your own proxy secret, containing the public certificate and private key pair, before deployment. To enable a successful Transport Layer Security (TLS) handshake, import the CA signer certificate into the keystore of any client application as a trusted source.

Before you begin

Note: If you deployed IBM Netcool Operations Insight on IBM Cloud Private V3.2.1, configure TLS encryption with the default certificate. For more information, see "Configuring TLS encryption with the default certificate" on page 191.

Before deploying Operations Management on IBM Cloud Private, you can create your own certificate key pair and create the proxy TLS secret by completing the following steps:

About this task

|Follow this procedure when the public certificate and private key have already been created and signed by an external CA. When creating the certificate, it is important to ensure that the subject Common Name (CN) field matches the following format:

```
proxy.{{ .Release.Name }}.{{global.cluster.fqdn}}
```

Where *{{ .Release.Name }}* is the name of the Helm release and *{{global.cluster.fqdn}}* is the fully qualified domain name (FQDN) of the cluster's master node. For more information about the FQDN, see Master node.

Procedure

1. Set the `global.tls.certificate.useExistingSecret` global property in the Helm chart.
2. Create the proxy TLS secret by running the following command:

   ```
   kubectl create secret tls {{ .Release.Name }}-proxy-tls-secret --cert=certificate.pem --
   key=key.pem [--namespace namespace]
   ```

   Where:

   - *{{ .Release.Name }}* is the unique name that is assigned to the deployment, for example, **noi**.
   - **certificate.pem** is the signed certificate returned by the CA.
   - **key.pem** is the private key corresponding to the signed certificate.
3. To establish a successful TLS connection, import the CA public certificate, which is used in step "2" on page 193. Complete the following steps:

   a. If necessary, create the keystore using one of the following commands:

      ```
      $NCHOME/bin/nc_ikeyman
      ```

      or

      ```
      •    $NCHOME/bin/nc_gskcmd -keydb -create -db "$NCHOME/etc/security/keys/omni.kdb" -pw
      password -stash -expire 366
      ```

      For more information about creating a keystore, see https://www.ibm.com/support/knowledgecenter/en/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/omnibus/wip/install/task/omn_con_ssl_creatingkeydbase.html.

   b. Import a privacy enhanced mail (PEM) encoded signer certificate by running one of the following commands:

      ```
      $NCHOME/bin/nc_ikeyman
      ```

or

```
•    $NCHOME/bin/nc_gskcmd -cert -add -file mycert.pem -db $NCHOME/etc/security/keys/
omni.kdb -stashed
```

For more information about adding certificates from CAs, see https://www.ibm.com/support/knowledgecenter/en/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/omnibus/wip/install/task/omn_con_ssl_addingcerts.html.

4. **Note:** To successfully complete the TLS handshake and establish a secure TLS connection, the ObjectServer address, which is specified in the `omni.dat` file, must exactly match the certificate subject Common Name (CN) value. Certificates that are manually created must have a subject CN value in the following format:

```
proxy.{{ .Release.Name }}.{{global.cluster.fqdn}}
```

Edit the client's `omni.dat` file to configure a Secure Sockets Layer (SSL) connection. Specify the SSL for each Object Server entry and add the server address and port number in the `omni.dat` file, as displayed in the following example:

```
[AGG_P]
{
Primary: proxy.{{ .Release.Name }}.{{global.cluster.fqdn}} ssl 3XXXX
}
[AGG_B]
{
Primary: proxy.{{ .Release.Name }}.{{global.cluster.fqdn}} ssl 3XXXX
}
```

For more information, see “Identifying the proxy listening port” on page 191.

5. Run the following command to generate the interfaces file:

```
$NCHOME/bin/nco_igen
```

What to do next

For more information, see “Preparing secrets for TLS encryption” on page 105.

Related tasks

Preparing secrets for TLS encryption

Operations Management on IBM Cloud Private provides the ability to automatically generate TLS certificates for customers who do not have their own Certificate Authority (CA). However, if you have your own CA and you want to deploy on your own site then you can manually create a certificate and use your own CA to sign the certificate.

Disabling TLS encryption

To disable Transport Layer Security (TLS) encryption, edit the proxy configmap.

Procedure

1. Open the proxy configmap for editing.

```
kubectl edit configmap helm-release-name-proxy-config -n namespace
```

Where

- *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment.
- *namespace* is the name of the namespace in which Operations Management on IBM Cloud Private is installed.

2. Find the tlsEnabled flag and set it to false. Save and exit the configmap.

```
tlsEnabled: "false"
```

3. Find the proxy pod using the command

```
kubectl get pods --namespace namespace  | grep proxy
```

Where *namespace* is the name of the namespace in which Operations Management on IBM Cloud Private is installed.

4. Restart the proxy pod.

```
kubectl delete pod proxy-pod -n namespace
```

Where

- *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment.
- *proxy-pod* is the name of the proxy pod in your Operations Management on IBM Cloud Private deployment.
- *namespace* is the name of the namespace in which Operations Management on IBM Cloud Private is installed.

Connecting with the ObjectServer NodePort

Learn how to connect to the IBM Tivoli Netcool/OMNIbus ObjectServer fail over pair from outside the IBM Cloud Private deployment.

Before you begin

The ObjectServer NodePorts do not support Transport Layer Security (TLS) encryption. If TLS encryption is required, connect with a proxy NodePort. For more information, see "Connecting with the proxy NodePort" on page 190.

About this task

Learn how to make a direct plain text Transmission Control Protocol (TCP) connection to the ObjectServer failover pair running in a IBM Netcool Operations Insight on IBM Cloud Private deployment. Typically clients such as Netcool/OMNIbus Probes and Gateways require this type of connection to write event data into the deployment. Connection to the ObjectServer pair is enabled with cluster NodePorts which are defined by the following services:

```
helm_release_name-objserv-agg-primary-nodeport
helm_release_name-objserv-agg-backup-nodeport
```

Where *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment.

Procedure

1. Describe the primary ObjectServer service by running the following command:

```
kubectl get service helm_release_name-objserv-agg-primary-nodeport -n namespace
```

Where

- *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment.
- *namespace* is the name of the namespace in which Operations Management on IBM Cloud Private is installed.

2. Describe the backup ObjectServer service by running the following command:

```
kubectl get service helm_release_name-objserv-agg-backup-nodeport -n namespace
```

Where

- *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment.

- *namespace* is the name of the namespace in which Operations Management on IBM Cloud Private is installed.

3. Identify the NodePorts from the command output, for example:

```
NAME                                                  TYPE       CLUSTER-IP   EXTERNAL-IP
PORT(S)                          AGE
helm_release_name-objserv-agg-primary-nodeport   NodePort   10.0.0.18    <none>
4100:32312/TCP,31581:31581/TCP   3h
```

and

```
NAME                                                  TYPE       CLUSTER-IP   EXTERNAL-IP
PORT(S)                          AGE
helm_release_name-objserv-agg-backup-nodeport   NodePort   10.0.0.162   <none>
4100:30404/TCP,30302:30302/TCP   3h
```

In the examples, the NodePort for the primary ObjectServer is 32312 and the NodePort for the backup ObjectServer is 30404.

4. Combine the NodePort values with the public domain name or IP address of the cluster to form the address of the ObjectServer. Edit the client's `omni.dat` file to add entries for the primary and backup ObjectServer, as described in the following example:

```
[AGG_P]
{
    Primary: mycluster.icp 32312
}
[AGG_B]
{
    Primary: mycluster.icp 30404
}
```

Where *mycluster.icp* is the public domain name of the cluster.

5. Run the following command to generate the interfaces file:

```
$NCHOME/bin/nco_igen
```

Connecting an on-premises IBM Tivoli Netcool/OMNIbus ObjectServer to Operations Management on IBM Cloud Private

Learn how to configure a connection between an on-premises ObjectServer and a Netcool Operations Insight deployment on IBM Cloud Private.

After you successfully deploy Operations Management on IBM Cloud Private, you can connect to an existing on-premises installation to create an event feed between your on-premises and cloud installations. A uni-directional gateway allows event data to flow in a single direction, either from on-premises to cloud, or from cloud to on-premises. A bidirectional gateway can be configured to allow event data to flow in both directions at the same time.

The following figure shows the architecture.

Figure 9. Architecture of an on-premises ObjectServer connected to an Operations Management on IBM Cloud Private deployment.

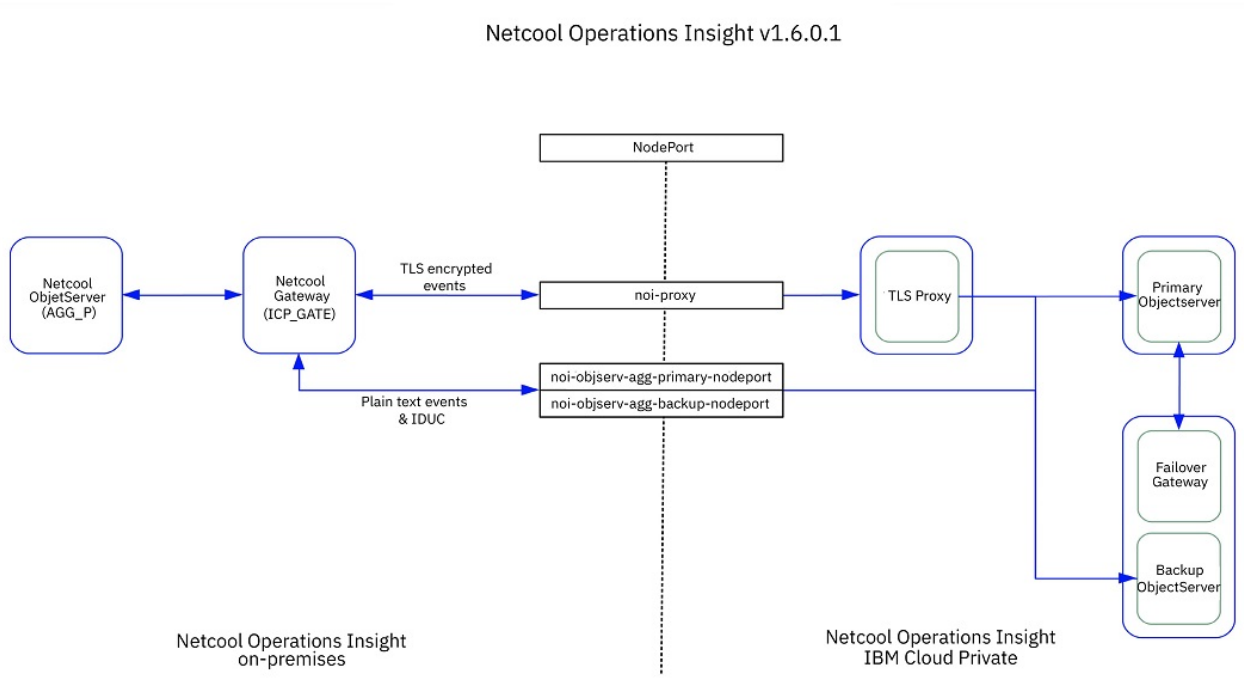

Configuring a uni-directional gateway

About this task

Learn how to configure a uni-directional gateway to create an event feed from an on-premises ObjectServer to a deployment of Operations Management on IBM Cloud Private. An on-premises primary aggregation ObjectServer must exist and be configured according to Installing a primary aggregation ObjectServer.

Procedure

1. Identify the NodePort details of the IBM Cloud Private deployment's ObjectServers.

   Follow the instructions in “Connecting with the ObjectServer NodePort” on page 195 to identify the primary and backup ObjectServer node port details.
2. On the on-premises host, edit the `omni.dat` file and add an entry for a new uni-directional gateway, `'ICP_GATE'`. Add entries for the source and destination ObjectServers, `AGG_P` and `ICP_AGG_V`. The `ICP_AGG_V` entry has the cloud ObjectServer node port details that you identified in step 1.

   Example `omni.dat` file:

   ```
   [AGG_P]
   {
       Primary: netcool1.onprem.fqdn 4100
   }
   [ICP_GATE]
   {
       Primary: netcool2.onprem.fqdn 4300
   }
   [ICP_AGG_V]
   {
       Primary: mycluster.icp 32312
       Backup: mycluster.icp 30404
   }
   ```
3. Run `$NCHOME/bin/nco_igen` to create the interfaces file.
4. Configure the uni-directional gateway table replication and mapping files, `ICP_GATE_tblrep.def` and `ICP_GATE.map`. These files control which ObjectServer tables and columns are replicated.

This example table replication file replicates the `alerts.status`, `alerts.journal`, and `alerts.details` tables.

```
cat << EOF > $NCHOME/omnibus/etc/ICP_GATE.tblrep.def
##############################################################################
# Netcool/OMNIbus Uni-directional ObjectServer Gateway 8.1.0
#
# ICP_GATE table replication definition file.
#
# Notes:
#
##############################################################################

REPLICATE INSERTS, UPDATES, FT_INSERTS
        FROM TABLE 'alerts.status'
        USING MAP 'StatusMap'
        ORDER BY 'Serial ASC';

REPLICATE INSERTS, UPDATES, FT_INSERTS FROM TABLE 'alerts.journal'
    USING MAP 'JournalMap';

REPLICATE INSERTS, UPDATES, FT_INSERTS FROM TABLE 'alerts.details'
    USING MAP 'DetailsMap';
EOF
```

This example map file maps the alerts.status, alerts.details, and alerts.journal tables.

```
cat << EOF > $NCHOME/omnibus/etc/ICP_GATE.map
##############################################################################
# Netcool/OMNIbus Uni-directional ObjectServer Gateway 8.1.0
#
# ICP_GATE Multitier map definition file.
#
# Notes:
#
# Fields that are marked as 'ON INSERT ONLY' will only be passed when an event
# is inserted for the first time. (ie. they will not be updated). The ordering
# of the fields is not important as the gateway will use named value insertion.
#
##############################################################################
CREATE MAPPING StatusMap
(
    'Identifier'        =     '@Identifier'         ON INSERT ONLY,
    'Node'              =    '@Node'                ON INSERT ONLY,
    'NodeAlias'         =    '@NodeAlias'           ON INSERT ONLY
                               NOTNULL '@Node',
    'Manager'           =    '@Manager'             ON INSERT ONLY,
    'Agent'             =    '@Agent'               ON INSERT ONLY,
    'AlertGroup'        =    '@AlertGroup'          ON INSERT ONLY,
    'AlertKey'          =    '@AlertKey'            ON INSERT ONLY,
    'Severity'          =    '@Severity',
    'Summary'           =    '@Summary',
    'StateChange'       =    '@StateChange',
    'FirstOccurrence'   =    '@FirstOccurrence'     ON INSERT ONLY,
    'LastOccurrence'    =    '@LastOccurrence',
    'InternalLast'      =    '@InternalLast',
    'Poll'              =    '@Poll'                ON INSERT ONLY,
    'Type'              =    '@Type'                ON INSERT ONLY,
    'Tally'             =    '@Tally',
    'ProbeSubSecondId'  =    '@ProbeSubSecondId',
    'Class'             =    '@Class'               ON INSERT ONLY,
    'Grade'             =    '@Grade'               ON INSERT ONLY,
    'Location'          =    '@Location'            ON INSERT ONLY,
    'OwnerUID'          =    '@OwnerUID',
    'OwnerGID'          =    '@OwnerGID',
    'Acknowledged'      =    '@Acknowledged',
    'Flash'             =    '@Flash',
    'EventId'           =    '@EventId'             ON INSERT ONLY,
    'ExpireTime'        =    '@ExpireTime'          ON INSERT ONLY,
    'ProcessReq'        =    '@ProcessReq',
    'SuppressEscl'      =    '@SuppressEscl',
    'Customer'          =    '@Customer'            ON INSERT ONLY,
    'Service'           =    '@Service'             ON INSERT ONLY,
    'PhysicalSlot'      =    '@PhysicalSlot'        ON INSERT ONLY,
    'PhysicalPort'      =    '@PhysicalPort'        ON INSERT ONLY,
    'PhysicalCard'      =    '@PhysicalCard'        ON INSERT ONLY,
    'TaskList'          =    '@TaskList',
    'NmosSerial'        =    '@NmosSerial',
    'NmosObjInst'       =    '@NmosObjInst',
```

```
    'NmosCauseType'     =    '@NmosCauseType',
    'NmosDomainName'    =    '@NmosDomainName',
    'NmosEntityId'      =    '@NmosEntityId',
    'NmosManagedStatus' =    '@NmosManagedStatus',
    'NmosEventMap'      =    '@NmosEventMap',
    'LocalNodeAlias'    =    '@LocalNodeAlias'      ON INSERT ONLY,
    'LocalPriObj'       =    '@LocalPriObj'         ON INSERT ONLY,
    'LocalSecObj'       =    '@LocalSecObj'         ON INSERT ONLY,
    'LocalRootObj'      =    '@LocalRootObj'        ON INSERT ONLY,
    'RemoteNodeAlias'   =    '@RemoteNodeAlias'     ON INSERT ONLY,
    'RemotePriObj'      =    '@RemotePriObj'        ON INSERT ONLY,
    'RemoteSecObj'      =    '@RemoteSecObj'        ON INSERT ONLY,
    'RemoteRootObj'     =    '@RemoteRootObj'       ON INSERT ONLY,
    'X733EventType'     =    '@X733EventType'       ON INSERT ONLY,
    'X733ProbableCause' =    '@X733ProbableCause'   ON INSERT ONLY,
    'X733SpecificProb'  =    '@X733SpecificProb'    ON INSERT ONLY,
    'X733CorrNotif'     =    '@X733CorrNotif'       ON INSERT ONLY,
    'URL'               =    '@URL'                 ON INSERT ONLY,
    'ExtendedAttr'      =    '@ExtendedAttr'        ON INSERT ONLY,
    'CollectionFirst'   =    '@CollectionFirst'     ON INSERT ONLY,

##############################################################################################
#
#    CUSTOM alerts.status FIELD MAPPINGS GO HERE
#
##############################################################################################



##############################################################################################

    'ServerName'        =    '@ServerName'          ON INSERT ONLY,
    'ServerSerial'      =    '@ServerSerial'        ON INSERT ONLY
);

CREATE MAPPING JournalMap
(
    'KeyField'  =    TO_STRING(STATUS.SERIAL) + ":" +
                TO_STRING('@UID') + ":" +
                TO_STRING('@Chrono')    ON INSERT ONLY,
    'Serial'    =   STATUS.SERIAL,
    'Chrono'    =   '@Chrono',
    'UID'       =   TO_INTEGER('@UID'),
    'Text1'     =   '@Text1',
    'Text2'     =   '@Text2',
    'Text3'     =   '@Text3',
    'Text4'     =   '@Text4',
    'Text5'     =   '@Text5',
    'Text6'     =   '@Text6',
    'Text7'     =   '@Text7',
    'Text8'     =   '@Text8',
    'Text9'     =   '@Text9',
    'Text10'    =   '@Text10',
    'Text11'    =   '@Text11',
    'Text12'    =   '@Text12',
    'Text13'    =   '@Text13',
    'Text14'    =   '@Text14',
    'Text15'    =   '@Text15',
    'Text16'    =   '@Text16'
);

CREATE MAPPING DetailsMap
(
    'KeyField'      =   '@Identifier' + '####' +
                TO_STRING('@Sequence')    ON INSERT ONLY,
    'Identifier'    =     '@Identifier',
    'AttrVal'       =     '@AttrVal',
    'Sequence'      =     '@Sequence',
    'Name'          =     '@Name',
    'Detail'        =     '@Detail'
);
EOF
```

5. Create a gateway properties file, `ICP_GATE.props`, by copying the default unidirectional gateway properties file `objserv_uni.props`.

```
cp $NCHOME/omnibus/gates/objserv_uni/objserv_uni.props $NCHOME/omnibus/etc/ICP_GATE.props
```

6. Configure the new gateway properties file, $NCHOME/omnibus/etc/ICP_GATE.props. Set the on-premises ObjectServer AGG_P as the source, set the cloud ObjectServer ICP_AGG_V as the destination, and set the resync type to 'UPDATE', as in the following example:

```
cat << EOF >> $NCHOME/omnibus/etc/ICP_GATE.props
Name                                : 'ICP_GATE'
Gate.MapFile                        : '$NCHOME/omnibus/etc/ICP_GATE.map'
Gate.Reader.TblReplicateDefFile     : '$NCHOME/omnibus/etc/ICP_GATE.tblrep.def'
Gate.Reader.Server                  : 'AGG_P'
Gate.Writer.Server                  : 'ICP_AGG_V'
Gate.Resync.Type                    : 'UPDATE'
Gate.Resync.LockType                : 'NONE'
Gate.Writer.Description             : 'collection_gate'
EOF
```

7. Start the new gateway with the following command:

```
$NCHOME/omnibus/bin/nco_g_objserv_uni -propsfile $NCHOME/omnibus/etc/ICP_GATE.props
```

Configuring a bidirectional gateway

About this task

Learn how to configure a bidirectional gateway to create an event feed from an on-premises ObjectServer to a deployment of Operations Management on IBM Cloud Private. An on-premises primary aggregation ObjectServer must exist and be configured according to Installing a primary aggregation ObjectServer.

Procedure

1. Follow steps 1 - 4 in “Configuring a uni-directional gateway” on page 197.
2. Create a gateway properties file, ICP_GATE.props, by copying the default bidirectional gateway properties file objserv_bi.props.

```
cp $NCHOME/omnibus/gates/objserv_bi/objserv_bi.props $NCHOME/omnibus/etc/ICP_GATE.props
```

3. Configure the new gateway properties file, $NCHOME/omnibus/etc/ICP_GATE.props. Set the on-premises ObjectServer AGG_P as the source, set the cloud ObjectServer ICP_AGG_V as the destination, and set the resync type to 'TWOWAYUPDATE', as in the following example:

```
cat << EOF >> $NCHOME/omnibus/etc/ICP_GATE.props
Name                                : 'ICP_GATE'
Gate.MapFile                        : '$OMNIHOME/etc/ICP_GATE.map'
Gate.ObjectServerA.Server           : 'AGG_P'
Gate.ObjectServerA.TblReplicateDefFile  : '$OMNIHOME/etc/ICP_GATE.tblrep.def'
Gate.ObjectServerB.Server           : 'ICP_AGG_V'
Gate.ObjectServerB.TblReplicateDefFile  : '$OMNIHOME/etc/ICP_GATE.tblrep.def'
Gate.Resync.Type                    : 'TWOWAYUPDATE'
Gate.ObjectServerA.Description      : 'failover_gate'
Gate.ObjectServerB.Description      : 'failover_gate'
```

4. Configure the IDUC hostname.

   ObjectServers in cloud deployments respond to IDUC clients with their local, in cloud hostnames. The /etc/hosts file of the gateway host must be updated to correctly resolve the IDUC host address and enable successful IDUC connections.

   a) Identify the IDUC hostname of the primary and backup ObjectServers on your cloud deployment with the following commands:

```
kubectl describe pod helm_release_name-ncoprimary-0 | grep IDUC_LISTENING_HOSTNAME
kubectl describe pod helm_release_name-ncobackup-0 | grep IDUC_LISTENING_HOSTNAME
```

   Where *helm_release_name* is your Operations Management on IBM Cloud Private helm release name.

   For example:

```
%> kubectl describe pod noi-ncoprimary-0 | grep IDUC_LISTENING_HOSTNAME
      NCO_IDUC_LISTENING_HOSTNAME:      noi-objserv-agg-primary-nodeport

%> kubectl describe pod noi-ncobackup-0 | grep IDUC_LISTENING_HOSTNAME
      NCO_IDUC_LISTENING_HOSTNAME:      noi-objserv-agg-backup-nodeport
```

b) Find the release FQDN, as in the following example:

```
%> helm get values --tls helm_release_name | grep fqdn
    fqdn: mycluster.icp
```

Where *helm_release_name* is your Operations Management on IBM Cloud Private helm release name.

c) Use the host command on the returned FQDN value to find the cluster IP address, as in the following example:

```
%> host mycluster.icp
mycluster.icp has address 1.2.3.4
```

d) Edit the `/etc/hosts` file on the gateway host to map the cluster IP address to the cluster FQDN.

```
sudo cat << EOF >> /etc/hosts
1.2.3.4 noi-objserv-agg-primary-nodeport noi-objserv-agg-backup-nodeport
EOF
```

5. Start the gateway with the following command:

```
$NCHOME/omnibus/bin/nco_g_objserv_bi -propsfile $NCHOME/omnibus/etc/ICP_GATE.props
```

Chapter 7. Getting started with Netcool Operations Insight

After you have installed or upgraded the product, log into the different components of Netcool Operations Insight.

About this task

Getting started with Netcool Operations Insight on premises

Log into the components of Operations Management, such as Web GUI and Operations Analytics - Log Analysis.

Getting started with Netcool Operations Insight

Use this information to start of the components of Operations Management for Operations Insight and to log in using a Web browser.

About this task

Tip:

You can create start-up scripts to automatically to start the various Netcool Operations Insight products. For instructions and an example of how to configure a start-up script, see the 'The Netcool Process Agent and machine start-up' section in the Netcool/OMNIbus*Best Practices Guide*, which can be found on the Netcool/OMNIbus best-practice Wiki: http://ibm.biz/nco_bps

You can configure Jazz not to prompt for user credentials when the stop command is run. After creating a backup, edit the following lines in the `/opt/IBM/JazzSM/profile/properties/soap.client.props` file:

```
com.ibm.SOAP.securityEnabled=true
com.ibm.SOAP.loginUserid=smadmin
com.ibm.SOAP.loginPassword=netcool
```

Run the following command to encrypt the embedded password within the file:

```
/opt/IBM/JazzSM/profile/bin/PropFilePasswordEncoder.sh \
/opt/IBM/JazzSM/profile/properties/soap.client.props \
com.ibm.SOAP.loginPassword
```

For more information, see the following technote: http://www.ibm.com/support/docview.wss?uid=swg21584635

Procedure

- Start the Dashboard Application Services Hub server using the `/opt/IMB/JazzSM/profiles/bin/startServer.sh` *`server name`* command.
- Log in at `http://`*`host.domain`*`:16310/ibm/console` or, for a secured environment, at `https://`*`host.domain`*`:16311/ibm/console`, where *host.domain* is the fully qualified host name or IP address of the Jazz for Service Management application server. 16310 and 16311 are the default ports for HTTP and HTTPS respectively.
- Assign roles to users. To give users access to the Event Analytics capability, assign the ncw_analytics_admin user.

What to do next

Useful information about how to configure your environment is in the *Netcool Operations Insight Example Guide*, which you can download at https://developer.ibm.com/itom/docs/noi/best-practices_noi/.

Related tasks

Getting started with Networks for Operations Insight

Getting started with Networks for Operations Insight

After the installation and configuration steps are completed, you can start the components of the Networks for Operations Insight feature and log in to the host using a Web browser.

Procedure

1. Ensure that Tivoli Netcool/OMNIbus ObjectServer is running, see http://www-01.ibm.com/support/knowledgecenter/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/omnibus/wip/admin/task/omn_con_startingobjserv.html?lang=en.
2. Start the Dashboard Application Services Hub server using the `/opt/IMB/JazzSM/profiles/bin/startServer.sh` *`server name`* command.
3. Source the environment variables. On the server where the Network Manager core components are installed, the script is `installation_directory/netcool/core/env.sh`. On the server where the Network Manager GUI components are installed, the script is `installation_directory/nmgui_profile.sh, for example, /opt/IBM/netcool/nmgui_profile.sh`.
4. Start the back-end processes for the products. Use the **`itnm_start`** command to start Network Manager. Do not use this command to start the ObjectServer. The ObjectServer is started and stopped by the Tivoli Netcool/OMNIbus commands. Start Netcool Configuration Manager separately by using the `itncm.sh start` script.
5. Log in at `http://`*`host.domain`*`:16310/ibm/console` or, for a secured environment, at `https://`*`host.domain`*`:16311/ibm/console`, where *host.domain* is the fully qualified host name or IP address of the Jazz for Service Management application server. 16310 and 16311 are the default ports for HTTP and HTTPS respectively. Use the supplied itnmadmin user with the password that you specified during the installation.
6. You can log in to the Netcool Configuration Manager - Base GUI at `http://`*`ncmserver-hostname`*`:`*`port-number`*, where *ncmserver-hostname* is the host name of the computer on which you installed Netcool Configuration Manager. *port-number* is the port that you specified during the installation. Use the default user name and password that you specified during the installation.

Related tasks

Starting Network Manager

Logging in

Launching Netcool Configuration Manager - Base

Related reference

itncm.sh script

Related information

Getting started with Network Manager

Getting started with Netcool Configuration Manager

Getting started with Operations Management on IBM Cloud Private

Log in to the components of Operations Management, such as Web GUI and Operations Analytics - Log Analysis.

Procedure

1. Retrieve the hostname for each component of Operations Management
   with the following command:

```
kubectl get ingress -n namespace
```

A Kubernetes ingress is a collection of rules that can be configured to give services externally reachable URLs. Each Operations Management component requires its own ingress. Run the `kubectl get ingress` command to retrieve the hostname allocated by the Kubernetes Ingress controller to satisfy each ingress. This command retrieves data similar to the following:

```
NAME
HOSTS                                          ADDRESS     PORTS     Age
helm_release_name-ibm-ea-ui-api-graphql
netcool.helm_release_name.global.cluster.fqdn  IP_address  80, 443   1d
helm_release_name-ibm-hdm-analytics-dev-backend-ingress
netcool.helm_release_name.global.cluster.fqdn  IP_address  80, 443   1d
helm_release_name-ibm-hdm-common-ui
netcool.helm_release_name.global.cluster.fqdn  IP_address  80, 443   1d
helm_release_name-impactgui
impact.helm_release_name.global.cluster.fqdn   IP_address  80, 443   1d
helm_release_name-nci-0
nci-0.helm_release_name.global.cluster.fqdn    IP_address  80, 443   1d
helm_release_name-nci-1
nci-1.helm_release_name.global.cluster.fqdn    IP_address  80, 443   1d
helm_release_name-scala
scala.helm_release_name.global.cluster.fqdn    IP_address  80, 443   1d
helm_release_name-webgui
netcool.helm_release_name.global.cluster.fqdn  IP_address  80, 443   1d
was-helm_release_name-webgui
was.helm_release_name.global.cluster.fqdn      IP_address  80, 443   1d
```

Where

- *namespace* is the name of the namespace into whichOperations Management on IBM Cloud Private is installed.
- *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment.
- *global.cluster.fqdn* is the fully qualified domain name (FQDN) of the cluster's master node. For more information about the FQDN, see Master node.

2. Construct the URL for each of the components of Operations Management. You can then copy and paste this URL directly into your browser.

```
https://hostname/ibm/console
```

Where *hostname* is the value from the HOSTS column in step 1 for the component that you want to log into. The hostname is made up of three elements:

- Component name; for example: `impact`.
- Release name: *helm_release_name*.
- Cluster name: *global.cluster.fqdn*

3. Ensure that the hostname in your URL resolves to an IP address. You can do this using one of the following methods:
   - Configure your `/etc/hosts` file with the hostname and IP address from step 1.
   - Query a DNS server on your network.
4. The following table lists the Operations Management components, together with the sample hostnames and URLs.

Component	Default hostname	Default URL
Netcool/Impact GUI	impact.*helm_release_name-global.cluster.fqdn*	https://impact.*helm_release_name.global.cluster.fqdn*/ibm/console

Component	Default hostname	Default URL
Operations Analytics - Log Analysis GUI	scala.*helm_release_name-global.cluster.fqdn*	https://scala.*helm_release_name.global.cluster.fqdn*/Unity
Web GUI	netcool.*helm_release_name-global.cluster.fqdn*	https://netcool.*helm_release_name.global.cluster.fqdn*/ibm/console
WebSphere Application Server	was.*helm_release_name-global.cluster.fqdn*	https://was.*helm_release_name.global.cluster.fqdn*/ibm/console
Netcool/Impact Primary	nci-0.*helm_release_name-global.cluster.fqdn*	https://nci-0.*helm_release_name.global.cluster.fqdn*/nameserver/services
Netcool/Impact Backup	nci-1.*helm_release_name-global.cluster.fqdn*	https://nci-1.*helm_release_name.global.cluster.fqdn*/nameserver/services

Related tasks

Loading a TLS certificate into Kubernetes

In a Netcool Operations Insight Kubernetes cluster environment, Operations Management components such as Web GUI, WebSphere Application Server, and Netcool/Impact are known as services. In order for your users to be able to use these Netcool Operations Insight services, the URL for each service requires its own Transport Layer Security (TLS) certificate. Operations Management on IBM Cloud Private ships with automatically generated certificates signed by the cluster Certificate Authority (CA).

Related information

Helm documentation: Helm ListClick here to view information on the Helm List command.

Chapter 8. Administering

Perform the following tasks to administer the solution.

About this task

Administering users on IBM Cloud Private

If you installed Operations Management on IBM Cloud Private, then you provide access to all of the Operations Management interfaces to users, based on default user group settings. You can also optionally create new users and groups.

About this task

You can manage users using a built-in LDAP server (openLDAP server), or using your organization's LDAP server. The mechanism available to you was configured at installation time.

Single sign-on

Single sign-on is preconfigured in Operations Management on IBM Cloud Private. Federated repositories to support authentication for single sign-on are configured by default.

The following table lists the federated repositories, and list the default users and groups within these repositories. These default users are enabled by default.

Table 35. Federated repositories for single sign-on

Repository	Default groups	Default users	Capability
`InternalFile Repository`	None	admin	
`NetcoolObjectServe r`	This repository links to default Netcool/ OMNIbus groups	root, ncoadmin, ncouser, nobody	Access individual Netcool Operations Insight components only
`ICP_LDAP`	icpadmins, icpusers, unityAdmins	icpadmin, icpuser, impactadmin, unityadmin	Perform launch-in context actions across Netcool Operations Insight components. Needed to access Event Analytics and Event Search functionality.

Users defined in any of these repositories can access individual Netcool Operations Insight components directly based on their role defined within the repository. For example, the `ncoadmin` user within the `NetcoolObjectServer` repository can log into Web GUI to perform tasks in the **Event Viewer**.

However, any user that needs to perform launch in context actions across Netcool Operations Insight components must be defined be in the `ICP_LDAP` repository and in either one of the groups `icpadmins` or `icpusers`. An example of this is where a user needs to access Event Search functionality, which involves launch-in context actions from the **Event Viewer** to Operations Analytics - Log Analysis.

Defining the user in the `ICP_LDAP` repository and in either one of the groups `icpadmins` or `icpusers` ensures that all of the relevant roles in Netcool/OMNIbus, Dashboard Application Services Hub, Netcool/ Impact, and Operations Analytics - Log Analysis are assigned to the user in order for launch in context actions to function properly.

Default users

The following table describes users that are present after installation, along with their groups.

Users and their groups

The following table describes users that are present after installation, along with their groups.

Note: impactadmin and unityadmin will not exist by default if you installed with *LDAP mode: proxy*, and must be manually created. For more information, see "Creating users on an external LDAP server" on page 210.

Table 36. Users present after installation

User name	Roles	Group	Password	Description
`icpadmin`	Inherited from the group	`icpadmins`	netcool	Sample administrator user for Operations Management on IBM Cloud Private.
`icpuser`	Inherited from the group	`icpusers`	netcool	Sample end user for Operations Management on IBM Cloud Private.
`impactadmin`	Netcool/Impact-specific roles: • `impactAdminUser` • `impactFullAccessUser` • `impactOpViewUser` • `impactMWMAdminUser` • `impactMWMUser` • `impactSelectedOpViewUser` • `impactUIDataProviderUser` • `impactWebServiceUser` • `ConsoleUser` • `WriteAdmin` • `ReadAdmin` • `impactRBAUser`	`icpadmins`	netcool	Administrator user for Netcool/Impact.
`unityadmin`	Operations Analytics - Log Analysis-specific roles: • `UnityAdmin` • `UnityUser`	`unityAdmins`	unityadmin	Administrator user for Operations Analytics - Log Analysis.

Table 36. Users present after installation (continued)				
User name	**Roles**	**Group**	**Password**	**Description**
admin	Dashboard Application Services Hub-specific roles: • iscadmins • chartAdministrator • samples • administrator		netcool	The administrator for Dashboard Application Services Hub. In a new installation, this user has permissions to administer users, groups, roles, and pages.
Default Netcool/ OMNIbus users	See Netcool/OMNIbus V8.1.0 documentation: users			
Default Web GUI users	See Netcool/OMNIbus V8.1.0 documentation: Web GUI users and groups			

Default groups

Use groups to organize users into units with common functional goals. Several Operations Management on IBM Cloud Private groups are created at installation.

Default user groups

The following groups are supplied with Operations Management on IBM Cloud Private. Users are assigned to these groups during installation.

Note: icpadmins and icpusers will not exist by default if you installed with *LDAP mode: proxy*, and must be manually created. For more information, see "Creating users on an external LDAP server" on page 210.

Table 37. Default user groups		
Name	**Description**	**Roles associated with the group**
icpadmins	Assign all IBM Cloud Private administrators to this group so that they have administrative permissions over all of the Netcool Operations Insight components deployed in IBM Cloud Private.	**Dashboard Application Services Hub-specific roles** chartViewer, suppressmonitor, monitor, configurator, iscadmins, chartAdministrator, ncw_user, samples, operator, ncw_analytics_admin, chartCreator, ncw_gauges_editor, ncw_admin, netcool_rw, ncw_dashboard_editor, administrator, netcool_ro, ncw_gauges_viewer **Netcool/Impact-specific roles** impactAdminUser, impactFullAccessUser, impactOpViewUser, impactMWMAdminUser, impactMWMUser, impactSelectedOpViewUser, impactUIDataProviderUser, impactWebServiceUser, ConsoleUser, WriteAdmin, ReadAdmin, impactRBAUser **Operations Analytics - Log Analysis-specific roles** icpuser

Table 37. Default user groups (continued)

Name	Description	Roles associated with the group
`icpusers`	Assign all IBM Cloud Private users and operators to this group so that they have permissions to use the Netcool Operations Insight components deployed in IBM Cloud Private.	**Dashboard Application Services Hub-specific roles** chartViewer, monitor, configurator, ncw_user, samples, operator, netcool_rw,netcool_ro, ncw_gauges_viewer **Netcool/Impact-specific roles** impactOpViewUser, impactMWMAdminUser, impactMWMUser, impactSelectedOpViewUser, impactUIDataProviderUser, impactWebServiceUser, ConsoleUser, WriteAdmin, ReadAdmin, impactRBAUser **Operations Analytics - Log Analysis-specific roles** icpuser
Default Netcool/ OMNIbus groups and roles	See Netcool/OMNIbus V8.1.0 documentation: default groups	
Default Web GUI groups and roles	See Netcool/OMNIbus V8.1.0 documentation: Web GUI users and groups	

Creating users on an external LDAP server

Certain LDAP entries must be created in the target LDAP server that is used by the Operations Management on IBM Cloud Private deployment if you installed with *LDAP mode:proxy*. If you installed with *LDAP mode:standalone* then these mandatory entries will already exist.

If the required LDAP entries are missing, then some pods do not start correctly. There are other recommended LDAP entries whose absence would not cause a failure to deploy, but which improve the organization of NOI entities in a deployment: NOI Users, Groups, and Roles. Before creating a NOI on ICP deployment, the LDAP server administrator must provide a base Distinguished Name (DN) value for the destination LDAP server. Additionally, the LDAP administrator must create the required LDAP entries at the base DN, and also review and optionally create the recommended LDAP entries.

All LDAP entries are described in the following sections along with their DN and requirement. In all cases, the LDAP_SUFFIX placeholder must be replaced with the base DN value that is provided by the LDAP administrator.

Organizational Units

Unit name	Distinguished Name	Requirement
`groups`	`ou=groups,LDAP_SUFFIX`	Required
`users`	`ou=users,LDAP_SUFFIX`	Required

Example LDIF to create organizational units

In all LDIF examples, LDAP_SUFFIX is replaced with `dc=myldap,dc=org`

```
dn: ou=groups,dc=myldap,dc=org
objectClass: organizationalUnit
objectClass: top
```

```
ou: groups

dn: ou=users,dc=myldap,dc=org
objectClass: organizationalUnit
objectClass: top
ou: users
```

Users

User Name	Distinguished Name	Requirement
icpadmin	uid=icpadmin,ou=users,LDAP_SUFFIX	Recommended
icpuser	uid=icpuser,ou=users,LDAP_SUFFIX	Recommended
impactadmin	uid=impactadmin,ou=users,LDAP_SUFFIX	Required
unityadmin	uid=unityadmin,ou=users,LDAP_SUFFIX	Required

Example LDIF for creating users

```
dn: uid=icpuser,ou=users,dc=myldap,dc=org
objectClass: top
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
cn: ICP User
uid: icpuser
givenName: ICP User
sn: icpuser
userPassword:: password

dn: uid=icpadmin,ou=users,dc=myldap,dc=org
objectClass: top
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
cn: ICP Admin
uid: icpadmin
givenName: ICP Admin
sn: icpadmin
userPassword:: password

dn: uid=unityadmin,ou=users,dc=myldap,dc=org
objectClass: top
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
cn: Unity Admin
uid: unityadmin
givenName: Unity
sn: unityadmin
userPassword:: password

dn: uid=impactadmin,ou=users,dc=myldap,dc=org
objectClass: top
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
cn: Impact Admin User
uid: impactadmin
givenName: Impact Admin User
sn: impactadmin
userPassword:: password
```

Groups

Group name	Distinguished Names	Members	Requirement
icpadmins	cn=icpadmins,ou=groups,LDAP_SUFFIX	icpadmin	Recommended
icpusers	cn=icpusers,ou=groups,LDAP_SUFFIX	icpadmin,icpuser	Recommended
unityadmins	cn=unityadmins,ou=groups,LDAP_SUFFIX	unityadmin	Recommended
impactadmins	cn=impactadmins,ou=groups,LDAP_SUFFIX	impactadmin	Recommended

Example LDIF for creating groups

```
dn: cn=icpadmins,ou=groups,dc=myldap,dc=org
cn: icpadmins
owner: uid=icpadmin,ou=users,dc=myldap,dc=org
description: ICP Admins group
objectClass: groupOfNames
member: uid=icpadmin,ou=users,dc=myldap,dc=org

dn: cn=icpusers,ou=groups,dc=myldap,dc=org
cn: icpusers
owner: uid=icpuser,ou=users,dc=myldap,dc=org
description: ICP Users group
objectClass: groupOfNames
member: uid=icpuser,ou=users,dc=myldap,dc=org
member: uid=icpadmin,ou=users,dc=myldap,dc=org

dn: cn=unityadmins,ou=groups,dc=myldap,dc=org
cn: unityadmins
owner: uid=unityadmin,ou=users,dc=myldap,dc=org
description: Unity Admins group
objectClass: groupOfNames
member: uid=unityadmin,ou=users,dc=myldap,dc=org
```

Managing users with LDAP

You can create users and perform other user management tasks by using LDAP if you selected the option to use the built-in LDAP server LDAP `mode:standalone` at installation time, or if you selected LDAP `mode:proxy` at installation time and the required users and groups exist in your external LDAP server. See “Creating users on an external LDAP server” on page 210

About this task

The open source OpenLDAP server is installed as part of the Operations Management on IBM Cloud Private installation. Access this server to manage users and groups.

Users and groups can also be managed through the WebSphere Application Server UI. For more information, see https://www.ibm.com/support/knowledgecenter/en/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/webtop/wip/task/web_adm_createuserwebsphere.html

Creating a user and adding a user to a group

Use LDAP to create a new user and add that user to an existing group. You can also add an existing user to an existing group.

Before you begin

During the installation of Operations Management on IBM Cloud Private you must communicate with the cluster from the command line, using the Kubernetes command line interface `kubectl`, You must configure the command line on your terminal to communicate with the cluster using the Kubernetes

command line interface `kubectl`. You do this by obtaining the cluster configuration details from the IBM Cloud Private GUI and pasting this text into a terminal on your local machine. For more details, see IBM Cloud Private documentation: Accessing your IBM Cloud Private cluster by using the kubectl command-line interface

Procedure

1. Run the following command to retrieve the identifier of the LDAP Proxy Server pod.

   ```
   kubectl get pods | grep helm_release_name-openldap-0
   ```

   Where *helm_release_name* is the name of your Operations Management on IBM Cloud Private Helm release.
2. Log in to the LDAP Proxy Server pod.

   ```
   kubectl exec -it openldap_pod_id /bin/bash
   ```

   Where *openldap_pod_id* is the identifier of the LDAP Proxy Server pod.

   Proceed as follows:

If you want to...	Then...
Create a new user and add it to a group	Go to the next step
Add an existing user to a group	Go to step 3

3. Create the new user.
   a) Create an LDAP Data Interchange Format file to define the new user.

      For example:

      ```
      vi newuser.ldif
      ```

   b) Define the contents of the LDIF file that you created by using a format similar to this example:

      ```
      dn: uid=icptester,ou=users,dc=mycluster,dc=icp
      objectClass: top
      objectClass: person
      objectClass: organizationalPerson
      objectClass: inetOrgPerson
      cn: ICP Test User
      uid: icptester
      givenName: ICP Test User
      sn: icptester
      userPassword: password
      ```

      Where:

      - `uid` is the user ID of the new user. For example, `icptester`.
      - *dc* is the domain components that were specified for the suffix and baseDN. By default the value of this parameter is `dc=mycluster,dc=icp`.
      - `userPassword` is the password for this user.

      All other attributes in the file can be defined as shown in the code sample.

   c) Run the following command to create the new user.

      ```
      ldapadd -c -x -w LDAP_BIND_PWD -D LDAP_BIND_DN -f filename.ldif
      ```

      Where:

      - *LDAP_BIND_PWD* is the password for the `ldap_bind` function, which asynchronously authenticates a client with the LDAP server. By default the value of this parameter is `admin`.

- *LDAP_BIND_DN* is an object in LDAP that can carry a password. In the example, the value is:

  ```
  "cn=admin,dc=mycluster,dc=icp"
  ```

- *filename* is the name of the LDAP Data Interchange Format file that is defined in step 2b. In the example used there, *filename* is newuser.

4. Add the user to an existing group.
   a) Create an LDAP Data Interchange Format file to add the user to a group.
      For example:

      ```
      vi addUsersToGroup.ldif
      ```

   b) Define the contents of the file by using a format similar to the following:

      ```
      dn: cn=icpadmins,ou=groups,dc=mycluster,dc=icp
      changetype: modify
      add: member
      member: uid=icptester,ou=users,dc=mycluster,dc=icp
      ```

   c) Run the following command to add the user to a group.

      ```
      ldapmodify -w LDAP_BIND_PWD -D LDAP_BIND_DN -f filename.ldif
      ```

      Where:

      - *LDAP_BIND_PWD* is the password for the ldap_bind function, which asynchronously authenticates a client with the LDAP server. By default the value of this parameter is admin.
      - *LDAP_BIND_DN* is an object in LDAP that can carry a password. In the example, the value is:

        ```
        "cn=admin,dc=mycluster,dc=icp"
        ```

      - *filename* is the name of the LDAP Data Interchange Format file that is defined in step 2b. In the example used there, *filename* is newuser.

5. Check that the users and groups were added to LDAP by running the following command.

   ```
   ldapsearch -x -LLL -H ldap:/// -b dc=mycluster,dc=icp
   ```

Removing a user from a group
Use LDAP to remove a user from a group.

Procedure

1. Create an LDIF file, such as the one in the example that removes the user *icptester* from the group *icpadmins*.

   ```
   $ cat remove-testuser-from-group.ldif
   dn: cn=icpadmins,ou=groups,dc=mycluster,dc=icp
   changetype: modify
   delete:member
   member: uid=icptester,ou=users,dc=mycluster,dc=icp
   ```

   Where

   - uid is the user ID of the user to be removed from the group.
   - cn is the group that the user is to be removed from.
   - *dc* is the domain components that were specified for the suffix and baseDN. By default the value of this parameter is dc=mycluster,dc=icp.

2. Run ldapmodify with the LDIF file that you created.

   ```
   $ ldapmodify -w mypassword -D 'cn=admin,dc=mycluster,dc=icp' -f remove-testuser-from-
   group.ldif
   ```

Deleting a user

Delete a user by using LDAP.

Procedure

Run the `ldapdelete` command to delete a user, as in the following example that deletes the user *icptester*.

```
ldapdelete -w mypassword -D 'cn=admin,dc=mycluster,dc=icp'
'uid=icptester,ou=users,dc=mycluster,dc=icp'
```

Where

- `uid` is the user ID of the user to be removed from the group.
- `cn` is the group that the user is to be removed from.
- *dc* is the domain components that were specified for the suffix and baseDN. By default the value of this parameter is `dc=mycluster,dc=icp`.
- `mypassword` is the password of the user to be deleted.

Backing up and restoring your system

By regularly backing up your system, you can restore service after failure of a component.

For more information, see http://ibm.biz/backup_restore.

Chapter 9. Event search

Event search applies the search and analysis capabilities of Operations Analytics - Log Analysis to events that are monitored and managed by Tivoli Netcool/OMNIbus.

Events are transferred from the ObjectServer through the Gateway for Message Bus to Operations Analytics - Log Analysis, where they are ingested into a datasource and indexed for searching. After the events are indexed, you can search every occurrence of real-time and historical events. The Tivoli Netcool/OMNIbus Insight Pack is installed into Operations Analytics - Log Analysis and provides custom apps that search the events based on various criteria. The custom apps can generate dashboards that present event information to show how your monitoring environment is performing over time. Keyword searches and dynamic drilldown functions allow you to go deeper into the event data for detailed information. The apps can be run from the Operations Analytics - Log Analysis. Tooling can be installed into the Web GUI that launches the apps from the right-click menus of the Event Viewer and the Active Event List. An "event reduction wizard" is also supplied that includes information and apps that can help you analyze and reduce volumes of events and minimize the "noise" in your monitored environment.

Required products and components

Event search requires the following products and components:

- Operations Analytics - Log Analysis V1.3.3 or V1.3.5. For the system requirements of this product, including supported operating systems, see the following topics:
  - Operations Analytics - Log Analysis V1.3.5: https://www.ibm.com/support/knowledgecenter/SSPFMY_1.3.5/com.ibm.scala.doc/install/iwa_hw_sw_req_scen_c.html.
  - Operations Analytics - Log Analysis V1.3.3: https://www.ibm.com/support/knowledgecenter/SSPFMY_1.3.3/com.ibm.scala.doc/install/iwa_hw_sw_req_scen_c.html

  Note: Operations Analytics - Log Analysis Standard Edition is included in Netcool Operations Insight. For more information about Operations Analytics - Log Analysis editions, search for "Editions" at the Operations Analytics - Log Analysis Knowledge Center, at https://www.ibm.com/support/knowledgecenter/SSPFMY.
- OMNIbusInsightPack_v1.3.1
- Gateway for Message Bus V8.0
- Netcool/OMNIbus core components V8.1.0 fix pack 7 or later
- Netcool/OMNIbus Web GUI V8.1.0 fix pack 5 or later

For the system requirements of the core components and Web GUI for Netcool/OMNIbus V8.1.0, see https://ibm.biz/BdRNaT,

Related tasks

Checking the version of the Insight Pack
To ensure compatibility between the versions of the Tivoli Netcool/OMNIbus Insight Pack, the Web GUI and the Operations Analytics - Log Analysis product, run the **pkg_mgmt** command to check which version of the Insight Pack is installed.

Related reference

On-premises scenarios for Operations Management

This topic presents the scenarios available in a deployment of Operations Management on premises together with the associated architectures.

Netcool/OMNIbus Insight Pack

The Netcool/OMNIbus Insight Pack enables you to view and search both historical and real time event data from Netcool/OMNIbus in the IBM Operations Analytics - Log Analysis product. This documentation is for Tivoli Netcool/OMNIbus Insight Pack V1.3.0.2.

The Insight Pack parses Netcool/OMNIbus event data into a format suitable for use by Operations Analytics - Log Analysis. The event data is transferred from Netcool/OMNIbus to Operations Analytics - Log Analysis by the IBM Tivoli Netcool/OMNIbus Gateway for Message Bus (**`nco_g_xml`**). For more information about the Gateway for Message Bus, see http://www-01.ibm.com/support/knowledgecenter/SSSHTQ/omnibus/gateways/xmlintegration/wip/concept/xmlgw_intro.html.

Content of the Insight Pack

The Insight Pack provides the following data ingestion artifacts:

- A Rule Set (with annotator and splitter) that parses Netcool/OMNIbus event data into Delimiter Separated Value (DSV) format.
- A Source Type that matches the event fields in the Gateway for Message Bus map file.
- A Collection that contains the provided Source Type.
- Custom apps, which are described in Table 38 on page 219.
- A wizard to help you analyze and reduce event volumes, which is described in “Event reduction wizard” on page 221. The wizard also contains custom apps, which are described in Table 39 on page 222.

Tip: The data that is shown by the custom apps originates in the alerts.status table of the Netcool/OMNIbus ObjectServer. For example, the Node identifies the entities from which events originate, such as hosts or device names. For more information about the columns of the alerts.status table, see IBM Knowledge Center at http://www-01.ibm.com/support/knowledgecenter/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/omnibus/wip/common/reference/omn_ref_tab_alertsstatus.html.

Custom apps

The following table describes the custom apps. The apps are all launched from the Operations Analytics - Log Analysis UI. Some apps can also be launched from event lists in the Netcool/OMNIbus Web GUI, that is, the Event Viewer or Active Event List (AEL). The configuration for launching the tools from the Web GUI is not included in this Insight Pack. To obtain this configuration, install the latest fix pack of the Web GUI V8.1.

Table 38. Custom apps in the Netcool/OMNIbus Insight Pack

Name and file name of app	Can also be launched from Web GUI event list	Description
OMNIbus Static Dashboard `OMNIbus_ Static_ Dashboard. app`	Yes	Opens a dashboard with charts that show the following event distribution information: • Event Trend by Severity • Event Storm by AlertGroup • Event Storm by Node • Hotspot by Node and AlertGroup • Severity Distribution • Top 5 AlertGroups Distribution • Top 5 Nodes Distribution • Hotspot by AlertGroup and Severity The app searches against the specified data source, a time filter specified by the operator when they launch the tool, and the Node of the selected events. The app then generates charts based on the events returned by the search. Charts supplied by the Tivoli Netcool/OMNIbus Insight Pack have changed in V1.3.0.2. The charts now specify a filter of `NOT PubType:U` which ensures that each event is counted once only, even if deduplications occur. The exception is the keyword search custom app which searches all events, including modified ones. In the Operations Analytics - Log Analysis UI, the app requires data from a search result before it can run. If you do not run search before you run the apps, an error is displayed. 1. To run a new search, click **Add search** and specify the string that you want to search for. 2. A list of corresponding events is displayed in the search results. 3. In the left panel, click **Search Dashboards** > **OMNIbusInsightPack** and double-click **Static Event Dashboard**.

Table 38. Custom apps in the Netcool/OMNIbus Insight Pack (continued)

Name and file name of app	Can also be launched from Web GUI event list	Description
OMNIbus Keyword Search `OMNIbus_ Keyword_ Search.app`	Yes	Uses information from the selected events to generate a keyword list with count, data source filter, and time filter in Operations Analytics - Log Analysis. The app generates the keyword list from the specified columns of the selected events. The default columns are Summary, Node, and AlertGroup. The app then creates the data source filter with the value specified by the event list tool and creates the time filter with the value that was selected when the tool was launched. In the Operations Analytics - Log Analysis UI, the app requires data from a search result before it can run. If you do not run search before you run the apps, an error is displayed. 1. To run a new search, click **Add search** and specify the string that you want to search for. 2. A list of corresponding events is displayed in the search results. Switch to the grid view and select the required entries. Click a column header to select the entire column. 3. In the left panel, click **Search Dashboards** > **OMNIbusInsightPack** and double-click **Keyword Search**. In the **Search Patterns** section, a list of keywords from the selected data is displayed. The event count associated with those keywords is in parentheses ().

Table 38. Custom apps in the Netcool/OMNIbus Insight Pack (continued)

Name and file name of app	Can also be launched from Web GUI event list	Description
OMNIbus Dynamic Dashboard `OMNIbus_ Dynamic_ Dashboard. app`	No	Searches the events in the "omnibus" data source over the last day and generates a dashboard with eight charts. The charts are similar to the charts generated by the OMNIbus Static Dashboard app but they also support drill down. You can double-click any data point in the chart to open a search workspace that is scoped to the event records that make up that data point. To open the dashboard in the Operations Analytics - Log Analysis user interface, click **Search Dashboards** > **OMNIbusInsightPack** > **Last_Day** > **Dynamic Event Dashboard**. This dashboard is not integrated with the event lists in the Web GUI.
OMNIbus_ Operational_ Efficiency `OMNIbus_ Operational_ Efficiency. app`	No	Searches the events from the "omnibus" data source over the last month and generates a dashboard with the following charts. • **Last Month - Top 10 AlertKeys**: Shows the AlertKeys that generated the most events, distributed by severity. • **Last Month - Top 10 AlertGroups**: Shows the AlertGroups that generated the most events, distributed by severity. • **Last Month - Top 10 Node**: Shows the Nodes that generated the most events, distributed by severity. • **Last Month - Hotspot by Node, Group, AlertKey**: Combines the three other charts to show the Nodes, AlertGroups, and AlertKeys that generated the most events in a tree map. To open the dashboard in the Operations Analytics - Log Analysis user interface, click **Search Dashboards** > **OMNIbusInsightPack** > **Last_Month** > **Operational Efficiency**. This dashboard is not integrated with the event lists in the Web GUI.

Event reduction wizard

The Event_Analysis_And_Reduction app is a guide to analyzing events in your environment and reducing event volumes. It consists of three sets of information and seven custom apps. The information is designed to help you understand the origin of high event volumes in your environment and create an action plan to reduce volumes. The information is in the first three nodes of the **Event_Analysis_And_Reduction** node on the UI: **OMNIbus_Analyze_and_reduce_event_volumes**,

OMNIbus_Helpful_links, and **OMNIbus_Introduction_to_the_Apps**. The seven custom apps analyze the origins of the high event volumes in your environment. They are described in the following table. For the best results, run the apps in the order that is given here. The wizard and the app that it contains can be run only from the Operations Analytics - Log Analysis UI.

Table 39. Custom apps in the Event_Analysis_And_Reduction wizard

Name and file name of app	**Description**
OMNIbus_Show_Event_1_Trend_Severity `OMNIbus_Show_Event_1_Trend_Severity.app`	Shows charts with five common dimensions for analyzing trends in event volumes over time: • Event trends by severity for the past hour, aggregated by minute. • Event trends by severity for the past day, aggregated by hour. • Event trends by severity for the past week, aggregated by day. • Event trends by severity for the past month, aggregated by week. • Event trends by severity for the past year, aggregated by month.
OMNIbus_Show_Event_2_HotSpots_Node `OMNIbus_Show_Events_2_HotSpots_Node.app`	Analyzes events by node, that is, the entities from which events originate. Examples include the source end point system, EMS or NMS, probe or gateway, and so on. You can modify this app to analyze the manager field, so that it shows the top event volumes by source system or integration. The app has the following charts: • The 20 nodes with the highest event counts over the past hour. • The 20 nodes with the highest event counts over the past day. • The 20 nodes with the highest event counts over the past week. • The 20 nodes with the highest event counts over the past month. • The 20 nodes with the highest event counts over the past year.
OMNIbus_Show_Event_3_HotSpots_AlertGroup `OMNIbus_Show_Events_3_HotSpots_AlertGroup.app`	Analyzes the origin of events by the classification that is captured in the AlertGroup field, for example, the type of monitoring agent, or situation. The app has the following charts: • The 20 AlertGroups with the highest event counts over the past hour. • The 20 AlertGroups with the highest event counts over the past day. • The 20 AlertGroups with the highest event counts over the past week. • The 20 AlertGroups with the highest event counts over the past month. • The 20 AlertGroups with the highest event counts over the past year.
OMNIbus_Show_Event_4_HotSpots_AlertKey `OMNIbus_Show_Event_4_HotSpots_AlertKey.app`	Analyzes the origin of events by the classification that is captured in the AlertKey field, for example, the type of monitoring agent or situation. The app has the following charts: • The 20 AlertKeys with the highest event counts over the past hour. • The 20 AlertKeys with the highest event counts over the past week. • The 20 AlertKeys with the highest event counts over the past month. • The 20 AlertKeys with the highest event counts over the past year.

Table 39. Custom apps in the Event_Analysis_And_Reduction wizard (continued)	
Name and file name of app	**Description**
OMNIbus_Show_Event_5_HotSpots_Node Severity OMNIbus_Show_Event_5_HotSpots_Node Severity.app	Shows the nodes with the highest event counts by event severity. The app has the following charts: • The 10 nodes with the highest event counts by event severity over the past hour. • The 10 nodes with the highest event counts by event severity over the past day. • The 10 nodes with the highest event counts by event severity over the past week. • The 10 nodes with the highest event counts by event severity over the past month. • The 10 nodes with the highest event counts by event severity over the past year.
OMNIbus_Show_Event_6_HotSpots_ NodeAlert Group OMNIbus_Show_Event_6_HotSpots_ NodeAlert Group.app	Shows the nodes with the highest event counts by the classification in the AlertGroup field, for example, the type of monitoring agent or situation. The app has the following charts: • The 10 nodes with the highest event counts from the top 5 AlertGroups over the past hour. • 10 nodes with the highest event counts from the top 5 AlertGroups over the past day. • The 10 nodes with the highest event counts from the top 5 AlertGroups over the past week. • The 10 nodes with the highest event counts from the top 5 AlertGroups over the past month. • The 10 nodes with the highest event counts from the top 5 AlertGroups over the past year.
OMNIbus_Show_Event_7_HotSpots_NodeAlert Key OMNIbus_Show_Event_7_HotSpots_NodeAlert Key. app	Shows the nodes with the highest event counts by the classification in the AlertKey field, for example, the monitoring agent or situation. The app has the following charts: • 10 nodes with the highest event counts from the top 5 AlertKeys over the past hour. • 10 nodes with the highest event counts from the top 5 AlertKeys over the past day. • 10 nodes with the highest event counts from the top 5 AlertKeys over the past week. • 10 nodes with the highest event counts from the top 5 AlertKeys over the past month. • 10 nodes with the highest event counts from the top 5 AlertKeys over the past year.

By default the custom apps include all events. To exclude certain events, for example, events that occur during maintenance windows, customise the search query used in the custom apps. For more information, see “Customizing the Apps” on page 178.

Configuring event search

This section describes how to configure the integration of the Netcool/OMNIbus and Operations Analytics - Log Analysis products. Events are forwarded from Netcool/OMNIbus to Operations Analytics - Log Analysis by the Gateway for Message Bus.

Before you begin

- Use a supported combination of product versions. For more information, see "Required products and components" on page 217. The best practice is to install the products in the following order:
  1. Netcool/OMNIbus V8.1.0 and the Web GUI
  2. Gateway for Message Bus. Install the gateway on the same host as the Netcool/OMNIbus product.
  3. Operations Analytics - Log Analysis, see one of the following links:
     - V1.3.5: see https://www.ibm.com/support/knowledgecenter/SSPFMY_1.3.5/com.ibm.scala.doc/install/iwa_install_ovw.html
     - V1.3.3: see https://www.ibm.com/support/knowledgecenter/SSPFMY_1.3.3/com.ibm.scala.doc/install/iwa_install_ovw.html

     https://ibm.biz/BdXeZk
  4. Netcool/OMNIbus Insight Pack, see "Installing the Tivoli Netcool/OMNIbus Insight Pack" on page 61.

  Tip: The best practice is to install the Web GUI and Operations Analytics - Log Analysis on separate hosts.

  Restriction: Operations Analytics - Log Analysis does not support installation in Group mode of IBM Installation Manager.
- Ensure that the ObjectServer that forwards event data to Operations Analytics - Log Analysis has the **NmosObjInst** column in the alerts.status table. **NmosObjInst** is supplied by default and is required for this configuration. You can use ObjectServer SQL commands to check for the column and to add it if it is missing, as follows.
  - Use the DESCRIBE command to read the columns of the alerts.status table.
  - Use the ALTER COLUMN setting with the ALTER TABLE command to add **NmosObjInst** to the alerts.status table.

  For more information about the alerts.status table, including the **NmosObjInst** column, see https://ibm.biz/BdXcBF. For more information about ObjectServer SQL commands, see https://ibm.biz/BdXcBX.
- Configure the Web GUI `server.init` file as follows:

  Note: The default values do not have to be changed on Web GUI V8.1.0 fix pack 5 or later.

  ```
  scala.app.keyword=OMNIbus_Keyword_Search
  scala.app.static.dashboard=OMNIbus_Static_Dashboard
  scala.datasource=omnibus
  scala.url=protocol://host:port
  scala.version=1.2.0.3
  ```

  Restart the server if you change any of these values. See https://ibm.biz/BdXcBc.
- Select and plan a deployment scenario. See "On-premises scenarios for Operations Management" on page 34. If your deployment uses the Gateway for Message Bus for forwarding events via the IDUC channel, you can skip step "5" on page 226. If you use the AEN client for forwarding events, complete all steps.
- Start the Operations Analytics - Log Analysis product.
- Familiarize yourself with the configuration of the Gateway for Message Bus. See https://ibm.biz/BdEQaD. Knowledge of the gateway is required for steps "1" on page 225, "5" on page 226, and "6" on page 226 of this task.

Procedure

The term *data source* has a different meaning, depending on which product you configure. In the Web GUI, a data source is always an ObjectServer. In the Operations Analytics - Log Analysis product, a data source is a source of raw data, usually log files. In the context of the event search function, the Operations Analytics - Log Analysis data source is a set of Netcool/OMNIbus events.

1. Configure the Gateway for Message Bus.

   At a high-level, this involves the following:

   - Creating a gateway server in the Netcool/OMNIbus interfaces file
   - Configuring the `G_SCALA.props` properties file, including specifying the `.map` mapping file.
   - Configuring the endpoint in the `scalaTransformers.xml` file
   - Configuring the SSL connection, if required
   - Configuring the transport properties in the `scalaTransport.properties` file

   For more information about configuring the gateway, see the Gateway for Message Bus documentation at https://ibm.biz/BdEQaD.

2. If you are ingesting data that is billable, and do not want data ingested into the Netcool Operations Insight data source to be included in the usage statistics, you need to set the Netcool Operations Insight data source as non-billable. Add the path to your data source (default is NCOMS, see following step) to a seed file and restart Operations Analytics - Log Analysis as described in one of the following topics:

   - V1.3.5: see https://www.ibm.com/support/knowledgecenter/SSPFMY_1.3.5/com.ibm.scala.doc/admin/iwa_nonbill_ds_t.html
   - V1.3.3: see https://www.ibm.com/support/knowledgecenter/SSPFMY_1.3.3/com.ibm.scala.doc/admin/iwa_nonbill_ds_t.html

   Note: Ensure you follow this step before you configure an "omnibus" data source for Netcool/OMNIbus events.

3. In Operations Analytics - Log Analysis, start the **Add Data Source** wizard and configure an "omnibus" data source for Netcool/OMNIbus events.

   Only a single data source is required. The event management tools in the Web GUI support a single data source only.

   a) In the **Select Location** panel, select **Custom** and type the Netcool/OMNIbus server host name.

   Enter the same host name that was used for the **`JsonMsgHostname`** transport property of the Gateway for Message Bus.

   b) In the **Select Data** panel, enter the following field values:

Field	Value
File path	`NCOMS`. This is the default value of the **`jsonMsgPath`** transport property of the Gateway for Message Bus. If you changed this value from the default, change the value of the **File path** field accordingly.
Type	This is the name of the data source type on which this data source is based. • To use the default data source type, specify `OMNIbus1100`. • To use a customized data source type, specify the name of the customized data source type; for example: `customOMNIbus`
Collection	`OMNIbus1100-Collection`

   c) In the **Set Attributes** panel, enter the following field values:

Field	Value
Name	`omnibus`. Ensure that the value that you type is the same as the value of the **`scala.datasource`** property in the Web GUI `server.init` file. If the **Name** field has a value other than `omnibus`, use the same value for the **`scala.datasource`** property.
Group	Leave this field blank.
Description	Type a description of your choice.

4. Configure access to the data source you set up in the previous step. This involves the following steps in the administrative settings for Operations Analytics - Log Analysis:
   a) Create a role using the **Roles** tab, for example, `noirole`, and ensure you assign the role permission to access the data source.
   b) Add a user, for example, `noiuser`, and assign the role you created that has permissions to access the data source (in this example, `noirole`).

   For information about creating and modifying users and roles in Operations Analytics - Log Analysis, see one of the following links:

   - V1.3.5: see https://www.ibm.com/support/knowledgecenter/SSPFMY_1.3.5/com.ibm.scala.doc/config/iwa_config_pinstall_userrole_ovw_c.html
   - V1.3.3: see https://www.ibm.com/support/knowledgecenter/SSPFMY_1.3.3/com.ibm.scala.doc/config/iwa_config_pinstall_userrole_ovw_c.html

   .

   Note: The contents of the Netcool/OMNIbus Insight Pack dashboards are empty unless you log in with a user that has a role assigned with permissions to access the data source.
5. Configure the Accelerated Event Notification (AEN) client:
   a) Configure AEN event forwarding in the Gateway for Message Bus.
   b) Configure the AEN channel and triggers in each ObjectServer by enabling the following postinsert triggers and trigger group:
      - scala_triggers
      - scala_insert
      - scala_reinsert

      These items are included in the default configuration of the ObjectServer, as well as the SQL commands to configure the AEN channel, but they are disabled by default. For more information about configuring the AEN client in an integration with the Operations Analytics - Log Analysis product, search for *Configuring event forwarding using AEN* in the Gateway for Message Bus documentation.
6. Start the Gateway for Message Bus in SCA-LA mode.

   The gateway begins sending Netcool/OMNIbus events to Operations Analytics - Log Analysis.
7. Install the Web GUI with the event search feature.

Results

After the configuration is complete, you can search for Netcool/OMNIbus events in Operations Analytics - Log Analysis. You can also use the Web GUI event management tools to launch into Operations Analytics - Log Analysis to display event data.

What to do next

- Install any available interim fixes and fix packs for the Operations Analytics - Log Analysis product, which are available from IBM Fix Central at http://www.ibm.com/support/fixcentral/.
- You can customize event search in the following ways:

  - Change the Operations Analytics - Log Analysis index configuration. For more information, see "Customizing events used in Event Search using the DSV toolkit" on page 176. If you change the index configuration, also change the map file of the Gateway for Message Bus. After the map file is changed, restart the gateway. For more information, search for *Map definition file* at https://ibm.biz/BdEQaD.
  - Customize the Operations Analytics - Log Analysis custom apps that are in the Insight Pack, or create new apps. For more information, see "Customizing the Apps" on page 178.
  - Customize the Web GUI event list tools. For more information, see "Customizing event management tools" on page 229.
- If the Web GUI and Operations Analytics - Log Analysis are on the same host, configure single sign-on to prevent browser sessions expiring. See "Configuring single sign-on for the event search capability" on page 227.

Related tasks

Customizing event management tools

Configuring triggers

Searching for events

Related reference

unity.sh command (for starting SmartCloud Analytics - Log Analysis V1.3.5)

unity.sh command (for starting SmartCloud Analytics - Log Analysis V1.3.3)

Related information

Gateway for Message Bus documentation

Configuring single sign-on for the event search capability

Configure single sign-on (SSO) between Web GUI and Operations Analytics - Log Analysis so that users can switch between the two products without having to log in each time.

Before you begin

Before performing this task, ensure that the following requirements are met:

- All server instances are in same domain; for example, *domain_name*.uk.ibm.com.
- LTPA keys are the same across all server instances.
- The LTPA cookie name that is used in Operations Analytics - Log Analysis must contain the string `ltpatoken`.

About this task

First create dedicated users in your LDAP directory, which must be used by both Web GUI and Operations Analytics - Log Analysis for user authentication, and then configure the SSO connection.

Table 40. Quick reference for configuring single sign-on

Step	Action	More information
1.	Create the dedicated users and groups in your LDAP directory. For example: 1. Create a new Organization Unit (OU) named NetworkManagement. 2. Under the NetworkManagement OU, create a new group named webguildap. 3. Under the NetworkManagement OU, create the following new users: webgui1, webgui2, webgui3, and webgui4. 4. Add the new users to the webguildap group.	The LDAP groups that you want to use in Web GUI must have roles that Web GUI recognizes. For more information, see the following topic: Configuring user authentication for Web GUI against an LDAP directory.
2.	In the Web GUI, assign the ncw_admin and ncw_user roles to the webguildap group that you created in step 1.	For more information see the following topics: • Assigning roles to Web GUI users and groups • Web GUI roles
3.	Configure Dashboard Application Services Hub and Operations Analytics - Log Analysis to use the same LDAP directory for authentication.	For more information on configuring these products to use LDAP, see the following topics: • Configuring Dashboard Application Services Hub to use LDAP • Configuring Operations Analytics - Log Analysis to use LDAP
4.	Configure Dashboard Application Services Hub for single sign-on. This enables users to access all of the applications running in Dashboard Application Services Hub by logging in only once.	For more information see the following topic: Configuring Dashboard Application Services Hub for single sign-on.
5.	Configure the SSO connection from the Operations Analytics - Log Analysis product to the Dashboard Application Services Hub instance in which the Web GUI is hosted. The following steps of the Operations Analytics - Log Analysis SSO configuration are important: • Export LTPA keys from the Jazz for Service Management server. • Update LA ldapRegistryHelper.properties file. • Run the LA ldapRegistryHelper.sh script. • Configure LTPA on the Liberty Profile for WAS (copy LTPA keys from Jazz)	For more information see the following topic: Configuring SSO for Operations Analytics - Log Analysis with Jazz for Service Management
6.	Assign Operations Analytics - Log Analysis roles to the users and groups that you created in step 1.	

Table 40. Quick reference for configuring single sign-on (continued)		
Step	**Action**	**More information**
7.	In the $*SCALAHOME*/wlp/usr/servers/Unity/server.xml/server.xml file, ensure that the <webAppSecurity> element has a httpOnlyCookies="false" attribute. Add this line before the closing </server> element. For example: `<webAppSecurity ssoDomainNames="hostname" httpOnlyCookies="false"/> </server>` Where The httpOnlyCookies="false" attribute disables the httponly flag on the cookie that is generated by Operations Analytics - Log Analysis and is required to enable SSO with the Web GUI	

Related tasks

Configuring user authentication for the Web GUI against LDAP directories

Assigning roles to Web GUI users and groups

Configuring SSO with Operations Analytics - Log Analysis V1.3.5

Configuring SSO with Operations Analytics - Log Analysis V1.3.3

Related reference

Web GUI roles

Troubleshooting event search

How to resolve problems with your event search configuration.

Customizing event management tools

The tools in the Event Viewer and AEL search for event data based on fields from the Netcool/OMNIbus ObjectServer. The fields are specified by URLs that are called when the Operations Analytics - Log Analysis product is started from the tools. You can change the URLs in the tools from the default event fields to search on fields of your choice.

For example, the **Search for similar events** > **15 minutes before event** tool filters events on the `AlertGroup`, `Type`, and `Severity` fields. The default URL is:

```
$(SERVER)/integrations/scala/Search?queryFields=AlertGroup,Type,Severity
&queryValuesAlertGroup={$selected_rows.AlertGroup}
&queryValuesType={CONVERSION($selected_rows.Type)}
&queryValuesSeverity={CONVERSION($selected_rows.Severity)}
&firstOccurrences={$selected_rows.FirstOccurrence}
&timePeriod=15
&timePeriodUnits=minutes
```

About this task

You can change the URLs in the following ways:

- Change the `scalaIntegration.xml` configuration file and apply the changes with the Web GUI **`runwaapi`** command that is included in the Web GUI Administration API (WAAPI) client.
- Change the tool configuration in the Web GUI **Administration** console page.

Procedure

As an example, the following steps show how to use each method to change the URLs in the **Search for similar events** > **15 minutes before event** tool to search on the `AlertKey` and `Location` event fields.

- To change the URLs in the `scalaIntegration.xml` configuration file:
  a) In *WEBGUI_HOME*`/extensions/LogAnalytics/scalaIntegration.xml`, or the equivalent XML file if you use a different file, locate the following <tool> element:

```
<tool:tool name="scalaSearchByEvent15Minutes">
```

  b) Change the URL in this element as follows.

  The changes are shown in **bold** text:

```
<tool:cgiurl foreach="true" windowforeach="false" target="_blank" method="GET"
 url="$(SERVER)/integrations/scala/Search?queryFields=AlertKey,Location
&queryValuesAlertKey={$selected_rows.AlertKey}
&queryValuesLocation={$selected_rows.Location}
&firstOccurrences={$selected_rows.FirstOccurrence}
&timePeriod=15
&timePeriodUnits=minutes">
</tool:cgiurl>
```

  c) Use the **runwaapi** command to reinstall the tools:

```
runwaapi -file scalaIntegration.xml
```

  d) Reinstall the following tool menus to the Event Viewer or AEL **alerts** menu item:

  – scalaStaticDashboard
  – scalaSimilarEvents
  – scalaEventByNode
  – scalaKeywordSearch

- To change the URLs on the **Administration** page:
  a) In the Web GUI, click **Administration** > **Event Management Tools** > **Tool Creation**. Then, on the **Tool Creation** page, locate the `scalaSearchByEvent15Minutes` tool.
  b) Change the URL as follows.

  The changes are shown in **bold** text:

```
$(SERVER)/integrations/scala/Search?queryFields=AlertKey,Location
&queryValuesAlertKey={$selected_rows.AlertKey}
&queryValuesLocation={$selected_rows.Location}
&firstOccurrences={$selected_rows.FirstOccurrence}
&timePeriod=15
&timePeriodUnits=minutes
```

  c) Refresh the Event Viewer or AEL so that the changes to the tool URL are loaded.

What to do next

- The Gateway for Message Bus uses a lookup table to convert the `Severity`, `Type`, and `Class` event field integer values to strings. After a tool is changed or created, use the `CONVERSION` function to change these field values to the strings that are required by Operations Analytics - Log Analysis.
- Change the other tools in the menu so that they search on the same field. It is more efficient to change the configuration file and then use the **runwaapi** command than to change each tool in the UI. The following table lists the names of the event management menu items and tools that are displayed in the **Tool Creation** and **Menu Configuration** pages.

Table 41. Web GUI menu and tool names

Menu item	Menu item name	Tool	Tool name
Search for events by node	scalaEventByNode	15 minutes before event	scalaSearch ByNode15Minutes
		1 hour before event	scalaSearch ByNode1Hour
		1 day before event	scalaSearch ByNode1Day
		1 week before event	scalaSearch ByNode1Week
		1 month before event	scalaSearch ByNode1Month
		1 year before event	scalaSearch ByNode1Year
		Custom ...	scalaSearch ByNodeCustom
Search for similar events	scalaSimilarEvents	15 minutes before event	scalaSearchByEvent15Minutes
		1 hour before event	scalaSearchByEvent1Hour
		1 day before event	scalaSearchByEvent1Day
		1 week before event	scalaSearchByEvent1Week
		1 month before event	scalaSearchByEvent1Month
		1 year before event	scalaSearchByEvent1Year
		Custom ...	scalaSearchByEventCustom
Show event dashboard by node	scalaStaticDashboard	15 minutes before event	scalaEventDistribution ByNode15Minutes
		1 hour before event	scalaEventDistribution ByNode1Hour
		1 day before event	scalaEventDistribution ByNode1Day
		1 week before event	scalaEventDistribution ByNode1Week
		1 month before event	scalaEventDistribution ByNode1Month
		1 year before event	scalaEventDistribution ByNode1Year
		Custom ...	scalaEventDistribution ByNodeCustom
Show keywords and event count	scalaKeywordSearch	15 minutes before event	scalaSetSearchFilter15Minutes

Table 41. Web GUI menu and tool names (continued)			
Menu item	**Menu item name**	**Tool**	**Tool name**
		1 hour before event	scalaSetSearchFilter1Hour
		1 day before event	scalaSetSearchFilter1Day
		1 week before event	scalaSetSearchFilter1Week
		1 month before event	scalaSetSearchFilter1Month
		1 year before event	scalaSetSearchFilter1Year
		Custom ...	scalaSetSearchFilterCustom

- The **Show event dashboard by node** and **Show keywords and event count** tools start the OMNIbus Static Dashboard and OMNIbus Keyword Search custom apps in Operations Analytics - Log Analysis. For more information about customizing the apps, see "Customizing the Apps" on page 178.

Related tasks

Searching for events

Adding custom apps to the Table View toolbar

To quickly launch custom apps, add them to the Table View toolbar of the Operations Analytics - Log Analysis UI. It is good practice to add the `OMNIbus_Keyword_Search.app` and `OMNIbus_Static_Dashboard.app` apps to the toolbar.

Procedure

- To add the `OMNIbus_Keyword_Search.app` app, use a configuration that is similar to the following example:

```
{
"url": "https://hostname:9987/Unity/CustomAppsUI?
name=OMNIbus_Keyword_Search&appParameters=[]",
"icon": "https://hostname:9987/Unity/images/keyword-search.png",
"tooltip": "OMNIbus Keyword Search"
}
```

  Where *hostname* is the fully qualified domain name of the Operations Analytics - Log Analysis host and *keyword-search* is the file name for a `.png` file that represents the app on the toolbar. Create your own `.png` file.

- To add the `OMNIbus_Static_Dashboard.app` app, use a configuration that is similar to the following example:

```
{
"url": "https://hostname:9987/Unity/CustomAppsUI?
name=OMNIbus_Static_Dashboard&appParameters=[]",
"icon": "https://hostname:9987/Unity/images/dashboard.png",
"tooltip": "OMNIbus Static Dashboard"
}
```

  Where *hostname* is the fully qualified domain name of the Operations Analytics - Log Analysis host and *dashboard* is the file name for a `.png` file that represents the app on the toolbar. Create your own `.png` file.

Related reference

Netcool/OMNIbus Insight Pack

The Netcool/OMNIbus Insight Pack enables you to view and search both historical and real time event data from Netcool/OMNIbus in the IBM Operations Analytics - Log Analysis product. This documentation is for Tivoli Netcool/OMNIbus Insight Pack V1.3.0.2.

Related information

Adding a shortcut to a Custom App to the Table view toolbar (Operations Analytics - Log Analysis V1.3.5)
Adding a shortcut to a Custom App to the Table view toolbar (Operations Analytics - Log Analysis V1.3.3)

Using Event Search

The event search tools can find the root cause of problems that are generating large numbers of events in your environment. The tools can detect patterns in the event data that, for example, can identify the root cause events that cause event storms. They can save you time that would otherwise be spent manually looking for the event that is causing problems. You can quickly pinpoint the most important events and issues.

The tools are built into the Web GUI event lists (AEL and Event Viewer). They run searches against the event data, based on default criteria, filtered over specific time periods. You can search against large numbers of events. You can change the search criteria and specify different time filters. When run, the tools start the Operations Analytics - Log Analysis product, where the search results are displayed.

Before you begin

- Set up the environment for event search. See "Configuring event search" on page 224.
- Familiarize yourself with the Operations Analytics - Log Analysis search workspace.
  - If you are using V1.3.5, then see http://www-01.ibm.com/support/knowledgecenter/SSPFMY_1.3.5/com.ibm.scala.doc/use/iwa_using_ovw.html.
  - If you are using V1.3.3, then see http://www-01.ibm.com/support/knowledgecenter/SSPFMY_1.3.3/com.ibm.scala.doc/use/iwa_using_ovw.html.
- To understand the event fields that are indexed for use in event searches, familiarize yourself with the ObjectServer alerts.status table. See http://www-01.ibm.com/support/knowledgecenter/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/omnibus/wip/common/reference/omn_ref_tab_alertsstatus.html.

Procedure

- To start using the event search tools, select one or more events from an event list and right-click. From the right-click menu, click **Event Search**, click a tool, and click a time filter.

  The tools are as follows:

Tool	Description
Show event dashboard by node	Searches for all events that originate from the same host name, service name, or IP address, which is equivalent to the Node field of the ObjectServer alerts.status table.
Search for similar events	Searches for all events that have the same failure type, type, and severity as the selected events. The failure type equates to the AlertGroup field of the alerts.status table. The type equates to the Type field. The severity equates to the Severity field.
Search for events by node	Searches for all events that originate from the same source, that is, host name, service name, or IP address. This is equivalent to the Node field of the alerts.status table. The results are displayed in a list in in the Operations Analytics - Log Analysis GUI.
Show keywords and event count	Extracts a list of keywords from the text of the event summary, event source, and failure type. The event summary text equates to the Summary field of the

Tool	Description
	alerts.status table. The event source equates to the Node field. The failure type equates to the AlertGroup field.

The time filters are calculated from the time stamp of the selected event or events. The Operations Analytics - Log Analysis time stamp is equivalent to the FirstOccurrence field of the ObjectServer alerts.status table. The default time filters are as follows. If you click **Custom** specify an integer and unit of time, such as 15 weeks.

- **15 minutes before event**
- **1 hour before event**
- **1 day before event**
- **1 week before event**
- **1 month before event**
- **1 year before event**
- **Custom ...**

If a single event is selected that has the time stamp 8 January 2014 08:15:26 AM, and you click **Search for events by node** > **1 hour before event**, the result is filtered on the following time range: (8 January 2014 07:15:26 AM) to (8 January 2014 08:15:26 AM).

If multiple events are selected, the time filter is applied from the earliest to the most recent time stamp. For three events that have the time stamps 1 January 2014 8:28:46 AM, 7 January 2014 8:23:20 AM, and 8 January 2014 8:15:26 AM, the **Search for events by node** > **1 week before event**, returns matching events in the following time range: (25 December 2013 08:28:46 AM) to (08 January 2014 08:15:26 AM).

Restriction: The Web GUI and Operations Analytics - Log Analysis process time stamps differently. The Web GUI recognizes hours, minutes, and seconds but Operations Analytics - Log Analysis ignores seconds. This problem affects the **Show event dashboard by node** and **Search for events by node**. If the time stamp 8 January 2014 07:15:26 AM is passed, Operations Analytics - Log Analysis interprets this time stamp as 8 January 2014 07:15 AM. So, the results of subsequent searches might differ from the search that was originally run.

The results are displayed differently depending on the tool. The time filter has no effect on how the results are displayed.

Tool	How search results are displayed
Show event dashboard by node	A dashboard is opened for the OMNIbus Static Dashboard custom app that shows the following information about the distribution of the matching events: – Event Trend by Severity – Event Storm by AlertGroup – Event Storm by Node – Hotspot by Node and AlertGroup – Severity Distribution – Top 5 AlertGroups Distribution – Top 5 Nodes Distribution – Hotspot by AlertGroup and Severity For more information about the OMNIbus Static Dashboard custom app, see “Netcool/OMNIbus Insight Pack” on page 218.
Search for similar events	The results are displayed in the search timeline, which shows the distribution of matching events over the specified time period. Below the timeline, the list of

Tool	How search results are displayed
and **Search for events by node**	results is displayed. Click **Table View** or **List View** to change how the results are formatted. Click **>** or **<** to move forward and back in the pages of results. Keywords that occur multiple times in the search results are displayed in the **Common Patterns** area of the navigation pane, with the number of occurrences in parentheses ().
Show keywords and event count	The keywords are displayed in the **Configured Patterns** area of the Operations Analytics - Log Analysis GUI. Each occurrence of the keyword over the time period is counted and displayed in parentheses () next to the keyword.

- After the results are displayed, you can refine them by performing further searches on the results in the search workspace.
  For example, click a keyword from the **Configured Patterns** list to add it to the **Search** field.

  Important: Because of the difference in handling seconds between the two products, if you run a further search against the keyword counts that result from the **Show keywords and event count** tool, you might see a difference in the count that was returned for a keyword under **Configured Patterns** and in the search that you run in the search workspace.

  Above the **Search** field, a sequence of breadcrumbs is displayed to indicate the progression of your search. Click any of the breadcrumb items to return the results of that search.

Example

The **Show keywords and event count** tool can examine what happened before a problematic event in your environment. Assume that high numbers of critical events are being generated in an event storm. A possible work flow is as follows:

- You select a number of critical events and click **Event search** > **Show keywords and event count** > **1 hour before event** so that you can identify any similarities between critical events that occurred in the last hour.
- The most recent time stamp (FirstOccurrence) of an event is 1 January 2014 8:28:00 AM. In the Operations Analytics - Log Analysis GUI, the search results show all keywords from the Summary, Node, and AlertGroup fields and the number of occurrences.
- You notice that the string "swt0001", which is the host name of a switch in your environment, has a high number of occurrences. You click **swt0001** and run a further search, which reduces the number of results to only the events that contain "swt0001".
- From this pared-down results list, you quickly notice that one event shows that switch is misconfigured, and that this problem is causing problems downstream in the environment. You can then return to the event list in the Web GUI and take action against this single event.

What to do next

Perform the actions that are appropriate for your environment against the events that are identified by the searches. See http://www-01.ibm.com/support/knowledgecenter/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/webtop/wip/task/web_use_jsel_manageevents.html? for the Event Viewer and http://www-01.ibm.com/support/knowledgecenter/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/webtop/wip/task/web_use_ael_managingevents.html? for the AEL.

Related concepts

Operations Management tasks
Use this information to understand the tasks that users can perform using Operations Management.

Related tasks

Configuring event search

Customizing event management tools

Related reference

Troubleshooting event search
How to resolve problems with your event search configuration.

Related information

Using Operations Analytics - Log Analysis V1.3.5 to search data
Using Operations Analytics - Log Analysis V1.3.3 to search data
Event search examples on IBM DeveloperWorks

Event search workflow for operators

A typical workflow to show operators how the event search tools can assist triaging and diagnostics from the event list.

Assume the following situation: An event storm has been triggered but the cause of the storm is unclear. For the past hour, large numbers of critical events have been generated. Run the event search tools against the critical events.

1. To gain an overview of what has happened since the event storm started, select the critical events. Then, right-click and click **Event search** > **Show event dashboard by node** > **1 hour before event**. The charts that are displayed show how the critical events break down, by node, alert group, severity, and so on.
2. Check whether any nodes stand out on the charts. If so, close the Operations Analytics - Log Analysis GUI, return to the event list and find an event that originates on that node. For example, type a filter in the text box on the Event Viewer toolbar like the following example that filters on critical events from the *mynode* node.

   ```
   SELECT * from alerts.status where Node = mynode; and Severity = 5;
   ```

   After the event list refreshes to show only matching events, select an event, right-click, and click **Event search** > **Search for events by node** > **1 hour before event**.
3. In the search results, check whether an event from that node stands out. If so, close the Operations Analytics - Log Analysis GUI, return to the event list, locate the event, for example, by filtering on the summary or serial number:

   ```
   SELECT * from alerts.status where Node = mynode; and Summary like
   “Link Down ( FastEthernet0/13 )”;
   ```

   ```
   SELECT * from alerts.status where Node = mynode; and Serial = 4586967;
   ```

   Action the event.
4. If nothing stands out that identifies the cause of the event storm, close the Operations Analytics - Log Analysis GUI and return to the event list. Select all the critical events again and click **Event search** > **Show keywords and event count** > **1 hour before event**.
5. From the results, look in the **Common Patterns** area on the navigation pane. Looks for keywords that are non generic but have a high occurrence, for instance host name or IP addresses.
6. Refine the search results by clicking relevant keywords to copy them to the **Search** field and running the search. All events in which the keyword occurs are displayed, and the **Common Patterns** area is updated.
7. If an event stands out as the cause of the event storm, close the Operations Analytics - Log Analysis GUI, return to the event list, and action the event. If not, continuously refine the search results by searching against keywords until a likely root cause event stands out.

For possible actions from the Event Viewer see http://www-01.ibm.com/support/knowledgecenter/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/webtop/wip/task/web_use_jsel_manageevents.html. For possible actions from the Active Event List, see http://www-01.ibm.com/support/knowledgecenter/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/webtop/wip/task/web_use_ael_managingevents.html. Other actions are possible, depending on the tools that are implemented in your environment.

Related information

Event search examples on IBM DeveloperWorks

Using Operations Analytics - Log Analysis V1.3.5 to search data

Using Operations Analytics - Log Analysis V1.3.3 to search data

Chapter 10. Event Analytics

You can deploy Event Analytics in your on premises environment or you can deploy Cloud Native Analytics in your IBM Cloud Private environment as part of Operations Management on IBM Cloud Private.

Event Analytics on premises

Use Event Analytics to help you analyze seasonal trends and related events, while you monitor and manage events.

Event Analytics overview

Event Analytics allows you to identify seasonal patterns of events and related events within their monitored environment.

Seasonal events

Event Analytics uses statistical analysis of IBM Tivoli Netcool/OMNIbus historical event data to determine the seasonality of events, such as when and how frequently events occur. The results of this analysis are output in both reports and graphs.

The data that is presented in the event seasonality report helps you to identify seasonal event patterns within their infrastructure. For example, an event that periodically occurs at an unscheduled specific time is highlighted. Seasonal Event Rules are grouped by state in the Seasonal Event Rules portlet. You can

- Use the View Seasonal Events UI to analyze seasonal events and associated related events.
- Deploy validated seasonal event rules, without writing new code. Rules that are generated in this way can have various actions applied to them.

Related events

Event Analytics uses statistical analysis of Tivoli Netcool/OMNIbus historical event data to determine which events have a statistical tendency to occur together. Event Analytics outputs the results of this statistical analysis as event groups, on a scheduled basis. You can:

- Use the related events UI to analyze these event groups.
- Deploy validated event groups as Netcool/Impact correlation rules with a single click, without the need to write any code. Correlation rules that are generated in this way act on real-time event data to show a single synthetic event for the events in the event group.
- Present all events in the group as children of this synthetic event. This view decreases the number of events displayed to your operations staff in the Event Viewer.
- Use the Related Event portlet to analyze patterns in groups and deploy correlation rules based on common event types between the groups.

The system uses the most actionable event in the group as the parent event to be set by the correlation rule. By default, the most actionable event in the group is the most ticketed or acknowledged event. Before you deploy the correlation rule, you can change the parent event setting. A synthetic event is created with some of the properties of the parent event, and all the related events are grouped under this synthetic event.

Event groups are generated by scheduled runs of related event configurations. A default related event configuration is provided. You can create your own configurations and specify which historical data to analyze. For example, you can specify a custom time range, an event filter, and schedule. For more information about related events, see “Related events” on page 314.

You can create a pattern based on the related event groups discovered by the related event analytic. The system can also suggest name patterns to you based on the related event groups discovered by the

related event analytic. You can use an event in the group as the parent event to be set by the correlation rule, or create a synthetic event as the parent. You can also test the performance of a pattern before it is created to check the number of related events groups and events returned for a pattern.

Installing and uninstalling Event Analytics

Read the following topics before you install or uninstall Event Analytics.

Prerequisites

Before you install Event Analytics you must complete the following preinstallation tasks.

Event Archiving

You must be running a database with archived events. Event Analytics supports the Db2 and Oracle databases. Event Analytics support of MS SQL requires a minimum of IBM Tivoli Netcool/Impact 7.1.0.1.

You can use a gateway to archive events to a database. In reporting mode, the gateway archives events to a target database. For more information, see http://www.ibm.com/support/knowledgecenter/SSSHTQ/omnibus/gateways/jdbcgw/wip/concept/jdbcgw_intro.html.

Note: The gateway can operate in two modes: audit mode and reporting mode. Event Analytics only supports reporting mode.

Browser Requirements

To display the Seasonal Event Graphs in Microsoft Internet Explorer, you must install the Microsoft Silverlight plug-in.

Reduced Memory

If you are not running Event Analytics on Solaris, remove the comment from or add the following entry in the `jvm.options` file:

```
#-Xgc:classUnloadingKickoffThreshold=100
```

Removing the comment from or adding that entry dynamically reduces memory requirements.

Installing Event Analytics

You can install Event Analytics with the IBM Installation Manager GUI or console, or do a silent installation. Event Analytics supports IBM Installation Manager 1.7.2 up to 1.8.4.

For more information about installing and using IBM Installation Manager, see the following IBM Knowledge Center:

http://www-01.ibm.com/support/knowledgecenter/SSDV2W/im_family_welcome.html

Netcool/Impact installation components

Select the components of Netcool/Impact that you want to install.

If you purchased IBM Netcool Operations Insight the Impact Server Extensions component is displayed in the list and is selected automatically. This component contains extra Impact Server features that work with IBM Netcool Operations Insight.

If you accept the default selection, both the GUI Server and the Impact Server are installed on the same computer. In a production environment, install the Impact Server and the GUI Server on separate computers. So, for example, if you already installed the Impact Server on another computer, you can choose to install the GUI Server alone.

The component Installation Manager is selected automatically on the system that is not already installed with Installation Manager.

Netcool/Impact does not support Arabic or Hebrew, therefore Event Analytics users, who are working in Arabic or Hebrew, see some untranslated English text.

Installing Event Analytics (GUI)
You can install Event Analytics with the IBM Installation Manager GUI.

Before you begin

- Determine which Installation Manager user mode you require.
- Ensure that the necessary user permissions are in place for your intended installation directories.
- Configure `localhost` on the computer where Event Analytics packages are to be installed.

About this task
The installation of Event Analytics requires you to install product packages for the following product groups:

- IBM Tivoli Netcool/Impact
- IBM Tivoli Netcool/OMNIbus
- IBM Netcool.

The steps for starting Installation Manager are different depending on which user mode you installed it in. The steps for completing the Event Analytics installation with the Installation Manager wizard are common to all user modes and operating systems.

Installation Manager takes account of your current umask settings when it sets the permissions mode of the files and directories that it installs. If you use Administrator mode or Non-administrator mode and your umask is 0, Installation Manager uses a umask of 22. If you use Group mode, Installation Manager ignores any group bits that are set and uses a umask of 2 if the resulting value is 0.

To install the packages and features, complete the following steps.

Procedure

1. Start Installation Manager. Change to the `/eclipse` subdirectory of the Installation Manager installation directory and use the following command to start Installation Manager:

   `./IBMIM`

   To record the installation steps in a response file for use with silent installations on other computers, use the `-record` *`response_file`* option. For example:

   ```
   ./IBMIM -record /tmp/install_1.xml
   ```

2. Configure Installation Manager to point to either a local repository or an IBM Passport Advantage repository, where the download packages are available. Within the IBM Knowledge Center content for Installation Manager, see the topic that is called *Installing packages by using wizard mode*.
3. In the main Installation Manager window, click **Install** and follow the installation wizard instructions to complete the installation.
4. In the **Install** tab select the following installation packages, and then click **Next**.
   - Packages for IBM Tivoli Netcool/Impact:

     IBM Tivoli Netcool/Impact GUI Server_7.1.0.17
     IBM Tivoli Netcool/Impact Server_7.1.0.17
     IBM Tivoli Netcool/Impact Server Extensions for Netcool Operations Insight_7.1.0.17
   - Packages for IBM Tivoli Netcool/OMNIbus:

     IBM Tivoli Netcool/OMNIbus_8.1.0.21
   - Packages for IBM Tivoli Netcool/OMNIbus Web GUI:

     IBM Tivoli Netcool/OMNIbus Web GUI_8.1.0.17

5. In the **Licenses** tab, review the licenses. If you are happy with the license content select **I accept the terms in the license agreements** and click **Next**.
6. In the **Location** tab, enter information for the `Installation Directory` and `Architecture` or progress with the default values, and click **Next**.
   - For `IBM Netcool`, the default values are `/opt/IBM/netcool` and `64-bit`.
   - For `IBM Netcool Impact`, the default values are `/opt/IBM/tivoli/impact` and `64-bit`.
7. In the **Features** tab, select the following features and then click **Next**. Other features are auto-selected.

Table 42. Available features

Feature	Description
IBM Tivoli Netcool/OMNIbus Web GUI 8.1.0.17 > **Install base features**	To install and run Event Analytics.
Netcool Operations Insight Extensions Web GUI 8.1.0.17 > **Install Event Analytics**	Contains the Event Analytics components.

8. In the **Summary** tab, review summary details. If you are happy with summary details click **Next**, but if you need to change any detail click **Back**.
9. To complete the installation, click **Finish**.

Results
Installation Manager installs Event Analytics.

What to do next

1. Configure the ObjectServer for Event Analytics, see "Configuring the ObjectServer " on page 277.
2. Connect to a valid database from within IBM Tivoli Netcool/Impact. To configure a connection to one of the Event Analytics supported databases, see the following topics:
   - Db2: "Configuring Db2 database connection within Netcool/Impact" on page 259
   - Oracle: "Configuring Oracle database connection within Netcool/Impact" on page 261
   - MS SQL: "Configuring MS SQL database connection within Netcool/Impact" on page 263
3. If you add a cluster to the Impact environment, you must update the data sources in IBM Tivoli Netcool/Impact For more information, see "Configuring extra failover capabilities in the Netcool/Impact environment" on page 279.
4. If you add a cluster to the Impact environment, you must update the data sources in IBM Tivoli Netcool/Impact. For more information, see "Configuring extra failover capabilities in the Netcool/Impact environment" on page 279.
5. You must set up a remote connection from the Dashboard Application Services Hub to Netcool/Impact. For more information, see "Netcool/Impact remote connection" on page 245.

Installing Event Analytics (Console)
You can install Event Analytics with the IBM Installation Manager console.

Before you begin

Obtain an IBM ID and an entitlement to download Event Analytics from IBM Passport Advantage. The packages that you are entitled to install are listed in Installation Manager.

Take the following actions:

- Determine which Installation Manager user mode you require.
- Ensure that the necessary user permissions are in place for the installation directories.

- Decide which features that you want to install from the installation packages and gather the information that is required for those features.
- Configure `localhost` on the computer where Event Analytics is to be installed.

About this task

The steps for starting Installation Manager are different depending on which user mode you installed it in. The steps for completing the Event Analytics installation with the Installation Manager console are common to all user modes and operating systems.

Installation Manager takes account of your current umask settings when it sets the permissions mode of the files and directories that it installs. If you use Administrator mode or Non-administrator mode and your umask is 0, Installation Manager uses a umask of 22. If you use Group mode, Installation Manager ignores any group bits that are set and uses a umask of 2 if the resulting value is 0.

Procedure

1. Change to the `/eclipse/tools` subdirectory of the Installation Manager installation directory.
2. Use the following command to start Installation Manager:
   - `./imcl -c` OR `./imcl -consoleMode`
3. Configure Installation Manager to download package repositories from IBM Passport Advantage:
   a) From the Main Menu, select **Preferences**.
   b) In the **Preferences** menu, select **Passport Advantage**.
   c) In the **Passport Advantage** menu, select **Connect** to **Passport Advantage**.
   d) When prompted, enter your IBM ID user name and password.
   e) Return to the Main Menu.
4. From the options that are provided on the installer, add the repository that you want to install.
5. From the Main Menu, select **Install**.

   Follow the installer instructions to complete the installation. The installer requires the following inputs at different stages of the installation:

   - Select Event Analytics
   - When prompted, enter an Installation Manager shared directory or accept the default directory.
   - When prompted, enter an installation directory or accept the default directory.
   - Clear the features that you do not require.
   - If required, generate a response file for use with silent installations on other computers. Enter the directory path and a file name with a `.xml` extension. The response file is generated before installation completes.
6. When the installation is complete, select **Finish**.

Results

Installation Manager installs Event Analytics.

What to do next

If you add a cluster to the Impact environment, you must update the data sources in IBM Tivoli Netcool/Impact 7.1. For more information, see “Configuring extra failover capabilities in the Netcool/Impact environment” on page 279.

Silently installing Event Analytics

You can install Event Analytics silently with IBM Installation Manager. This installation method is useful if you want identical installation configurations on multiple workstations. Silent installation requires a response file that defines the installation configuration.

Before you begin

Take the following actions:

- Create or record an Installation Manager response file.

  You can specify a local or remote IBM Tivoli Netcool/OMNIbus package and a Netcool Operations Insight Extensions Web GUI package with a repository in the response file. You can also specify that Installation Manager downloads the packages from IBM Passport Advantage. For more information about specifying authenticated repositories in response files, search for the *Storing credentials* topic in the Installation Manager information center:

  http://www-01.ibm.com/support/knowledgecenter/SSDV2W/im_family_welcome.html

  A default response file is included in the Event Analytics installation package in `responsefiles/`*platform*, where *platform* can be `unix` or `windows`.

  When you record a response file, you can use the `-skipInstall` argument to create a response file for an installation process without performing the installation. For example:

  – Create or record a skipInstall:

  ```
  IBMIM.exe -record C:\response_files\install_1.xml -skipInstall
   C:\Temp\skipInstall
  ```

- Determine which Installation Manager user mode you require.
- Read the license agreement. The license agreement file, `license.txt`, is stored in the `/native/license_version.zip` archive, which is contained in the installation package.
- Ensure that the necessary user permissions are in place for your intended installation directories.
- Configure `localhost` on the computer where Event Analytics is to be installed.

Procedure

1. Change to the `/eclipse/tools` subdirectory of the Installation Manager installation directory.
2. To encrypt the password that is used by the administrative user for the initial log-in to Dashboard Application Services Hub, run the following command:

   - `./imutilsc encryptString` *password*

   Where *password* is the password to be encrypted.
3. To install Event Analytics, run the following command:

   - `./imcl -input` *response_file* `-silent -log /tmp/install_log.xml -acceptLicense`

   Where *response_file* is the directory path to the response file.

Results

Installation Manager installs Event Analytics.

What to do next

If you add a cluster to the Impact environment, you must update the data sources in IBM Tivoli Netcool/Impact 7.1. For more information, see “Configuring extra failover capabilities in the Netcool/Impact environment” on page 279.

Post-installation tasks
You must perform these tasks following installation of Event Analytics.

Netcool/Impact remote connection
Db2 is the default archive database. To use Oracle or MS SQL as the archive database, you must set up a remote connection to Netcool/Impact.

About this task
The IBM Dashboard console Netcool/Impact connection should use HTTPS.

Procedure

1. Log in to the IBM Dashboard console. If you fail to connect to the Dashboard console, ensure that the firewall on your computer is disabled.
2. Click the **Console** icon.
3. Select **Connections**.
4. Click the **Create new remote provider** icon.
5. Enter the Netcool/Impact UI server *Host name*, *Port*, *Name* and *Password*.
6. Click **Search**.
7. Select **Impact_NCICLUSTER** as your data provider.
8. Click **OK**.

Uninstalling Event Analytics
You can uninstall Event Analytics with the IBM Installation Manager GUI or console, or do a silent uninstall.

For more information about installing and using IBM Installation Manager, see the following IBM information center:

http://pic.dhe.ibm.com/infocenter/install/v1r7/index.jsp

Uninstalling Event Analytics
Use IBM Installation Manager to remove Event Analytics .

Before you begin

Take the following actions:

- Stop all Event Analytics processes.
- Back up any data or configuration files that you want to retain.
- To do a silent removal, create or record an Installation Manager response file.

  Use the `-record` *`response_file`* option.

  To create a response file without installing the product, use the `-skipInstall` option. For example:

  1. Create or record a skipInstall:

     ```
     IBMIM.exe -record C:\response_files\install_1.xml -skipInstall C:\Temp
     \skipInstall
     ```

  2. To create an uninstall response file, using the created skipInstall:

     ```
     IBMIM.exe -record C:\response_files\uninstall_1.xml -skipInstall C:\Temp
     \skipInstall
     ```

About this task

Note: To uninstall Tivoli Netcool/OMNIbus Web GUI 8.1.0.17, you must first uninstall the Netcool Operations Insight Extensions for IBM Tivoli Netcool/OMNIbus Web GUI_8.1.0.17, including the Event Analytics feature.

Procedure

GUI removal

1. To remove Event Analytics with the Installation Manager GUI:
   a) Change to the `/eclipse` subdirectory of the Installation Manager installation directory.
   b) Use the following command to start the Installation Manager wizard:

      ```
      ./IBMIM
      ```

   c) In the main Installation Manager window, click **Uninstall**.
   d) Select the offerings that you want to remove and follow the Installation Manager wizard instructions to complete the removal.

Console removal

2. To remove Event Analytics with the Installation Manager console:
   a) Change to the `/eclipse/tools` subdirectory of the Installation Manager installation directory.
   b) Use the following command to start the Installation Manager:

      ```
      ./imcl -c
      ```

   c) From the Main Menu, select **Uninstall**.
   d) Select the offerings that you want to remove and follow the Installation Manager instructions to complete the removal.

Silent removal

3. To silently remove Event Analytics:
   a) Change to the `/eclipse/tools` subdirectory of the Installation Manager installation directory.
   b) Use the following command to start the Installation Manager:

      ```
      ./imcl -input response_file -silent -log /tmp/install_log.xml -
      acceptLicense
      ```

      Where *response_file* is the directory path to the response file that defines the removal configuration

Results

Installation Manager removes the files and directories that it installed.

What to do next

Files that Installation Manager did not install, and configuration files that were changed, are left in place. Review these files and remove them or back them up as appropriate.

Upgrading Event Analytics

Follow these instructions to upgrade Event Analytics to the latest version.

Upgrading Event Analytics

You can upgrade the IBM Netcool Operations Insight packages for Event Analytics by applying the latest fix packs.

About this task

To perform the upgrade, use the IBM Installation Manager **Update** functions to locate update packages, and update your environment with the following product update packages:

- Packages for IBM Tivoli Netcool/Impact:

IBM Tivoli Netcool/Impact GUI Server_7.1.0.17
IBM Tivoli Netcool/Impact Server_7.1.0.17
IBM Tivoli Netcool/Impact Server Extensions for Netcool Operations Insight_7.1.0.17

- Packages for IBM Tivoli Netcool/OMNIbus:

  IBM Tivoli Netcool/OMNIbus_8.1.0.21

- Packages for IBM Tivoli Netcool/OMNIbus Web GUI:

  IBM Tivoli Netcool/OMNIbus Web GUI_8.1.0.17
  Netcool Operations Insight Extensions for IBM Tivoli Netcool/OMNIbus Web GUI_8.1.0.17

Procedure

The product update packages must be updated individually. Complete steps 1 - 3 for each product update package.

1. Start Installation Manager. Change to the `/eclipse` subdirectory of the Installation Manager installation directory and enter the following command to start Installation Manager:

   `./IBMIM`

2. Configure Installation Manager to point to either a local repository or an IBM Passport Advantage repository, where the download package is available. Within the IBM Knowledge Center content for Installation Manager, see the topic that is called *Installing packages by using wizard mode*.

   See the following URL within the IBM Knowledge Center content for Installation Manager:

   http://www-01.ibm.com/support/knowledgecenter/SSDV2W/im_family_welcome.html

3. In the main Installation Manager window, click **Update** and complete the following type of installation wizard instructions to complete the installation of your update package:
   a) In the **Update Packages** tab, select the product group to find related update packages, and click **Next**. A list of the available update packages displays.
   b) From the list of available update packages, select one update package that you want to install, and click **Next**. Remember you can install only one update package at a time.
   c) In the **Licenses** tab, review the licenses. Select **I accept the terms in the license agreements** and click **Next**.
   d) In the **Features** tab, select the features for your update package, and click **Next**.
   e) Complete the configuration details, and click **Next**.
   f) In the **Summary** tab, review summary details. If you need to change any detail click **Back**, but if you are happy with summary details click **Update** and wait for the installation of the update package to complete.
   g) When the installation of the update package completes, the window updates with details of the installation. Click **Finish**.
4. To ensure that the seasonal event reports that were created before upgrading are visible, you must run the `SE_CLEANUPDATA` policy as follows.
   a) Login as the administrator to the server where IBM Tivoli Netcool/Impact is stored and running.
   b) Navigate to the policies tab and search for the `SE_CLEANUPDATA` policy.
   c) To open the policy, double-click on the policy.
   d) To run the policy, select the run button on the policy screen toolbar.
5. To view the event configurations in the View Seasonal Events portlet, rerun the configurations. For more information about running event configurations, see the "Configuring analytics" on page 287 topics.

What to do next

1. Verify that the correct packages are installed. After you update each package, and to ensure that you have the correct environment, verify that the following packages are installed.

- Packages for IBM Tivoli Netcool/Impact:

  IBM Tivoli Netcool/Impact GUI Server_7.1.0.17
  IBM Tivoli Netcool/Impact Server_7.1.0.17
  IBM Tivoli Netcool/Impact Server Extensions for Netcool Operations Insight_7.1.0.17

- Packages for IBM Tivoli Netcool/OMNIbus:

  IBM Tivoli Netcool/OMNIbus_8.1.0.21

- Packages for IBM Tivoli Netcool/OMNIbus Web GUI:

  IBM Tivoli Netcool/OMNIbus Web GUI_8.1.0.17
  Netcool Operations Insight Extensions for IBM Tivoli Netcool/OMNIbus Web GUI_8.1.0.17

2. Configure the ObjectServer for Event Analytics. For more information about configuring the ObjectServer for Event Analytics, see "Configuring the ObjectServer " on page 277.
3. Connect to a valid database from within IBM Tivoli Netcool/Impact. To configure a connection to one of the Event Analytics supported databases, see the following topics:
   - Db2: "Configuring Db2 database connection within Netcool/Impact" on page 259
   - Oracle: "Configuring Oracle database connection within Netcool/Impact" on page 261
   - MS SQL: "Configuring MS SQL database connection within Netcool/Impact" on page 263
4. If you add a cluster to the Impact environment, you must update the data sources in IBM Tivoli Netcool/Impact 7.1. For more information, see "Configuring extra failover capabilities in the Netcool/Impact environment" on page 279.
5. If you want to make use of the pattern generalization feature in Event Analytics, you must configure the type properties used for event pattern creation in IBM Tivoli Netcool/Impact. For more information about configuring the type properties used for event pattern creation in IBM Tivoli Netcool/Impact, see "Configuring event pattern processing" on page 270.

Upgrading Event Analytics from stand-alone installations of Netcool/OMNIbus and Netcool/Impact
If you have stand-alone installations of Netcool/OMNIbus with Web GUI and Netcool/Impact, you can perform the upgrade.

Before you begin
Ensure that the following product packages are already installed:

- IBM Tivoli Netcool/Impact packages:

  IBM Tivoli Netcool/Impact GUI Server_7.1.0.17
  IBM Tivoli Netcool/Impact Server_7.1.0.17

- IBM Tivoli Netcool/OMNIbus packages:

  IBM Tivoli Netcool/OMNIbus_8.1.0.21

- IBM Netcool packages:

  IBM Tivoli Netcool/OMNIbus Web GUI_8.1.0.17

About this task
This upgrade scenario is for users who already use Tivoli Netcool/OMNIbus and Netcool/Impact but do not have the Netcool Operations Insight packages that are needed for the Event Analytics function, and now want the Event Analytics function.

For this upgrade scenario, you must use IBM Installation Manager to **Install** the product packages that are required for the Event Analytics function, then **Update** the product packages. The **Install** of product packages locates and installs the following two packages:

IBM Tivoli Netcool/Impact Server Extensions for Netcool Operations Insight_7.1.0.17
Netcool Operations Insight Extensions for IBM Tivoli Netcool/OMNIbus Web GUI_8.1.0.17

To perform the upgrade, complete the following steps.

Procedure

1. Start Installation Manager. Change to the `/eclipse` subdirectory of the Installation Manager installation directory and enter the following command to start Installation Manager:

   `./IBMIM`
2. Configure Installation Manager to point to either a local repository or an IBM Passport Advantage repository, where the download package is available. Within the IBM Knowledge Center content for Installation Manager, see the topic that is called *Installing packages by using wizard mode*.

   See the following URL within the IBM Knowledge Center content for Installation Manager:

   http://www-01.ibm.com/support/knowledgecenter/SSDV2W/im_family_welcome.html
3. To install your packages in the main Installation Manager, click **Install** and complete the steps in the installation wizard to complete the installation of your packages:
   a) In the **Install** tab, select the following product groups and product installation packages, and click **Next**.
      - IBM Tivoli Netcool/Impact Server Extensions for Netcool Operations Insight_7.1.0.17
      - Tivoli Netcool/OMNIbus Web GUI Version 8.1.0.17
      - Netcool Operations Insight Extensions for IBM Tivoli Netcool/OMNIbus Web GUI_8.1.0.17

   b) In the **Licenses** tab, review the licenses. When you are happy with the license content select **I accept the terms in the license agreements** and click **Next**.
   c) In the **Location** tab, use the existing package group and location.
   d) In the **Features** tab, select the features for your packages, and click **Next**.
   e) In the **Summary** tab, review summary details. If you need to change any detail click **Back**, but if you are happy with summary details click **Install** and wait for installation of the package to complete.
   f) When installation of the packages completes, the window updates with details of the installation. Click **Finish**.
4. Migrate the rollup configuration. For more information about updating the rollup configuration, see "Adding columns to seasonal and related event reports" on page 265.

What to do next

1. Verify that the correct packages are installed. After you update each package, and to ensure that you have the correct environment, verify that the following packages are installed.
   - Packages for IBM Tivoli Netcool/Impact:

     IBM Tivoli Netcool/Impact GUI Server_7.1.0.17
     IBM Tivoli Netcool/Impact Server_7.1.0.17
     IBM Tivoli Netcool/Impact Server Extensions for Netcool Operations Insight_7.1.0.17
   - Packages for IBM Tivoli Netcool/OMNIbus:

     IBM Tivoli Netcool/OMNIbus_8.1.0.21
   - Packages for IBM Tivoli Netcool/OMNIbus Web GUI:

     IBM Tivoli Netcool/OMNIbus Web GUI_8.1.0.17
     Netcool Operations Insight Extensions for IBM Tivoli Netcool/OMNIbus Web GUI_8.1.0.17
2. Configure the ObjectServer for Event Analytics. For more information about configuring the ObjectServer for Event Analytics, see "Configuring the ObjectServer " on page 277.
3. Connect to a valid database from within IBM Tivoli Netcool/Impact. To configure a connection to one of the Event Analytics supported databases, see the following topics:
   - Db2: "Configuring Db2 database connection within Netcool/Impact" on page 259
   - Oracle: "Configuring Oracle database connection within Netcool/Impact" on page 261

- MS SQL: "Configuring MS SQL database connection within Netcool/Impact" on page 263

4. If you add a cluster to the Impact environment, you must update the data sources in IBM Tivoli Netcool/Impact. For more information, see "Configuring extra failover capabilities in the Netcool/Impact environment" on page 279.
5. If you want to make use of the pattern generalization feature in Event Analytics, you must configure the type properties used for event pattern creation in IBM Tivoli Netcool/Impact. For more information about configuring the type properties used for event pattern creation in IBM Tivoli Netcool/Impact, see "Configuring event pattern processing" on page 270.

Migration of rollups in Netcool/Impact v7.1.0.13

Fix pack v7.1.0.13 uses a new format for the creation of rollups that is different to previous versions of Netcool/Impact. A migration script is automatically executed during the install or upgrade process to convert pre-existing rollups to the v7.1.0.13 format. Run the Event Analytics configuration wizard after the upgrade to v7.1.0.13 to verify and save your configuration (see **note** below).

- In **v7.1.0.12 (or earlier)** the rollup display names are free-form text with no formatting applied. For example:

```
reevent_rollup_1_column_name=ORIGINALSEVERITY
reevent_rollup_1_type=MAX
reevent_rollup_1_display_name=MaxSeverity
```

  In this scenario, Netcool Operations Insight creates a new column in the database called `MaxSeverity`. The display name in Dashboard Application Services Hub will be `MaxSeverity`, or whatever is defined in the translation directory in the Netcool/Impact `uiprovider` directory.

  With the introduction of the Event Analytics configuration wizard in v7.1.0.13, it was necessary to apply a new format to rollup display names.

- In **v7.1.0.13** a format of <*column_name*>_<*type*> is applied to rollup display names: For example:

```
reevent_rollup_1_column_name=ORIGINALSEVERITY
reevent_rollup_1_type=MAX
reevent_rollup_1_display_name=ORIGINALSEVERITY_MAX
```

  Using the wizard, you can apply any display name to a column, in any language. Because creating database columns in any language could have been error prone, the format of the names for rollup database columns is now set to <*column_name*>_<*type*>. The display name is stored in the *translation* directory and files in the `uiproviderconfig` directory. This format makes it is possible to change display names using the Event Analytics configuration wizard. For this reason, a migration script is executed during install/upgrade to transform all pre-existing rollups to the v7.1.0.13 format.

Note: You must run the Event Analytics configuration wizard after upgrading to Netcool/Impact v7.1.0.13. The following artifacts will be changed as a result of the rollup migration script in Netcool/Impact v7.1.0.13:

- Stored metadata for rollups in configuration
- Database columns (renamed)
- Output parameters for policies
- Properties files
- Translated properties files

Complete the steps of the wizard as described in "Configuring Event Analytics using the wizard" on page 251 after upgrading to v7.1.0.13 to verify and save any customizations to your configuration. Backup files containing previous customizations are stored in `$IMPACT_HOME/backup/install/gui_backup/<pre-FP13 fp name>/uiproviderconfig/`.

Configuring the system

Perform these tasks to configure and optionally customize the system prior to use.

Configuring Event Analytics

Configure Event Analytics prior to use.

Configuring Event Analytics using the wizard

In Netcool/Impact V7.1.0.13 it is recommended to use the Event Analytics configuration wizard instead of the `./nci_trigger` command to edit properties in the NOI Shared Configuration properties file. The setup wizard guides you through the Event Analytics configuration process. You must run the Event Analytics configuration wizard after upgrading to Netcool/Impact V7.1.0.13 to verify and save your configuration.

Note: If you are running the wizard within a load-balancing environment, then first perform the following steps on each Netcool/Impact UI server:

1. Edit the `$IMPACT_HOME/etc/server.props` file.
2. Set the `impact.noi.ui.hostname` variable to the hostname of one of the Netcool/Impact UI servers.

To launch the wizard, click  **Insights** and select **Event Analytics Configuration**.

The Event Analytics Configuration wizard consists of two parts:

1. Configuring access to the following databases:
   - The Tivoli Netcool/OMNIbus Historical Event Database, containing historical event data used to analyze historical events for Event Analytics.
   - The Tivoli Netcool/OMNIbus ObjectServer, containing live event data to be enriched based on insights derived from Event Analytics processing.
2. Configuring settings to control Event Analytics processing.

Configuring the historical event database

Configure access to the Tivoli Netcool/OMNIbus historical event database that contains the data used to analyze historical events for Event Analytics. On the historical event database window, you specify the database type, connection details, table name, and timestamp format.

Before you begin

If you have custom fields in your Historical Event database, then before doing this task you must first map the custom field names to the corresponding standard field names in Netcool/Impact by creating a database view, as described in "Mapping customized field names" on page 279. In the appropriate step in the procedure below, you must specify that database view name instead of the Historical Event database `reporter_status` table name.

You can set up an Oracle database as your historical event database with a custom URL. You can also configure a system identifier (SID) or Service Name in order to connect the database. The SID or Service Name configuration settings are not available in the Event Analytics Configuration wizard. Instead refer to the backend configuration instructions: "Configuring Oracle database connection within Netcool/Impact" on page 261.

Procedure

1. Specify the database type used for the historical event database:
   - **Db2**
   - **Oracle**
   - **MS SQL Server**
2. Enter the connection details for the database in the fields provided.

Hostname
Enter the name of the server hosting the database.

Port
Enter the port number to be used to connect to the server that hosts the database.

Username
Enter the username for connecting to the database.

Password
Enter the password for the specified username.

The remaining fields in this section differ depending on the database type selected in step 1.

- If you selected a database type of **Db2** or **MS SQL Server**, then complete the following field:

  Database Name
  In the Database Name field enter the name of the database you want to access. For example REPORTER.

- If you selected a database type of **Oracle**, then complete the following fields:

  Communication method
  Specify whether to use a custom URL, Oracle SID (alphanumeric system identifier), or Oracle service name.

  Custom URL
  If you set **Communication method** to Custom URL, then specify the URL in the *<hostname>*:*<port>*:*<server>* format.

  Note: For Real Application Clusters (RAC) servers, see additional information here: https://www.ibm.com/support/knowledgecenter/SSSHYH_7.1.0.16/com.ibm.netcoolimpact.doc/user/rac_cluster_support.html

  SID
  If you set **Communication method** to SID, then specify the SID.

  Service name
  If you set **Communication method** to Service name, then specify the service name.

3. Click **Connect** to validate your connection to the historical event database.
4. For the table you want to query, select a **Database schema** and **History table** or view from the drop-down lists provided.
   a) The options available under **Database schema** are based on the username provided to connect to the historical event database.
   b) The options available under **History table** are based on the selected Database schema.

   Important: If you have custom fields in your Historical Event database, and you created a database view to map these fields, as described in "Mapping customized field names" on page 279, then you must select that database view from the **History table** drop-down list.
5. Specify the timestamp field used in the historical event database to store the first occurrence of an event.

Specifying the primary and backup ObjectServer
On the ObjectServer window, enter the hostname, port, and user credentials to connect to the primary and backup ObjectServers.

Procedure

1. In the fields provided, enter the connection details to the primary ObjectServer:

   Hostname
   Enter the hostname where the primary ObjectServer is installed.

   Port
   Specify the port number that the primary ObjectServer will use.

Username
Enter the username to access the ObjectServer.

Password
Enter the password for the specified username.

2. To enable a backup ObjectServer, select the **Enable backup ObjectServer** check box and enter the connection details:

   Note: Selecting **Enable backup ObjectServer** enables the fail back option when Impact cannot connect to the database. Deselecting this option disables the backup.

   Hostname
   Enter the hostname where the backup ObjectServer is installed.

   Port
   Specify the port number that the backup ObjectServer will use.

   The Username and Password are same as the credentials specified for the primary ObjectServer.

   Click **Connect** to connect to the ObjectServer.

Adding report fields
Select report fields to add additional information to seasonal and related event reports, historical event reports, and instance reports.

Before you begin
If you add any custom columns to any of the reports defined in the Event Analytics Setup Wizard and define a column title for this column, then, if you want translations of the column title to appear in the relevant Event Analytics report, you must edit the relevant translation files. For more information about translating column labels, see http://ibm.biz/impact_trans.

If you do not do this, then when the column title is displayed in the relevant Event Analytics report, it will appear in English, regardless of the locale of the browser.

Procedure

1. Specify the **Aggregate fields**.

   You can add aggregate fields to the seasonal and related event reports, by applying predefined aggregate functions to selected fields from the Historical Event Database. These fields are displayed in the seasonal and related event reports in the same order as they appear below.

   Example: To display the maximum severity of all the events that make up a related event group select the SEVERITY field, then apply the Max aggregate function, and click **Include in reports: Related**.

   Note: Once you have saved these settings, you must rerun the relevant configuration scans in order for the changes to become visible in your seasonal and related event reports.

2. Specify the **Historical report** fields.

   The fields specified here are displayed in the historical event report, in the same order as they appear below. The historical event report is shown when you drill in from a seasonal event to its contributing historical events.

   Note: When you next open the historical event report, wait 20 seconds for your changes to appear.

3. Specify the **Instance report** fields.

   The fields specified here are displayed as additional fields in the instance report for a related event group, in the same order as they appear below. The instance report is shown when you drill into event details from a related event group, to show the instances of that group.

   Note: When you next open the instance event report, wait 20 seconds for your changes to appear.

4. Click **Save** to save your changes.

Configuring event suppression
Some events might not be important with respect to monitoring your network environment. For events that do not need to be viewed or acted on, event suppression is available as an action when creating a seasonal event rule.

About this task

For seasonal event rules, specify the ObjectServer fields to use for suppressing and unsuppressing events.

Procedure

1. To suppress an event, select a **Suppression field** and **Suppression field value** from the drop-down lists provided. The field and value that you define here are used to mark the event for suppression when the incoming event matches the seasonal event rule with event suppression selected as one of its actions.
2. To unsuppress an event, select an **Unsuppression field** and **Unsuppression field value** from the drop-down lists provided. The field and value that you define here are used to unsuppress an event when the incoming event matches the seasonal event rule with event suppression selected as one of its actions.
3. Click **Save** to save your changes.

Configuring event pattern processing
Configure how patterns are derived from related events using this example-driven wizard panel.

Before you begin
To configure event pattern processing, you must specify Historical Event Database columns to use for settings such as event type, event identity, and resource, or accept the columns specified as default. If you want to use custom columns, then you must first configure the Impact Event Reader to read these custom fields, as described in the following topic: Netcool/Impact Knowledge Center: OMNIbus event reader service.

About this task

An event pattern is a set of events that typically occur in sequence on a network resource. For example, on a London router LON-ROUTER-1, the following sequence of events might frequently occur: FAN-FAILURE, POWER®-SUPPLY-FAILURE, DEVICE-FAILURE, indicating that the router fan needs to be changed. Using the related event group feature, Event Analytics will discover this sequence of events as a related event group on LON-ROUTER-1.
Using the event pattern feature, Event Analytics can then detect this related event group on any network resource. In the example above, the related event group FAN-FAILURE, POWER-SUPPLY-FAILURE, DEVICE-FAILURE detected on the London router LON-ROUTER-1 can be stored as a pattern and that pattern can be detected on any other network resource, for example, on a router in Dallas, DAL-ROUTER-5.

Procedure

1. Select the appropriate Historical Event Database column(s) for the following **Global settings**:

   Default event type
   An event type is a category of event, for example: FAN-FAILURE, POWER-SUPPLY-FAILURE and DEVICE-FAILURE are event types. By default event type information is stored in the following Historical Event Database column: ALERTGROUP. If you have another set of events that you categorize in a different way, then you can specify additional event type columns in section 2 below.

   Default event identity
   The event identity uniquely identifies an event on a specific network resource. By default the event identity is stored in the following Historical Event Database column: IDENTIFIER.

Resource
A resource identifies a network resource on which events occur. In the example, LON-ROUTER-1 and DAL-ROUTER-5 are examples of resources on which events occur. By default this resource information is stored in the following Historical Event Database column: NODE.

2. If you have another set of events that you categorize in a different way, you can add them as **Additional event types**.
   a) Select the check box to enable **Additional event types**.
   b) Click **Add new**. Add a row for each distinct set of events.
   c) Specify the filters and fields listed below for each set of events. Event Analytics uses these settings to determine event patterns for a set of events. Filters are applied from top to bottom, in the order that they appear in the table. You can change the order by using the controls at the end of the row.

      Type name
      Specify the type name.

      Database filter
      Specify the filter that matches this set of historical events in the Historical Event Database.

      ObjectServer filter
      Specify the filter that matches the corresponding set of live events in the ObjectServer. The ObjectServer filter should be semantically identical to the Database filter, except that you should specify ObjectServer column syntax for the columns.

      Event type field
      An event type is a category of event, for example: FAN-FAILURE, POWER-SUPPLY-FAILURE, and DEVICE-FAILURE are event types. For this set of events, specify the Historical Event Database column that stores event type information.

      Event identity field(s)
      The event identity uniquely identifies an event on a specific network resource. For this set of events, specify the Historical Event Database column or columns that stores event identity information.

Reviewing the configuration
On the **Summary** window, review your settings. You can also save the settings here or click **Back** to make changes to the settings that you configured.

Procedure

1. Review the settings on the **Summary** window.

   Click **Back** or any of the navigation menu links to modify the settings as appropriate.
2. When you are satisfied with the configuration settings, click **Save**.

Exporting the Event Analytics configuration
Use the **nci_trigger** command to export a saved Event Analytics configuration to another system.

Procedure

1. To generate a properties file from the command-line interface, use the following command:

```
nci_trigger server <UserID>/<password> NOI_DefaultValues_Export
 FILENAME directory/filename
```

   Where:

   SERVER
   The server where Event Analytics is installed.

   <UserID>
   The user name of the Event Analytics user.

<password>
The password of the Event Analytics user.

directory
The directory where the file is stored.

filename
The name of the properties file.

For example:

```
./nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Export
 FILENAME
 /tmp/eventanalytic.props
```

2. To import the modified properties file into Netcool/Impact, use the following command:

```
nci_trigger SERVER <UserID>/<password> NOI_DefaultValues_Configure
 FILENAME
directory/filename
```

For example:

```
./nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Configure
 FILENAME
 /tmp/eventanalytic.props
```

Generated properties file
Overwritten property values can be updated in the generated properties file.

You can edit the generated properties file to set up and customize Netcool/Impact for seasonal events and related events. The following properties are in the generated properties file.

```
#
##############################################################################################
####                          NOI Shared Configuration                                  ####
##############################################################################################
#
# If you are updating the Rollup configuration, go to
# The end of the file
# Following section holds the configuration for accessing
# Alerts historical information and storing results
# history_datasource_name Contains the Impact datasourcename
# history_datatype_name Contains the Impact datatype name
# history_database_type Contains the Impact datasource type (Db2, Oracle, MSSQL)
# history_database_table Contains the database table and if required, the schema,
to access the event history
# results_database_type Contains the database type for storing results.
#
# Most likely you do not have to change this configuration
#
history_datasource_name=ObjectServerHistoryDb2ForNOI
history_datatype_name=AlertsHistoryDb2Table
history_database_table=Db2INST1.REPORTER_STATUS
history_database_type=Db2
#
#
#
results_database_type=DERBY
#
# Column name for the analysis
#
history_column_names_analysis=SUMMARY
#
# The column name where the timestamp associated with the records is stored
#
history_column_name_timestamp=FIRSTOCCURRENCE
#
#
#
history_db_timestampformat=yyyy-MM-dd HH:mm:ss.SSS
configuration_db_timestampformat=yyyy-MM-dd HH:mm:ss.SSS
#
##############################################################################################
####                     Seasonality Only Configuration                                 ####
##############################################################################################
```

```
#
# Will only save and process events of this confidence level or higher
#
save_event_threshold=.85
#
# Used in determining the confidentialiy level ranges, by determining the
threshold values.
# level_threshold_high Level is high, when confidentiality is greater than
or equal to
# level_threshold_medium Level is medium, when confidentiality is greater
than or equal to
# level_threshold_low Level is low, when confidentiality is greater than or
equal to
# If the confidentiality doesn't meet any of these conditions, level will be
set to unknown.
#
level_threshold_high=99
level_threshold_medium=95
level_threshold_low=0
#
# Rollup configuration adds additional information to the Seasonal Report data
# number_of_rollup_configuration Contains the number of additional rollup
configuration
# rollup_ <number where its 1 to n >_column_nameContains the column name from
which the data is retreived
# rollup_ <number where its 1 to n >_type Contains the type value
# rollup_ <number where its 1 to n >_display_name A name that needs to be
defined in the UI
#  Types can be defined as follows :
#      MAX, MIN, SUM, NON_ZERO, DISTINCT and EXAMPLE
#      MAX: The maximum value observed for the column, if no value is ever seen
this will default to Integer.MIN_VALUE
#      MIN: The minimum value observed for the column, if no value is ever seen
this will default to Integer.MAX_VALUE
#      SUM: The sum of the values observed for the column.
#      NON_ZERO: A counting column, that counts "Non-Zero"/"Non-Blank"
occurrences of events, this can be useful to
#                track the proportion of events that have been actioned,
or how many events had a ticket number associated
#                with them.
#      DISTINCT: The number of distinct values that have been seem for this key,
value pair
#      EXAMPLE: Show the first non-blank "example" of a field that contained this
key, useful when running seasonality on a
#               field that can't be accessed, such as ALERT_IDENTIFIER, and you want an
example human readable
#               SUMMARY to let you understand the type of problem
#
number_of_rollup_configuration=2
rollup_1_column_name=SEVERITY
rollup_1_type=MIN
rollup_1_display_name=MINSeverity
rollup_2_column_name=SEVERITY
rollup_2_type=MAX
rollup_2_display_name=MAXSeverity
#
################################################################################
####                  Related Events Only Configuration                     ####
################################################################################
#
# Rollup configuration adds additional information to the Related Events data
# reevent_number_of_rollup_configuration Contains the number of additional
rollup configuration
# reevent_rollup_ <number where its 1 to n >_column_nameContains the column
name from which the data is retrieved
# reevent_rollup_ <number where its 1 to n >_type Contains the type value
# reevent_rollup_ <number where its 1 to n >_display_name A name that needs
to be defined in the UI
# reevent_rollup_ <number where its 1 to n >_actionable Numeric only column
that determines the weight for probable root cause
#  Types can be defined as follows :
#      MAX, MIN, SUM, NON_ZERO, DISTINCT and EXAMPLE
#      MAX: The maximum value observed for the column, if no value is ever
seen this will default to Integer.MIN_VALUE
#      MIN: The minimum value observed for the column, if no value is ever
seen this will default to Integer.MAX_VALUE
#      SUM: The sum of the values observed for the column.
#      NON_ZERO: A counting column, that counts "Non-Zero"/"Non-Blank"
occurrences of events, this can be useful to
#                track the proportion of events that have been actioned, or
how many events had a ticket number associated
#                with them.
```

```
#      DISTINCT: The number of distinct values that have been seem for this
key, value pair
#      EXAMPLE: Show the first non-blank "example" of a field that contained
this key, useful when running Seasonality on a
#               field that can't be accessed, such as ALERT_IDENTIFIER, and you want
an example human readable
#               SUMMARY to let you understand the type of problem
#
reevent_number_of_rollup_configuration=3
reevent_rollup_1_column_name=ORIGINALSEVERITY
reevent_rollup_1_type=MAX
reevent_rollup_1_display_name=MAXSeverity
reevent_rollup_1_actionable=true
reevent_rollup_2_column_name=ACKNOWLEDGED
reevent_rollup_2_type=NON_ZERO
reevent_rollup_2_display_name=Acknowledged
reevent_rollup_2_actionable=true
reevent_rollup_3_column_name=ALERTGROUP
reevent_rollup_3_type=EXAMPLE
reevent_rollup_3_display_name=AlertGroup
reevent_rollup_3_actionable=false
#
# Group Information adds additional group information under the Show Details ->
Group More Information portion of the UI
# reevent_num_groupinfo Contains the number of group information columns to
display
# reevent_groupinfo_ <number where its 1 to n >_columnContains the column
name from which the data is retrieved
#  The following columns are allowed :
#      PROFILE, EVENTIDENTITIES, INSTANCES, CONFIGNAME, TOTALEVENTS,
UNIQUEEVENTS, REVIEWED, GROUPTTL
#      PROFILE: The relationship profile, or strength of the group.
#      EVENTIDENTITIES: A comma separated list that creates the event identity.
#      INSTANCES: The total number of group instances.
#      CONFIGNAME: The configuration name the group was created under.
#      TOTALEVENTS: The total number of events within the group.
#      UNIQUEEVENTS: The total number of unique events within the group.
#      REVIEWED: Whether the group has been reviewed by a user or not.
#      GROUPTTL: The number of seconds the group will stay active after the
first event occurs.
#
reevent_num_groupinfo=3
reevent_groupinfo_1_column=PROFILE
reevent_groupinfo_2_column=EVENTIDENTITIES
reevent_groupinfo_3_column=INSTANCES
#
# Event Information adds additional event information under the Show Details ->
Event More Information portion of the UI
# reevent_num_eventinfo Contains the number of event information columns to
display
# reevent_eventinfo_ <number where its 1 to n >_columnContains the column name
from which the data is retrieved
#  The following columns are allowed :
#      PROFILE, INSTANCES, EVENTIDENTITY, EVENTIDENTITIES, CONFIGNAME, and GROUPNAME
#      PROFILE: The relationship profile, or strength of the related event.
#      INSTANCES: Total number of instance for the related event.
#      EVENTIDENTITY: The unique event identity for the related event.
#      EVENTIDENTITIES: A comma separated list that creates the event identity.
#      CONFIGNAME: The configuration name the related event was created under.
#      GROUPNAME: The group name the related event is created under.
#
reevent_num_eventinfo=1
reevent_eventinfo_1_column=INSTANCES
#
##################################################################################################
# The following properties are used to configure event pattern creation         ##
# type.resourcelist=<columns include information. Comma separated list >        ##
# type.servername.column=<SERVERNAME column name if different than default>     ##
# type.serverserial.column=<SERVERSERIAL column name if different than default> ##
# type.default.eventid=<default event identities when there is no mactch        ##
#                       found in the types configuration. Comma separated list  ##
#                       The id should not include a timestamp component.        ##
# type.default.eventtype=<default event type when there is no match             ##
#                         found in the types configuration.                     ##
# type index starts with 0                                                      ##
# type_number_of_type_configurations=number of type  configurations             ##
# type.index.eventid=event identity column name                                 ##
# type.index.eventtype=event column includes the type to use                    ##
# type.index.filterclause=History DB filter to filter events to find the types  ##
# type.index.osfilterclause=ObjectServer filter tp filter matching events types ##
#                                                                               ##
#                                                                               ##
```

```
# NOTE : It is recommended to create database index(s) on the reporter status   ##
#         table for the fields used in the filtercaluse to speed the query(s).  ##
#         Example to create an index:                                           ##
#         create index types_index on db2inst1.reporter_stuats (Severity)       ##
#                                                                               ##
#                                                                               ##
#Use the following as an example creating one type only                         ##
#                                                                               ##
#type_number_of_type_configurations=1                                           ##
#type.0.eventid=NODE,SUMMARY,ALERTGROUP                                         ##
#type.0.eventtype=ACMEType                                                      ##
#type.0.filterclause=Vendor =( 'ACME' )                                         ##
#type.0.osfilterclause=Vendor = 'ACME'                                          ##
###############################################################################################
type.resourcelist=NODE
type.default.eventid=IDENTIFIER
type.default.eventtype=ALERTGROUP
type.servername.column=SERVERNAME
type.serverserial.column=SERVERSERIAL

type_number_of_type_configurations=1
type.0.eventid=SUMMARY
type.0.eventtype=ALERTGROUP
type.0.filterclause=( Severity >=3 )
type.0.osfilterclause=Severity >=3
#
###############################################################################################
# The following properties are used to configure Name Similarity NS             ##
###############################################################################################

name_similarity_feature_enable=true
name_similarity_default_pattern_enable=false
name_similarity_default_threshold=0.9
name_similarity_default_lead_restriction=1
name_similarity_default_tail_restriction=0
```

Configuring Event Analytics using the command line

You can configure Event Analytics from the command line using the `./nci_trigger` utility.

Before you begin

In Netcool/Impact v7.1.0.13 it is recommended to use the Event Analytics configuration wizard instead of the `./nci_trigger` command to edit properties in the NOI Shared Configuration properties file. For more information, see “Configuring Event Analytics using the wizard” on page 251.

Configuring the historical event database

Configure access to the Tivoli Netcool/OMNIbus historical event database that contains the data used to analyze historical events for Event Analytics.

Before you begin

In Netcool/Impact v7.1.0.13 it is recommended to use the Event Analytics configuration wizard instead of the `./nci_trigger` command to edit properties in the NOI Shared Configuration properties file. For more information, see “Configuring Event Analytics using the wizard” on page 251.

Configuring Db2 database connection within Netcool/Impact

You can configure a connection to a valid Db2 database from within IBM Tivoli Netcool/Impact.

Before you begin

In Netcool/Impact v7.1.0.13 it is recommended to use the Event Analytics configuration wizard instead of the `./nci_trigger` command to edit properties in the NOI Shared Configuration properties file. For more information, see “Configuring Event Analytics using the wizard” on page 251.

About this task

Users can run seasonality event reports and related event configurations, specifying the time range and name with Db2. Complete the following steps to configure the ObjectServer data source or data type.

Procedure

1. Log in to the Netcool/Impact UI.

   ```
   https://impacthost:port/ibm/console
   ```

2. Configure the ObjectServer data source and data type.
   a) In the Netcool/Impact UI, from the list of available projects, select the **NOI project**.
   b) Select the **Data Model** tab and select **ObjectServerForNOI**.
      1) Click **Edit** and enter information for *<username>*, *<password>*, *<host name>*, *<port>*.
      2) To save the Netcool/Impact data source, click **Test Connection**, followed by the **Save** icon.
   c) Edit the data type. Expand the data source and edit the data type to correspond to the ObjectServer event history database type.
   For example, `AlertsForNOITable`
   d) For `Base Table`, select *<database table>*.
   e) To update the schema and table, click **Refresh** and then click **Save**.
   f) Select the **Data Model** tab and select **ObjectServerHistoryDb2ForNOI**.
      1) Click **Edit** and enter information for *<username>*, *<password>*, *<host name>*, *<port>*.
      2) To save the Netcool/Impact data source, click **Test Connection**, followed by the **Save** icon.
   g) Edit the data type. Expand the **ObjectServerHistoryDb2ForNOI** data source and edit `AlertsHistoryDb2Table`.
   h) For `Base Table`, select *<database name>* and *<database table name>*
   i) To update the schema and table, click **Refresh** and then click **Save**.
   j) Select the **Services** tab and ensure that the following services are started.

      ```
      ProcessRelatedEvents
      ProcessSeasonalityEvents
      ProcessRelatedEventConfig
      ```

3. Configure the Db2 database connection within Netcool/Impact if it was previously configured for Oracle or MSSQL. The following steps configure the report generation to use the Db2 database. Export the default properties, change the default configuration, and update the properties.
   a) Generate a properties file, go to the *<Impact install location>*`/bin` directory to locate the `nci_trigger`, and run the following command from the command-line interface.

      ```
      nci_trigger <server> <username>/<password> NOI_DefaultValues_Export
       FILENAME directory/filename
      ```

      where

      <server>
      The server where Event Analytics is installed.

      <user name>
      The user name of the Event Analytics user.

      <password>
      The password of the Event Analytics user.

      directory
      The directory where the file is stored.

      filename
      The name of the properties file.

      For example:

      ```
      ./nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Export
       FILENAME /tmp/seasonality.props
      ```

b) Update the properties file. Some property values are overwritten by the generated properties file, you might need to update other property values in the generated properties file. For a full list of effected properties, see "Generated properties file" on page 256.

- If you do not have the following parameter values, update your properties file to reflect these parameter values.

```
history_datasource_name=ObjectServerHistoryDb2ForNOI
history_datatype_name=AlertsHistoryDb2Table
history_database_table=<database table name>
history_database_type=Db2
```

c) Import the modified properties file into Netcool/Impact, enter the following command.

```
nci_trigger <Server> <username>/<password> NOI_DefaultValues_Configure
 FILENAME directory/filename
```

For example:

```
./nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Configure
 FILENAME /tmp/seasonality.props
```

Related tasks

Installing Netcool/OMNIbus and Netcool/Impact

Related information

Netcool/Impact documentation: nci_triggerUse the nci_trigger tool to run a policy from the command line.

Configuring Oracle database connection within Netcool/Impact

You can configure a connection to a valid Oracle database from within IBM Tivoli Netcool/Impact.

Before you begin

In Netcool/Impact v7.1.0.13 it is recommended to use the Event Analytics configuration wizard instead of the `./nci_trigger` command to edit properties in the NOI Shared Configuration properties file. For more information, see "Configuring Event Analytics using the wizard" on page 251.

To use Oracle as the archive database, you must set up a remote connection to Netcool/Impact. For more information, see "Netcool/Impact remote connection" on page 245.

About this task

Users can run seasonality event reports and related event configurations, specifying the time range and name with Oracle. Complete the following steps to configure the ObjectServer data source or data type.

Procedure

1. Log in to the Netcool/Impact UI.

```
https://impacthost:port/ibm/console
```

2. Configure the ObjectServer data source and data type.
   a) In the Netcool/Impact UI, from the list of available projects, select the **NOI project**.
   b) Select the **Data Model** tab, and select **ObjectServerForNOI**.
      1) Click **Edit** and enter information for *<username>*, *<password>*, *<host name>*, and *<port>*.
      2) Save the Netcool/Impact data source. Click **Test Connection**, followed by the **Save** icon.
   c) Edit the data type. Expand the data source **ObjectServerForNOI** and edit the data type to correspond to the ObjectServer event history database type.
   For example, `AlertsForNOITable`.
   d) For `Base Table`, select *<database table>*.

e) To update the schema and table, click **Refresh** and then click **Save**.

f) Select the **Data Model** tab, and select **ObjectServerHistoryOrclForNOI**.

1) Click **Edit** and enter information for *<username>*, *<password>*, *<host name>*, *<port>*, and *<sid>*.

2) Save the Netcool/Impact data source. Click **Test Connection**, followed by the **Save** icon.

g) Edit the data type. Expand the data source **ObjectServerHistoryOrclForNOI** and edit the `AlertsHistoryOrclTable` data type.

h) For `Base Table`, select *<database name>* and *<database table name>*.

i) To update the schema and table, click **Refresh** and then click **Save**.

j) Edit the data type. Expand the data source **ObjectServerHistoryOrclForNOI** and edit the `SE_HISTORICALEVENTS_ORACLE` data type.

k) For `Base Table`, select *<database name>* and *<database table name>*.

l) To update the schema and table, click **Refresh** and then click **Save**.

m) Select the **Services** tab and ensure that following services are started:

```
ProcessRelatedEvents
ProcessSeasonalityEvents
ProcessRelatedEventConfig
```

3. Configure the report generation to use the Oracle database. Export the default properties, change the default configuration, and update the properties.

a) Generate a properties file. Go to the *<Impact install location>*/`bin` directory to locate the `nci_trigger` utility, and run the following command from the command-line interface:

```
nci_trigger <server> <username>/<password> NOI_DefaultValues_Export
 FILENAME directory/filename
```

Where

<server>
The server where Event Analytics is installed.

<user name>
The user name of the Event Analytics user.

<password>
The password of the Event Analytics user.

directory
The directory where the properties file is stored.

filename
The name of the properties file.

For example:

```
./nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Export
 FILENAME /tmp/seasonality.props
```

b) You need to modify the property values that are overwritten by the generated properties file. For a full list of properties, see "Generated properties file" on page 256.

- If you do not have the following values for these properties, update your properties file to reflect these property values:

```
history_datasource_name=ObjectServerHistoryOrclForNOI
history_datatype_name=AlertsHistoryOrclTable
history_database_table=<database table name>
history_database_type=Oracle
```

- Enter the following value, which is the Oracle database timestamp format from the policy, to the `history_db_timestampformat` property:

```
history_db_timestampformat=yyyy-mm-dd hh24:mi:ss
```

Note: The `history_db_timestampformat` property delivers with the properties file with a default value of `yyyy-MM-dd HH:mm:ss.SSS`. This default timestamp format for the `history_db_timestampformat` property does not work with Oracle. Thus, you need to perform the previous step to change the default value to the Oracle database timestamp format from the policy (`yyyy-mm-dd hh24:mi:ss`).

c) Import the modified properties file into Netcool/Impact using the following command:

```
nci_trigger <Server> <username>/<password> NOI_DefaultValues_Configure
 FILENAME directory/filename
```

For example:

```
./nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Configure
 FILENAME /tmp/seasonality.props
```

Related information

Netcool/Impact documentation: nci_triggerUse the nci_trigger tool to run a policy from the command line.

Configuring MS SQL database connection within Netcool/Impact

You can configure a connection to a valid MS SQL database from within IBM Tivoli Netcool/Impact.

Before you begin

In Netcool/Impact v7.1.0.13 it is recommended to use the Event Analytics configuration wizard instead of the `./nci_trigger` command to edit properties in the NOI Shared Configuration properties file. For more information, see “Configuring Event Analytics using the wizard” on page 251.

MS SQL support requires, at minimum, IBM Tivoli Netcool/Impact 7.1.0.1.

To use MS SQL as the archive database, you must set up a remote connection to Netcool/Impact. For more information, see “Netcool/Impact remote connection” on page 245.

About this task

Users can run seasonality event reports and related event configurations, specifying the time range and name with MS SQL. Complete the following steps to configure the ObjectServer data source and data type.

Procedure

1. Log in to the Netcool/Impact UI.

```
https://impacthost:port/ibm/console
```

2. Configure the ObjectServer data source and data type.
   a) In the Netcool/Impact UI, from the list of available projects, select the **NOI project**.
   b) Select the **Data Model** tab and select **ObjectServerForNOI**.
      1) Click **Edit** and enter the following information *<username>, <password>, <host name>, <port>*.
      2) Save the Netcool/Impact data source. Click **Test Connection**, followed by the **Save** icon.
   c) Edit the data type, expand the data source and edit the data type to correspond to the ObjectServer event history database type.
   For example, `AlertsForNOITable`
   d) For `Base Table`, select *<database table>*.

e) To update the schema and table, click **Refresh** and then click **Save**.

f) Select the **Data Model** tab and select **ObjectServerHistoryMSSQLForNOI**.

1) Click **Edit** and enter the following information *<username>*, *<password>*, *<host name>*, *<port>*, *<sid>*.

2) Save the Netcool/Impact data source. Click **Test Connection**, followed by the **Save** icon.

g) Edit the data type. Expand the data source **ObjectServerHistoryMSSQLForNOI** and edit `AlertsHistoryMSSQLTable`.

h) For `Base Table`, select *<database table name>*.

i) To update the schema and table, click **Refresh** and then click **Save**.

j) Select the **Services** tab and ensure that the following services are started.

```
ProcessRelatedEvents
ProcessSeasonalityEvents
ProcessRelatedEventConfig
```

3. Configure the report generation to use the MS SQL database.

a) Generate a properties file, go to the *<Impact install location>*`/bin` directory to locate the `nci_trigger` and in the command-line interface enter the following command.

```
nci_trigger <server> <username>/<password> NOI_DefaultValues_Export
 FILENAME directory/filename
```

<server>
The server where Event Analytics is installed.

<user name>
The user name of the Event Analytics user.

<password>
The password of the Event Analytics user.

directory
The directory where the file is stored.

filename
The name of the properties file.

For example, `./nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Export FILENAME /tmp/seasonality.props`.

b) Update the properties file. Some property values are overwritten by the generated properties file, you might need to update other property values in the generated properties file. For a full list of effected properties, see “Generated properties file” on page 256.

- If you do not have the following parameter values, update your properties file to reflect these parameter values.

```
history_datasource_name=ObjectServerHistoryMSSQLForNOI
history_datatype_name=AlertsHistoryMSSQLTable
history_database_table=<database table name>
history_database_type=MSSQL
```

c) Import the modified properties file into Netcool/Impact, enter the command.

```
nci_trigger <Server> <username>/<password> NOI_DefaultValues_Configure
FILENAME directory/filename
```

For example,

```
./nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Configure
FILENAME /tmp/seasonality.props
```

Related information

Netcool/Impact documentation: nci_triggerUse the nci_trigger tool to run a policy from the command line.

Adding columns to seasonal and related event reports

You can add columns to seasonal event reports and related events reports by updating the rollup configuration.

Before you begin

In Netcool Operations Insight v1.4.1.2 and later (corresponding to Netcool/Impact v7.1.0.13 and later) it is recommended to use the Event Analytics Configuration Wizard instead of the `./nci_trigger` command to edit properties in the NOI Shared Configuration properties file.

About this task

To update the rollup configuration, complete the following steps.

Procedure

1. Generate a properties file containing the latest Event Analytics system settings.
   a) Navigate to the directory `$IMPACT_HOME/bin`.
   b) Run the following command to generate a properties file containing the latest Event Analytics system settings.

```
nci_trigger server_name username/password NOI_DefaultValues_Export
 FILENAME directory/filename
```

   Where:

   - *server_name* is the name of the server where Event Analytics is installed.
   - *user name* is the user name of the Event Analytics user.
   - *password* is the password of the Event Analytics user.
   - `NOI_DefaultValues_Export` is a Netcool/Impact policy that performs an export of the current Event Analytics system settings to a designated properties file.
   - *directory* is the directory where the properties file is stored.
   - *filename* is the name of the properties file.

   For example:

```
nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Export
 FILENAME /tmp/properties.props
```

2. Update the properties file that you generated.
   a. Specify the number of columns you want to add to the reports:

      Increase the value of the `number_of_rollup_configuration=2` parameter for seasonal events.

      Increase the value of the `reevent_number_of_rollup_configuration=2` parameter for related events.

   For example, to add one column to the reports, increase the parameter value by one from 2 to 3.

   b. For a new rollup column, add property information.

   - For a new Seasonal Event reports column, add the following properties.

```
rollup_<rollup number>_column_name=<column name>
rollup_<rollup number>_display_name=<column_name>_<type>
rollup_<rollup number>_type=<type>
```

   - For a new Related Events reports column, add the following properties.

```
reevent_rollup_<rollup number>_column_name=<column name>
reevent_rollup_<rollup number>_display_name=<column_name>_<type>
reevent_rollup_<rollup number>_type=<type>
reevent_rollup_<rollup number>_actionable=<true/false>
```

<rollup number>
Specifies the new column rollup number.

<column name>
Specifies the new column name. The column name must match the column name in the history table.

<display name>
Specifies the new column display name. The display name must match the column name in the report.

<type>
Specifies one of the following types:

MAX
The maximum value observed for the column. If no value is observed, the value defaults to the minimum value of an integer.

MIN
The minimum value observed for the column. If no value is observed, the value defaults to the maximum value of an integer.

SUM
The sum of all of the values observed for the column.

NON_ZERO
A counting column that counts *nonzero* occurrences of events. This column can be useful to track the proportion of actioned events, or how many events had an associated ticket number.

DISTINCT
The number of distinct values that are seen for this key-value pair.

EXAMPLE
Displays the first non-blank *example* of a field that contained this key. The EXAMPLE type is useful when you are running seasonality on a field that can't be accessed, such as ALERT_IDENTIFIER, and you want an example human readable SUMMARY to demonstrate the type of problem.

Note: You cannot change the <type> property of a rollup column once the configuration has been updated. You must add a new rollup column and specify a different <type> (with a new <display name> if you are keeping the old rollup).

actionable=<true/false>
If this property is set to `true` for a rollup, the rollup is used to determine the probable root cause of a correlation rule. This root cause determination is based on the rollup that has the most actions that are taken against it. For example, if `Acknowledge` is part of your rollup configuration and has a property value of `actionable=true`, then the event with the highest occurrence of `Acknowledge` is determined to be the probable root cause. Probable root cause determination uses the descending order of the actionable rollups, that is, the first actionable rollup is a higher priority than the second actionable rollup. Only four of the possible *<type>* keywords are valid for root cause: `MAX, MIN, SUM, NON_ZERO`.

If this property is set to `false` for a rollup, the rollup is not used to determine the probable root cause of a rule. If all rollup configurations have a property value of `actionable=false`, the first event that is found is identified as the parent.

To manually change a root cause event for a correlation rule, see "Selecting a root cause event for a correlation rule" on page 327.

3. Import the modified properties file into Event Analytics.
   a) Ensure you are in the directory `$IMPACT_HOME/bin`.

b) Run the following command to perform an import of Event Analytics system settings from a designated properties file.

```
nci_trigger server_name username/password NOI_DefaultValues_Configure
 FILENAME directory/filename
```

Where:

- *server_name* is the name of the server where Event Analytics is installed.
- *user name* is the user name of the Event Analytics user.
- *password* is the password of the Event Analytics user.
- NOI_DefaultValues_Configure is a Netcool/Impact policy that performs an import of Event Analytics system settings from a designated properties file.
- *directory* is the directory where the properties file is stored.
- *filename* is the name of the properties file.

For example:

```
nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Configure
 FILENAME /tmp/properties.props
```

Results

The rollup configuration is updated.

Example

Example 1. To add a third column to the Seasonal Event report, change the rollup configuration value to 3, and add the properties.

```
number_of_rollup_configuration=3
rollup_1_column_name=SEVERITY
rollup_1_display_name=SEVERITY_MIN
rollup_1_type=MIN
rollup_2_column_name=SEVERITY
rollup_2_display_name=SEVERITY_MAX
rollup_2_type=MAX
rollup_3_column_name=TYPE
rollup_3_display_name=TYPE_MAX
rollup_3_type=MAX
```

Example 2. The configuration parameters for a default Related Events report.

```
reevent_rollup_1_column_name=ORIGINALSEVERITY
reevent_rollup_1_display_name=ORIGINALSEVERITY_MAX
reevent_rollup_1_type=MAX
reevent_rollup_1_actionable=true
reevent_rollup_2_column_name=ACKNOWLEDGED
reevent_rollup_2_display_name=ACKNOWLEDGED_NON_ZERO
reevent_rollup_2_type=NON_ZERO
reevent_rollup_2_actionable=true
reevent_rollup_3_column_name=ALERTGROUP
reevent_rollup_3_display_name=ALERTGROUP_EXAMPLE
reevent_rollup_3_type=EXAMPLE
reevent_rollup_3_actionable=false
```

What to do next

To add columns to the Seasonal Event reports, Historical Event portlet, Related Eventreports, or Related Event Details portlet complete the following steps:

1. Log in to the Tivoli Netcool/Impact UI.
2. Go to the **Policies** tab.
3. Open the policy that you want to modify. You can modify one policy at a time.
   - For Historical Events, open the **SE_GETHISTORICALEVENTS** policy.

- For Seasonal Events, open the **SE_GETEVENTDATA** policy.
- For related events groups, open one of the following policies.

  ```
  RE_GETGROUPS_ACTIVE
  RE_GETGROUPS_ARCHIVED
  RE_GETGROUPS_EXPIRED
  RE_GETGROUPS_NEW
  RE_GETGROUPS_WATCHED
  ```

  Note: Each policy is individually updated. To update two or more policies, you must modify each policy individually.
- For related events, open one of the following policies.

  ```
  RE_GETGROUPEVENTS_ACTIVE
  RE_GETGROUPEVENTS_ARCHIVED
  RE_GETGROUPEVENTS_EXPIRED
  RE_GETGROUPEVENTS_NEW
  RE_GETGROUPEVENTS_WATCHED
  ```

  Note: Each policy is individually updated. To update two or more policies, you must modify each policy individually.
- For related events details group instances table, open the following policy:

  ```
  RE_GETGROUPINSTANCEV1
  ```

4. Click the **Configure Policy Settings** icon.
5. Under **Policy Output Parameters**, click **Edit**.
6. To create a custom schema definition, open the **Schema Definition Editor** icon.
7. To create a new field, click **New**.
8. Specify the new field name and format.

   The new field name must match the display name in the configuration file.

   The format must match the format in the **AlertsHistory** Table.

   The format must be appropriate for the rollup type added. For example, for numerical types such as SUM or NON_ZERO use a numeric format. Use String for DISTINCT, if the base column is String.

   Refresh the **SE_GETHISTORICALEVENTS_Db2** table, or other database model, before you run Event Analytics with the added Historical Event table fields.

   For the `RE_GETGROUPS_` policies, only rollup columns with a *<type>* value of `MAX, MIN, SUM, NON_ZERO` are supported. Therefore, add only numeric fields to the schema.
9. To complete the procedure, click **Ok** on each of the open dialog boxes, and **Save** on the **Policies** tab.

   Note: Columns that are created for the Related Event Details before Netcool Operations Insight release 1.4.0.1 are displayed as designed. Configurations and groups that are created after you upgrade to Netcool Operations Insight release 1.4.0.1, display the events from the historical event. By adding columns to the Related Event Details, you can display additional information such as the **`owner ID`** or **`ticketnumber`**.

You can also add the following columns for group instances in the Related Event Details:

- `SERVERSERIAL`
- `SERVERNAME`
- `TALLY`
- `OWNERUID`

By default, the previously listed columns are hidden for group instances in the Related Event Details. To display these columns in the Related Event Details, you need to edit the `Policy_RE_GETGROUPINSTANCEV1_RE_GETGROUPINSTANCEV1.properties` file, which is located in the following directory: `$IMPACT_HOME/uiproviderconfig/properties`.

Specifically, set the following properties in the `Policy_RE_GETGROUPINSTANCEV1_RE_GETGROUPINSTANCEV1.properties` file from their default values of `true` to the values `false` (or comment out the field or fields):

```
SERVERSERIAL.hidden=true
SERVERNAME.hidden=true
TALLY.hidden=true
OWNERUID.hidden=true
```

For example,

```
OWNERUID.hidden=false
```

Or, for example,

```
#OWNERUID.hidden=true
```

Configuring event suppression
Event suppression is available as an action when creating a seasonal event rule. You can configure event suppression by modifying the `NOI_DefaultValues` properties file.

Before you begin

In Netcool Operations Insight v1.4.1.2 and later (corresponding to Netcool/Impact v7.1.0.13 and later) it is recommended to use the Event Analytics Configuration Wizard instead of the `./nci_trigger` command to edit properties in the NOI Shared Configuration properties file.

For more information on the relevant section of the wizard, see "Configuring event suppression" on page 254.

About this task
To add details about suppressing and unsuppressing events, you must modify the `NOI_DefaultValues` properties file in the `$IMPACT_HOME/bin` directory.

Procedure

1. Log in to the server where IBM Tivoli Netcool/Impact is stored and running.
2. Generate a properties file containing the latest Event Analytics system settings.
   a) Navigate to the directory `$IMPACT_HOME/bin`.
   b) Run the following command to generate a properties file containing the latest Event Analytics system settings.

      ```
      nci_trigger server_name username/password NOI_DefaultValues_Export
       FILENAME directory/filename
      ```

      Where:

      - *`server_name`* is the name of the server where Event Analytics is installed.
      - *`user name`* is the user name of the Event Analytics user.
      - *`password`* is the password of the Event Analytics user.
      - `NOI_DefaultValues_Export` is a Netcool/Impact policy that performs an export of the current Event Analytics system settings to a designated properties file.
      - *`directory`* is the directory where the properties file is stored.
      - *`filename`* is the name of the properties file.

      For example:

      ```
      nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Export
       FILENAME /tmp/properties.props
      ```

3. Add the following lines of text to the properties file:

```
seasonality.suppressevent.column.name=SuppressEscl
seasonality.suppressevent.column.type=NUMERIC
seasonality.suppressevent.column.value=4
seasonality.unsuppressevent.column.name=SuppressEscl
seasonality.unsuppressevent.column.type=NUMERIC
seasonality.unsuppressevent.column.value=0
```

4. Import the modified properties file into Event Analytics.
   a) Ensure you are in the directory `$IMPACT_HOME/bin`.
   b) Run the following command to perform an import of Event Analytics system settings from a designated properties file.

```
nci_trigger server_name username/password NOI_DefaultValues_Configure
 FILENAME directory/filename
```

   Where:

   - *server_name* is the name of the server where Event Analytics is installed.
   - *user name* is the user name of the Event Analytics user.
   - *password* is the password of the Event Analytics user.
   - `NOI_DefaultValues_Configure` is a Netcool/Impact policy that performs an import of Event Analytics system settings from a designated properties file.
   - *directory* is the directory where the properties file is stored.
   - *filename* is the name of the properties file.

   For example:

```
nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Configure
 FILENAME /tmp/properties.props
```

Configuring event pattern processing
You can configure how patterns are derived from related events by editing properties in the generated NOI Shared Configuration properties file.

Before you begin

In Netcool Operations Insight v1.4.1.2 and later (corresponding to Netcool/Impact v7.1.0.13 and later) it is recommended to use the Event Analytics Configuration Wizard instead of the `./nci_trigger` command to edit properties in the NOI Shared Configuration properties file.

For more information on the relevant section of the wizard, see “Configuring event pattern processing” on page 254.

Note:

- You should perform this configuration task prior to running any related events configurations that use the global type properties associated with event pattern creation. It is expected that you will perform this configuration task only when something in your environment changes that affects where type information is found in events.
- Avoid configuring multiple types for the same event. By default, `Identifier` is used to identify the same events. This can be overridden, but assuming the default, you should setup the type properties so that events identified by the same `Identifier` only have one type value. For example, if there are 10 events with `Identifier=xxx` and you want to use a `type=ALERTGROUP` then the events should have the same ALERTGROUP. If events for the same `Identifier` have many alert group values, the first one will be picked.

The default NOI Shared Configuration properties file is divided into sections, where each section contains a number of properties that allow you to instruct how Netcool/Impact handles a variety of operations,

such as how it should handle event pattern creation. There are three categories of event pattern creation properties defined in the NOI Shared Configuration properties file:

- Properties related to configuring which table columns in the Historical Event Database that Netcool/Impact should use in performing the event pattern analysis.
- Properties related to configuring the default unique event identifier and event type in the Historical Event Database that you want Netcool/Impact to use when there is no match in the event type index related properties.
- Properties related to configuring one or more event identity and event type indexes.

Table 43 on page 271 describes the event pattern creation properties defined in the NOI Shared Configuration properties file. Use these descriptions to help you configure the values appropriate for your environment.

Table 43. Event pattern creation properties

Global type	Description	Example
Properties related to configuring table columns in the Historical Event Database		
`type.resourcelist`	Specifies the name of the table column or columns in the Historical Event Database that Netcool/Impact should use in performing the event pattern analysis.	The NOI Shared Configuration properties file that you generate with the `nci_trigger` command provides the following default value: `type.resourcelist=NODE` **Note:** You should use the default value, `NODE`.
`type.servername.column`	Specifies the name of the table column in the Historical Event Database that contains the name of the server associated with any particular event that arrives in the Historical Event Database.	The NOI Shared Configuration properties file that you generate with the `nci_trigger` command provides the following default value: `type.servername.column= SERVERNAME` **Note:** You should use the default value, `SERVERNAME`, where possible.
`type.serverserial.column`	Specifies the name of the table column in the Historical Event Database that contains the server serial number associated with any particular event that arrives in the Historical Event Database. Note that the server serial number should be unique.	The NOI Shared Configuration properties file that you generate with the `nci_trigger` command provides the following default value: `type.serverserial.column= SERVERSERIAL` **Note:** You should use the default value, `SERVERSERIAL`, where possible.
Properties related to configuring the default unique event identifier and event type in the Historical Event Database		

Table 43. Event pattern creation properties (continued)

Global type	Description	Example
`type.default.eventid`	This property contains the database field in the Historical Event Database that you want to specify as the default Event Identity. An Event Identity is a database field that identifies a unique event in the Historical Event Database. When you configure a related events configuration, you select database fields for the Event Identity from a drop-down list of available fields. In the User Interface, you perform this from the **Advanced** tab when you want to override the settings in the configuration file. Netcool/Impact uses the database field specified in this property as the default Event Identity when there is no match in the value specified in the `type.`*`index`*`.eventid` property. **Note:** The database field specified for this property should not contain a timestamp component.	The NOI Shared Configuration properties file that you generate with the `nci_trigger` command provides the following default value: `type.default.eventid= IDENTIFIER`

Table 43. Event pattern creation properties (continued)		
Global type	**Description**	**Example**
`type.default.eventtype`	Specifies the default related events type to use when creating an event pattern to generalize. Netcool/Impact uses this default related events type when there is no match in the`type.`*`index`*`.eventtype` property. **Note:** You choose the related events type values based on the fields for which you want to create a generalized pattern. For example, if you want to create a pattern and generalize it based on the EVENTID for an event, you would specify that value in this property. When the related events configuration completes and you create a pattern for generalization, the pattern generalization screen will contain a drop down menu that lists all of the EVENTIDs found in the Historical Event Database. You can then create a pattern/rule that will be applied to all EVENTIDs selected for that pattern. This means that you can expand the definition of the pattern to include all types, not just the types in the Related Events Group.	The NOI Shared Configuration properties file that you generate with the `nci_trigger` command provides the following default value: `type.default.eventtype= EVENTID`
Properties related to configuring one or more event identity and event type indexes. You should specify values for each of the properties described in this section.		

Table 43. Event pattern creation properties (continued)

Global type	Description	Example

`type_number_of_ type_configurations`	Specifies the number of types to use in the NOI Shared Configuration properties file for the global type configuration. There is no limit on how many types you can configure.	The following example specifies two types for the global type configuration:
		`type_number_of_ type_configurations=2`
		Thus, you would define the other `type.`*`index`* related properties as follows. Note that the index numbering starts with 0 (zero).
		`type.0.eventid=Identifier`
		`type.0.eventtype=ACMEType`
		`type.0.filterclause= Vendor='ACME'`
		`type.0.osfilterclause= Vendor='ACME'`
		`type.1.eventid=SUMMARY, NODE`
		`type.1.eventtype= TAURUSType`
		`type.1.filterclause= Vendor = 'TAURUS'`
		`type.1.osfilterclause= Vendor = 'TAURUS'`

`type.`*`index`*`.eventid`	Specifies the database field in the Historical Event Database that you want to specify as the Event Identity. Multiple fields are separated by commas.	The following shows an example of a database field used as the Event Identity:
		`type.0.eventid=SUMMARY`
		The following shows an example of multiple database fields used as the Event Identity:
		`type.0.eventid=NODE, SUMMARY, ALERTGROUP`

`type.`*`index`*`.eventtype`	Specifies the event type to return for pattern generalization. **Note:** The returned event types display in the event type drop down menu in the pattern generalization screen.	The following example shows an event type to return for pattern generalization: `type.0.eventtype=EVENTID`

Table 43. Event pattern creation properties (continued)

Global type	Description	Example
type.*index*.filterclause	Specifies an Historical Event Database filter that defines a set of events. For the set of events defined by this filter, the event type will be found in the table column or columns in the type.*index*.eventtype property. **Note:** It is recommended that you create one or more database indexes on the reporter status table for the fields used in the type.*index*.filterclause to speed up the query.	`type.0.filterclause= Vendor = 'ACME'`
type.*index*.osfilterclause	Specifies an ObjectServer filter to filter matching event types. **Note:** The filter that you specify for the type.*index*.osfilterclause property should be semantically identical to the filter that you specify for the type.*index*.filterclause property, except for this property you use the ObjectServer syntax.	`type.0.osfilterclause= Vendor = 'ACME'`

About this task

To configure the event pattern creation properties that Netcool/Impact uses for generalization, you must modify the default NOI Shared Configuration properties file in the <*Impact_install_location*>/bin directory.

Procedure

1. Log in to the server where IBM Tivoli Netcool/Impact is stored and running.
2. Generate a properties file containing the latest Event Analytics system settings.
   a) Navigate to the directory $IMPACT_HOME/bin.
   b) Run the following command to generate a properties file containing the latest Event Analytics system settings.

```
nci_trigger server_name username/password NOI_DefaultValues_Export
 FILENAME directory/filename
```

   Where:

   - *server_name* is the name of the server where Event Analytics is installed.
   - *user name* is the user name of the Event Analytics user.
   - *password* is the password of the Event Analytics user.
   - NOI_DefaultValues_Export is a Netcool/Impact policy that performs an export of the current Event Analytics system settings to a designated properties file.
   - *directory* is the directory where the properties file is stored.

- *filename* is the name of the properties file.

For example:

```
nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Export
 FILENAME /tmp/properties.props
```

3. Go to the directory where you generated the NOI Shared Configuration properties file and open it for editing.
4. Create a backup copy of the generated NOI Shared Configuration properties file.
5. Using the editor of your choice open the generated NOI Shared Configuration properties file for editing.
6. Using the information about the event pattern creation properties described in Table 43 on page 271, specify values appropriate to your environment. Remember that the following properties have default values that you should not change:
   - `type.resourcelist`
   - `type.servername.column`
   - `type.serverserial.column`
7. After specifying appropriate values to the event pattern creation properties, write and then quit the NOI Shared Configuration properties file.
8. Import the modified properties file into Event Analytics.
   a) Ensure you are in the directory `$IMPACT_HOME/bin`.
   b) Run the following command to perform an import of Event Analytics system settings from a designated properties file.

```
nci_trigger server_name username/password NOI_DefaultValues_Configure
 FILENAME directory/filename
```

Where:

- *server_name* is the name of the server where Event Analytics is installed.
- *user name* is the user name of the Event Analytics user.
- *password* is the password of the Event Analytics user.
- `NOI_DefaultValues_Configure` is a Netcool/Impact policy that performs an import of Event Analytics system settings from a designated properties file.
- *directory* is the directory where the properties file is stored.
- *filename* is the name of the properties file.

For example:

```
nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Configure
 FILENAME /tmp/properties.props
```

Example

The following example sets the Event Identity, defines a set of events, and finds the type information in the specified table column or columns in the Historical Event Database:

```
type_number_of_type_configurations=1
type.0.eventid=NODE,SUMMARY,ALERTGROUP
type.0.eventtype=ACMEType
type.0.filterclause=( Vendor = 'ACME' )
type.0.osfilterclause=Vendor = 'ACME'
```

More specifically, the examples shows that if there is an event and the value for `Vendor` for that event is `ACME`, then look in the table column called `ACMEType` to find the event type.

The following example expands on the previous example by showing two configurations (as indicated by the value 2 in the `type_number_of_type_configurations` property:

```
type_number_of_type_configurations=2
type.0.eventid=NODE
type.0.eventtype=ACMEType
type.0.filterclause=( Vendor = 'ACME' )
type.0.osfilterclause=Vendor = 'ACME'
type.1.eventid=NODE,SUMMARY,ALERTGROUP
type.1.eventtype=TAURUSType
type.1.filterclause=( Vendor = 'TAURUS' )
type.1.osfilterclause=Vendor = 'TAURUS'
```

Note: Netcool/Impact attempts to match each event to the filter defined in configuration `0` first. If the event matches the filter defined in configuration `0`, then Netcool/Impact defines the event's type as defined in the filter. If the event does not match the filter defined in configuration `0`, Netcool/Impact attempts to match the event to the filter defined in configuration `1`. If the event matches the filter defined in configuration `1`, then Netcool/Impact defines the event's type as defined in the filter. Netcool/Impact continues this processing sequence for as many configuration types you define.

If no events match the filters defined in the defined configuration types you define, Netcool/Impact uses the default configuration to determine where type and identity are to be found.

Other configuration tasks
Perform these configuration tasks to further configure your system.

Configuring the ObjectServer
Prior to deploying rules based on related event events or patterns you must run SQL to update the ObjectServer. This SQL introduces relevant triggers into the ObjectServer to enable to rules to be fully functional.

About this task
The SQL provides commands for creating and modifying ObjectServer objects and data. Complete the following steps to run the SQL to update the ObjectServer.

Procedure

1. Copy the SQL file `IMPACT_HOME/add-ons/RelatedEvents/db/relatedevents_objectserver.sql` from Netcool/Impact into the `tmp` directory on your ObjectServer.
2. Run the SQL against your ObjectServer, enter the following command.

   On Windows, enter the command `%OMNIHOME%\..\bin\redist\isql -U <username> -P <password> -S <server_name> < C:\tmp\relatedevents_objectserver.sql`
   On Linux and UNIX, enter the command `$OMNIHOME/bin/nco_sql -user <username> -password <password> -server <server_name> < /tmp/relatedevents_objectserver.sql`
3. If you have not previously configured the Event Analytics ObjectServer, you must enter the following command.

   On Windows, enter the command `%OMNIHOME%\..\bin\redist\isql -U <username> -P <password> -S <server_name> < C:\tmp\relatedevents_objectserver.sql`
   On Linux and UNIX, enter the command `$OMNIHOME/bin/nco_sql -user <username> -password <password> -server <server_name> < /tmp/relatedevents_objectserver.sql`
4. All users must run the SQL against your ObjectServer, enter the following command.

   On Windows, enter the command `%OMNIHOME%\..\bin\redist\isql -U <username> -P <password> -S <server_name> < C:\tmp\relatedevents_objectserver_update_fp5.sql`
   On Linux and UNIX, enter the command `$OMNIHOME/bin/nco_sql -user <username> -password <password> -server <server_name> < /tmp/relatedevents_objectserver_update_fp5.sql`

What to do next
Event correlation for the related events function in Event Analytics, uses a `ParentIdentifier` column that is added the ObjectServer. If the size of this identifier field changes in your installation, you must change the value of the `ParentIdentifier` column within the ObjectServer SQL file that creates the event grouping automation `relatedevents_objectserver.sql`, to ensure that both values are the same. The updated SQL is automatically picked up.
Related tasks
Installing Netcool/OMNIbus and Netcool/Impact

Additional configuration for Netcool/Impact server failover for Event Analytics
The following additional configuration is required if a seasonality report is running during a Netcool/Impact node failover. Without this configuration the seasonality report might hang in the processing state following the failover. This will give the impression that the report is running, however it will remain stuck in the phase and percentage complete level that is displayed following the failover. Any queued reports will also not run. This is due to a limitation of the derby database. Use the workaround in this section to avoid this problem.

Procedure

1. Locate the **jvm.options** file in `<impact_home>/wlp/usr/servers/<Impact_server _name>/`.
2. Uncomment the following line in all the nodes of the failover cluster:

```
#-Xgc:classUnloadingKickoffThreshold=100
```

3. Restart all nodes in the failover Netcool/Impact cluster.

Results
Following these changes, any currently running seasonality report will terminate correctly during the cluster failover and any queued reports will continue running after the failover has completed.

Configuring extra failover capabilities in the ObjectServer
Related events use standard ObjectServer components to provide a high availability solution. These ObjectServer components require extra configuration to ensure high availability where there is an ObjectServer pair and the primary ObjectServer goes down before the cache on the Netcool/Impact node refreshes.

In this scenario, if you deploy a correlation rule, the rule is picked up if you have replication setup between the ObjectServer tables. Otherwise, the new rule is not picked up and this state continues until you deploy another new rule. Complete the following steps to setup replication between the ObjectServer tables.

- In the `.GATE.map` file, add the following lines.

```
CREATE MAPPING RE_CACHEMAP
(
'name'   = '@name' ON INSERT ONLY,
'updates' = '@updates'
);
```

- If your configuration does not use the standard `StatusMap` file, add the following line to the `StatusMap` file that you use to control `alerts.status`, you can find the `StatusMap` file in the `.tblrep.def` file.

```
   'ParentIdentifier' =    '@ParentIdentifier'
```

- In the `.tblrep.def` file, add the following lines.

```
REPLICATE ALL FROM TABLE 'relatedevents.cacheupdates'
        USING map 'RE_CACHEMAP';
```

For more information about adding collection ObjectServers and displaying ObjectServers to your environment, see the following topics within IBM Knowledge Center for IBM Tivoli Netcool/OMNIbus, Netcool/OMNIbus v8.1.0 Welcome page.

Setting up the standard multitiered environment
Configuring the bidirectional aggregation ObjectServer Gateway
Configuring the unidirectional primary collection ObjectServer Gateway

Configuring extra failover capabilities in the Netcool/Impact environment
Configure extra failover capabilities in the Netcool/Impact environment by adding a cluster to the Netcool/Impact environment. To do this you must update the data sources in IBM Tivoli Netcool/Impact.

About this task
Update the following data sources when you add a cluster to the Netcool/Impact environment.

`seasonalReportDataSource`
`RelatedEventsDatasource`
`NOIReportDatasource`

Complete the following steps to update the data sources.

Procedure

1. In Netcool/Impact, go to the **Database Failure Policy**.
2. Select *Fail over* or *Fail back* depending on the high availability type you want. For more information, see the failover and failback descriptions.
3. Go to **Backup Source**.
4. Enter the secondary Impact Server's Derby *Host Name*, *Port*, and *Database* information.

   Standard failover
   Standard failover is a configuration in which an SQL database DSA switches to a secondary database server when the primary server becomes unavailable and then continues by using the secondary until Netcool/Impact is restarted.

   Failback
   Failback is a configuration in which an SQL database DSA switches to a secondary database server when the primary server becomes unavailable and then tries to reconnect to the primary at intervals to determine whether it returned to availability.

What to do next

If you encounter error ATKRST132E, see details in “Troubleshooting Event Analytics (on premises)” on page 521

If you want your Netcool/Impact cluster that is processing events to contain the same cache and update the cache, at or around the same time, you must run the file `relatedevents_objectserver.sql` with `nco_sql`. The `relatedevents_objectserver.sql` file contains the following commands.

```
create database relatedevents;
create table relatedevents.cacheupdates persistent (name varchar (20) primary key,
updates integer);
insert into relatedevents.cacheupdates (name, updates) values ('RE_CACHE', 0);
```

Mapping customized field names
Any customized field names in the Historical Event Database must be mapped to the corresponding standard field names in Netcool/Impact. You do this by creating a database view to map to the standard field names.

Before you begin
You must perform this task if you have customized table columns in your Historical Event Database. For example, if you have defined columns called SUMMARYTXT and `IDENTIFIERID` instead of the default

names SUMMARY and IDENTIFIER, you must perform this task. You create a database view and map back to the actual field names.

About this task

The steps documented here are for a Db2 database. The procedure is similar for an Oracle database.

To map customized columns in your Historical Event Database, complete the following steps.

Procedure

1. Use the following statement to create the view and point the data types to the new view.

```
DROP VIEW REPORTER_STATUS_STD;
CREATE VIEW REPORTER_STATUS_STD AS SELECT SUMMARYTXT AS SUMMARY,
IDENTIFIERID AS IDENTIFIER, * FROM REPORTER_STATUS;
```

2. Change the data types from REPORTER_STATUS to REPORTER_STATUS_STD. The data types for Db2 are AlertsHistoryDb2Table and SE_HISTORICALEVENTS_Db2 under **ObjectServerHistoryDb2ForNOI** data source.
3. Delete RELATEDEVENTS.RE_MAPPINGS records from the table:

```
DELETE FROM RELATEDEVENTS.RE_MAPPINGS WHERE TRUE;
```

4. Run the Event Analytics Configuration wizard to configure the Netcool/Impact properties to use for Event Analytics.
5. On the **Historical event database** configuration screen, connect to the database and then select the **REPORTER_STATUS_STD** (view) from the **History table** drop-down menu as the **Table name** for Event Analytics.
6. When using any other columns that were mapped in the view, for example Summary for SUMMARYTXT, use the new value in any of the wizard screens. In this case use Summary.
   For example, when adding fields to the report in the **Configure report fields** screen, use the values mapped in the view (Identifier or Summary).
7. Save the Event Analytics configuration. You can now use the mapped fields for Event Analytics.

Customizing tables in the Event Analytics UI

You can use the uiproviderconfig files to customize the tables in the Event Analytics UI.

To customize how tables are displayed in the Event Analytics UI, you can update the properties and translation files that are specific to the policy or data type that you want to update. These files are stored in the *properties* and *translation* directories in the $IMPACT_HOME/uiproviderconfig/ directory.

If you want to update the properties and translation files, make a backup of the files before you do any updates.

Configuring columns to display in the More Information panel

You can configure the columns that you want to display in the **More Information** panel.

About this task

The **More Information** panel can be started from within the Related Event Details portlet, when you click the hyperlink for either the Group Name or the Pivot Event, and the panel provides more details about the Group Name or the Pivot Event. The Event Analytics installation installs a default configuration of columns that display in the **More Information** panel, but you can change the configuration of columns that display. Complete the following steps to configure columns to display in the **More Information** panel.

Procedure

1. Generate a properties file containing the latest Event Analytics system settings.
   a) Navigate to the directory $IMPACT_HOME/bin.
   b) Run the following command to generate a properties file containing the latest Event Analytics system settings.

```
nci_trigger server_name username/password NOI_DefaultValues_Export
 FILENAME directory/filename
```

Where:

- *server_name* is the name of the server where Event Analytics is installed.
- *user name* is the user name of the Event Analytics user.
- *password* is the password of the Event Analytics user.
- NOI_DefaultValues_Export is a Netcool/Impact policy that performs an export of the current Event Analytics system settings to a designated properties file.
- *directory* is the directory where the properties file is stored.
- *filename* is the name of the properties file.

For example:

```
nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Export
 FILENAME /tmp/properties.props
```

2. Update the properties file with properties for columns you want to display in the **More Information** panel.
   - For columns related to the Group Name in the **More Information** panel, the following properties are the default properties in the properties file. You can add, remove, and change the default properties.

```
reevent_num_groupinfo=3
reevent_groupinfo_1_column=PROFILE
reevent_groupinfo_2_column=EVENTIDENTITIES
reevent_groupinfo_3_column=INSTANCES
```

reevent_num_groupinfo=3
This property represents the number of group information columns to display. The default value is 3 columns. The value can be any number between 1 and 8, as eight columns are allowed.

reevent_groupinfo_1_column=PROFILE
Enter this property line item for each column. The variables in this property line item are *1* and *PROFILE*.

1 denotes that this column is your first column. This value can increment up to 8 per property line item, as eight columns are allowed.

PROFILE represents the column. The following eight columns are allowed.

PROFILE
Specifies the relationship profile, or strength of the group.

EVENTIDENTITIES
Specifies a comma-separated list that creates the event identity.

INSTANCES
Specifies the total number of group instances.

CONFIGNAME
Specifies the configuration name under which the group was created.

TOTALEVENTS
Specifies the total number of events within the group.

UNIQUEEVENTS
Specifies the total number of unique events within the group.

REVIEWED
Specifies the review status of a group by a user.

GROUPTTL
Specifies the number of seconds the group will stay active after the first event occurs.

- For columns related to the `Pivot Event` in the **More Information** panel, the following properties are the default properties in the properties file. You can add, remove, and change the default properties.

```
reevent_num_eventinfo=1
reevent_eventinfo_1_column=INSTANCES
```

reevent_num_eventinfo=1
This property represents the number of group information columns to display. The default value is 1 column. The value can be any number between 1 and 6, as six columns are allowed.

reevent_eventinfo_1_column=INSTANCES
Enter this property line item for each column. The variables in this property line item are *1* and *INSTANCES*.

1 denotes that this column is your first column. This value can increment up to 6 per property line item, as six columns are allowed.

`INSTANCES` represents the column. The following six columns are allowed:

INSTANCES
Specifies the total number of instances for the related event.

PROFILE
Specifies the relationship profile, or strength of the related event.

EVENTIDENTITY
Specifies the unique event identity for the related event.

EVENTIDENTITIES
Specifies a comma-separated list that creates the event identity.

CONFIGNAME
Specifies the configuration name under which the related event was created.

GROUPNAME
Specifies the group name under which the related event was created.

3. Import the modified properties file into Event Analytics.
   a) Ensure you are in the directory `$IMPACT_HOME/bin`.
   b) Run the following command to perform an import of Event Analytics system settings from a designated properties file.

```
nci_trigger server_name username/password NOI_DefaultValues_Configure
 FILENAME directory/filename
```

Where:

- *server_name* is the name of the server where Event Analytics is installed.
- *user name* is the user name of the Event Analytics user.
- *password* is the password of the Event Analytics user.
- `NOI_DefaultValues_Configure` is a Netcool/Impact policy that performs an import of Event Analytics system settings from a designated properties file.
- *directory* is the directory where the properties file is stored.
- *filename* is the name of the properties file.

For example:

```
nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Configure
 FILENAME /tmp/properties.props
```

Configuring name similarity

Configure the name similarity feature by tuning parameters that govern how the system processes multiple resources when performing pattern matching.

About this task

Pattern matching enables Event Analytics to identify types of events that tend to occur together on a specific network resource. The name similarity feature extends pattern matching by enabling it to identify types of events that tend to occur together on more than one resource, where the resources within the pattern have a similar name. For examples of similar resource names that might be discovered by the name similarity feature, see "Examples of name similarity" on page 333.

Depending on how name similarity is configured, pattern matching will see these resource names as similar and will create a single pattern including events from all of these resource names.

Similarity threshold value: Algorithms are used to determine name similarity. First, an edit distance is calculated by a third-party algorithm. The edit distance is the minimum number of operations needed to transform one string into the other, where an operation is defined as an insertion, deletion, or substitution of a single character, or a transposition of two adjacent characters. Then, the algorithm calculates a normalized similarity distance, which lies in the range 0.0 to 1.0. In this range, 0.0 means that the strings are identical and 1.0 means that the strings are completely different. The normalized similarity distance is calculated by using a contribution of the edit distance weighted according to the first string length, the second string length, and the number of transpositions. Finally, the name similarity algorithm calculates a normalized threshold value (in the range 0.0 to 1.0) by subtracting the normalized similarity distance from the value 1.0. A threshold value of 0.0 means strings can be completely different. A threshold value of 1.0 means that strings must match exactly.

By default name similarity is configured with values that should enable it to work effectively in most environments. Use this procedure to change these settings.

Note: Only change name similarity settings if you understand the underlying algorithm.

Procedure

1. Generate a properties file containing the latest Event Analytics system settings.
   a) Navigate to the directory `$IMPACT_HOME/bin`.
   b) Run the following command to generate a properties file containing the latest Event Analytics system settings.

```
nci_trigger server_name username/password NOI_DefaultValues_Export
 FILENAME directory/filename
```

   Where:

   - *server_name* is the name of the server where Event Analytics is installed.
   - *user name* is the user name of the Event Analytics user.
   - *password* is the password of the Event Analytics user.
   - `NOI_DefaultValues_Export` is a Netcool/Impact policy that performs an export of the current Event Analytics system settings to a designated properties file.
   - *directory* is the directory where the properties file is stored.
   - *filename* is the name of the properties file.

   For example:

```
nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Export
 FILENAME /tmp/properties.props
```

2. Edit the properties file that you generated in the previous step.

   For example:

```
vi /tmp/properties.props
```

3. Find the section of the properties file that reads as follows:
   This code snippet shows the default values of the name similarity parameters.

```
##############################################################################################
# The following properties are used to configure Name Similarity NS                          ##
##############################################################################################

name_similarity_feature_enable=true
name_similarity_default_pattern_enable=false
name_similarity_default_threshold=0.9
name_similarity_default_lead_restriction=1
name_similarity_default_tail_restriction=0
```

4. Update one or more of the name similarity settings.
   The following table describes each of these settings.

Table 44. Name similarity settings

Parameter	Description	Values
`name_similarity_feature_enable`	Boolean that switches the name similarity feature on or off. **Note:** This is a global flag that governs all name similarity functionality. For example, if you set this flag to `false`, then no aspect of name similarity will be enabled, and none of the other flags in this table will have any effect.	Possible values: • `true`: Name similarity is switched on. • `false`: Name similarity is switched off. Default value: `true`
`name_similarity_default_pattern_enable`	Boolean that specifies whether to apply name similarity processing to historical patterns, meaning patterns that were created before name similarity was introduced into the Netcool Operations Insight solution. Name similarity was introduced into Netcool Operations Insight in V1.5.0, which corresponds to Netcool/Impact fix pack 14.	Possible values: • `true`: Apply name similarity processing to historical patterns. • `false`: Do not apply name similarity processing to historical patterns. Default value: `false`
`name_similarity_default_threshold`	String comparison threshold value, where `0` equates to completely dissimilar strings, and `1` equates to identical strings. The value specified in the `name_similarity_default_threshold` parameter, is used to determine whether two strings are similar. **Note:** The string similarity test is also governed by the *lead* and *tail* restriction parameters described below.	Possible values: `0` to `1` inclusive Default value: `0.9`

Table 44. Name similarity settings (continued)

Parameter	Description	Values
`name_similarity_default_lead_restriction`	Number of characters at the beginning of the strings being compared that must be identical. **Important:** If this number of characters is not identical then the strings automatically fail the similarity test.	Default value: `1` **Note:** This default setting assumes that the front end of the strings being compared is usually different.
`name_similarity_default_tail_restriction`	Number of characters at the end of the strings being compared that must be identical. **Important:** If this number of characters is not identical then the strings automatically fail the similarity test.	Default value: `0` **Note:** This default setting assumes that the tail end of the strings being compared is usually the same; for example `".com"`.

5. Import the modified properties file into Event Analytics.
   a) Ensure you are in the directory `$IMPACT_HOME/bin`.
   b) Run the following command to perform an import of Event Analytics system settings from a designated properties file.

```
nci_trigger server_name username/password NOI_DefaultValues_Configure
 FILENAME directory/filename
```

   Where:

   - *server_name* is the name of the server where Event Analytics is installed.
   - *user name* is the user name of the Event Analytics user.
   - *password* is the password of the Event Analytics user.
   - `NOI_DefaultValues_Configure` is a Netcool/Impact policy that performs an import of Event Analytics system settings from a designated properties file.
   - *directory* is the directory where the properties file is stored.
   - *filename* is the name of the properties file.

   For example:

```
nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Configure
 FILENAME /tmp/properties.props
```

Related information

Netcool/Impact documentation: nci_triggerUse the nci_trigger tool to run a policy from the command line.

Configuring multiple resource columns

Learn how to configure multiple resource columns for one or more pattern resource definitions.

About this task

The `repattern_multiresource_correlation_logic` parameter is configured with the OR value by default. Use this procedure to change the setting. Only change the `repattern_multiresource_correlation_logic` setting if you understand the underlying algorithm. When OR logic is specified, it correlates two events by resource as soon as the criteria is met

for just one pattern resource definition. When AND logic is specified, only EQUALITY resource matching is used and the criteria must be met for all of the pattern resource definitions.

Note: If one resource field is selected per event type, the resource fields for each event type can be different. In this case AND logic is the same as OR logic. If more than one resource field is selected, the resource fields for each event type must be the same.

Note: Suggested patterns only use one resource field. They are never generated with multiple resources.

Procedure

1. Generate a properties file containing the latest Event Analytics system settings.
   a) Navigate to the directory `$IMPACT_HOME/bin`.
   b) Run the following command to generate a properties file containing the latest Event Analytics system settings.

      ```
      nci_trigger server_name username/password NOI_DefaultValues_Export
       FILENAME directory/filename
      ```

      Where:

      - *server_name* is the name of the server where Event Analytics is installed.
      - *user name* is the user name of the Event Analytics user.
      - *password* is the password of the Event Analytics user.
      - `NOI_DefaultValues_Export` is a Netcool/Impact policy that performs an export of the current Event Analytics system settings to a designated properties file.
      - *directory* is the directory where the properties file is stored.
      - *filename* is the name of the properties file.

      For example:

      ```
      nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Export
       FILENAME /tmp/properties.props
      ```

2. Edit the properties file that you generated in the previous step.

   For example:

   ```
   vi /tmp/properties.props
   ```

3. Find the section of the properties file that reads as follows:

   This code snippet shows the default values of the name similarity parameters.

   ```
   repattern_multiresource_correlation_logic=OR
   ```

4. Update the `repattern_multiresource_correlation_logic` setting, as described in the following table.
5. Import the modified properties file into Event Analytics.
   a) Ensure you are in the directory `$IMPACT_HOME/bin`.
   b) Run the following command to perform an import of Event Analytics system settings from a designated properties file.

      ```
      nci_trigger server_name username/password NOI_DefaultValues_Configure
       FILENAME directory/filename
      ```

      Where:

      - *server_name* is the name of the server where Event Analytics is installed.
      - *user name* is the user name of the Event Analytics user.
      - *password* is the password of the Event Analytics user.

- `NOI_DefaultValues_Configure` is a Netcool/Impact policy that performs an import of Event Analytics system settings from a designated properties file.
- *`directory`* is the directory where the properties file is stored.
- *`filename`* is the name of the properties file.

For example:

```
nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Configure
 FILENAME /tmp/properties.props
```

Configuring analytics

You can create and run configuration scans on demand or on a scheduled basis to generate analytics based on your event data.

Configure Analytics window

The Configure Analytics window contains a list of existing event configurations, or reports. Use this window to view, create, modify, delete, run, or stop event configurations.

Note: To access the Configure Analytics window, users must be assigned the `ncw_analytics_admin` role.

You can use the Configure Analytics window to determine whether an event recurs and when it recurs most frequently. For example, an event occurs frequently at 9 a.m. every Monday. Knowledge of the type of event and the patterns of recurrence can help to determine the actions that are required to reduce the number of events.

The Configure Analytics table displays the following columns of information for each event configuration.

Name
Specifies the unique event configuration name.

Event Identity
Specifies the database fields that identify a unique event in the database. Event seasonality runs on all events selected from the **Event Identity** drop-down list. If the **Event Identity** value is **Using Global Settings**, the Event Identity is set up in the configuration file.

Seasonality Enabled
Specifies whether the event configuration has seasonality event analytics enabled. This column displays one of the following values:

- `True`: Seasonality analytics is enabled.
- `False`: Seasonality analytics is not enabled.

The column displays the value `true` if seasonality analytics is enabled.

Related Event Enabled
Specifies whether the event configuration has related event analytics enabled. This column displays one of the following values:

- `True`: Related event analytics is enabled.
- `False`: Related event analytics is not enabled.

Seasonality Status
Specifies the status of the seasonality event configuration. The column can display one of the following status icons: Waiting, Running, Completed, or Error.

Related Event Status
Specifies the status of the related event configuration. The column can display one of the following status icons: Waiting, Running, Completed, or Error.

Start Time
Specifies the inclusive start date of historical data for the event configuration.

End Time
Specifies the inclusive end date of historical data for the event configuration.

Seasonality Phase
Specifies the phase of the seasonality event configuration run. In total, this column displays five phases during the run of the seasonality event configuration. For example, when the seasonality event configuration completion phase occurs, the value `Completed` displays in the column.

Seasonality phase progress
Displays the progress of the seasonality event phase, expressed in terms of percentage. For example, when the seasonality event configuration completion phases finishes, the value `100%` displays in the column.

Related Event Phase
Specifies the phase of the related event configuration run. In total, this column displays five phases during the run of the related event configuration. For example, when the related event configuration completion phase occurs, the value `Completed` displays in the column.

Related Event Phase Progress
Displays the progress of the related event phase, expressed in terms of percentage. For example, when the related event configuration completion phases finishes, the value `100%` displays in the column.

Scheduled
Indicates whether the event configuration run is scheduled to run every x number of days, weeks, or months. The value `Yes` displays in the column if the event configuration run is scheduled.

Relationship profile
The entry in this column specifies the strength of the relationship between the events in the event groups that are determined by the algorithm. One value is specified.

`Strong` Represents a high confidence level for relationships between events, but fewer events and fewer event groups. Your event groups might be missing weaker related events.

`Medium` Represents medium confidence level for relationships between events, average number of events and fewer event groups. Your event groups might be missing weakly related events.

`Weak` You see more events and more event groups but potentially more false positives.

Important: Netcool/Impact does not support Arabic or Hebrew. Event Analytics users who are working in Arabic or Hebrew see some untranslated English text.

Setting the Impact data provider and other portlet preferences

If there are multiple connections with the Impact data provider, you must specify which Impact data provider to use.

About this task

If a single connection with the Impact data provider exists, then that connection is used to compile the list of seasonal reports and display them in a table. If there are multiple connections with the Impact data provider, you must edit your portlet preferences to select one of the options.

Procedure

1. To edit your portlet preferences, or as an administrator to edit the portlet defaults:
   - To edit your portlet preferences, click **Page Actions** > **Personalize Page** > **Widget** > **Personalize**.
   - To edit the portlet defaults of all users, click **Page Actions** > **Edit Page** > **Widget** > **Edit**.

   The Configure Analytics dialog box is displayed.
2. Select a data provider from the **Data Provider** drop-down list.
3. In the **Bidi Settings** tab, specify the settings for the display of bidirectional text.

Component direction
Select the arrangement of items in the portlet, left-to-right, or right-to-left. The default setting uses the value that is defined for the page or the console. If the page and console both use the default setting, the locale of your browser determines the layout.

Text direction
Select the direction of text on the portlet. The default settings use the value that is defined for the page or the console. If the page and console both use the default setting, the locale of your browser determines the text direction. The `Contextual Input` setting displays text that you enter in the appropriate direction for your globalization settings.

4. To save your changes, complete the following steps.
   a) Select **Save** in the Configure Analytics dialog box. The Configure Analytics dialog box is closed.
   b) Select **Save**, which is in the upper right of the window.

Viewing current analytics configurations

An Administrator can see the list of current analytics configurations and some basic information about the analytics (related events and seasonal events) related to those configurations.

About this task

Event Analytics includes a default analytics configuration for you to run a basic configuration with default values. You can run, modify, or delete this analytics configuration. To view your analytics configurations, complete the following steps. This task assumes that you have logged into the Dashboard Application Services Hub as a user with the `ncw_analytics_admin` role.

Procedure

1. Start the Configure Analytics portlet.
   a) In the Dashboard Application Services Hub navigation menu, go to the **Insights** menu.
   b) Select **Configure Analytics**.
2. In the Configure Analytics portlet, a table presents a list of analytics configurations that are already configured. Scroll down the list to view all analytics configurations. The table automatically refreshes every 60 seconds and displays information for select column headings.
   - To view configuration parameters for a specific analytics configuration, select a configuration and then select the **Modify Selected Configuration** icon.
   - To view progress of the latest action that is taken for an analytics configuration, look at the content that is displayed in the following columns:

     Seasonality Status
     Specifies the status of the seasonality event configuration. The column can display one of the following status icons: Waiting, Running, Completed, or Error.

     Related Event Status
     Specifies the status of the related event configuration. The column can display one of the following status icons: Waiting, Running, Completed, or Error.

     Start Time
     Specifies the inclusive start date of historical data for the event configuration.

     End Time
     Specifies the inclusive end date of historical data for the event configuration.

     Seasonality Phase
     Specifies the phase of the seasonality event configuration run. In total, this column displays five phases during the run of the seasonality event configuration. For example, when the seasonality event configuration completion phase occurs, the value `Completed` displays in the column.

     Seasonality Phase Progress
     Displays the progress of the seasonality event phase, expressed in terms of percentage. For example, when the seasonality event configuration completion phases finishes, the value `100%` displays in the column.

Related Event Phase
Specifies the phase of the related event configuration run. In total, this column displays five phases during the run of the related event configuration. For example, when the related event configuration completion phase occurs, the value `Completed` displays in the column.

Related Event Phase Progress
Displays the progress of the related event phase, expressed in terms of percentage. For example, when the related event configuration completion phases finishes, the value `100%` displays in the column.

- To view other details about the analytics configuration, look at the information displayed in other columns:

Name
Specifies the unique event configuration name.

Event identity
Specifies the database fields that identify a unique event in the database. Event seasonality runs on all events selected from the **Event Identity** drop-down list.

Scheduled
Advises whether the analytics configuration is scheduled to query the historical events database.

Seasonality Enabled
Specifies whether the event configuration has seasonality event analytics enabled. This column displays one of the following values:

– `True`: Seasonality analytics is enabled.

– `False`: Seasonality analytics is not enabled.

Related Event Enabled
Specifies whether the event configuration has related event analytics enabled. This column displays one of the following values:

– `True`: Related event analytics is enabled.

– `False`: Related event analytics is not enabled.

Relationship Profile
The entry in this column specifies the strength of the relationship between the events in the event groups that are determined by the algorithm. One value is specified.

`Strong` Represents a high confidence level for relationships between events, but fewer events and fewer event groups. Your event groups might be missing weaker related events.

`Medium` Represents medium confidence level for relationships between events, average number of events and fewer event groups. Your event groups might be missing weakly related events.

`Weak` You see more events and more event groups but potentially more false positives.

Creating a new or modifying an existing analytics configuration
An Administrator can create a new analytics configuration or modify an existing analytics configuration. You choose the analytics type (related events, seasonal events, or both) you want to run during the create new analytics configuration operation.

Before you begin
When you modify an existing analytics configuration, you cannot change the following parameter fields in the dialog box:

- **Name:**
- **Name**
- **Analytics Type**
- **Event identity**
- **Event identity:**
- **Seasonal event analytics**

- **Related event analytics**

Procedure

1. Start the Configure Analytics portlet. See "Viewing current analytics configurations" on page 289.
2. Select the **Create New Configuration** icon to create a new analytics configuration, or highlight an existing analytics configuration and select the **Modify Selected Configuration** icon to modify an existing analytics configuration. The UI displays a dialog box that contains parameter fields for the new or existing analytics configuration.
3. Populate the parameter fields in the **General** tab of the dialog box with the details applicable to the analytics configuration.

   Name
   Enter the name of the analytics configuration. The name should reflect the type of analytics configuration you are creating.

   For example, `TestSeasonality1` and `TestRelatedEvents1` might be names you assign to analytics configurations for seasonality events and related events. The name for an analytics configuration must be unique and not contain certain invalid characters.

   The invalid character list is the list of characters listed in the *webgui_home*`/etc/illegalChar.prop` file.

   Analytics Type
   Select **Seasonal event analytics**, **Related event analytics**, or both.

   Event identity
   From the drop-down list, select the database fields that identify a unique event in the database. Event seasonality runs on all events that are selected from the **Event Identity** drop-down list. For information about how to change the fields in the drop-down list, see "Changing the choice of fields for the Event Identity" on page 295.

   Date Range
   Select either **RelativeFixed date range** or **FixedRelative date range**

   Relative: Enter the time frame that is to be included in the analytics configuration. The relative time frame is measured in **Months**, **Weeks**, or **Days**.

   Fixed: The **Start date** and **End date** parameter fields are active. Enter an inclusive **Start date** and **End date** for the analytics configuration.

   Fixed date range: The **Start date** and **End date** parameter fields are active. Enter an inclusive **Start date** and **End date** for the analytics configuration.

   Relative date range: Enter the time frame that is to be included in the analytics configuration.

   Run every
   To schedule the analytics configuration to run at specific time intervals, enter how frequently the configuration is to run. When you enter a value greater than zero, the analytics configuration becomes a scheduled configuration.

   Note: This option applies to the relative date range only. You cannot apply this option to the fixed date range.

   Filter
   Detail any filters that are applicable to the analytics configuration. For example, enter `Summary NOT LIKE '%maintenance%'`.

   *** Select the analytics type you want to run**
   Select **Seasonal event analytics**, **Related event analytics**, or both.
4. Populate the parameter fields in the **Related Events** tab of the dialog box with the details applicable to the analytics configuration.

   Relationship Profile
   Select the strength of the relationship between the events in an analytics configuration. If this value is set to `Strong`, there is more confidence in the result and less number of groups produced.

Automatically deploy rules discovered by this configuration
Select this option to automatically deploy rules that are discovered by this analytics configuration.

Relationship Profile Select the strength of the relationship between the events in an analytics configuration. If this value is set to `Strong`, there is more confidence in the result and less number of groups produced.

Automatically deploy rules discovered by this configuration Select this option if you want to automatically deploy rules that are discovered by this analytics configuration.

5. Populate the parameter field in the **Advanced** tab of the dialog box with the details applicable to the analytics configuration. You can use the defined event identities, select the **Override global event identity** to identify unique events in the database.

Override global event identity
Select this option to enable the **Event identity** drop-down list.

When the **Override global event identity** check box is selected, you cannot create a pattern from a configuration. However, you can deploy a related events group.

Event identity
From the **Event identity** drop-down list, select the database fields that identify a unique event in the database. Event seasonality runs on all events that are selected from the **Event Identity** drop-down list. For information about how to change the fields in the drop-down list, see "Changing the choice of fields for the Event Identity" on page 295.

Override global event type
Select this option to enable the **Event type** drop-down list.

When the **Override global event type** check box is selected, you cannot create a pattern from a configuration. However, you can deploy a related events group.

Event type
From the **Event type** drop-down list, select the database fields that identify an event type in the database.

Override global resource
Select this option to enable the **Resource** drop-down list.

When the **Override global resource** check box is selected, you cannot create a pattern from a configuration. However, you can deploy a related events group.

Resource
From the **Resource** drop-down list, select the database fields that identify a resource in the database.

6. Click either **Save** to save the report without running, or click **Save & Run** to save and run the report. You can also cancel the operation by clicking **Cancel**.

Results

- If no errors are found by the system validation of the analytics configuration content, the new or updated analytics configuration and its parameters are displayed in the table.
- If errors are found by the system validation of the analytics configuration content, you are prevented from saving the configuration and you are asked to reset the invalid parameter.

Scope-based groups

Events in a scope-based group are grouped together because they share a common attribute, such as a resource.

Create an event policy to set `ScopeID` for events that match your defined filter. Scope-based event grouping is activated for events that match the filter, based on the `ScopeID` that you specify. For more information about creating a scope-based grouping policy, see OMNIbus documentation: About scope-based grouping events with Event Analytics.

Manually running an unscheduled analytics configuration

An Administrator can manually run an unscheduled analytics configuration at any stage. You choose the analytics type (related events or seasonal events) you want to run during the run unscheduled analytics operation.

Before you begin

The analytics configuration that you try to manually run cannot have a **Related Event Status** or **Seasonality Status** of `Running`. If you try to manually run an analytics configuration that is already running, the GUI displays a warning message.

Note: Sequentially running reports can take longer to complete than parallel running reports in previous releases.

Procedure

1. Start the Configure Analytics portlet. See "Viewing current analytics configurations" on page 289.
2. Within the list of analytics configurations that are displayed, select one configuration.
3. From the toolbar, click the **Run Selected Configuration** icon. Some columns are updated for your selected analytics configuration.

   The icon in the **Seasonality Status** or **Related Event Status** column changes to a time glass icon.

   The text in the **Seasonality Phase** or **Related Event Phase** column changes to `Waiting to Start`.

   The percentage in the **Seasonality Phase Progress** or **Related Event Phase Progress** column starts at `0%` and changes to reflect the percentage complete for the phase.

Results

The analytics configuration is put into the queue for the scheduler to run. As the analytics configuration is running, the following columns are updated to communicate the progress of the run:

- **Seasonality Status** or **Related Event Status**
- **Seasonality Phase** or **Related Event Phase**
- **Seasonality Phase Progress** or **Related Event Phase Progress**

What to do next

If you want to stop an analytics configuration that is in `Running` status, from the toolbar click the **Stop Selected Configuration** icon.

Stopping an analytics configuration

You can stop an analytics configuration that is running.

About this task

If you create an analytics configuration and select to run the configuration, you might realize that some configuration values are incorrect while the configuration is still running. In this situation you can choose to stop the analytics configuration instead of deleting the configuration or waiting for the configuration run to complete. To stop a running analytics configuration, complete the following steps.

Procedure

1. Start the related events configuration portlet, see "Viewing current analytics configurations" on page 289.
2. Within the list of analytics configurations that are displayed, select the running configuration.
3. From the toolbar, click the **Stop Selected Configuration** icon.

Deleting an analytics configuration

Analytics configurations can be deleted individually, regardless of their status.

Procedure

1. Start the Configure Analytics portlet, see "Viewing current analytics configurations" on page 289.
2. Select the name of the analytics configuration that you want to delete and from the toolbar click the **Delete Selected Configuration** icon.
3. Within the confirmation dialog that is displayed, select **OK**.

   If you attempt to delete an analytics configuration with one or more rules created for it, a text warning dialog box appears with the current rules status for that analytics configuration. The following example illustrates text that the warning dialog box can contain:

   ```
   Configuration EventAnalytics_Report_1 contains the following rules:

   Seasonality Rules:
   0 watched rules, 1 active rules, 0 expired rules and 0 archived

   Related Event Rules:
   0 watched rules, 0 active rules, 0 expired rules and 0 archived

   Delete the rules manually before deleting the configuration.
   ```

   As the message indicates, manually delete the rule or rules associated with the specified analytics configuration before deleting the analytics configuration. In the example, the one active rule associated with the analytics configuration called `EventAnalytics_Report_1` would need to be deleted first.

Results

- The table of analytics configurations refreshes, and the deleted configuration no longer appears in the list of analytics configurations.
- Deleting the analytics configuration does not delete the related results if the results are in the **Deployed** or **Expired** state. However, deleting the analytics configuration does delete the related results that are in the **New** or **Archived** state.
- You are unable to reuse the name of a deleted analytics configuration until all related event groups that contain the name of the deleted configuration are deleted from the system.

Changing the expiry time for related events groups

You can modify the expiry time for `Active` related events groups. When the expiry time is reached, the expired groups and related events display in the View Related Events portlet, within the **Expired** tab.

About this task

Groups that contain an expired group or pattern continue to correlate. The system administrator should review the group and expired group or event.

By default the related events expiry time is 6 months. Complete the following steps to change the related events expiry time.

Note: `Watched` related events groups do not expire.

Procedure

1. Log in to the Netcool/Impact UI.
2. Select the **Related Events project**.
3. Select the **Policies** tab.
4. Within the **Policies** tab, select to edit the **RE_CONSTANTS** policy.
5. Within the **RE_CONSTANTS** policy, change the value for the `RE_EXPIRE_TIME` constant. Enter your new value in months.

6. Save the policy.

Results
This change takes effect only with newly discovered related event groups in the Active tabs.

What to do next
If you want to configure the expiry time so that deployed groups never expire, change the value for the RE_EXPIRE_TIME constant to 0 and save the policy for this change to take effect. You do not need to restart the Impact Server.

If you want to enable the expiry time at any stage, set this variable back to a value greater than 0.

Changing the choice of fields for the Event Identity
You can change which fields, from your event history database, are available for selection as the Event Identity.

About this task
An Event Identity is a database field that identifies a unique event in the event history database. When you configure a related events configuration, you select database fields for the Event Identity from a drop-down list of available fields. Through configuration of an exception list within Netcool/Impact, you can change the fields available for selection in the drop-down list. Fields included in the exception list do not appear in the Configure Analytics portlet.

The Netcool/Impact design displays the following default fields in the Event Identity drop-down list

```
Alert Group
Alert Key
Node
Summary
Identifier
LOCALNODEALIAS
LOCALPRIOBJ
LOCALROOTOBJ
LOCALSECOBJ
REMOTENODEALIAS
REMOTEPRIOBJ
REMOTEROOTOBJ
REMOTESECOBJ
```

If you have other database fields that are not in the exception list, these other fields also appear in the drop-down list. Complete the following steps to modify the exception list.

Procedure

1. Log in to Netcool/Impact.
2. From the list of available projects, select the **RelatedEvents** project.
3. Select the **Policies** tab. Within this tab, select and edit the **RE_CONSTANTS** policy.
4. Update the RE_OBJECTSERVER_EXCLUDEDFIELDS variable. Add or remove fields from the static array. Case sensitivity does not matter.
5. Save the policy.
6. Run the policy. If there is an error, check your syntax.

Results
The changes occur when the policy is saved. No restart of Netcool/Impact is needed.

Changing the discovered groups

You can modify the discovered groups for related events groups.

About this task

By default all unallocated groups are shown in the group sources for a configuration. Complete the following steps to change the discovered groups.

Procedure

1. Log in to the Netcool/Impact UI.
2. Select the **Related Events project**.
3. Select the **Policies** tab.
4. Within the **Policies** tab, select to edit the **NOI_CONSTANTS** policy.
5. Within the **NOI_CONSTANTS** policy, change the value for the *RE_DISCOVERED_GROUPS* variable.
6. Save the policy.
7. Run the `noi_derby_upgradefp15.sql` file to upgrade or re-run the configuration to change the existing data in the views. The default value after upgrade is: `Unallocated Groups`

Validating analytics and deploying rules

Review and validate the analytics to create rules to apply to the live event stream that the operators see in the Netcool/OMNIbus Event Viewer.

View Seasonal Events portlet

The View Seasonal Events portlet contains a list of configurations, a list of seasonal events, and seasonal event details.

In addition to viewing the seasonal events, you can mark events as reviewed and identify the events that were reviewed by others.

The View Seasonal Events portlet displays the following default columns in the group table:

Configuration
Displays a list of the seasonal event configurations.

Event Count
Displays a count of the number of seasonal events for each seasonal event configuration.

Node
Displays the managed entity from which the seasonal event originated. The managed entity could be a device or host name, service name, or other entity.

Summary
Displays the description of the seasonal event.

Alert Group
Displays the Alert Group to which the seasonal event belongs.

Reviewed by
Displays the list of user names of the users who reviewed the seasonal event.

Confidence Level
Displays icons and text based on the level of confidence that is associated with the seasonal event. The confidence level is displayed as high, medium, or low, indicating that an event has a high, medium, or low seasonality.

Maximum Severity
Displays the maximum severity of the events that contribute to the seasonality of the selected seasonal event.

Rule Created
Displays the name of the seasonal event rule that was created for the seasonal event.

Related Events Count
Displays a count of the number of related events for each seasonal event.

First Occurrence
Displays the date and time when the seasonal event first occurred. The time stamp is configurable by users and is displayed in the following format:

```
YYYY-MM-DD HH:MM:SS
```

For example:

```
2012-10-12 09:52:58.0
```

Viewing a list of seasonal event configurations and events
You can view a list of the seasonal event configurations and seasonal events in the View Seasonal Events portlet.

Before you begin
To access the View Seasonal Events portlet, users must be assigned the `ncw_analytics_admin` role.

Procedure

To view a list of the seasonal event configurations and seasonal events, complete the following steps.

1. Open the **View Seasonal Events** portlet.
2. By default, the Seasonal Event configurations are listed in the **Configuration** table.
3. To view a list of seasonal events associated with the configurations, select one of the following options.
   a) Select **All** to view of list of the seasonal events for all of the configurations.
   b) Select a specific configuration to view a list of the seasonal events for that configuration.

   The seasonal events are listed in the **Summary** table.

Results
The seasonal event configurations and associated seasonal events are listed in the View Seasonal Events portlet.

Reviewing a seasonal event
You can mark or unmark a seasonal event as reviewed.

About this task
The **Reviewed by** column in the View Seasonal Events portlet displays the user name of the reviewer.

Procedure

To update the review status of an event, complete the following steps.

1. Open the **View Seasonal Events** portlet.
2. Select a specific configuration or **ALL** in the configuration table.
3. Select a seasonal event from the events table.
4. Right-click the seasonal event and select **Mark as Reviewed** or **Unmark as Reviewed**.

   Each seasonal event can be reviewed by multiple users. The reviewers are listed in the **Reviewed by** column.

Results
The selected seasonal event is marked or unmarked as **Reviewed**. The **Reviewed by** column is updated to display the user name of the reviewer.

Sorting columns in the View Seasonal Events portlet

You can sort the columns in the View Seasonal Events portlet to organize the displayed data.

Before you begin

To access the View Seasonal Events portlet, users must be assigned the `ncw_analytics_admin` role.

About this task

The rows in the View Seasonal Events portlet are sorted by the configuration name. You can change the order of the rows by using the columns to sort the data.

Sorted columns are denoted by an upwards-pointing arrow or downwards-pointing arrow in the column header, depending on whether the column is sorted in ascending or descending order.

Procedure

To sort the rows by column, complete the following steps.

1. Open the **View Seasonal Events** portlet.
2. To sort single columns, complete the following steps.
   a) To sort a column, click the column header once. The rows are sorted in ascending order.
   b) To sort in descending order, click the column header again.
   c) To unsort the column, click the column header a third time.
3. To sort multiple columns, complete the following steps.
   a) Sort the first column as a single column.
   b) Move the mouse pointer over the column header of the next column you want to sort. Two icons are displayed. One is a standard sorting icon and the other is a nested sorting icon. The nested sorting icon has a number that represents how many columns are sorted as a result of selecting the option. For example, if this is the second column that you want to sort the number 2 is displayed.
   c) Click the nested sorting icon. The column is sorted with regard to the first sorted column.

      Tip: When you move the mouse pointer over the nested sorting icon, the hover help indicates that it is a nested sorting option. For example, the hover help for the icon displays "Nested Sort - Click to sort Ascending". The resulting sort order is ascending with regard to the previous columns on which a sorting order was placed.
   d) To reverse the order of the nested sort, click the nested sorting icon again. The order is reversed and the nested sorting icon changes to the remove sorting icon.
   e) To remove nested sorting from a column, move the mouse pointer over the column header and click the **Do not sort** icon.

   Note: In any sortable column after nested sorting is selected, when you click the standard sorting icon, it becomes the only sorted column in the table and any existing sorting, including nested is removed.

Results

Sorted columns are marked with an upwards-pointing arrow or a downwards-pointing arrow in the column header to indicate whether the column is sorted in ascending or descending order. The sorting is temporary and is not retained.

Exporting all seasonal events for a specific configuration to Microsoft Excel

You can export all seasonal events for a specific configuration to a Microsoft Excel spreadsheet from a supported browser.

Before you begin

You view seasonal events for one or more configurations in the **View Seasonal Events** portlet. To access the **View Seasonal Events** portlet, users must be assigned the `ncw_analytics_admin` role.

Procedure

To export all seasonal events for a specific configuration to a Microsoft Excel spreadsheet, complete the following steps.

1. Open the **View Seasonal Events** portlet.
2. Select a specific configuration from the configuration table.
3. Click the **Export Seasonal Events** button in the toolbar.
   After a short time, the **Download export results** link displays.
4. Click the link to download and save the Microsoft Excel file.

Results

The Microsoft Excel file contains a spreadsheet with the following tabs:

- `Report Summary`: This tab contains a summary report of the configuration that you selected.
- `Seasonal Events`: This tab contains the seasonal events for the configuration that you selected.
- `Export Comments`: This tab contains any comments relating to the export for informational purposes (for example, if the spreadsheet headers are truncated, or if the spreadsheet rows are truncated).

Exporting selected seasonal events for a specific configuration to Microsoft Excel

You can export selected seasonal events for a specific configuration to a Microsoft Excel spreadsheet from a supported browser.

Before you begin

You view seasonal events for one or more configurations in the **View Seasonal Events** portlet. To access the **View Seasonal Events** portlet, users must be assigned the `ncw_analytics_admin` role.

Procedure

To export selected seasonal events for a specific configuration to a Microsoft Excel spreadsheet, complete the following steps.

1. Open the **View Seasonal Events** portlet.
2. Select a specific configuration from the configuration table.
3. Select multiple seasonal events by using the `Crtl` key and select method. (You can also select multiple seasonal events by using the click and drag method.)
4. After selecting multiple seasonal events, right click on one of the selected seasonal events and select the **Export Selected Events** button in the toolbar.
   After a short time, the **Download export results** link displays.
5. Click the link to download and save the Microsoft Excel file.

Results

The Microsoft Excel file contains a spreadsheet with the following tabs:

- `Report Summary`: This tab contains a summary report of the configuration that you selected.
- `Seasonal Events`: This tab contains the seasonal events for the configuration that you selected.
- `Export Comments`: This tab contains any comments relating to the export for informational purposes (for example, if the spreadsheet headers are truncated, or if the spreadsheet rows are truncated).

Seasonal Event Rules

You can use seasonal event rules to apply an action to specific events.

You can choose to apply actions to a selected seasonal event, or to a seasonal event and some or all of its related events.

You can use seasonal event rules to apply actions to suppress and unsuppress an event, to modify or enrich an event, or to create an event if the selected event does not occur when expected.

Creating a seasonal event rule

You can create a watched or deployed seasonal event rule from the View Seasonal Events portlet.

Before you begin

To access the View Seasonal Events portlet, users must be assigned the `ncw_analytics_admin` role.

Procedure

To create a seasonal event rule in the View Seasonal Events portlet, complete the following steps.

1. Open the **View Seasonal Events** portlet.
2. Select a specific configuration or **ALL** in the configuration table.
3. Select a seasonal event from the events table.
4. Right-click the seasonal event and select **Create Rule**.
5. Input a unique rule name in the **Create Rule** window.
6. Input the following rule criteria for events and actions in the **Create Rule** window.
   a) To apply rule actions to an event and time condition, see "Applying rule actions to an event and time condition" on page 300.
   b) To apply actions when an event occurs, see "Applying actions when an event occurs" on page 301.
   c) To Applying actions when an event does not occur, see "Applying actions when an event does not occur" on page 302.
7. To save the seasonal event rule, choose one of the following criteria.
   a) Select **Watch** to monitor the rule's performance before it is deployed.
   b) Select **Deploy** to activate the rule.

Results

A seasonal event rule is created. To view a list of current seasonal event rules, open the **Seasonal Event Rules** portlet.

Applying rule actions to an event and time condition

To create a seasonal event rule, you must specify the selected events or time conditions, or both in the Create Rule or **Modify Existing Rule** window.

Before you begin

Create or modify an existing seasonal event rule. To create a seasonal event rule, see "Creating a seasonal event rule" on page 300. To modify an existing seasonal event rule, see "Modifying an existing seasonal event rule" on page 305.

About this task

The seasonal event that is selected by default in the **Event Selection** pane is the seasonal event from which the Create Rule or **Modify Existing Rule** window was opened.

Note: A seasonal event rule suppresses events when they occur for a deployed related event group. The seasonal rule actions do not apply to the synthetic parent event that is created.

Note: You can create a seasonal event rule to unsuppress an event or alarm. This rule has no actions if there are no suppressed alarms.

Procedure

To specify the selected events and time conditions, complete the following steps in the Create Rule window.

1. To select the all, or one or more of the events that are related to the seasonal event, complete the following steps.
   a) To select all of the related events, select the **Select all related events** check box.

b) To edit or select one or more of the related events, click **Edit Selection**, select one or more related events.

2. To save the related eventselection, click **OK**
3. To select a time condition, complete the following steps.
   a) Select one of the following time condition filter conditions.

      AND
      Select **AND** to apply rule actions to each of the selected the time conditions.

      OR
      Select **OR** to apply rule actions to individual time conditions.

   b) Select **Minute of the Hour**, **Hour of the Day**, **Day of Week**, or **Day of Month** from the drop-down menu.

   c) Select **Is** or **Is Not** from the drop-down menu.

   d) Select the appropriate minute, hour, day, or date from the drop-down menu. You can select multiple values from this drop-down menu.

      Note: High, medium, and low seasonality labels are applied to this time selection drop-down menu to indicate the seasonality of the events occurring at that time.

4. Click the add button to add another time condition.
5. To save the event selection and time conditions, choose one of the following criteria.
   a) Select **Watch** to monitor the rule's performance before it is deployed.

   b) Select **Deploy** to activate the rule.

Results
The seasonal event rule conditions are applied to the selected events, time conditions, or both.

Applying actions when an event occurs
You can apply specific actions to occur when an event occurs in a specific time window.

Before you begin

Create or modify an existing seasonal event rule. To create a seasonal event rule, see “Creating a seasonal event rule” on page 300. To modify an existing seasonal event rule, see “Modifying an existing seasonal event rule” on page 305.

Note: To suppress or unsuppress events you must update the `noi_default_values` file. For more information about the `noi_default_values` file, see “Configuring event suppression” on page 269.

About this task

The events that are selected in the **Event Selection** pane are the events to which the action is applied when the event occurs in a specific time window. For more information about selecting events, see “Applying rule actions to an event and time condition” on page 300.

You can suppress events that do not require you to take any direct action, and unsuppress the events after a specified time period.

You can set a column value on an event occurrence and again set it again after a specified time period.

Procedure

To specify the actions to apply when an event occurs, complete the following steps in the **Actions When Event(s) Occurs in Specified Time Window(s)** pane.

1. To suppress an event so that no action is taken when it occurs, complete the following steps.
   a) Select the **Suppress event(s)** check box.

   b) (Optional) To select a column value, see step 3 below.

2. To unsuppress an event after an action occurs, complete the following steps.
   a) To select the time after the action occurs to unsuppress the event, select a number from the **Perform Action(s) After** list, or type an entry in the field. Select **Seconds**, **Minutes**, or **Hours** from the **Perform Action(s) After** drop-down list.
   b) Select the **Unsuppress event(s)** check box.
   c) (Optional) To select a column value, see step 4 below.
3. To set the column value after an action occurs, complete the following steps.
   a) Select the **Set Column Values** check box and click the **Set Column Value** button for **Perform Action(s) on Event Occurrence**.
   b) In the **Set Column Value** page, input values for the `ObjectServer` columns.
   c) To save the column values, click **Ok**.
4. To reset the column value after a specified time period, complete the following steps.
   a) To specify a time period, select a number from the **Perform Action(s) After** list, or type an entry in the field. Select **Seconds**, **Minutes**, or **Hours** from the **Perform Action(s) After** drop-down list.
   b) Select the **Set Column Values** check box and click the **Set Column Value** button for **Perform Action(s) After**.
   c) In the **Set Column Value** page, input values for the `ObjectServer` columns.
   d) To save the column values, click **Ok**.
5. To save the seasonal event rule, choose one of the following options.
   a) Select **Watch** to monitor the rule's performance before it is deployed.
   b) Select **Deploy** to activate the rule.

Results
The action to be applied to a rule that occurs in a specific time window is saved.

Applying actions when an event does not occur
You can apply specific actions to occur when an event does not occur in a specific time window.

Before you begin
Create or modify an existing seasonal event rule. To create a seasonal event rule, see "Creating a seasonal event rule" on page 300. To modify an existing seasonal event rule, see "Modifying an existing seasonal event rule" on page 305.

About this task
The events that are selected in the **Event Selection** pane are the events to which the action is applied if the event does not occur in a specific time window. For more information about selecting events, see "Applying rule actions to an event and time condition" on page 300.

Procedure

To specify the actions to apply when an event does not occur, complete the following steps in the **Actions When Event(s) Does Not Occur in Specified Time Window(s)** pane.

1. To select the time after which the event does not occur to apply the action, complete the following steps.
   a) Select a number from the **Perform Action(s) After** list, or type an entry in the field.
   b) Select **Seconds**, **Minutes**, or **Hours** from the **Perform Action(s) After** drop-down list.
2. To create a synthetic event on a non-occurrence, select the **Create event** check box and click **Create event**.
3. To define the event, complete the fields in the new **Create Event** window.
4. To save the synthetic event, click **Ok**.
5. To save the seasonal event rule, choose one of the following options.

a) Select **Watch** to monitor the rule's performance before it is deployed.

b) Select **Deploy** to activate the rule.

Results
The action to be applied to a rule that does not occur in a specific time window is saved.

Seasonal event rule states
Seasonal event rules are grouped by state in the Seasonal Event Rules portlet.

Seasonal event rule states

Seasonal event rules are grouped in the following states.

Watched
A watched seasonal event rule is not active.

You can watch a seasonal event rule to monitor how the rule performs before you decide whether to deploy it.

Watched seasonal event rules take no actions on events. It is used to collect statistics for rule matches for incoming events.

Active
A deployed seasonal event rule is active. Active seasonal event rules take defined actions on live events.

Expired
An expired seasonal event rule remains active. If triggered, the seasonal event rule takes defined actions on live events. The default expiry time is **6 MONTHS**. To ensure that seasonal event rules are valid, regularly review the state and performance of the rules. You can customize the expiry time of a seasonal event rule. For more information, see "Modifying the default seasonal event rule expiry time" on page 303.

Archived
An archived seasonal event rule is not active. You can choose to archive a watched, active, or expired seasonal event rule. To delete a seasonal event rule, it must first be archived.

For more information about changing the state of a seasonal event rule, see "Modifying a seasonal event rule state" on page 306.

Modifying the default seasonal event rule expiry time
You can change the default seasonal event rules expiry time to a specific time or choose no expiry time to ensure that a seasonal event rule does not expire.

About this task
To ensure that seasonal event rules are valid, you should regularly review and update the state of the rules.

Procedure

To modify or remove the default seasonal event rules expiry time, complete the following steps.

1. To generate a properties file from the command line interface, use the following command:

```
./nci_trigger SERVER <UserID>/<Password> NOI_DefaultValues_Export FILENAME
 directory/filename
```

Where

SERVER
The server where Event Analytics is installed.

<UserID>
The user name of the Event Analytics user.

<Password>
The password of the Event Analytics user.

directory
The directory where the file is stored.

filename
The name of the properties file.

For example:

```
./nci_trigger NCI impactadmin/impact NOI_DefaultValues_Export FILENAME
 /space/noi_default_values
```

2. To modify the default seasonal event rules expiry time, edit the default values of the following parameters.

seasonality.rules.expiration.time.value=*6*
The number of days, hours, or months after which the seasonal event rule expires. The default value is *6*.

seasonality.rules.expiration.time.unit=*MONTH*
The seasonal event rules expiry time frequency. The default frequency is *MONTH*. The following time units are supported:

- HOUR
- DAY
- MONTH

3. To import the modified properties file into IBM Tivoli Netcool/Impact, use the following command:

```
./nci_trigger SERVER <UserID>/<Password> NOI_DefaultValues_Configure FILENAME
 directory/filename
```

For example:

```
./nci_trigger NCI impactadmin/impact NOI_DefaultValues_Configure FILENAME
 /space/noi_default_values
```

Results
The default seasonal event rules expiry time is modified.

Viewing performance statistics for seasonal event rules
You can view performance statistics for seasonal event rules in the Seasonal Event Rules portlet, within the **Watched**, **Active**, or **Expired** tabs of the group table.

Columns in the group table

The group table in the View Seasonal Events portlet displays the seasonal event rules that you have created. The left-side of the group table has these columns:

Configuration: Displays the list of configuration names for which a seasonal event rule has been created.

Rule Count: Displays the number of seasonal event rules created for each particular configuration. This number also indicates the total number of seasonal event rules created for all configurations under the **All** item.

Rule Name: Displays the name of the seasonal event rule.

Last Run: Displays the date and time when the seasonal event rule was last executed. If the column is blank, the seasonal event rule has not been executed.

Deployed: Displays the date and time when the seasonal event rule was deployed. The term deployed means that the seasonal event rule is available for use, is actively accumulating rule statistics, and any actions applied to the rules are being performed.

Note: For the **Last Run** and **Deployed** columns, the date is expressed as *month*, *day*, *year*. Likewise, the time is expressed as *hours*:, *minutes*:, *seconds*. The time also indicates whether AM or PM. For example: `Apr 13, 2015 4:45:17 PM`.

Performance statistics in the group table

The performance statistics are displayed in the following columns of the group table in the **Watched**, **Active**, or **Expired** tabs in the Seasonal Event Rules portlet.

Suppressed Events: Displays the total number of events that the seasonal event rule suppressed since the rule was deployed.

Unsuppressed Events: Displays the total number of events that the seasonal event rule unsuppressed since the rule was deployed.

Enriched/Modified Events: Displays the total number of events that the seasonal event rule enriched or modified since the rule was deployed.

Generated Events on Non-occurrence: Displays the total number of events that the seasonal event rule generated since the rule was deployed for events that do not meet the event selection criteria (that is, for those matching events that fall outside of the event selections condition for the rule).

Reset performance statistics

You can reset performance statistics to zero for a group in the **Watched**, **Active**, or **Expired** tabs. To reset performance statistics, right-click on the seasonal event rule name (from the **Rule Name** column) and from the menu select **Reset performance statistics**. A message displays indicating that the operation will reset statistics data for the selected seasonal event rule. The message also indicates that you will not be able to retrieve this data. Click OK to continue with the operation or `Cancel` to stop the operation. A success message displays after you select OK.

Resetting performance statistics to zero for a seasonal event rule also causes the following columns to be cleared: **Last Run** and **Deployed**. Note that performance statistics are not collected for the **Archived** tab. When a rule is moved between states, the performance statistics are reset. Every time an action is triggered by the rule the performance statistics increase.

Modifying an existing seasonal event rule

You can modify an existing seasonal event rule to update or change the event selection criteria or actions.

Before you begin

To access the **Seasonal Event Rules** portlet, users must be assigned the `ncw_analytics_admin` role.

Procedure

1. Open the Seasonal Event Rules portlet.

   The **Seasonal Event Rules** portlet lists the seasonal event rules configuration in the table on the left side, and the seasonal event rules are listed in the table in the right side.
2. Select the rule that you want to modify from the rule table.
3. Right click and select **Edit Rule**.
4. Modify the event selection criteria or actions in the **Modify Existing Rule** window.
5. To save the seasonal event rule, choose one of the following criteria.
   a) Select **Watch** to monitor the rule's performance before it is deployed.
   b) Select **Deploy** to activate the rule.

Results

The seasonal event rule is modified. To view a list of current seasonal event rules, open the **Seasonal Event Rules** portlet.

Viewing seasonal event rules grouped by state

You can view seasonal event rules grouped by state in the **Seasonal Event Rules** portlet.

Before you begin

To access the **Seasonal Event Rules** portlet, users must be assigned the `ncw_analytics_admin` role.

Procedure

To view seasonal event rules grouped by state, complete the following steps.

1. Open the Seasonal Event Rules portlet.

   The **Seasonal Event Rules** portlet lists the seasonal event rules configuration in the table on the left side, and the seasonal event rules are listed in the table in the right side.
2. Select the seasonal event rule state that you want to view from the status tabs.

   The seasonal event rules are stored in tabs that relate to their status. For example, to view a list of the active seasonal event rules configurations and rules, select the **Active** tab.

Results

The seasonal event rules configurations and rules for the chosen status are listed in the **Seasonal Event Rules** portlet.

Modifying a seasonal event rule state

You can change the state of a seasonal event rule to watched, active, or archived from the Seasonal Event Rules portlet.

Before you begin

To access the **Seasonal Event Rules** portlet, users must be assigned the `ncw_analytics_admin` role.

About this task

The seasonal event rules are stored in tabs that relate to their state. The total number of rules is displayed on the tabs. For example, when you **Archive** a **Watched** rule, the rule moves from the **Watched** tab to the **Archived** tab in the Seasonal Event Rules portlet and the rules total is updated.

Performance statistics about the rule are logged. You can use performance statistics to verify that a deployed rule is being triggered and that a monitored rule is collecting statistics for rule matches for incoming events. Performance statistics are reset when you change the state of a seasonal event rule.

Procedure

To change the state of a seasonal event rule in the Seasonal Event Rules portlet, complete the following steps.

1. Open the Seasonal Event Rules portlet.

   The **Seasonal Event Rules** portlet lists the seasonal event rules configuration in the table on the left side, and the seasonal event rules are listed in the table in the right side.
2. To change the state of seasonal event rule, complete one of the following actions:
   a) To change the state of a watched seasonal event rule, select the **Watched** tab. Select a rule from the rule table. To change the state of the rule right-click the rule and select **Deploy** or **Archive**.
   b) To change the state of an active seasonal event rule, select the **Active** tab. Select a rule from the rule table. To change the state of the rule right-click the rule and select **Watch** or **Archive**.
   c) To change the state of an expired seasonal event rule, select the **Expired** tab. Select a rule from the rule table. To change the state of the rule right-click the rule and select **Validate**, **Watch**, or **Archive**.
   d) To change the state of an archived seasonal event rule, select the **Archived** tab. Select a rule from the rule table. To change the state of the rule right-click the rule and select **Watch** or **Deploy**.

Results
The seasonal event rule state is changed from its current state to its new state. The rule totals are updated to reflect the new seasonal event rule state.

Applying rule actions to a list of events
You can apply defined actions to a list of events while you create a seasonal event rule.

Before you begin
To access the View Seasonal Events portlet, users must be assigned the `ncw_analytics_admin` role.

About this task

One of the events in the list on the **Related Event Selection** window is the seasonal event from which you launched the Create Rule dialog box. When the rule you created is fired, the rule is fired on the seasonal event and the related events that you selected. Because the rule is fired on the seasonal event, it is not possible for you to deselect this seasonal event from the list of related events displayed in the **Related Event Selection** window.

Procedure

To select a list of events to which the defined action applies, complete the following steps.

1. Open the **View Seasonal Events** portlet.
2. Select a specific configuration or **ALL** in the configuration table.
3. Select a seasonal event from the events table.
4. Right-click the seasonal event and select **Create Rule**.
5. To choose all related events:
   a) In the **Event Selection** pane of the **Create Rule** page, click the **Select all related events** checkbox.
6. Or, to choose one or more related events:
   a) In the **Event Selection** pane of the **Create Rule** page, click the **Edit Selection...** control button.
      The **Related Event Selection** window displays. Note that the seasonal event from which you launched the Create Rule dialog box has a check mark that you cannot deselect.
   b) Select one or more related events from the list displayed in the **Related Event Selection** window.
   c) Click **OK**.
7. To save your changes, choose one of the following options:
   a) Select **Watch** to monitor the rule's performance before it is deployed.
   b) Select **Deploy** to activate the rule.

Results
The updated seasonal event rule is saved and the defined actions are applied to the selected related events.

Setting the column value for an event
You can set the column value for an event when you set the actions for a rule.

Before you begin
To access the View Seasonal Events portlet, users must be assigned the `ncw_analytics_admin` role.

Procedure

To set the column value, complete the following steps.

1. Open the **View Seasonal Events** portlet.
2. Select a specific configuration or **ALL** in the configuration table.
3. Select a seasonal event from the events table.

4. Right-click the seasonal event and select **Create Rule**.
5. In the **Actions When Event(s) Occurs in Specific Time Window(s)** pane, select from the following options.
   a) To set the column value to suppress an event, select the **Set Column Values** check box and click the **Set Column Value** button for **Perform Action(s) on Event Occurrence**.
   b) To set the column value to unsuppress an event, select the **Set Column Values** check box and click the **Set Column Value** button for **Perform Action(s) After**.
6. In the **Set Column Value** page, input values for the `ObjectServer` columns.
   a) You can add or remove columns by using the **plus** and **minus** buttons.
7. To save the column values, click **Ok**.
8. To save the seasonal event rule, choose one of the following options.
   a) Select **Watch** to monitor the rule's performance before it is deployed.
   b) Select **Deploy** to activate the rule.

Results
The seasonal event rule that modifies the column values is saved.

Seasonal Event Graphs
The seasonal event graphs display bar charts and confidence level event thresholds for seasonal events.

The **Seasonal Event Graphs** portlet consists of four charts:

Minute of the hour
The minute or minutes of the hour that the event occurs.

Hour of the day
The hour or hours of the day that the event occurs.

Day of the week
The day or days of the week that the event occurs.

Day of the month
The date or dates of the month that the event occurs.

The confidence level of the data in the charts is displayed in three ways:

1. The overall distribution score of each chart is displayed as high (red), medium (orange), or low (green) seasonality at the top of each chart.
2. The degree of deviation of the events is indicated by the high (red) and medium (orange) seasonality threshold lines on the charts.
3. The maximum confidence level of each bar is displayed as high (red), medium (orange), or low (green).

The default confidence level thresholds are as follows:

- High: 99-100%
- Medium: 95-99%
- Low: 0-95%

To modify the default confidence level thresholds of the charts, see “Editing confidence thresholds of Seasonal Event Graphs” on page 311.

Understanding graphs

The four seasonal event graphs illustrate event seasonality. The graphs depict independent observations. For example, if the **Hour of the day** graph indicates a high confidence level for 5 p.m., and the **Minute of the hour** graph indicates a high confidence level for minute 35, it does not necessarily mean that the events all occur at 5:35 p.m. The 5 p.m. value can contain other minute values.

Note: In some instances, **Minute of the hour** is indicated as having a high confidence level but the overall confidence level of seasonality is low. This is due to the high-level statistic that does not include minute of the hour due to poll cycle of monitors.

Note: In some instances, the overall confidence level of a chart is indicated as high although none of the bars in the graph are in the red zone. An example of this is a system failure due to high load and peak times, with no failure outside of these times.

The seasonal event graphs **Count** refers to the number of observations that are recorded in each graph. There is a maximum of one observation for each minute, hour, day, and date range. Therefore, the count for each of the graphs can differ. For example, if an event occurs at the following times:

10:31 a.m., 1 June 2013
10:31 a.m., 2 June 2013
10:35 a.m., 2 June 2013

There is a count of two observations for 10 a.m., two observations for minute 31, and one observation for minute 35.

Viewing seasonal event graphs for a seasonal event
You can view seasonal event graphs for the seasonal events that are displayed in the View Seasonal Events portlet.

Before you begin
To access the View Seasonal Events portlet, users must be assigned the `ncw_analytics_admin` role.

Procedure

To view seasonal event graphs for a seasonal event, complete the following steps.

1. Open the **View Seasonal Events** portlet.
2. Select a specific configuration or **ALL** in the configuration table.
3. Select a seasonal event from the events table.
4. Right-click the seasonal event and select **Show Seasonal Event Graphs**.

Results
The Seasonal Event Graphs portlet displays the bar charts and confidence levels for the selected seasonal event. For more information about charts and threshold levels, see the "Seasonal Event Graphs" on page 308 topic.

Viewing historical events from seasonality graphs
You can view a list of historical events from seasonality graphs.

Before you begin
To access the View Seasonal Events portlet, users must be assigned the `ncw_analytics_admin` role.

Procedure

To view a list of historical events from seasonality graphs, complete the following steps.

1. Open the **View Seasonal Events** portlet.
2. Select a specific configuration or **ALL** in the configuration table.
3. Select a seasonal event from the events table.
4. Right-click the seasonal event and select **Show Seasonal Event Graphs**.
5. In the Seasonal Event Graphs tab, you can choose to view all of the historical events for a seasonal event, or filter the historical events by selecting bars in a graph.
   a) To view all of the historical events for a seasonal event, select **Show Historical Events** in the **Actions** drop-down list.

b) To view the historical events for specific times, hold down the **Ctrl** key and click the specific bars in the graphs. Select **Show Historical Events for Selected Bars** in the **Actions** drop-down list.

Multiple bars that are selected from one chart are filtered by the OR condition. For example, if you select the bars for 9am or 5pm in the **Hour of the Day** graph, all of the events that occurred between 9am and 10am and all events that occurred between 5pm and 6pm are displayed in the Historical Event portlet.

Multiple bar that are selected from more than one graph are filtered by the AND condition. For example, if you select the bar for 9am in the **Hour of the Day** graph and Monday in the **Day of the Week** graph, all of the events that occurred between 9am and 10am on Mondays are displayed in the Historical Event portlet.

Results

The historical events are listed in the **Historical Event** portlet.

Exporting seasonal event graphs for a specified seasonal event to Microsoft Excel

You can export seasonal event graphs for a specified seasonal event to a Microsoft Excel spreadsheet from a supported browser.

Before you begin

You view seasonal event graphs for the seasonal events that are displayed in the **View Seasonal Events** portlet. To access the **View Seasonal Events** portlet, users must be assigned the `ncw_analytics_admin` role.

About this task

In addition to exporting seasonal event graphs to a Microsoft Excel spreadsheet, you also export the historical event data and seasonal event data and confidence levels for the seasonal event that you selected. Currently, there is no way to export only the seasonal event graphs.

Procedure

To export seasonal event graphs for a specified seasonal event to a Microsoft Excel spreadsheet, complete the following steps.

1. Open the **View Seasonal Events** portlet.
2. Select a specific configuration from the configuration table.
3. Select a seasonal event from the events table.
4. Right-click the seasonal event and select **Show Seasonal Event Graphs**.
5. From the **Actions** menu, select **Export Seasonal Event Graphs**.

   After a short time, the **Download export results** link displays.
6. Click the link to download and save the Microsoft Excel file.

Results

The Microsoft Excel file contains a spreadsheet with the following tabs:

- `Seasonal Data`: This tab contains the seasonal event data and confidence levels for the seasonal event that you selected.
- `Seasonality Charts`: This tab contains the seasonal event graphs for the seasonal event that you selected.
- `Historical Events`: This tab contains the historical event data for the seasonal event that you selected.
- `Export Comments`: This tab contains any comments relating to the export for informational purposes (for example, if the spreadsheet headers are truncated, or if the spreadsheet rows are truncated).

For more information about charts and threshold levels, see the "Seasonal Event Graphs" on page 308 topic.

Editing confidence thresholds of Seasonal Event Graphs

You can edit the default confidence level thresholds of the Seasonal Event Graphs.

About this task

The confidence level of the data in the charts is displayed in two ways:

1. The overall distribution score of each chart is displayed as high (red), medium (orange), or low (green) seasonality at the top of each chart.
2. The degree of deviation of the events is indicated by the high (red) and medium (orange) seasonality threshold lines on the charts.
3. The maximum confidence level of each bar is displayed as high (red), medium (orange), or low (green).

The default confidence level thresholds are as follows:

- High: 99-100%
- Medium: 95-99%
- Low: 0-95%

To modify the default confidence level thresholds of the charts, see "Editing confidence thresholds of Seasonal Event Graphs" on page 311.

Procedure

To edit the default confidence level threshold, complete the following steps:

1. To generate a properties file from the command-line interface, use the following command:

   ```
   nci_trigger server <UserID>/<password> NOI_DefaultValues_Export
    FILENAME directory/filename
   ```

   where

   SERVER
   The server where Event Analytics is installed.

   <UserID>
   The user name of the Event Analytics user.

   <password>
   The password of the Event Analytics user.

   directory
   The directory where the file is stored.

   filename
   The name of the properties file.

   For example:

   ```
   ./nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Export
    FILENAME
    /tmp/seasonality.props
   ```

2. To modify the confidence level thresholds, edit the default values of the following parameters:

   - `level_threshold_high = 99`
   - `level_threshold_medium = 95`
   - `level_threshold_low = 0`

   Note: Other property values are overwritten by the generated properties file. You might need to update other property values. For a full list of properties, see "Generated properties file" on page 256.
3. To import the modified properties file into Netcool/Impact, use the following command:

```
nci_trigger SERVER <UserID>/<password> NOI_DefaultValues_Configure
 FILENAME
directory/filename
```

For example:

```
./nci_trigger NCI impactadmin/impactpass NOI_DefaultValues_Configure
 FILENAME
 /tmp/seasonality.props
```

Historical events

You can view a list of historical events for one or more seasonal events in the table that displays in the **Historical Event** portlet. You can also export the data associated with the list of historical events for the associated seasonal events to a spreadsheet.

The **Historical Event** portlet displays a table with the following default columns:

Summary
Displays the description of the historical event.

Node
Displays the managed entity from which the historical event originated. The managed entity could be a device or host name, service name, or other entity.

Severity
Displays the severity of the historical event. The following list identifies the possible values that can display in the **Severity** column:

- 0: Clear
- 1: Indeterminate
- 2: Warning
- 3: Minor
- 4: Major
- 5: Critical

FirstOccurrence
Displays the date and time in which the historical event was created or first occurred. The date is expressed as *month*, *day*, *year*. The time is expressed as *hours*`:`, *minutes*`:`, *seconds*. The time also indicates whether AM or PM. For example: `Apr 13, 2015 4:45:17 PM`.

LastOccurrence
Displays the date and time in which the historical event was last updated. The date is expressed as *month*, *day*, *year*. The time is expressed as *hours*`:`, *minutes*`:`, *seconds*. The time also indicates whether AM or PM. For example: `Jun 2, 2015 5:54:49 PM`.

Acknowledged
Indicates whether the historical event has been acknowledged:

- 0: No
- 1: Yes

The historical event can be acknowledged manually or automatically by setting up a correlation rule.

Tally
Displays an automatically maintained count of the number of historical events associated with a seasonal event.

Viewing historical events for a seasonal event

You can view a list of historical events for a seasonal event in the table that displays in the **Historical Event** portlet.

Before you begin

You view seasonal events for which you want a list of historical events in the **View Seasonal Events** portlet. To access the **View Seasonal Events** portlet, users must be assigned the `ncw_analytics_admin` role.

Procedure

To view a list of historical events for a seasonal event, complete the following steps.

1. Open the **View Seasonal Events** portlet.
2. Select a specific configuration or **ALL** from the configuration table.
3. Select a seasonal event from the events table.
4. Right-click the seasonal event and select **Show Historical Events**.

Results

The historical events are listed in the table that displays in the **Historical Event** portlet.

Exporting historical event data

You can export historical event data to a spreadsheet from `Firefox` or `Internet Explorer`.

Before you begin

You first view seasonal events for which you want a list of historical events in the **View Seasonal Events** portlet. To access the **View Seasonal Events** portlet, users must be assigned the `ncw_analytics_admin` role.

Procedure

To export historical event data, complete the following steps.

1. Open the **View Seasonal Events** portlet.
2. Select a specific configuration or **ALL** from the configuration table.
3. Select a seasonal event from the events table.
4. Right-click the seasonal event and select **Show Historical Events**.

   The historical events are listed in the table that displays in the **Historical Event** portlet.
5. Select one or more historical events from the table that displays in the **Historical Event** portlet.
6. To copy the selected historical events:
   a) In `Firefox`, to copy the data from the displayed clipboard click **Ctrl+C** followed by **Enter**.
   b) In `Internet Explorer`, to copy the data from the displayed clipboard right-click on the selected historical event and select **Copy Ctrl+C** from the drop down menu.
7. Paste the historical event data to your spreadsheet.

Exporting historical event data for a specified seasonal event to Microsoft Excel

You can export historical event data for a specified seasonal event to a Microsoft Excel spreadsheet from a supported browser.

Before you begin

You first view seasonal events for which you want a list of historical events in the **View Seasonal Events** portlet. To access the **View Seasonal Events** portlet, users must be assigned the `ncw_analytics_admin` role.

About this task

In addition to exporting historical event data to a Microsoft Excel spreadsheet, you also export the seasonal event charts and seasonal event data and confidence levels for the seasonal event that you selected. Currently, there is no way to export only the historical event data.

Procedure

To export historical event data to a Microsoft Excel spreadsheet, complete the following steps.

1. Open the **View Seasonal Events** portlet.
2. Select a specific configuration from the configuration table.
3. Select a seasonal event from the events table.
4. Right-click the seasonal event and select **Show Seasonal Event Graphs**.
5. From the **Actions** menu, select **Export Seasonal Event Graphs**.

   After a short time, the **Download export results** link displays.
6. Click the link to download and save the Microsoft Excel file.

Results

The Microsoft Excel file contains a spreadsheet with the following tabs:

- `Seasonal Data`: This tab contains the seasonal event data and confidence levels for the seasonal event that you selected.
- `Seasonality Charts`: This tab contains the seasonal event graphs for the seasonal event that you selected.
- `Historical Events`: This tab contains the historical event data for the seasonal event that you selected.
- `Export Comments`: This tab contains any comments relating to the export for informational purposes (for example, if the spreadsheet headers are truncated, or if the spreadsheet rows are truncated).

Related events

Use the related events function to identify and show events that are historically related and to deploy chosen correlation rules, which are derived from related events configurations. You can create a pattern based on a related events group. The pattern applies the events in the group, which are specific to a resource, to any resource.

The related events function is accessible through three portlets.

- The Configure Analytics portlet. Use this portlet to create, modify, run, and delete related events configurations.
- The View Related Events portlet. Use this portlet to review the events and event groups that are derived from a related events configuration and to deploy correlation rules.
- The Related Event Details portlet. Use this portlet to access more detail about an event or an event group.

To access the View Related Events portlet, users must be assigned the `ncw_analytics_admin` role.

The related events function uses an algorithm with the event database columns you select to determine relationships between events.

Related events find signatures and patterns that occur together in the historic event stream. This discovery allows subject matter experts to easily review the detected signatures and derive correlation rules from related events configurations, without having to write correlation triggers or policies.

This diagram shows the relationship between the components for the related events functions.

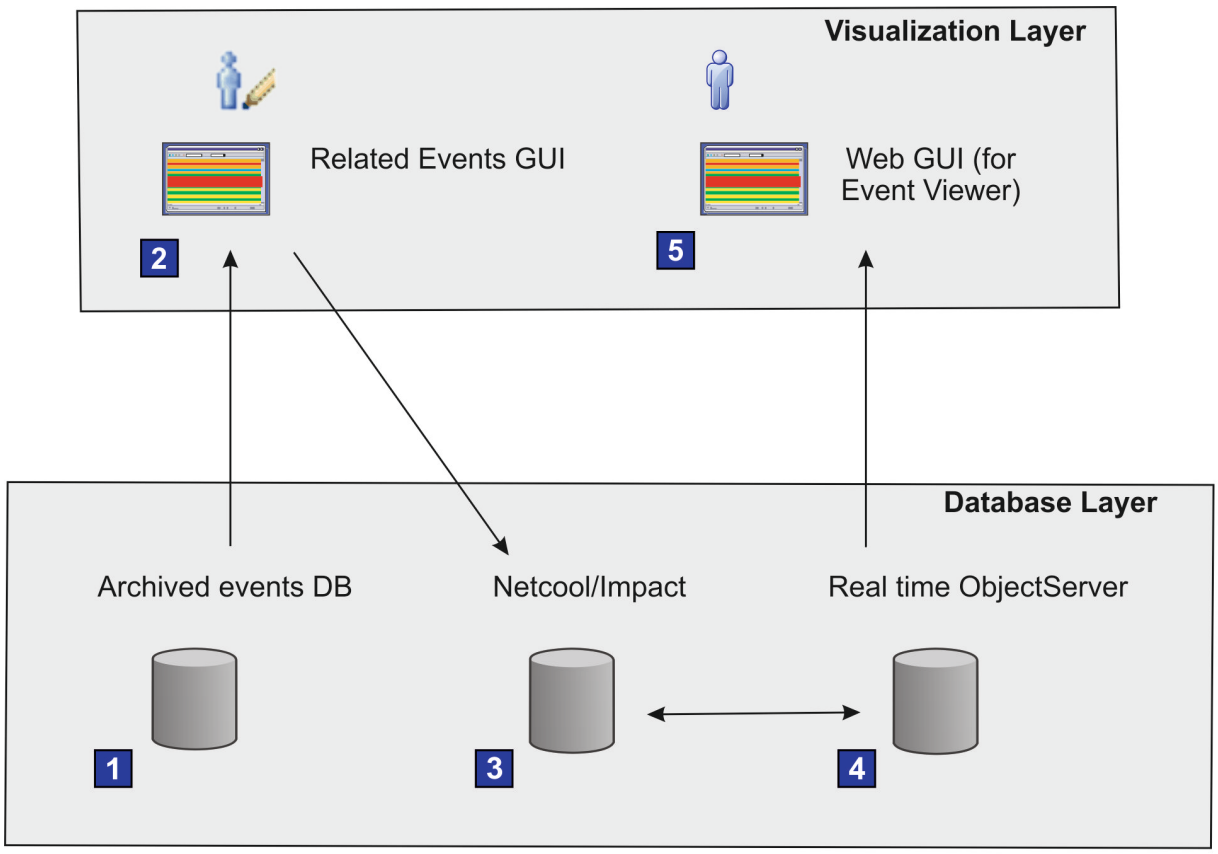

1. *Netcool/OMNIbus continually archives real-time events to an archived events database.*
2. *The Administrator creates a related events configuration. The configuration identifies and groups related events from the archive database and derives correlation rules. The Administrator watches and deploys the rules or configures the configuration to automatically deploy the rules.*
3. *Netcool/Impact policies are automatically created from the deployed correlation rules.*
4. *Netcool/Impact policies take action on real time events and group child events under a synthetic parent event.*
5. *The Operator is presented with a reduced number of events in the Event Viewer.*

Figure 10. Related events architecture overview

Work with related events

Use the View Related Events portlet to work with related events and related event groups that are derived from your related events configuration.

To access the View Related Events portlet, users must be assigned the `ncw_analytics_admin` role.

In the Configuration, Group, or Event tables you can right-click on a group, a configuration, or the `All` container and a menu is displayed. The menu lists some of the following actions for you to select.

Watch For more information about this action, see "Watching a correlation rule" on page 327.
Deploy For more information about this action, see "Deploying a correlation rule" on page 328.
Archive For more information about this action, see "Archiving related events" on page 324.
Delete This action is only available from within the **Archived** tab. If you want to delete event groups from the system, choose this action.

Reset performance statistics For more information about this action, see "Viewing performance statistics for a correlation rule" on page 329.

New This action is only available from within the **Archived** tab. If you choose this action, your selected row reinstates into the **New** tab.

Copy Choose this action if you want to copy a row, which you can then paste into another document.

Within the View Related Events portlet, in the **New**, **Watched**, **Active**, **Expired**, or **Archived** tabs, four tables display information about your related events.

Configuration table
Displays a list of the related event configurations.

Group Sources table
Displays the source information for related event groups based on the configuration and created patterns.

Groups table
Displays the related event groups for a selected configuration.

Events table
Displays the related events for a selected configuration or a selected group.

A performance improvement implemented in V1.6.0.1 ensures that the View Related Events portlet displays Events, Groups, and Groups Sources more quickly once an item is selected. As part of this update, each tab in the View Related Events portlet now lists all configurations in the panel on the left of the portlet following the successful run of a configuration. Configurations are displayed in the panel even if there are no events or groups in a particular state for a given configuration. If no data exists for a particular state, the panels will display a **No items to display** message. The configuration will be listed in all five tabs, **New**, **Watched**, **Active**, **Expired**, and **Archived**.

Right-click on a configuration in the **Configuration** table to display a list of menu items. You can select the following actions from the menu list.

Watch For more information about this action, see "Watching a correlation rule" on page 327.

Deploy For more information about this action, see "Deploying a correlation rule" on page 328.

Archive For more information about this action, see "Archiving related events" on page 324.

Copy Choose this action if you want to copy a row, which you can then paste into another document.

Right-click on a pattern in the **Group Sources** table to display a list of menu items. You can select the following actions from the menu list.

Edit Pattern For more information about this action, see "Editing an existing pattern" on page 340.

Delete Pattern For more information about this action, see "Deleting an existing pattern" on page 340.

Copy Choose this action if you want to copy a row, which you can then paste into another document.

Right-click on a group name in the **Groups** table to display a list of menu items. You can select the following actions from the menu list.

Show details For more information about this action, see "Viewing related events details for a seasonal event" on page 317.

Create Pattern For more information about this action, see "Creating patterns" on page 331.

Unmark as reviewed For more information about this action, see "Marking a related events group as reviewed" on page 318.

Mark as reviewed For more information about this action, see "Marking a related events group as reviewed" on page 318.

Watch For more information about this action, see "Watching a correlation rule" on page 327.

Deploy For more information about this action, see "Deploying a correlation rule" on page 328.

Archive For more information about this action, see "Archiving related events" on page 324.

Delete This action is only available from within the **Archived** tab. If you want to delete event groups from the system, choose this action.

Reset performance statistics For more information about this action, see "Viewing performance statistics for a correlation rule" on page 329.

New This action is only available from within the **Archived** tab. If you choose this action, your selected row reinstates into the **New** tab.

Copy Choose this action if you want to copy a row, which you can then paste into another document.

Right-click on an event in the **Events** table to display a list of menu items. You can select the following actions from the menu list.

Show details For more information about this action, see "Viewing related events details for a seasonal event" on page 317.

Copy Choose this action if you want to copy a row, which you can then paste into another document.

Within the View Related Events portlet, you can also complete the following types of tasks.

- View related events.
- View related events by group.
- Sort a related events view.
- View performance statistics for a deployed correlation rule.

Within the Related Event Details portlet, you can also complete the following types of tasks.

- Change the pivot event.
- Work with correlation rules and related events.
- View events that form a correlation rule.
- Select a root cause event for a correlation rule

Viewing related events

In the View Related Events portlet, you can view a listing of related events as determined by related events configurations that ran.

Procedure

1. Log in to the Dashboard Application Services Hub as a user with the `ncw_analytics_admin` role.
2. In the Dashboard Application Services Hub navigation menu, go to the **Insights** menu.
3. Under **View Analytics**, select **View Related Events**.
4. By default, within the View Related Events portlet the **New** tab opens, this tab lists related events with a status of New.

What to do next

If you want to see related events with another status, select the relevant toolbar button within the View Related Events portlet toolbar.

Viewing related events details for a seasonal event

You can view related event details for a seasonal event in the Related Event Details portlet.

Before you begin

To access the Seasonal Event Rules portlet, users must be assigned the `ncw_analytics_admin` role.

Procedure

To view a list of historical events for a seasonal event, complete the following steps.

1. Open the **View Seasonal Events** portlet.
2. Select a specific configuration or **ALL** in the configuration table.
3. Select a seasonal event in the events table.
4. Right-click the seasonal event and select **Show Related Event Details**.

Results
The Related Event Details portlet displays the related event details.

Viewing related events by group
From the full list of related events, you can view only the related events that are associated to a specific group.

About this task

A related events configuration can contain one or many related events groups. A related events group is determined by a related events configuration and a related events group can be a child of one or more related events configurations.

Note:

- Discovered Groups and any "Suggested patterns" on page 339 are displayed in the **Group Sources** table of the View Related Events portlet. Any groups that are covered by a suggested pattern will not appear under the list of groups associated with Discovered Groups. A group that is a member of a suggested patten will only show up under the events of Discovered Groups once the suggested patten is deleted.
- You might see a different number of **Unique Events** in a related events group when a **Relationship Profile** of `Strong` has been selected for the events in the configuration. This is caused by the same events being repeated more than once.

Procedure

1. Start the View Related Events portlet, see "Viewing related events" on page 317.
2. Within any tab, in the **Configuration** table, expand the root node `All`. The list of related events configurations display.
3. In the **Configuration** table, select a related events configuration. The list of related events groups display in the **Group Sources** and **Groups** tables and the related events display in the **Events** table.
4. In the **Group** table, select a group. The **Event** table updates and displays only the events that are associated to the selected group.

Viewing related events in the Event Viewer
To see the grouping of related events in the **Event Viewer**, you must apply a view that uses the **IBM Related Events** relationship.

Procedure

1. Open the **Event Viewer**.
2. Click **Edit Views**.
3. Select the **Relationships** tab.
4. Select **IBM Related events** from the drop-down menu.
5. Click **Save**.

Results
This relationship is used to present the results of correlations generated by the Related Events Analytics functionality.

Marking a related events group as reviewed
The review status for a related events group can be updated in the View Related Events portlet.

About this task

In the View Related Events portlet, within the group table you can modify the review status for a related events group. The review status values that are displayed indicates to Administrators whether related events groups are reviewed or not.

In the View Related Events portlet, in the **Groups** table you can modify the review status for a related events group. The review status values that are displayed indicates to Administrators the review status of the related events groups.

Related events groups can display these review status values.

Yes. The group is reviewed.
No. The group is not reviewed.

To mark a related events group as reviewed or not reviewed, complete the following steps.

Procedure

1. View related events, see "Viewing related events" on page 317.
2. In the View Related Events portlet, within the group table, select a line item, which represents a group, and right-click. A menu is displayed.
3. From the menu, select **Mark as Reviewed** or **Unmark as Reviewed**. A success message in a green dialog box displays.

Results

The values in the Reviewed column are updated, to one of the following values Yes, No.

When you enable sorting for the group table, you can sort on the Yes or No values.

Sorting a related events view
Within a related events view, it is possible to sort the information that is displayed.

Before you begin
Within the View Related Events portlet, select the tab view where you want to apply the sorting.

About this task

Sorting by single column or multiple columns is possible within either the **Configuration**, **Group** or **Event** table. Sorting within the **Group** or **Event** table can be done independently or in parallel by using the sorting arrows that display in the table column headings. When you apply sorting within the **Configuration** table the configuration hierarchy disappears, but the configuration hierarchy reappears when you remove sorting. For more details about rollup information, see "Adding columns to seasonal and related event reports" on page 265.

Sorting by single column or multiple columns is possible within the **Configuration**, **Group Sources**, **Groups** or **Event** table. Sorting within the **Groups** or **Event** table can be done independently or in parallel by using the sorting arrows that display in the table column headings. When you apply sorting within the **Configuration** table the configuration hierarchy disappears, but the configuration hierarchy reappears when you remove sorting. For more details about rollup information, see "Adding columns to seasonal and related event reports" on page 265.

Procedure

1. In either the **Configuration**, **Group** or **Event** table, hover the mouse over a column heading. Arrows are displayed, hover the mouse over the arrow, one of the following sort options is displayed.

   Click to sort Ascending
   Click to sort Descending
   Do not sort this column
2. In either the **Configuration**, **Group Sources**, **Groups** or **Event** table, hover the mouse over a column heading. Arrows are displayed, hover the mouse over the arrow, one of the following sort options is displayed.

   Click to sort Ascending

Click to sort Descending
Do not sort this column

3. Left-click to select and apply your sort option, or left click a second or third time to view and apply one of the other sort options.
4. For sorting by multiple column, apply a sort option to other column headings. Sorting by multiple columns is not limited, as sorting can be applied to all columns.

Results
The ordering of your applied sort options, is visible when you hover over column headings. The sorting options that you apply are not persistent across portlet sessions when you close the portlet the applied sorting options are lost.

Filtering related events
Filtering capability is possible on the list of related events within the View Related Events portlet.

Procedure

1. Start the View Related Events portlet, see "Viewing related events" on page 317
2. Within the toolbar, in the filter text box, enter the filter text that you want to use. Filtering commences as you type.

Results
The event list is reduced to list only the events that match the filter text in at least one of the displayed columns.

What to do next
To clear the filter text, click the **x** in the filter text box. After you clear the filter text, the event list displays all events.

Exporting related events for a specific configuration to Microsoft Excel
You can export related events for a specific configuration to a Microsoft Excel spreadsheet from a supported browser.

Before you begin
You view related events for one or more configurations in the **View Related Events** portlet. To access the **View Related Events** portlet, users must be assigned the `ncw_analytics_admin` role.

Procedure

To export related events for a specific configuration to a Microsoft Excel spreadsheet, complete the following steps.

1. Open the **View Related Events** portlet.
2. Select a specific configuration from the configuration table.
3. Click the **Export Related Events** button in the toolbar.
   After a short time, the **Download export results** link displays.
4. Click the link to download and save the Microsoft Excel file.

Results

The Microsoft Excel file contains a spreadsheet with the following tabs:

- `Report Summary`: This tab contains a summary report of the configuration that you selected.
- `Groups Information`: This tab contains the related events groups for the configuration that you selected.

- `Groups Instances`: This tab contains a list of all the related events instances for all of the related events groups for the configuration that you selected.
- `Group Events`: This tab contains a list of all the events that occurred in the related events groups for the configuration that you selected.
- `Instance Events`: This tab contains a list of all the events that occurred in all of the related events instances for all the related events groups for the configuration that you selected.
- `Export Comments`: This tab contains any comments relating to the export for informational purposes (for example, if the spreadsheet headers are truncated, or if the spreadsheet rows are truncated).

Exporting selected related events groups to Microsoft Excel
You can export related events groups for a specific configuration to a Microsoft Excel spreadsheet from a supported browser.

Before you begin
You view related events for one or more configurations in the **View Related Events** portlet. To access the **View Related Events** portlet, users must be assigned the `ncw_analytics_admin` role.

Procedure

To export related events groups for a specific configuration to a Microsoft Excel spreadsheet, complete the following steps.

1. Open the **View Related Events** portlet.
2. Select a specific configuration from the configuration table.
3. Select multiple related event groups by using the `Ctrl` key and select method. (You can also select multiple related events groups by using the click and drag method.)
4. After selecting multiple related event groups, right click on one of the selected groups and select the **Export Selected Groups** button in the toolbar.

   After a short time, the **Download export results** link displays.
5. Click the link to download and save the Microsoft Excel file.

Results

The Microsoft Excel file contains a spreadsheet with the following tabs:

- `Report Summary`: This tab contains a summary report of the configuration that you selected.
- `Groups Information`: This tab contains the related events groups for the configuration that you selected.
- `Groups Instances`: This tab contains a list of all the related events instances for all of the related events groups for the configuration that you selected.
- `Group Events`: This tab contains a list of all the events that occurred in the related events groups for the configuration that you selected.
- `Instance Events`: This tab contains a list of all the events that occurred in all of the related events instances for all the related events groups for the configuration that you selected.
- `Export Comments`: This tab contains any comments relating to the export for informational purposes (for example, if the spreadsheet headers are truncated, or if the spreadsheet rows are truncated).

Expired related events
An active related event group expires if no live events matching that related event group come into the Object Server for a period of six months. This six-month period is known as the *automated expiry time*. When the automated expiry time is reached for an active related events group, the group and its related events are moved to the **Expired** tab within the View Related Events portlet.

Even though the expired groups and related events are visible in the **Expired** tab, you must acknowledge that the group is expired. In the **Expired** tab, right-click on the group that you want to acknowledge. A menu is displayed, from the menu select **Validate**. The default automated expiry time for an active

configuration is six months. To change the expiry time see “Changing the expiry time for related events groups” on page 294.

To perform other actions on related events within the **Expired** tab, right-click on a group or event and a menu is displayed. From the menu, select the action that you want to take.

Impacts of newly discovered groups on existing groups
An existing related events group can be replaced by newly discovered related event group, or the newly discovered group can be ignored.

The following bullet points describe the Event Analytics functions for management of newly discovered groups and existing groups.

- If a newly discovered group is a subset or same as an existing group within `Watched`, `Active`, `Expired`, or `Archived`, then Event Analytics ignores the newly discovered group.
- If a newly discovered group is a superset of an existing group within `New`, then Event Analytics deletes the existing group and displays the newly discovered group in `New`. Otherwise, no changes occur with the existing group.
- If a newly discovered group is a superset of an existing group within `Watched`, `Active`, or `Expired`, then Event Analytics moves the existing group to `Archived`, and displays the newly discovered group in `Watched`, `Active`, or `Expired`.
- If a newly discovered group is a superset of an existing group within `Archived`, then Event Analytics adds the newly discovered group in to `New` and leaves the `Archived` group where it is.

Extra details about related events
Use the Related Event Details portlet to access extra details about related events.

Within the Related Event Details portlet, you can complete the following types of tasks.

- View the occurrence time of an event.
- Switch between tabulated and charted event information.
- Remove an event from a related events group.

Only one instance of the Related Event Details portlet can be open at any stage. If you select **Show Details** for an event or an event group in the View Related Events portlet and the Related Event Details portlet is already open, the detail in the Related Event Details portlet refreshes to reflect your selected event or event group.

Viewing the occurrence time of an event
You can get details about the time that an event occurred.

About this task
When you look at the occurrence time of an event, you might be able to relate this event to some other events that occurred around the same time. Within a particular event group, the same event might occur multiple times. For example, if the group occurs 10 times within the time period over which the related events report is run, then there are 10 instances of the group. The event might occur in each of those group instances, resulting in 10 occurrence times for that event. Events in strong related events groups appear in all group instances, but events in medium or weak related events groups might appear in a subset of the group instances. This information is visible in the Related Event Details portlet by switching between different instances in the **Event Group Instance Table** as explained in the following procedure.

Procedure

1. Start the View Related Events portlet, see “Viewing related events” on page 317.
2. Within the View Related Events portlet, in the **Events** table select an event or in the **Group** table select an event group, and right-click. A menu is displayed.
3. From the menu, select **Show Details** and a Related Event Details portlet opens.
4. Within the Related Event Details portlet, in the **Events** tab, two tables are displayed.

- **Event Group Instance Table**: This table lists each instance of the event group and the time at which the instance occurred. The time of the group instance is set to the occurrence time of the event that you selected.
  - The `Unique Events` column shows the number of unique events for each group instance.
- **Events Table**: This table lists the events for a selected group instance and the occurrence time of each event.
  - The `Offset` column displays `Not Applicable` if the pivot event is not in the selected group instance. However, if the pivot event is in the selected group instance the `Offset` column displays an offset time that is related to the pivot event. For more information about the pivot event, see "Changing the pivot event" on page 325.
  - The `Instances` column shows the number of group instances each event participates in. If an event is deleted and then reoccurs, it has a different server serial value. For a strong profile, all events occur in all instances, for example, a value of 3/3 in the `Instances` column denotes that the event occurred in all three instances.

What to do next
Within the **Event Group Instance Table**, select another instance of the event group. The **Events Table** now displays the events in the newly selected group instance.

Switching between tabulated and charted event information
Event information is visible in table or chart format, within the Related Event Details portlet.

About this task
A pivot event is an event that acts as a pivot around which you can extrapolate related event occurrences, in relation to the pivot event occurrence. To view the event distribution of a pivot event, complete the following steps to switch from tabulated event information to charted event information within the Related Event Details portlet.

Procedure

1. Start the View Related Events portlet, see "Viewing current analytics configurations" on page 289.
2. Within the View Related Events portlet, in the **Events** table select an event or in the **Group** table select an event group, and right-click. A menu is displayed.
3. From the menu, select **Show Details** and a Related Event Details portlet opens.
4. Within the Related Event Details portlet, in the **Events** tab, two tables are displayed.

   Event Group Instance Table: This table lists each instance of the event group and the time at which the instance occurred. The time of the group instance is set to the occurrence time of the event that you selected.
   Events Table: This table lists the events for a selected group instance.
5. From the **Events** tab toolbar, select **Timeline**. The event information displays in chart format.

   For information about understanding the timeline chart, see "Understanding the timeline chart" on page 324

   Note: The timeline chart scale is displayed as **seconds** (**s**), **minutes** (**m**), or **hours** (**h**). If many timelines are displayed, users might need to scroll down to view all of the timelines. The timeline chart scale is anchored in place at the top of the **Timeline** view.

Results

The timeline chart shows the event distribution for each event type in the group, relative to the pivot event. Each comb in the timeline chart represents an event type and the teeth represent the number of instances of the event type. The pivot event is not represented by a comb, but the pivot event instance is always at zero seconds, minutes, or hours. For a selected event group instance, the highlighted tooth on each comb denotes the event type instance relative to the pivot event, in seconds, minutes, or hours.

The long summary labels under the combs in the timeline chart are truncated. Move the mouse cursor over a truncated summary label to see the tooltip that shows the full summary label.

What to do next

- If there are many event types to view, use the pagination function in addition to scrolling. In the pagination toolbar, select the page to view and the page size.
- If you want to change the pivot event, see "Changing the pivot event" on page 325.
- If you want to revert to the tabulated event information, select **Events** from the **Events** tab toolbar.

Understanding the timeline chart
Event information in the Related Event Details portlet is available in chart format.

You can use the Related Event Details portlet to view more information about related events. For example, you can view charted event information on a timeline chart. For more information about how to view the charted event information, see "Switching between tabulated and charted event information" on page 323

The timeline chart shows the event distribution for each event type in the group, relative to the pivot event. The pivot event is always at zero seconds, minutes, or hours.

Each comb in the timeline chart represents an event type and the teeth represent the number of instances of the event type. The blue event markers represent all the times the event occurred relative to the pivot event. The red event markers indicate the time that the event occurred in the selected group instance.

Removing an event from a related events group
You can remove an event from a related events group.

About this task
When you believe that an event is no longer related to other events in the related events group, you can remove the event from that events group. When you remove an event from a related events group, the event is hidden from the UI and the correlation process but the event is not deleted from the system. Complete the following steps to remove an event from a related events group.

Procedure

1. View event groups and events in the View Related Events portlet, see "Viewing related events" on page 317 and "Viewing related events by group" on page 318.
2. Within the View Related Events portlet, in the **Group** table select an event group or in the **Event** table select an event, and right click. A menu is displayed.
3. From the menu, select **Show Details** and a Related Event Details portlet opens.
4. Within the Related Event Details portlet, in either the **Events** tab on the events table or in the **Correlation Rule** tab, right-click on an event. A menu is displayed.
5. From the menu, select **Remove Event**.
6. A confirmation message is displayed, select **Yes** or **No**.

Results
The event is removed from the group and no longer appears in the event list in either the **Events** tab and the **Correlation Rule** tab.

Archiving related events
You can archive related events by archiving the related events group.

Before you begin
View event groups and events in the View Related Events portlet, see "Viewing related events" on page 317 and "Viewing related events by group" on page 318.

About this task
When you believe that events within a related events group are no longer relevant, you can archive that group. Complete the following steps to archive a related events group.

Procedure

- To archive a related events group within the View Related Events portlet, complete the following steps.
  1. Within the View Related Events portlet, select the **New**, **Watched**, **Active**, or **Expired** tab.
  2. In your chosen tab, within the **Group** table select an event group and right click. A menu is displayed.
  3. From the menu, select **Archive**.
- To archive a related events group within the Related Event Details portlet, complete the following steps.
  1. Within the View Related Events portlet, select the **New**, **Watched**, **Active**, or **Expired** tab.
  2. In your chosen tab, within the **Group** table select an event group or within the **Event** table select an event, and right click. A menu is displayed.
  3. From the menu, select **Show Details**, a Related Event Details portlet opens.
  4. Within the Related Event Details portlet, in either the **Events** or **Correlation Rule** tab, select **Archive**. A success message is displayed.

Results
The related events group moves into the **Archived** tab in the View Related Events portlet.

What to do next
Within the **Archived** tab, from the list of archived groups you can select a group and right click. A menu is displayed with a choice of tasks for your selected group.

- If you want to move a group out of the **Archived** tab and into the **New** tab, from the menu select **New**. A number of actions can be performed with groups and events within the **New** tab, see "Work with related events" on page 315.
- If you want to delete a related events group from the system, from the menu select **Delete**. This is the only way to delete a related events group from the system.

Changing the pivot event
You can change a pivot event to view related events that are of interest.

About this task
Use a pivot event as a baseline to determine a sequence of events in the group. A pivot event displays in the Related Event Details portlet. Within the Related Event Details portlet, a pivot event can be changed. Also, a pivot event history, of the 20 most recent pivot events, is available for you to revisit.

- When you open the Related Event Details portlet from an event in the View Related Events portlet, that event becomes the pivot event within the Related Event Details portlet.
- When you open the Related Event Details portlet from a group in the View Related Events portlet, one of the events from that group becomes the pivot event within the Related Event Details portlet. The pivot event is not always the parent event.

Complete the following steps to change the pivot event.

Procedure

1. Within the View Related Events portlet, right-click on an event or a group. A menu is displayed.
2. From the menu, select **Show Details**. The Related Event Details portlet opens.
3. Within the Related Event Details portlet, information about the pivot event is displayed.

- In the Event Group Instance table, the `Contains Pivot Event` column reports if a group instance has a pivot event, or not. Some groups might not have a Pivot Event set because the event identity is different for these events.
- In the Events table, the pivot event is identifiable by a red border.
- Next to the `Group Name` entry, a **Pivot Event** link displays. To see more details about the pivot event, click the **Pivot Event** link and a **More Information** widow opens displaying details about the pivot event.

4. In the Related Event Details portlet, within the **Events** tab, in the **Events** table, right-click on the event you want to identify as the pivot event. A menu is displayed.
5. From the menu, select **Set as Pivot Event**.

Results
Your selected event becomes the pivot event with a red border. Data updates in the timeline chart, in the **Pivot Event** link, in the **Event Group Instance** table and in the **Events** table.

What to do next
Within the Related Event Details portlet, you can reselect one of your 20 recent pivot events as your current pivot event. From the **Events** tab toolbar, select either the forward arrow or back arrow to select one of the 20 recent pivot events.

Correlation rules and related events
A correlation rule is a mechanism that enables automatic action on real-time events that are received by the ObjectServer, if a trigger condition is met. The result is fewer events in the Event Viewer for the operator to troubleshoot.

Writing a correlation rule in code is complex but the related events function removes the need for administrators to code a correlation rule. Instead, the related events function derives a correlation rule from your related events configuration and deploys a correlation rule, all through the GUI. After the correlation rule is deployed in a live environment and if the trigger condition is met, then automatic action occurs.

- The trigger condition is the occurrence of one or more related event types, from an event group, on the Tivoli Netcool/OMNIbus ObjectServer. Only one event must be the parent event. Related event types are derived from your related events configuration.
- The automatic action is the automatic creation of a synthetic event with some of the properties of the parent event, and the automatic grouping of the event group events under this synthetic event.

Viewing events that form a correlation rule
You can view the related events that form a correlation rule in the Related Event Details portlet.

About this task
Administrators can view related events that form a correlation rule to understand associations between events. Complete the following steps to view related events that form a correlation rule.

Procedure

1. Start the View Related Events portlet, see "Viewing current analytics configurations" on page 289.
2. Within the View Related Events portlet, in the **Events** table select an event or in the **Group** table select an event group, and right-click. A menu is displayed.
3. From the menu, select **Show Details** and a Related Event Details portlet opens.
4. In the Related Event Details portlet, select the **Correlation Rule** tab.

Results
A table displays with a list of the related events that make up the correlation rule.

Selecting a root cause event for a correlation rule
You can select the root cause event for the correlation rule

About this task
When you select the root cause event for the correlation rule, the selected event becomes a parent event. A parent synthetic event is created with some of the properties from the parent event and a parent-child relationship is created between the parent synthetic event and the related events. When these events occur in a live environment, they display in the Event Viewer within a group as child events of the parent synthetic event. With this view of events, you can quickly focus on the root cause of the event, rather than looking at other related events.

To select the root cause event for the correlation rule, complete the following steps. If you want to see automated suggestions about the root cause event for a group, see configuration details in "Adding columns to seasonal and related event reports" on page 265.

Procedure

1. View all events that form a correlation rule, see "Viewing events that form a correlation rule" on page 326
2. In the Related Event Details portlet, within the **Correlation Rule** tab, right-click an event and select `Use Values in Parent Synthetic Event`.

Results
The table in the **Correlation Rule** tab refreshes and the `Use Values in Parent Synthetic Event` column for the selected event updates to `Yes`, which indicates this event is now the parent event.

For a related events group, if all of the children of a parent synthetic event are cleared in the Event Viewer, then the parent synthetic event is also cleared in the Event Viewer. If another related event comes in for that same group, the parent synthetic event either reopens or re-creates in the Event Viewer, depending on the status of the parent synthetic event.

Watching a correlation rule
You can watch a correlation rule and monitor the rule performance before you deploy the rule for the rule to correlate live data.

Before you begin
Complete your review of the related events and the parent event that form the correlation rule. If necessary, change the correlation rule or related events configuration.

About this task
When you are happy with the correlation rule, you can choose to **Watch** the correlation rule.

When you choose to **Watch** the correlation rule, the rule moves out of its existing tab and into the **Watched** tab within the View Related Events portlet. While the rule is in **Watched**, the rule is not creating synthetic events or correlating but does record performance statistics. You can check the rule's performance before you deploy the rule for the rule to correlate live data.

Note: On rerun of a related events configuration scan, a warning message is displayed if any new groups are discovered that conflict with groups on which an existing watched rule is based.

- Click **OK** to accept the warning and continue watching the existing rule. The new groups are ignored.
- Click **Cancel** to ignore the warning and replace the existing watched rule with a new rule based on the newly discovered group.

Note: any NEW groups which conflict with existing non-NEW groupsany already existing patterns cannot be edited. Only newly discovered patterns can be edited.

Complete the following steps to **Watch** the correlation rule.

Procedure

- Within the View Related Events portlet, perform the following steps:
  a) View related events by group, see "Viewing related events by group" on page 318.
  b) In the View Related Events portlet, within the group table, select either a related events group or a related events configuration and right click. A menu is displayed.
  c) From the menu, select **Watch**.
- Within the Related Event Details portlet for a group or an event, perform the following steps:
  a) View related events or related event groups, see "Viewing related events" on page 317 and "Viewing related events by group" on page 318.
  b) Select an event or a related events group.
     – In the View Related Events portlet, within the group table, select a related events group and right click. A menu is displayed.
     – In the View Related Events portlet, within the event table, select an event and right click. A menu is displayed.
  c) From the menu, select **Show Details**. The Related Event Details portlet opens.
  d) In the Related Event Details portlet, within any tab, select **Watch**.

Results
The rule displays in the **Watched** tab.

What to do next
Within the **Watched** tab, monitor the performance statistics for the rule. When you are happy with the performance statistics consider "Deploying a correlation rule" on page 328.

Deploying a correlation rule
You can deploy a correlation rule, for the rule to correlate live data.

Before you begin
Complete your review of the related events and the parent event that form the correlation rule. If necessary, change the correlation rule or related events configuration.

About this task
When you are happy with the correlation rule, you can choose to **Deploy** the correlation rule.

When you choose to **Deploy** the correlation rule, the rule moves out of its existing tab and into the **Active** tab within the View Related Events portlet. **Active** rule algorithm works to identify the related events in the live incoming events and correlates them so the operator knows what event to focus on. Performance statistics about the rule are logged which you can use to verify whether the deployed rule is being triggered.

Note: On rerun of a related events configuration scan, a warning message is displayed if any new groups are discovered that conflict with groups on which an existing deployed rule is based.

- Click **OK** to accept the warning and continue deploying the existing rule. The new groups are ignored.
- Click **Cancel** to ignore the warning and replace the existing deployed rule with a new rule based on the newly discovered group.

Complete the following steps to **Deploy** the correlation rule.

Procedure

- Within the View Related Events portlet.
  a) View related events by group, see "Viewing related events by group" on page 318.

b) In the View Related Events portlet, within the groups table, select either a related events group or a related events configuration and right click. A menu is displayed.

c) From the menu, select **Deploy**.

- Within the Related Event Details portlet for a group or an event.

a) View related events or related event groups, see "Viewing related events" on page 317 and "Viewing related events by group" on page 318.

b) Select an event or a related events group.

– In the View Related Events portlet, within the groups table, select a related events group and right click. A menu is displayed.

– In the View Related Events portlet, within the events table, select an event and right click. A menu is displayed.

c) From the menu, select **Show Details**. The Related Event Details portlet opens.

d) In the Related Event Details portlet, within any tab, select **Deploy**.

Results
The rule moves out of the **New** tab and into the **Active** tab within the View Related Events portlet.

What to do next
When you establish confidence with the rules and groups that are generated by a related events configuration, you might want all the generated groups to be automatically deployed in the future. If you want all the generated groups to be automatically deployed, return to "Creating a new or modifying an existing analytics configuration" on page 290 and within the **Configure Related Events** window, tick the option `Automatically deploy rules discovered by this configuration`.

Viewing performance statistics for a correlation rule
You can view performance statistics for a correlation rule in the View Related Events portlet, within the **Watched**, **Active**, or **Expired** tabs.

Performance statistics in the group table

Times Fired: The total number of times the rule ran since the rule became active.
Times Fired in Last Month: The last month time period is counted as 30 days instead of a calendar month. The total number of times that the rule is fired in the current 30 days. Time periods are calculated from the creation date of the group.
Last Fired: The last date or time that the rule was fired.
Last Occurrence I: The percentage of events that occurred from the group, in the last fired rule.
Last Occurrence II: The percentage of events that occurred from the group in the second last fired rule.
Last Occurrence III: The percentage of events that occurred from the group in the third last fired rule.

Performance statistics in the event table

Occurrence: The number of times the event occurred, for all the times the rule fired.

Reset performance statistics

You can reset performance statistics to zero for a group in the **Watched**, **Active**, or **Expired** tabs. To reset performance statistics, right-click on the group name and from the menu select **Reset performance statistics**. A message displays indicating that the operation will reset statistics data for the selected correlation rule. The message also indicates that you will not be able to retrieve this data. Click `Yes` to continue with the operation or `No` to stop the operation. A success message displays after you select `Yes`.

Resetting performance statistics to zero for a group also causes the following columns to be cleared: **Times Fired**, **Times Fired in Last Month**, and **Last Fired**. Note that performance statistics are not

collected for the **Archived** tab. When a rule is moved between states, the performance statistics are reset. Every time an action is triggered by the rule the performance statistics increase.

Related Events statistics
When sending some events a synthetic event is created, but the statistics can appear not to be updated.

This is because there are delays in updating the related events statistics. These delays are due to the time window during which the related event groups are open, so that events can be correlated.

The statistics (Times Fired, Times Fired in last month, last fired) are updated only when the **Group Time to Live** has expired. The sequence is; synthetic event is triggered, action is done, and the statistics are calculated later.

Take the following query as an example:

```
SELECT GROUPTTL FROM RELATEDEVENTS.RE_GROUPS WHERE GROUPNAME = 'XXX';
```

There was an occurrence of the GROUPTTL being equal to 82800000 milliseconds, this is 23 hours. In this instance an update to the statistics wouldn't be visible to the user for 23 hours. If GROUPTTL is reduced to 10 seconds by running the following command:

```
UPDATE RELATEDEVENTS.RE_GROUPS SET GROUPTTL = 10000 WHERE GROUPNAME = 'XXX';
```

Subsequent tests will show that the statistics are updated promptly.

An algorithm creates GROUPTTL based on historical occurrences of the events. There is no default value for GROUPTTL and no best practice recommendation. GROUPTTL should be determined and set on a per case basis.

Data is displayed for the Times Fired, Times Fired in Last Month, and Last Fired columns for groups allocated to patterns. In previous versions, the group statistics were only updated for unallocated groups and not updated for groups allocated to a pattern. The new statistics represent the total occurrences for events with an active, watched or expired status.

The Times Fired value increments every time there is an event matching the Event Identifier of a deployed group. This statistic is suppressed by default. It can be enabled by turning on the **timesfired_group_stats_enable** property.

The Times Fired in Last Month value is the sum of events received in the last 30 days. This statistic is suppressed by default. It can be enabled by turning on the **timesfired_group_stats_enable** property. When the 30 days time period has passed and a new event comes in, this value resets back to 0. When an event is not firing for two or more months, the value persists and is not reset. In this case, the value in the **Times Fired in Last Month** column refers to the 30 days before the timestamp in the Last Fired column.

The Last Fired value is the timestamp of the last time such an event was received. This statistic is suppressed by default. It can be enabled by turning on the **timesfired_group_stats_enable** property.

Performance statistics in the group table

The statistics in the **Times Fired** column in the **Group Sources** panel represent the sum of incoming events, which match a given pattern. The value increments every time there is an incoming event that matches a pattern in a watched, active, or expired state. This statistic is active by default, and can't be turned off.

When a group is part of a pattern, the statistics in the **Times Fired** column in the **Groups** panel represent the sum of incoming events that are part of the group and match a given pattern. The value increments every time there is an event, which is part of an active group, that matches an active pattern. This functionality is controlled with the **timesfired_group_stats_enable** property. By default, this statistic is disabled. To enable this functionality, complete the following steps:

1. Export the current default IBM Netcool Operations Insight export property values to a file by running the following command:

```
./nci_trigger <NCI_Cluster_name>  <impactadmin_user_name>/<impactadmin_user_pwd>
NOI_DefaultValues_Export FILENAME <path/filename>
```

2. In the new *<path/filename>* file, change the **timesfired_group_stats_enable** value from false to true.
3. Import the updated properties file by running the following command:

```
./nci_trigger <NCI_Cluster_name> <impactadmin_user_name>/<impactadmin_user_pwd>
NOI_DefaultValues_Configure FILENAME <path/filename>
```

When a group is not part of a pattern, the statistics in the **Times Fired** column in the **Groups** panel represent the sum of incoming events that are part of a group in active status. The value increments every time there is an event, which is part of an active group. This statistic is active by default, and can't be turned off

Creating patterns

Groups of related events are discovered using Related Event analytics. Automatically discovered groups in the View Related Events portlet can be used to create patterns.

About patterns

Use this information to understand how patterns are created and how they differ from related event groups.

Patterns and related event groups

The use of patterns allows events with different Event Identifier fields to be grouped in the Event Viewer.

Discovered groups can be deployed independently of a pattern. In this case, incoming events are matched by using the Event Identifier field and grouped in the Event Viewer by Resource Field.

In contrast, the use of patterns allows events with different Event Identifier fields to be grouped in the Event Viewer. In this case, incoming events are matched by using the Event Type field. Matching events for deployed patterns are grouped by Resource in the Event Viewer. To allow a group of related events, with *different* Resource field values, to be allocated to a pattern, use name similarity or specify a regular expression in the pattern. To allow deployed patterns to group events across multiple resources, use name similarity or specify a regular expression in the pattern.

Event Types

By default, the Event Type field is set to the `AlertGroup` column in the Object Server `alerts.status` table. You can configure the system to use a different column or combination of columns by using the Event Analytics configuration wizard, as described in “Configuring event pattern processing” on page 254.

When a pattern is manually created, at least one Event Type must be selected in the drop-down list. There is no maximum number of Event Types. Whichever related event groups are allocated to the pattern must have all the Event Types specified (and no extra ones). For example, if a pattern is created with Event Type set to `NmosEventType` and `ITNMMonitor` then only groups with related events that have both these event types can be allocated. In this example, if a group contains three related events with `AlertGroup` values of, in turn, `NmosEventType`, `ITNMMonitor`, and some other value, then this group is not a candidate for allocation to the pattern. It is not a candidate because its related events contain an Event Type that is not part of the pattern.

Resources

A resource can be a hostname (“server name”), or an IP address.

By default the Resource field is set to the `Node` column in the Object Server `alerts.status` table. You can configure the system to use a different column by using the Event Analytics configuration wizard, as described in “Configuring event pattern processing” on page 254.

If name similarity is disabled and no regular expression is specified for the pattern, then all the resource values for the events within a related events group must be the same.

Note: The check on the Resource field is performed by using the historical event data, not the related event data. However, most of the time the Resource value in the historical event data is the same as the Resource value in the related event data.

If many related events groups are allocated to a pattern, then the Resource value does not have to match across groups; however, the Resource value must match *within* a group. Name similarity is enabled by default in Netcool Operations Insight V1.5.0 and higher. When name similarity is enabled, the Resource values in the related events group (by default, the contents of the Node column) must be sufficiently similar based on the name similarity settings. The default name similarity settings require the lead characters to be the same and the text to be 90% similar.

Warning: There is an important exception to this scenario. If the Node column contains IP addresses, then the different IP address values must match down to the subnet value; that is, the first, second, and third segment of the IP address must be the same. For example:

- 123.456.789.10 matches 123.456.789.11.
- 123.456.789.10 does *not* match 123.456.788.10.

Example

For example, assume that Link up and Link down regularly occurs on Node A. Analytics detects the occurrence in the historical data and generates a specific grouping of those two events for Node A. Likewise, if Link up and Link down also regularly occurs on Node B, a grouping of those two events is generated but specifically for Node B.

With generalization, the association of such events is encapsulated by the system as a pattern: Link up / Link down on any Node. In generalization terms, Link Up / Link Down represents the event type and Node* represents the resource.

Advantages of patterns

A created pattern has the following advantages over a related event group:

- For any instance of a pattern, not all of the events in the definition must occur for the pattern to apply. This is dependent on the Trigger Action settings. For more information about Trigger Action setting, see “Creating an event pattern” on page 335.
- The pattern definition encompasses groups of events with the defined event types.
- A single pattern can capture the occurrence of events on any resource. For example, with discovered groups, analytics only found historical events that occurred on a specific hostname, and created groups for each host name. If real time events happen on different host names in the future, the discovered groups will not capture them. However, patterns will discover the events because the event type is the same.
- A pattern can encompass event groupings that were not previously seen in the event history. An event group that did not previously occur on a specific resource is identified by the pattern, as the pattern is not resource dependent, but event type specific. **Note:** when selecting an event type (during the event type configuration), the column that identifies the event type should be unique across multiple event groups.
- A single pattern definition can encompass multiple event groups. Patterns will act on event types for different host names which might have occurred historically (discovered groups) or will happen in future real time events. For example, an event type could be "*Server Shutting Down*", "*Server Starting Up*", "*Interface Ping Failure*", and so on. Each group is resource specific, but an event pattern is event type specific. Therefore, an environment might have multiple groups for different resources, and an event pattern will encompass all of those different groups since their event type is the same.

Extending patterns

You can extend pattern matching functionality using regular expressions and using name similarity.

You can extend pattern matching functionality in the following two ways to enable the discovery of a pattern instance on more than one resource:

- Using *regular expressions*: you can define a regular expression to apply to the contents of the resource field or fields during pattern matching. Resource names that match the regular expressions are candidates to be included in a single pattern. You can optionally specify a regular expression when you create a pattern.
- Using *name similarity*: this feature uses a string comparison algorithm to determine whether the resource names contained in two resource fields are similar. Name similarity is enabled by default. If enabled, name similarity is applied at two points in the process:

  1. When patterns are suggested, as described in "Suggested patterns" on page 339.
  2. When live events are correlated to identify pattern instances, as described in "Examples of name similarity" on page 333.

  The end result is that during pattern processing of live events it is possible to have more than one resource name in a single pattern instance if the string comparison algorithm determines that the resource names are similar.

Note: Name similarity and regular expression functionality are not mutually exclusive. If name similarity is configured, you can also define regular expressions. Pattern matching is processed in the following order:

1. Exact match
2. Regular expression
3. Name similarity

Examples of name similarity

This topic presents examples of similar resource names that might be discovered by using the default name similarity settings.

By default, name similarity is configured with the following default settings. For more information on these configuration parameters, see "Configuring name similarity" on page 283.

Parameter	Description	Values
`name_similarity_default_threshold`	String comparison threshold. For example, a similarity threshold value of 0.9 means that strings must match by at least that value to be considered similar. For more information about the threshold value, see Similarity threshold value. • 1 equates to identical strings. • 0 equates to completely dissimilar strings.	0.9
`name_similarity_default_lead_restriction`	Lead restriction. Number of characters at the beginning of the string that must be identical.	1
`name_similarity_default_tail_restriction`	Tail restriction. Number of characters at the end of the string that must be identical.	0

Based on these settings, the following event snippets present an example for resource name similarity analysis. Note that the resource name is stored in the resource column. `NODE` is the default resource column, but can be changed for the pattern event type.

```
1     NODE                                          SUMMARY
ALERTKEY             Similar?
2     acme.env1.base.adm_chk_probeCheck             System Alert SEV2 ABC
```

```
adm_probe            No
3     cnz.env2.base.adm_chk_reports              System Alert SEV2 ABC
adm_report           No
4     abc.lyf.base.logs1                         System Alert SEV2 DEF
logs1                No
5     abc.gbs.stato.dotnetcore                   System Alert SEV2 ABC
runtime_down         No
6     caripa.env1.stato.dotnetcore               System Alert SEV2 GHI
runtime_down         Yes
7     caripa.env1.stato.TNT                      System Alert SEV2 GHI
runtime_down         Yes
8     caripa.env1.stato.TNT                      System Alert SEV2 GHI
runtime_down         Yes
9     emperor.env3.stato.pythonRuntime           System Alert SEV2 ABC
runtime_down         No
10    abc.env5.base.bash.total_cpu_noncore       System Alert SEV2 DEF
bash_cpu_noncore     No
11    abc.cio.base.total_cpu_noncore             System Alert SEV2 ABC
bash_cpu_noncore     No
12    banca.env1.base.bosh.jobstate.console      System Alert SEV2 GHI
job_fail             No
```

As a result of this similarity analysis, the resource names in the NODE column for the events listed in rows 6, 7, and 8 are considered similar. The reasons for this include the following:

- All of the resource names other than those in the NODE column of rows 2, 6, 7, and 8, start with a letter other than c, hence they are rejected automatically, because the lead restriction is set to 1 character.
- The resource name in the NODE column of row 2 fails the similarity threshold of 0.9 because it is very different to the resource names in rows 6, 7, and 8.
- The tail restriction is set to 0, so this allows the resource name in row 6 to pass overall similarity, even though the final letters of its resource name are different to the final letters of the resource names in rows 7 and 8.

Starting the Events Pattern portlet

The administrator can start the Events Pattern portlet from a number of locations on the Event Analytics UI.

Before you begin

To access the View Related Events, Related Event Details, and Events Pattern portlets, users must be assigned the `ncw_analytics_admin` role.

About this task

You can start the Events Pattern portlet from the View Related Events portlet or the Related Event Details portlet. Starting the Events Pattern portlet directly from the Related Event Details portlet ensures that you do not need to return to the View Related Events portlet to start the Events Pattern portlet after you review the details of a group.

Procedure

You can start the Events Pattern portlet from the View Related Events portlet or the Related Event Details portlet.

1. To start the Events Pattern portlet from the View Related Events portlet, start the View Related Events portlet and select one of the following options. For more information about starting the View Related Events portlet, see “Viewing related events” on page 317.
   a) To create a pattern, complete the following steps.
      1) Select a related events group in the **Groups** table.
      2) Right-click the related events group and select **Create Pattern**.
   b) To edit an existing pattern, complete the following steps.
      1) Select a pattern in the **Group Sources** table.
      2) Right-click the pattern and select **Edit Pattern**.

2. To start the Events Pattern portlet from the Related Event Details portlet, complete the following steps.
   a) Start the Related Event Details portlet. For more information, see "Viewing related events details for a seasonal event" on page 317
   b) Click **Create Pattern**.

      Note: The Events Pattern portlet is updated with each newly selected group.

What to do next

Input the details of the pattern in the Events Pattern portlet. For more information about completing the Events Pattern portlet, see "Creating an event pattern" on page 335

Creating an event pattern

You can create a pattern based on the automatically discovered groups.

Before you begin

To access the View Related Events and Events Pattern portlets, users must be assigned the `ncw_analytics_admin` role.

About this task

A related events configuration automatically discovers groups of events that apply to specific managed resources. You can create an event pattern that is not specific to resources based on an automatically discovered group.

Procedure

1. Start the Events Pattern portlet for a group. For more information about starting the portlet, see "Starting the Events Pattern portlet" on page 334.
2. Complete the parameter fields in the **Pattern Criteria** tab of the Events Pattern portlet.

   Merge into
   Merge a Related Event Group into an existing pattern or select **NONE** to create new pattern. To merge a group into a pattern, select from the list of patterns with one or more event types in common. **NONE** is the default option.

   Name
   The name of the pattern. The name must contain alphanumeric characters. Special characters are not permitted.

   Pattern Filter
   The ObjectServer SQL filters that are applied to the pattern. This filter is used to restrict the events to which the pattern is applied. For example, enter `Summary NOT LIKE '%maintenance%'`.

   Time between first and last event
   The maximum time that can elapse between the occurrence of the first event and the last event in this pattern, which is measured in minutes. The default value is determined by the Related Events Group on which the pattern is based. Events that occur outside of this time window are not considered part of this group.

   Trigger Action
   Select the **Trigger Action** check box to group the live events when the selected event comes into the ObjectServer. When an event with the selected event type occurs, the grouping is triggered to start. The created grouping includes events that contain all of the selected event types.

   For example, if the following three event types are part of the pattern criteria, `A`, `B`, and `C`, with only the **Trigger Action** check box for event `C` selected, the grouping only occurs when an event with event type `C` occurs. The grouping contains events that contain all three event types.

   Note: A group will be triggered even if only one event with the triggering event type occurs. In this case a group will be created in the **Event Viewer** made up of either of a synthetic and the triggering

event as a child event, or of the triggering event as both parent and child event, depending on how you configure the **Parent Event** tab of the Events Pattern portlet, in step 3.

Event Type
The event type or types that are included in the pattern. The **Event Type** is prepopulated with existing event types for the selected pattern, and can be modified.

Note: Origin of event type

Triangle, circle, and square icons signify where the event types originate from, when a group is merged into an existing pattern.

- Triangle: Common to both the existing pattern and the group.
- Circle: Part of the group.
- Square: Part of the existing pattern.

Resource Column(s)
The resource or resources to which the action is applied. The **Resource Column(s)** is prepopulated with existing event type resources for the selected pattern, and can be modified. To modify the selection, click the drop-down list arrow and select one or more columns from the checklist.

If you specify multiple resource columns, then by default these columns will be combined using OR logic. You can configure whether multiple resource columns should be combined using AND or OR logic. For more information, see "Configuring multiple resource columns" on page 285.

- OR logic: correlates two events by resource as soon as the criteria are met for just one pattern resource definition.
- AND logic: correlates two events by resource only once criteria are met for all of the pattern resource definitions.

Note: If you specify AND logic, then you cannot specify a regular expression for matching the resource information from the multiple selected resource columns.

Note: Duplicate **Event Type** and **Resource Columns** pairs are not permitted.

Regular Expression

(Optional) Click the regular expression icon  to specify a regular expression pattern for matching the resource information from the selected resource column.

To match a string, add the asterisk symbol * before and after the characters. For example, to match the resource information in the "`the application` *`abc`* `on` *`myhost.1xyz.com`* `encountered an unrecoverable error.`" event, use the following regular expression:

```
.*[0-9]*xyz.*
```

Resource names that match the regular expression are identified when the Events Pattern is created.

Note: A regular expression can only be specified under the following conditions:

- One column has been selected for the resource.
- Multiple columns with OR logic have been selected for the resource. OR logic is the default.

A regular expression cannot be specified when multiple columns with AND logic have been selected for the resource.

For more information about creating and editing regular expressions, see "Applying a regular expression to the pattern criteria" on page 338.

3. In the **Parent Event** tab of the Events Pattern portlet, select one of the following parent event options.

Most Important Event by Type
The system checks the events as they occur. The events are ranked based on the order defined in the UI. The highest ranking event is the parent. The parent event changes if a higher ranking event

occurs after a lower ranking event. To prevent a dynamically changing parent event, select **Synthetic Event**.

You can manually reorder the ranking by selecting an event and clicking the **Move Up** and **Move Down** arrows.

Synthetic Event
Create an event to act as the parent event or select **Use Selected Event as Template** to use an existing event as the parent event.

To create or modify a synthetic event, populate the following parameter fields, as required. All of the synthetic event fields are optional.

Node
The managed entity from which the event originated. Displays the managed entity from which the seasonal event originated.

Summary
The event description.

Severity
The severity of the event. Select one of the following values from the **Severity** drop-down list.

Critical
Major
Minor
Warning
Indeterminate
Clear

Alert Group
The Alert Group to which the event belongs.

Add additional fields
Select the **Add additional fields** check box to add more fields to the synthetic parent event.

4. In the **Test** tab of the Events Pattern portlet, you can run a test to display the existing auto-discovered groups that match the pattern criteria. The test displays the types of events that are matched by the chosen criteria. To run the test, select **Run Test**. To cancel the test at any time, select **Cancel Test**.
5. To save, watch, or deploy the pattern, select one of the following options.
   - Select **Save** to save the pattern details to the View Related Events **New** tab.
   - Select **Watch** to add the pattern to the View Related Events **Watched** tab.
   - Select **Deploy** to add the pattern to the View Related Events **Active** tab.

Results
The events pattern is created and displayed in the **Group Sources** table in the View Related Events portlet.

Note:

- If the patterns display 0 group and 0 events, the pattern creation process might not be finished. To confirm that the process is running,
  1. Append the policy name to the policy logger file from the **Services** tab, **Policy Logger** service. For more information about configuring the Policy logger, see https://www.ibm.com/support/knowledgecenter/SSSHYH_7.1.0.16/com.ibm.netcoolimpact.doc/user/policy_logger_service_window.html.
  2. Check the following log file.

     ```
     $IMPACT_HOME/logs/<serverName>_policylogger_PG_ALLOCATE_PATTERNS_GROUPS.log
     ```

  If the process is not running, see the Event Analytics troubleshooting information.

- After creating a new pattern, the allocation of groups to the pattern happens in the background, via a policy. If the new pattern does not have any groups allocated (this is determined by the data set) then the new pattern will be deleted. For more information, see the following technote: http://www.ibm.com/support/docview.wss?uid=swg22012714.
- A pattern will not have any groups allocated under the following conditions:
  - Name similarity has been switched off. By default it is on.
  - No regular expressions have been associated with the pattern.
  - Resource names identified in any potential groups are different.

Related reference

Troubleshooting Event Analytics (on premises)
Use the following troubleshooting information to resolve problems with Event Analytics.

Applying a regular expression to the pattern criteria

You can apply a new regular expression to the pattern criteria, or edit an existing regular expression, to match resource information from unstructured data in the selected resource column.

Before you begin

The Resource field is used for grouping events in the Event Viewer. Event matching occurs in the following order:

1. Exact match
2. Regular expression
3. Name similarity

To access the View Related Events and Events Pattern portlets, users must be assigned the `ncw_analytics_admin` role.

About this task

The regular expression is used to match specific information from unstructured data in the selected resource column.

Note: A regular expression can only be specified under the following conditions:

- One column has been selected for the resource.
- Multiple columns with OR logic have been selected for the resource. OR logic is the default.

A regular expression cannot be specified when multiple columns with AND logic have been selected for the resource.

You can configure whether multiple resource columns should be combined using AND or OR logic. For more information, see “Configuring multiple resource columns” on page 285.

Procedure

1. Start the Events Pattern portlet for a group. For more information about starting the portlet, see “Starting the Events Pattern portlet” on page 334.
2. Proceed as follows:

   - To apply a new regular expression, click the regular expression icon  in the **Pattern Criteria** tab of the Events Pattern portlet.
   - To modify an existing regular expression, click the Confirm icon  in the **Pattern Criteria** tab of the Events Pattern portlet.

   The **Regular Expression** dialog box is displayed.

3. Insert or edit the regular expression in the **Expression** field.

   To match a string, add the asterisk symbol * before and after the characters. For example, to match the resource information in the "`the application` *`abc`* `on` *`myhost.1xyz.com`* `encountered an unrecoverable error.`" event, use the following regular expression:

   ```
   .*[0-9]*xyz.*
   ```

   Resource names that match the regular expression are identified when the Events Pattern is created.
4. To change or select the event type to which the regular expression is applied, select an event type from the drop-down list in the **Test Data** field.

   Note: A regular expression works on one resource field only. It does not work if multiple resource fields are selected. If a pattern has two or more event types, they must all use the same Resource field.
5. To test the regular expression, select **Test**. The test results are displayed in the **Result** field.

   Note: If there are multiple matches for the given regular expression, the matches are displayed in the **Result** field as a comma-separated list.
6. To save and apply the regular expression, select **Save**. The **Regular Expression** dialog box is closed. A confirm symbol is displayed beside the Resource Column.

Viewing related event details in the Events Pattern portlet

You can view the related event details for a selected related events group in the Events Pattern portlet, when you create a new pattern for the group.

Before you begin

To access the View Related Events and Events Pattern portlets, users must be assigned the `ncw_analytics_admin` role.

Procedure

To view the related event details in the Events Pattern portlet, complete the following steps.

1. Open the View Related Events portlet.
2. Select a related events group in the **Groups** table.
3. Right-click the related events group and select **Create Pattern**. The Events Pattern portlet is displayed.

Results

The related event details are displayed in the **Group instances** and **Events** tables in the **Pattern Criteria** tab of the Events Pattern portlet.

Note: The related event details columns in the **Pattern Criteria** tab of the Events Pattern portlet matches the Related Event Details portlet columns.

Suggested patterns

Suggested Patterns are automatically created during a Related Events Configuration.

With generalization, the association of events is encapsulated by the system as a pattern. Any Suggested Patterns that are generated can be viewed in the **Group Sources** table of the **View Related Events** portlet. For more information, see “Viewing related events by group” on page 318.

Note: Patterns are not created when the **Override global event identity** option is selected in the **Configure Analytics portlet**.

Right-click on a suggested pattern in the **Group Sources** table to display a list of menu items. You can select the following actions from the menu list.

Edit Pattern For more information about this action, see “Editing an existing pattern” on page 340.

Delete Pattern For more information about this action, see “Deleting an existing pattern” on page 340.

Watch For more information about this action, see “Watching a correlation rule” on page 327.

Deploy For more information about this action, see “Deploying a correlation rule” on page 328.

Archive For more information about this action, see “Archiving related events” on page 324.

Copy Choose this action if you want to copy a row, which you can then paste into another document.

Editing an existing pattern

You can edit an existing pattern to modify the pattern criteria.

Before you begin

To access the View Related Events and Events Pattern portlets, users must be assigned the `ncw_analytics_admin` role.

Note: On rerun of a related events configuration scan, any already existing patterns cannot be edited. Only newly discovered patterns can be edited.

Procedure

1. Start the View Related Events portlet. For more information about starting the View Related Events portlet, see “Viewing related events” on page 317.
2. Select a pattern in the **Group Sources** table.
3. Right-click the pattern and select **Edit Pattern**.
4. Modify the parameter fields in the **Pattern Criteria** and **Parent Event** tabs. For more information about modifying the parameters, see “Creating an event pattern” on page 335.
5. To save, watch, or deploy the pattern, select one of the following options.
   - Select **Save** to save the pattern details to the View Related Events **New** tab.
   - Select **Watch** to add the pattern to the View Related Events **Watched** tab.
   - Select **Deploy** to add the pattern to the View Related Events**Active** tab.

Deleting an existing pattern

You can delete an existing pattern to remove it from the **Group Sources** table.

Before you begin

To access the View Related Events and Events Pattern portlets, users must be assigned the `ncw_analytics_admin` role.

Procedure

1. Start the View Related Events portlet. For more information about starting the View Related Events portlet, see “Viewing related events” on page 317.
2. Select the pattern you want to delete in the **Group Sources** table.
3. Right-click the pattern and select **Delete Pattern**.
4. To delete the pattern, select **Yes** in the confirmation dialog window.

Results

The selected pattern is deleted.

Exporting pattern generalization test results to Microsoft Excel

You can export pattern generalization test results for a specific configuration to a Microsoft Excel spreadsheet from a supported browser.

Before you begin

To access the View Related Events, Related Event Details, and Events Pattern portlets, users must be assigned the `ncw_analytics_admin` role.

About this task

You can start the Events Pattern portlet from the View Related Events portlet or the Related Event Details portlet. Starting the Events Pattern portlet directly from the Related Event Details portlet ensures that you do not need to return to the View Related Events portlet to start the Events Pattern portlet after you review the details of a group.

Procedure

To export pattern generalization test results for a specific configuration to a Microsoft Excel spreadsheet, complete the following steps.

1. Open the **View Related Events** portlet.
2. Select a specific configuration from the configuration table.
3. Enter the pattern criteria and navigate to the **Test** tab of the Events Pattern portlet and select **Run Test**.
4. Click the **Export Generalization Test Results** button in the toolbar.

   After a short time, the **Download export results** link displays.
5. **Note:** The user is restricted to exporting the first 100 groups of the pattern test results to provide a sample of the results in the exported file.

   Click the link to download and save the Microsoft Excel file.

Results

The Microsoft Excel file contains a spreadsheet with the following tabs:

- `Groups Information`: This tab contains the related events groups for the configuration that you selected.
- `Groups Instances`: This tab contains a list of all the related events instances for all of the related events groups for the configuration that you selected.
- `Group Events`: This tab contains a list of all the events that occurred in the related events groups for the configuration that you selected.
- `Instance Events`: This tab contains a list of all the events that occurred in all of the related events instances for all the related events groups for the configuration that you selected.
- `Export Comments`: This tab contains any comments relating to the export for informational purposes (for example, if the spreadsheet headers are truncated, or if the spreadsheet rows are truncated).

Cloud Native Analytics on IBM Cloud Private

Cloud Native Analytics allows you to identify seasonal patterns of events and temporal event groups while running on IBM Cloud Private.

Figure 11. Event viewer

Use the **Example_IBM_CloudAnalytics** view to see the **Grouping** and **Seasonal** columns in the **Event Viewer**. The **Default** view does not display the grouping and seasonal columns. To add these columns to the **Default** view, use the **Edit Views** icon to add the **`CEACorrelationDetails`** and **`CEASeasonalDetails`** columns to the Event list view.

Seasonal events

Cloud Native Analytics uses statistical analysis of IBM Tivoli Netcool/OMNIbus historical event data to determine the seasonality of events, such as when and how frequently events occur. Use the **Example_IBM_CloudAnalytics** view to see the **Seasonal** column in the **Event Viewer**. Seasonal events are also highlighted using an icon in the **Incident** panel.

Temporal event grouping

Cloud Native Analytics uses statistical analysis of Tivoli Netcool/OMNIbus historical event data to determine which events have a statistical tendency to occur together. Cloud Native Analytics outputs the results of this statistical analysis as event groups, on a scheduled basis. Use the **Example_IBM_CloudAnalytics** view to see the **Grouping** column in the **Event Viewer**. Click **Investigate** to open the **Incident Viewer**. View details of the events and quickly determine the temporal or scope-based groups to which each event belongs.

The **Example_IBM_CloudAnalytics** view displays parent events, with one or more child events hidden from view. You can toggle a parent event in the **Event Viewer** to display the child events.

Roles

Three roles are available in Cloud Native Analytics:

- noi_lead - The noi_lead role can perform all operations on the UI. With the noi_lead role, you can manage policies in **Temporal policies**. This feature is not available to other roles.
- noi_engineer - The noi_engineer role can perform all operations on the UI, except for managing policies.
- noi_operator - The noi_operator role can open the **Incident Viewer** from the **Event Viewer**, but cannot click-through on the seasonality and grouping icons in the **Incident Viewer**. The **Temporal group** or **Seasonal event** panels are not available with this role. The **See more info** option is not available on individual events with this role. Also, policies cannot be approved or rejected with this role.

Incidents

The **Incident Viewer** in Cloud Native Analytics shows the events that make up the incident, together with the seasonal, temporal or scope-based group to which each event belongs. An event can belong to more than one group.

Drill-down to an individual event to see the event details and timeline.

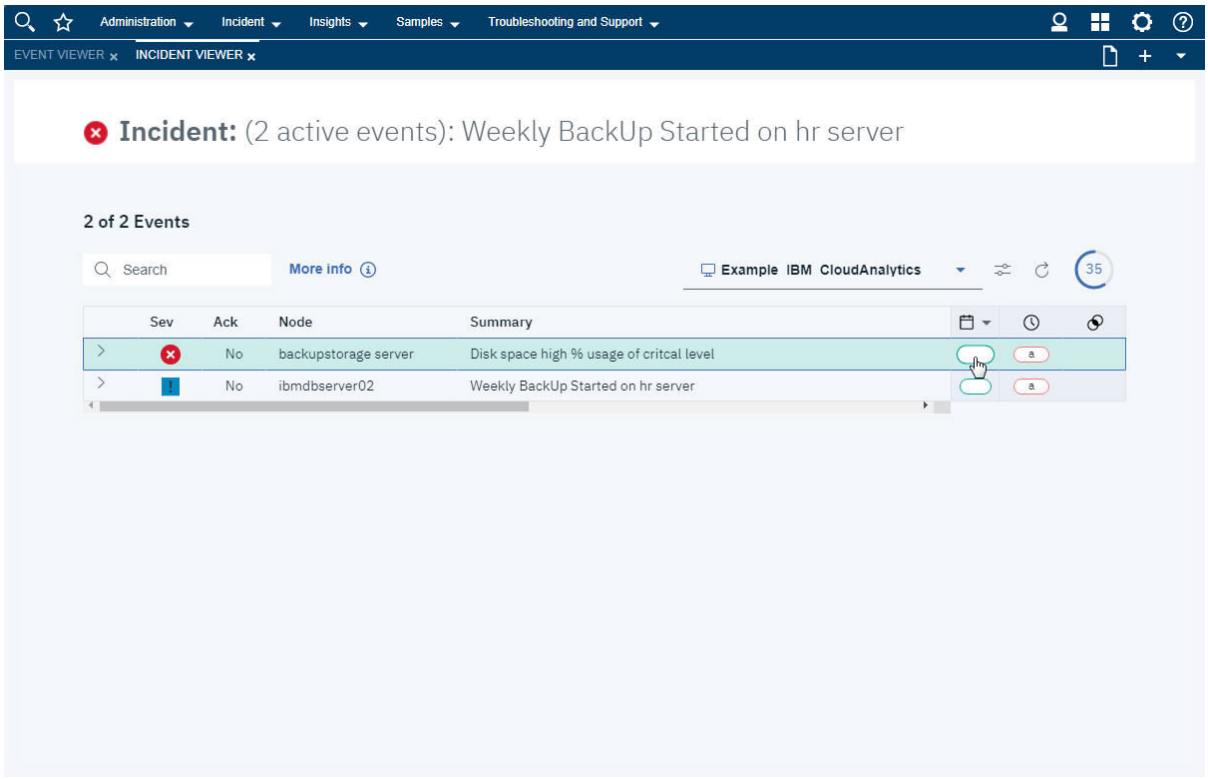

Figure 12. Incident viewer

Investigate further by selecting a group.

Figure 13. Temporal group in Incident viewer

Figure 14. Seasonal group in Incident viewer

Seasonal events

Cloud Native Analytics allows you to identify seasonal events within your monitored environment. Seasonal events tend to occur at specific times.

The seasonality details panel provides a calendar view of all of the historical events that contributed to the selected seasonal event. One or more seasonal time windows are listed in the top left pane. Select a seasonal time window to filter the view to show just historical events that contributed to that time window.

Figure 15. Seasonal event

Temporal groups

Cloud Native Analytics performs temporal correlation to identify groups of events that tend to occur together within your monitored environment.

The temporal group panel shows the events that make up a temporal group, together with all of the historical instances of the group. You can navigate into these instances to view detailed timelines of each instance. You can also view event details.

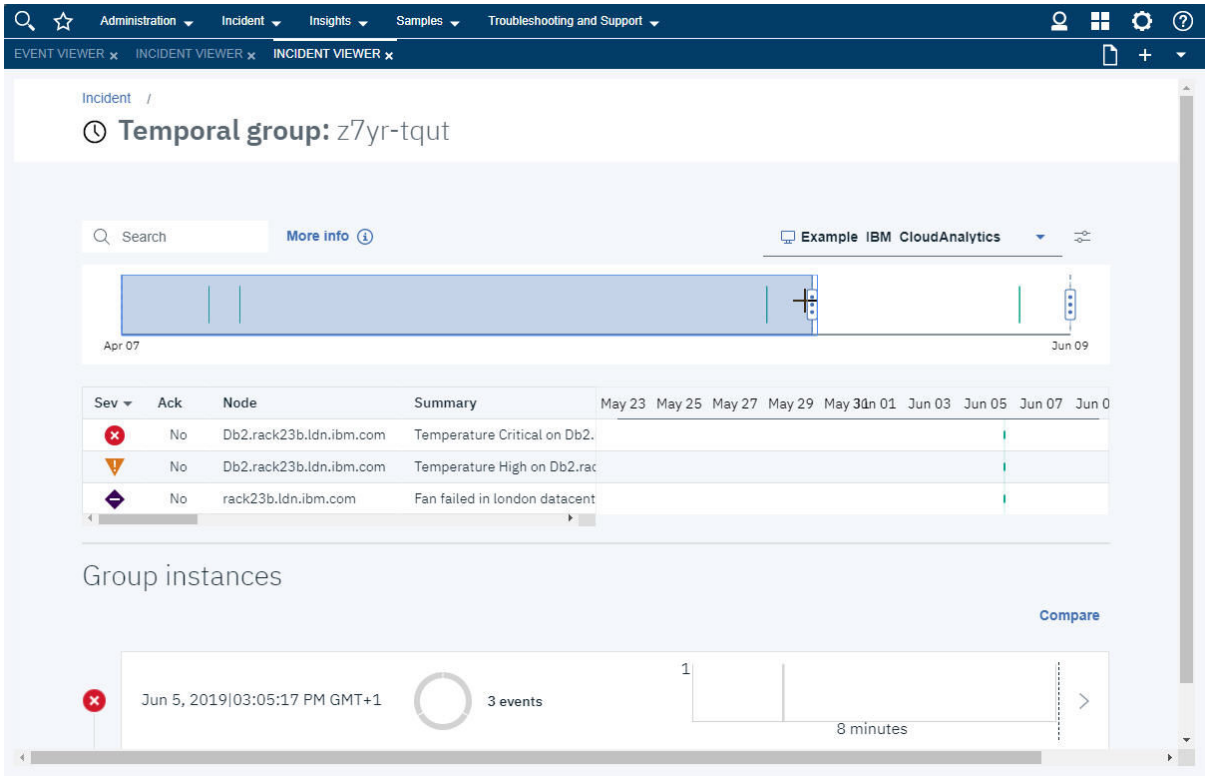

Figure 16. Temporal group

Click the new window icon to open the temporal group panel in a new browser window. You can copy the group URL from the new window to share the temporal group with others. If the temporal group is deleted or updated in a training run, a new URL is required.

Scope-based groups

Events in a scope-based group are grouped together because they share a common attribute, such as a resource.

Create an event policy to set `ScopeID` for events that match your defined filter. Scope-based event grouping is activated for events that match the filter, based on the `ScopeID` that you specify. For more information about creating a scope-based grouping policy, see OMNIbus documentation: Creating a scope-based grouping policy for Event Analytics.

Manage policies

Cloud Native Analytics allows you to manage policies within your monitored environment. Open **Manage Policies** from the **Insights** menu.

You can rank, approve, or reject policies from the **Temporal policies** panel. You can also drill-down to find out more information about individual policies.

The **Live** tab lists live temporal policies. If you are running in autodeploy mode, then the neighboring **Suggested** tab is always empty and all policies are listed in the **Live** tab.

Figure 17. Temporal policies

Autodeploy mode

When you install Cloud Native Analytics, you can select an option to automatically deploy all temporal group policies. For more information, see the *autodeploy* parameter in "Configuring Installation Parameters for Operations Management on IBM Cloud Private" on page 125.

Loading data

To learn about Cloud Native Analytics, you can install data from your local system, migrate historical data from a reporter database, or use the sample data set, which is provided with the Operations Management on IBM Cloud Private installation. Load data, train the system, and see the results.

Loading sample data: scenario for Operations Management on IBM Cloud Private

To learn about Operations Management on IBM Cloud Private, you can install a sample data set. Learn how to install and load sample data, train the system, and see the results.

Before you begin

Before you complete these steps, complete the following prerequisite items:Before you complete these steps, complete the following prerequisite items:

- The **ea-events-tooling** container is installed with the helm chart. It is not started as a pod and contains scripts to install data on the system, which can be run with the `kubectl run` command.
- Determine *image*, the location of the **ea-events-tooling** container image on your system, by running the following command: **kubectl get image ea-events-tooling -o json**. You can also run `helm status` *helm_release_name* `--tls`, where *helm_release_name* is the Cloud Native Analytics Helm release name. In the Sample Data output of this command there is an example command to create a data ingestion job, and the value specified for the **--image** parameter for that command is the **ea-events-tooling** container image location. Launch the **ea-events-tooling** container image with a `kubectl run` command.
- Determine *image_tag,* the image version number of the **ea-events-tooling** container, by launching the IBM Cloud Private UI. Select **Menu** > **Container images** and click the **ea-events-tooling** container to display the tag details.
- Install a Kubernetes client and configure it to connect to the cluster. For more information, see https://www.ibm.com/support/knowledgecenter/SSBS6K_3.2.0/manage_cluster/install_kubectl.html.

- Create a secrete which maps to the **ea-events-tooling** container image location. For more information, see "Configuring passwords and secrets" on page 111. Attach image pull secrets to the default service account. For example, if you have a secret called `my-docker-repository-secret`, run the following command:

```
kubectl patch serviceaccount default -p '{"imagePullSecrets": [{"name": "my-docker-repository-secret"}]}'
```

Note: As an alternative to patching the default service account with image pull secrets, you can add the following option to each **kubectl run** command that you issue:

```
--overrides='{ "apiVersion": "v1", "spec": { "imagePullSecrets":
      [{"name": "my-docker-repository-secret"}] } }'
```

About this task

You can use scripts in the **ea-events-tooling** container to install sample data on the system. Run the `loadSampleData.sh` script to load data to the ingestion service, train it, create a scope-based policy and load the data into IBM Netcool Operations Insight. This script loads prebuilt data into the ingestion service and ObjectServer and trains the system for seasonality and related events.

To access the secrets, which control access to the ObjectServer, Web GUI and policy administration, the `loadSampleData.sh` script needs to run as a job. For more information, see "Configuring passwords and secrets" on page 111.

Procedure

1. Run the `loadSampleData.sh` script from a job, so that it has access to the secrets that control the passwords. Complete the following steps:
   a. Use the `-j` option with the script to generate a YAML file, such as `loadSampleJob.yaml` in the following example:

```
kubectl delete pod ingesnoi3
kubectl run ingesnoi3  --restart=Never --env=LICENSE=accept --image-pull-policy=Always  \
--env=CONTAINER_IMAGE=image:image_tag  \
-i --image=image:image_tag \
loadSampleData.sh -- -r helm_release_name -j  > loadSampleJob.yaml
```

   Where:

   - *image* is the location of the **ea-events-tooling** container. The image value can be found from the helm status *release_name* --tls command, as described earlier.
   - *CONTAINER_IMAGE* is an environment variable, which is the same as the value you pass to the **--image** parameter in the Kubernetes command. This variable allows the container to populate the image details in the YAML file output that it creates.
   - *image_tag* is the image version tag, as described earlier.
   - *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment.

   b. Create a job using the generated YAML file, such as `loadSampleJob.yaml` in the following example:

```
kubectl create -f loadSampleJob.yaml -n <namespace>
```

   Where *<namespace>* is the name of the namespace in which Operations Management on IBM Cloud Private is installed.

   Note: If the default service account does not have access to the image repository, uncomment the image pull secrets section in the `loadSampleData.yaml` file and set the **imagePullSecrets.name** parameter to your Docker secret name before running the **kubectl create** command.

A job called `-loadSampleData` is created. You can view the job output with the pod logs created by the job.

2. View the sample data.
   a. Connect to Web GUI. The URL is displayed in the command output when you do a helm install of Operations Management on IBM Cloud Private. For more information, see "Installing Operations Management on IBM Cloud Private" on page 119.
   b. Select **Incident** > **Event Viewer**. The list of all events are displayed.
   c. Select the **Example_IBM_CloudAnalytics** view to see how the events from the sample data are grouped.

If you have installed Operations Management on IBM Cloud Private in manual deploy mode, running the `loadSampleData.sh` script switches your configuration to auto-deploy mode. For more information about manually deploying policies, see step #unique_146/unique_146_Connect_42_autodeploy in the *Installing Operations Management on IBM Cloud Private* topic.

To disable auto-deploy mode, re-run a portion of the training by running the `runTraining.sh` script.

Get the start and end times of the sample data as in the following example:

```
kubectl delete pod ingesnoi3;
kubectl run ingesnoi3 -i --restart=Never --env=LICENSE=accept --image=image:image_tag
getTimeRange.sh samples/demoTrainingData.json.gz
pod "ingesnoi3" deleted
{"minfirstoccurence":{"epoc":1552023064603,"formatted":"2019-03-08T05:31:04.603Z"},
"maxlastoccurrence":{"epoc":1559729860924,"formatted":"2019-06-05T10:17:40.924Z"}}
```

Where:

- *image* is the location of the **ea-events-tooling** container. The image value can be found from the helm status *release_name* --tls command, as described earlier.
- *image_tag* is the image version tag, as described earlier.

Re-run the training with seasonality disabled, as in the following example:

```
kubectl delete pod ingesnoi3;
kubectl run ingesnoi3 -i --restart=Never --env=LICENSE=accept --image-pull-policy=Always --
image=image:image_tag runTraining.sh -- -r test-install -a SEASONALITY -s 1552023064603 -e
1559729860924 -d false
```

Where:

- *image* is the location of the **ea-events-tooling** container. The image value can be found from the helm status *release_name* --tls command, as described earlier.
- *image_tag* is the image version tag, as described earlier.
- *-s* and *-e* are the start and end times returned by the **getTimeRange()** command
- *-d* disables live policy updates.

Loading local data: scenario for Operations Management on IBM Cloud Private
To learn about Operations Management on IBM Cloud Private, you can install a local data set. Learn how to install and load your local data, train the system, and see the results.

Before you begin
Before you complete these steps, complete the following prerequisite items:

- The **ea-events-tooling** container is installed with the helm chart. It is not started as a pod and contains scripts to install data on the system, which can be run with the `kubectl run` command.
- Determine *image*, the location of the **ea-events-tooling** container image on your system, by running the following command: **kubectl get image ea-events-tooling -o json**. You can also run `helm status helm_release_name --tls`, where *helm_release_name* is the Cloud Native Analytics Helm release name. In the Sample Data output of this command, there is an example command to create a data ingestion job. The value that is specified for **--image** for that command is the **ea-**

events-tooling container image. Start the **ea-events-tooling** container image with a `kubectl run` command.

- Determine *image_tag*, the image version number of the **ea-events-tooling** container, by starting the IBM Cloud Private UI. Select **Menu** > **Container images** and click the **ea-events-tooling** container to display the tag details.
- Determine the HTTP username and password from the secret which has **systemauth** in the name, by running the following command:

```
kubectl get secret -systemauth-secret -o yaml
```

- Install a Kubernetes client and configure it to connect to the cluster. For more information, see https://www.ibm.com/support/knowledgecenter/SSBS6K_3.2.0/manage_cluster/install_kubectl.html.
- Create a secret that maps to the **ea-events-tooling** container image location. For more information, see "Configuring passwords and secrets" on page 111. Attach image pull secrets to the default service account. For example, if you have a secret called `my-docker-repository-secret`, run the following command:

```
kubectl patch serviceaccount default -p '{"imagePullSecrets": [{"name": "my-docker-repository-
secret"}]}'
```

About this task

You can use scripts in the **ea-events-tooling** container to install local data on the system. To complete this task, run the `filetoingestionservice.sh`, `getTimeRange.sh`, `runTraining.sh`, `createPolicy.sh`, and `filetonoi.sh` scripts. The scripts load local data to the ingestion service, train the system for seasonality and temporal events, create live seasonal policies and suggest temporal policies, and load the data into Operations Management on IBM Cloud Private.

Procedure

1. Send local data to the ingestion service. Run the `filetoingestionservice.sh` script:

```
export HELM_RELEASE=helm_release_name
export HTTP_PASSWORD=$(kubectl get secret $HELM_RELEASE-systemauth-secret -o jsonpath --
template '{.data.password}' | base64 --decode)
export HTTP_USERNAME=$(kubectl get secret $HELM_RELEASE-systemauth-secret -o jsonpath --
template '{.data.username}' | base64 --decode)

cat mydata.json.gz | kubectl run ingesthttp -i --restart=Never --image=image:image_tag --
env=INPUT_FILE_NAME=stdin --env=LICENSE=accept --env=HTTP_USERNAME=$HTTP_USERNAME --
env=HTTP_PASSWORD=$HTTP_PASSWORD filetoingestionservice.sh $HELM_RELEASE
```

   Where:

   - *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment.
   - *mydata.json.gz* is the path to your local compressed data file.
   - *image* is the location of the **ea-events-tooling** container. The image value can be found from the helm status *release_name* `--tls` command, as described earlier.
   - *image_tag* is the image version tag, as described earlier.
   - You can override the user name and password by using `HTTP_USERNAME` and `HTTP_PASSWORD`.

   Note: If you specify the `--env=INPUT_FILE_NAME=stdin` parameter, you can send your local data to the scripts by using the `-i` option with the **kubectl run** command. This option links the `stdin` parameter on the target pod to the `stdout` parameter.
2. Use the `getTimeRange.sh` script to calculate the training time range. If no time range is specified, the trainer trains against all rows for the tenant ID. Instead of using all data associated with the tenant ID to train the system, run the following command to find the start and end time stamps of the data:

```
cat mydata_json.gz | kubectl run ingesnoi3 -i --restart=Never --env=LICENSE=accept --image-
pull-policy=Always
```

```
--image=image:image_tag  getTimeRange.sh stdin
```

Output similar to the following example is displayed:

```
{"minfirstoccurence":{"epoc":1540962968226,"formatted":"2018-10-31T05:16:08.226Z"},
"maxlastoccurrence":{"epoc":1548669553896,"formatted":"2019-01-28T09:59:13.896Z"}}
```

3. Train the system with the new data. Run the `runTraining.sh` script:

```
kubectl run trainer -it --command=true --restart=Never --env=LICENSE=accept
--image=image:image_tag runTraining.sh -- -r helm_release_name [-t tenantid] [-a algorithm]
[-s start-time] [-e end-time] [-d auto-deploy]
```

Where:

- *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment.
- *image* is the location of the **ea-events-tooling** container. The image value can be found from the helm status *release_name* `--tls` command, as described earlier.
- *image_tag* is the image version tag, as described earlier.
- *algorithm* is either **related-events** or **seasonal-events**. If not specified, defaults to **related-events**.
- Optional: *tenantid* is the tenant ID associated with the data that is ingested, as specified by the `global.common.eventanalytics.tenantId` parameter in the `values.yaml` file that is associated with the Helm chart.
- Optional: *start-time* and *end-time* are the start and end times to train against. These values are provided in the command output from step "2" on page 350. You can specify the start or end time, neither, or both. If neither are specified, the current time is used as the end time and the start time is 93 days before the end time. You can either specify the start and end times with an integer Epoch time format in milliseconds, or with the default date string formatting for the system. Run the **./runTraining.sh -h** command to determine the default date formatting.
- Optional: *auto-deploy* Set to true to deploy policies immediately. Set to false to review policies before deployment.

4. Create a policy. A policy can be created through the UI, or you can specify a policy by running the `createPolicy.sh` script:

```
export ADMIN_PASSWORD=$(kubectl get secret helm_release_name-systemauth-secret
-o jsonpath --template '{.data.password}'  | base64 --decode)
kubectl run createpolicy  --restart=Never --image=image:image_tag
--env=LICENSE=accept --env=ADMIN_PASSWORD=${ADMIN_PASSWORD} createPolicy.sh
-- -r helm_release_name
```

Where:

- *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment.
- *image* is the location of the **ea-events-tooling** container. The image value can be found from the helm status *release_name* `--tls` command, as described earlier.
- *image_tag* is the image version tag, as described earlier.

Note: This creates a policy that maps the node to `resource/name` by default. If you want to map to `resource/hostname` or `resource/ipaddress` instead, specify the `--env=CONFIGURATION_PROPERTIES=resource/hostname|ipaddress` parameter.

5. Send local data to the ObjectServer. Run the `filetonoi.sh` script:

```
export HELM_RELEASE=helm_release_name
export OMNIBUS_ROOT_PASSWORD=$(kubectl get secret  test-install-omni-secret
-o jsonpath --template '{.data.OMNIBUS_ROOT_PASSWORD}' | base64 --decode)

cat mydata.tar.gz | kubectl run ingesnoi -i --restart=Never --image=image:image_tag
--env=LICENSE=accept --env=JDBC_PASSWORD=$OMNIBUS_ROOT_PASSWORD
```

```
--env=INPUT_FILE_NAME=stdin filetonoi.sh $HELM_RELEASE
```

Where:

- *helm_release_name* is the helm release name of your Operations Management on IBM Cloud Private deployment.
- *image* is the location of the **ea-events-tooling** container. The image value can be found from the helm status *release_name* --tls command, as described earlier.
- *image_tag* is the image version tag, as described earlier.

You can specify a user ID and password by using JDBC_USERNAME and JDBC_PASSWORD. These parameters correspond to the user ID and password of the ObjectServer.

Note:

You can view the available overrides and their default values by using the **--help** command option.

If you are unable to replay some events, include the following parameters:

```
--env=INPUT_REPORTERDATA=false
--env=EVENT_REPLAY_TEMPORALITY_PRIMARY_TIMING_FIELD=LastOccurrence
--env=EVENT_REPLAY_PREEMPTIVE_DELETIONS=false
--env=EVENT_REPLAY_SKIP_DELETION=false
--env=INPUT_TAG_TOOL_GENERATED_EVENTS=false
```

If you encounter more data integrity issues, include the following parameters:

```
--env=EVENT_REPLAY_STRICT_DELETIONS=false
--env=EVENT_REPLAY_STRICT_PRIMARY_CONTINUITY=false
```

6. View the data.
   a. Connect to Web GUI. The URL is displayed in the command output when you do a helm installation of Operations Management on IBM Cloud Private. For more information, see "Installing Operations Management on IBM Cloud Private" on page 119.
   b. Select **Incident** > **Events** > **Event Viewer**. Select the **All Events** filter. The list of all events is displayed.
   c. Select the **Example_IBM_CloudAnalytics** view to see how the events from the local data are grouped.

Migrating historical data from a reporter database: scenario for Operations Management on IBM Cloud Private

To learn about Operations Management on IBM Cloud Private, you can install a historical data set. Learn how to install and load historical data from a reporter database and train the system.

Before you begin

Before you complete these steps, complete the following prerequisite items:

- The **ea-events-tooling** container is installed with the helm chart. It is not started as a pod and contains scripts to install data on the system, which can be run with the `kubectl run` command. Determine the location of the **ea-events-tooling** container image on your system by running the following command: **kubectl get image ea-events-tooling -o json** You can also run the `helm status` *`helm_release_name`* `--tls` command. The **--image** parameter specifies the container image location. Launch the image with a `kubectl run` command.
- Determine the image version number *<image_tag>* of the **ea-events-tooling** container, launch IBM Cloud Private. Select **Menu** > **Container images** and click the **ea-events-tooling** container to display tag details.
- Install a Kubernetes client and configure it to connect to the cluster. For more information, see https://www.ibm.com/support/knowledgecenter/SSBS6K_3.2.0/manage_cluster/install_kubectl.html.

About this task

You can use a script in the **ea-events-tooling** container to install historical data on the system. Run the `jdbctoingestionservice.sh` script to load historical data to the ingestion service. This script loads historical data from your reporter database into the ingestion service. The script runs inside the Kubernetes cluster. The virtual machines, which host the IBM Cloud Private worker nodes, must have access to the JDBC port on your reporter database server.

Procedure

1. Run the `jdbctoingestionservice.sh` script:

```
export HELM_RELEASE=<helm-release>
export WAS_PASSWORD=$(kubectl get secret $HELM_RELEASE-was-secret
-o jsonpath --template '{.data.WAS_PASSWORD}'  | base64 --decode)
 kubectl run ingesthttp -i --restart=Never --image=<values:global:image:repository>
/ea-events-tooling:<image_tag>  --image-pull-policy=Always
--env=INPUT_JDBC_HOSTNAME=<JDBC server host>
--env=INPUT_JDBC_PORT=<JDBC server port> --env=INPUT_JDBC_USERID=$JDBC_USER
--env=INPUT_JDBC_PASSWORD=$JDBC_PASSWORD  --env=LICENSE=accept --env=HTTP_PASSWORD=
$WAS_PASSWORD
jdbctoingestionservice.sh -- -r $HELM_RELEASE -s "start-time" -e "end-time"
```

   Where:

   - *$HELM_RELEASE* is the release of the helm chart and corresponds to the **NAME** field in the output from the `helm list` command.
   - *<values:global:image:repository>* is the location of the **ea-events-tooling** container, as described earlier.
   - *<image_tag>* is the image version tag, as described earlier.
   - *<JDBC server host>* and *<JDBC server port>* are the host name and port of the server that runs the Db2 instance and hosts the reporter database.
   - *$JDBC_USER* and *$JDBC_PASSWORD* are the user ID and password of a user with read access to the reporter database.
   - Optional: *start-time* and *end-time* are the start and end times to train against. You can specify the start or end time, neither, or both. If neither are specified, the current time is used as the end time and the start time is 93 days before the end time. You can either specify the start and end times with an integer Epoch time format in milliseconds, or with the default date string formatting for the system. Run the **./runTraining.sh -h** command to determine the default date formatting.
2. Train the system with the new data. Run the `runTraining.sh` script:

```
kubectl run trainer -it --command=true --restart=Never --env=LICENSE=accept
--image=<values:global:image:repository>/ea-events-tooling:<image_tag>
runTraining.sh -- -r $HELM_RELEASE [-t tenantid] [-a algorithm]
[-s start-time] [-e end-time]
```

   Where:

   - *algorithm* is either **related-events** or **seasonal-events**. If not specified, defaults to **related-events**.
   - *tenantid* is the tenant ID associated with the data that is ingested, as specified by the `global.common.eventanalytics.tenantId` parameter in the `values.yaml` file that is associated with the Helm chart.
   - *start-time* and *end-time* are the start and end times to train against. Specify the same values that you passed to the `jdbctoingestionservice.sh` in step “1” on page 353.

1.6.0.1 Cloud Native Analytics Service Monitoring

A self-monitoring policy can be enabled to provide assurance that Cloud Native Analytics is processing events. This policy is disabled by default.

About this task

When enabled, the Cloud Native Analytics self-monitoring policy causes OMNIbus to create an event every minute, with an *Identifier* field value of `Event Analytics Service Monitoring`. This heartbeat event follows the usual pathway of an event through the backend services of Cloud Native Analytics. At the end of the pathway, the *Grade* field value of the event is set to the time at which the event was processed by the Cloud Native Analytics self-monitoring policy. The heartbeat event is visible in the Event Viewer, and when Cloud Native Analytics is functioning correctly the time-stamp value in its *Grade* field increases every minute. If Cloud Native Analytics self-monitoring is enabled and the time-stamp in the Grade field of the heartbeat event is not incrementing after a few minutes, then there might be a failure in one of the Cloud Native Analytics services.

To enable Cloud Native Analytics self-monitoring, use the following procedure.

Procedure

1. Retrieve the password that is required to create your new policy.

   ```
   kubectl get secret helm_release_name-systemauth-secret -o=jsonpath='{.data.password}'
   ```

   Where *helm_release_name* is the name of your Cloud Native Analytics deployment.
2. Decode the password.

   ```
   encoded_password | base64 --decode
   ```

   Where *encoded_password* is the output from step 1.
3. Create a file called `selfMonitoring.json` with the following content:

   ```
   {
   "groupid": "self_monitoring",
   "type": "enrich",
   "dynamic": true,
   "configuration": {
   "deployed": true,
   "properties": [
   "/details/ScopeID"
   ]
   },
   "metadata": {
   "model": {
   "analytic": "self_monitoring",
   "type": "analytic"
   }
   },
   "resolver": {
   "stub": "com.ibm.itsm.inference.resolver.SelfMonitoringResolver",
   "version": "1.0.1"
   }
   }
   ```

4. Find the ingress point for the policy registry by running the following command, and finding the value of *spec.rules.host* from the output.

   ```
   kubectl get ingress | grep backend-ingress
   ```

5. Run the following command

   ```
   curl -u system:password --insecure -X POST "https://ingress-point/api/policies/system/v1/ea-
   generic-tenant/policies/system" -H "accept: application/json" -H "Content-Type: application/
   json" -d @selfMonitoring.json
   ```

   Where

- *password* is the decoded password output from step 2.
- *ingress-point* is the ingress point for the policy registry, as found in the previous step.

Note: Self-monitoring events are not archived.

1.6.0.1 Disabling Cloud Native Analytics Service Monitoring

Cloud Native Analytics self-monitoring can be disabled by removing the self-monitoring policy.

About this task

Procedure

1. Retrieve the password that is required to access the policy registry.

   ```
   kubectl get secret helm_release_name-systemauth-secret -o=jsonpath='{.data.password}'
   ```

   Where *helm_release_name* is the name of your Cloud Native Analytics deployment.
2. Decode the password.

   ```
   encoded_password | base64 --decode
   ```

   Where *encoded_password* is the output from step 1.
3. Find the ingress point for the policy registry by running the following command, and finding the value of *spec.rules.host* from the output.

   ```
   kubectl get ingress | grep backend-ingress
   ```

4. Run the following command to list all the policies. Then, find the policy ID of the policy that has a *group_id* value of `self-monitoring`.

   ```
   curl -u system:password --insecure -X GET "https://ingress-point/api/policies/system/v1/ea-
   generic-tenant/policies/system" -H "accept: application/json" -H "Content-Type: application/
   json"
   ```

   Where

   - *password* is the decoded password output from step 2.
   - *ingress-point* is the ingress point for the policy registry, as found in step 3.
5. Delete the self-monitoring policy with the following command:

   ```
    curl -u system:password --insecure -X DELETE "https://ingress-point/api/policies/system/v1/
   ea_generic_tenant/policies/batch/system" -H "accept: application/json" -H "Content-Type:
   application/json" -d "["policyid"]"
   ```

   Where

   - *password* is the decoded password output from step 2.
   - *ingress-point* is the ingress point for the policy registry, as found in step 3.
   - *policy_id* is the ID of the self-monitoring policy that you want to delete, as found in step 4.

1.6.0.1 Topology Analytics on IBM Cloud Private

If IBM Agile Service Manager is installed, and integration with it is enabled, then you are able to see topological context for your Cloud Native Analytics events, where there is an associated topology.

Prerequisites

For the Agile Service Manager integration to be enabled, *Enable ASM integration* must be set to true when you install Operations Management on IBM Cloud Private. For more information about this parameter, see “Configuring Installation Parameters for Operations Management on IBM Cloud Private” on page 125.

To use the Topology Analytics capability, the column *CEAAsmStatusDetails* must be selected as a **Display Column** in the Web GUI view that is used in the **Event Viewer** and **Incident Viewer**. This column toggles the display of the **Topology** column for incidents and events.

Event Viewer

Cloud Native Analytics groups events into incidents where a temporal or scope-based correlation exists. (For more information, see “Scope-based groups” on page 346 and “Temporal groups” on page 345). Incidents are shown in the **Event Viewer**, and an incident can be expanded to show its composite events.

If Topology Analytics matches an event to a topological resource, then a topology icon is displayed in the **Topology** column for that event. If an event is not associated with any resources, then a topology icon is not shown in the **Topology** column for that event. If any of an incident's events have associated topology, then a topology icon is displayed in the **Topology** column for the incident.

Click an event's topology icon to launch directly to a full Agile Service Manager **Topology Viewer** seeded with the resource that the event occurred on. The topology is displayed as it existed at the time that the event status was observed.

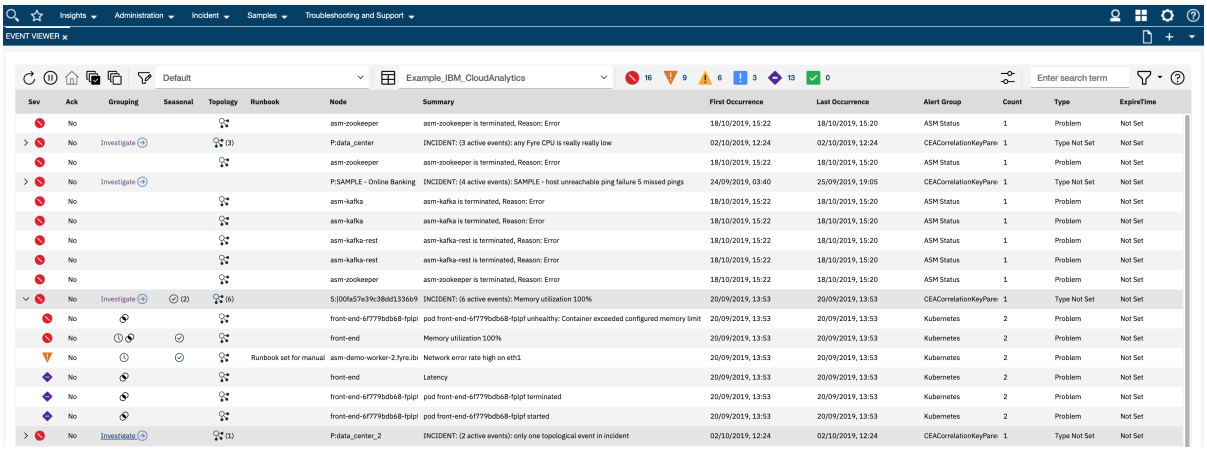

Figure 18. Event viewer

Incident Viewer

To open the **Incident Viewer** click **Investigate** on an incident in the **Event Viewer**. The **Incident Viewer** has an oblong pill icon in the **Topology** column if an incident or an event has associated topology. Click an event that has topological resources associated with it to launch an embedded one hop view of the topology of the resource that the event occurred on.

Context is shared between the Incident Viewer events and the one hop view to facilitate exploration of the topology of the event. The selection of an event causes associated resources to be highlighted in the one hop view. The selection of resources in the one hop view highlights any events that are associated with that resource in the event listing.

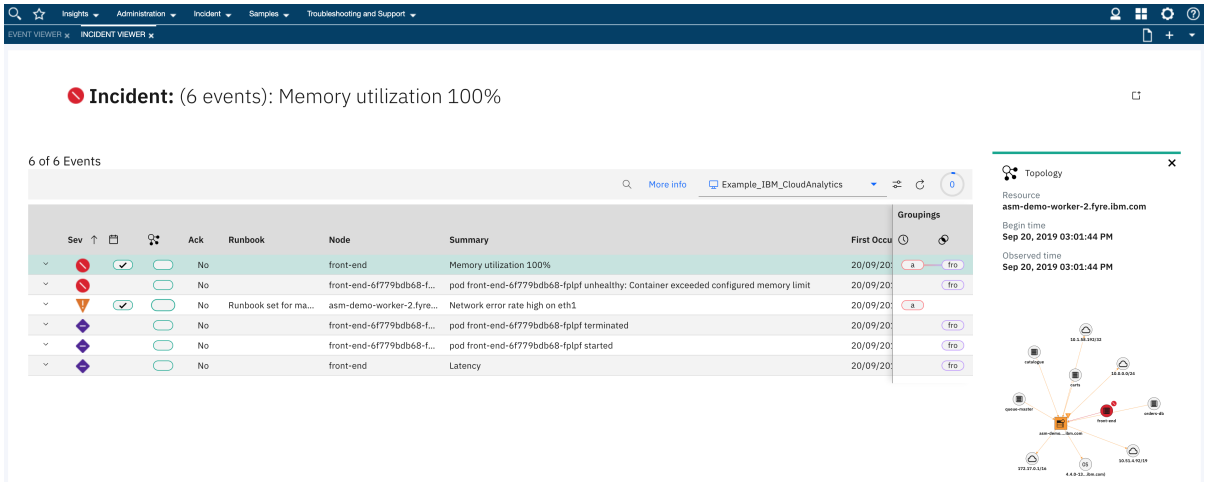

Figure 19. Incident viewer

To launch directly to a full Agile Service Manager Topology Viewer, seeded with the event's associated resource at the observation time, click **More Info**.

Launch to full Topology Viewer

Agile Service Manager's **Topology Viewer** can be used to explore the topological context of the resource that the event occurred on, and to view the history of the topology around the time of the event. For more information, see https://www.ibm.com/support/knowledgecenter/SS9LQB_1.1.6/Using/t_asm_viewatopologyhistory.html.

Right-clicking on a resource with status displays the associated event.

Chapter 11. Enabling topology search

The topology search capability is an extension of the Networks for Operations Insight feature. It applies the search and analysis capabilities of Operations Analytics - Log Analysis to give insight into network performance. Events that have been enriched with network data are analyzed by the Network Manager Insight Pack and are used to calculate the lowest-cost routes between two endpoints on the network topology over time. The events that occurred along the routes over the specified time period are identified and shown by severity. The topology search requires the Networks for Operations Insight feature to be installed and configured. The topology search capability can plot the lowest-cost route across a network between two end points and display all the events that occur on the devices on the routes.

The Network Manager IP Edition product enriches all the event data that is generated by the devices on the network topology. It is stored in Tivoli Netcool/OMNIbus, so that the Operations Analytics - Log Analysis product can cross-reference devices and events. The Gateway for Message Bus is used to pass event data from Tivoli Netcool/OMNIbus to Operations Analytics - Log Analysis. Also, the Network Manager Insight Pack reads topology data from the NCIM database in Network Manager IP Edition to identify the paths in the topology between the devices.

The scope of the topology search capability is that of the entire topology network, which includes all NCIM domains. To restrict the topology search to a single domain, you can configure a properties file that is included in the Network Manager Insight Pack.

After the Insight Pack is installed, you can run the apps from the Operations Analytics - Log Analysis UI. With Network Manager IP Edition installed and configured, the apps can also be run as right-click tools from the Network Views. With Tivoli Netcool/OMNIbus Web GUI installed and configured, the apps can be run as right-click tools from the Event Viewer and Active Event List (AEL).

The custom apps use the network-enriched event data and the topology data from the Network Manager IP Edition NCIM database. They plot the lowest-cost routes across the network between two nodes (that is, network entities) and count the events that occurred on the nodes along the routes. You can specify different time periods for the route and events. The algorithm uses the speed of the interfaces along the routes to calculate the routes that are lowest-cost. That is, the fastest routes from start to end along which a packet can be sent. The network topology is based on the most recent discovery. Historical routes are not accounted for. If your network topology is changeable, the routes between the nodes can change over time. If the network is stable, the routes stay current.

Before you begin

Ensure that you have a good knowledge of your network before you implement the topology search capability. Over large network topologies, the topology search can be performance intensive. It is therefore important to determine which parts of your network you want to use the topology search on. You can define those parts of the network into a single domain. Alternatively, implement the cross-domain discovery function in Network Manager IP Edition to create a single aggregation domain of the domains that you want to search. You can restrict the scope of the topology search to that domain or aggregation domain. For more information about deploying Network Manager IP Edition to monitor networks of small, medium, and larger networks, see https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/overview/concept/ovr_deploymentseg.html. For more information about the cross-domain discovery function, see https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/disco/task/dsc_configuringcrossdomaindiscoveries.html.

Related concepts

Network Manager Insight Pack

Supported products and components

The topology search capability is supported on a specific combination of products and components. Ensure that your environment has the requisite support before you enable topology search. These requirements apply to both new and upgraded environments.

The topology search capability requires the following products and components.

- Operations Analytics - Log Analysis V1.3 or later with the OMNIbusInsightPack_v1.3.1 and the NetworkManagerInsightPack_V1.3.0.0.
- Tivoli Netcool/OMNIbus Core V8.1.0.2 and Tivoli Netcool/OMNIbus Web GUI V8.1.0.2 or later. Install the **Install tools and menus for event search with IBM SmartCloud Analytics - Log Analysis** feature as part of the Web GUI installation.
- Gateway for Message Bus package version 6.0 or later. Earlier package versions do not include the configurations that are required for the topology search capability.
- Network Manager V4.1.1.1 or later. The topology search capability requires that the NCIM database for the network topology is IBM Db2 9.7 or 10.1. Oracle 10g or 11g is also supported, but requires more configuration than Db2. Although the Network Manager product supports other databases for storing the topology, the topology search capability is supported only on these databases.

Related tasks

Installing Netcool Operations Insight
Plan the installation and complete any pre-installation tasks before installing Netcool Operations Insight.

Network Manager Insight Pack

The Network Manager Insight Pack reads event data and network topology data so that it can be searched and visualized in the IBM Operations Analytics - Log Analysis product.

The Network Manager IP Edition product enriches all the event data that is generated by the devices on the network topology. It is stored in Tivoli Netcool/OMNIbus, so that the Operations Analytics - Log Analysis product can cross-reference devices and events. The Gateway for Message Bus is used to pass event data from Tivoli Netcool/OMNIbus to Operations Analytics - Log Analysis. Also, the Network Manager Insight Pack reads topology data from the NCIM database in Network Manager IP Edition to identify the paths in the topology between the devices.

The scope of the topology search capability is that of the entire topology network, which includes all NCIM domains. To restrict the topology search to a single domain, you can configure a properties file that is included in the Network Manager Insight Pack.

Related concepts

About cross-domain discoveries

Related tasks

Enabling topology search

Configuring cross-domain discoveries

Content of the Insight Pack

The data ingestion artifacts that are included in the Network Manager Insight Pack.

- Custom apps, which are described in Table 46 on page 361.

A rule set, source type, and collection are provided in the OMNIbusInsightPack_v1.3.1, which the Network Manager Insight Pack uses.

Custom apps

The following table describes the custom apps in the Insight Pack. After the Insight Pack is installed, you can run the apps from the Operations Analytics - Log Analysis UI. With Network Manager IP Edition

installed and configured, the apps can also be run as right-click tools from the Network Views. With Tivoli Netcool/OMNIbus Web GUI installed and configured, the apps can be run as right-click tools from the Event Viewer and Active Event List (AEL).

The custom apps use the network-enriched event data and the topology data from the Network Manager IP Edition NCIM database. They plot the lowest-cost routes across the network between two nodes (that is, network entities) and count the events that occurred on the nodes along the routes. You can specify different time periods for the route and events. The algorithm uses the speed of the interfaces along the routes to calculate the routes that are lowest-cost. That is, the fastest routes from start to end along which a packet can be sent. The network topology is based on the most recent discovery. Historical routes are not accounted for. If your network topology is changeable, the routes between the nodes can change over time. If the network is stable, the routes stay current.

The apps count the events that occurred over predefined periods of time, relative to the current time, or over a custom time period that you can specify. For the predefined time periods, the current time is calculated differently, depending on which product you run the apps from. Network Manager IP Edition uses the current time stamp. The Tivoli Netcool/OMNIbus Web GUI uses the time that is specified in the FirstOccurrence field of the events.

Table 46. Custom apps that are included in the Network Manager Insight Pack

Custom app name and file name	**Description**
Find alerts between two nodes on layer 2 topology `NM_Show_Alerts_Between_Two_Nodes_Layer2.app`	This app shows the distribution of alerts on the least-cost routes between two network end points in a layer 2 topology. Charts show the alert distribution by severity and alert group for each route over the specified time period. The ObjectServer field for the alert group is AlertGroup. A list of the routes is displayed from which you can search the events that occurred on each route over the specified time period. In the Operations Analytics - Log Analysis UI, the app requires search results before you can run it. In the search results, select the **NmosObjInst** column and then run the app. The app finds the events between the 2 nodes on which each selected event originated.
Find alerts between two nodes on layer 3 topology `NM_Show_Alerts_Between_Two_Nodes_Layer3.app`	This app shows the distribution of alerts on the least-cost routes between two network end points in a layer 3 topology. Charts show the alert distribution by severity and alert group for each route over the specified time period. The ObjectServer field for the alert group is AlertGroup. A list of the routes is displayed from which you can search the events that occurred on each route over the specified time period. In the Operations Analytics - Log Analysis UI, the app requires search results before you can run it. In the search results, select the **NmosObjInst** column and then run the app. The app finds the events between the 2 nodes on which each selected event originated.

Configuring topology search

Before you can use the topology search capability, configure the Tivoli Netcool/OMNIbus core and Web GUI components, the Gateway for Message Bus and Network Manager IP Edition.

Before you begin

Set up the environment for each product as follows:

- Configure the event search capability, including the Gateway for Message Bus. See "Configuring event search" on page 224. The topology search capability requires that the Gateway for Message Bus is configured to forward event data to Operations Analytics - Log Analysis.
- If your Operations Analytics - Log Analysis is upgraded from a previous version, migrate the data to your V1.3 instance. See one of the following topics:
  - Operations Analytics - Log Analysis V1.3.5: https://www.ibm.com/support/knowledgecenter/SSPFMY_1.3.5/com.ibm.scala.doc/admin/iwa_admin_backup_restore.html
  - Operations Analytics - Log Analysis V1.3.3: https://www.ibm.com/support/knowledgecenter/SSPFMY_1.3.3/com.ibm.scala.doc/admin/iwa_admin_backup_restore.html
- Ensure that the ObjectServer that forwards event data to Operations Analytics - Log Analysis has the **NmosObjInst** column in the alerts.status table. **NmosObjInst** is supplied by default and is required for this configuration. You can use ObjectServer SQL commands to check for the column and to add it if it is missing, as follows.
  - Use the DESCRIBE command to read the columns of the alerts.status table.
  - Use the ALTER COLUMN setting with the ALTER TABLE command to add **NmosObjInst** to the alerts.status table.

  For more information about the alerts.status table, including the **NmosObjInst** column, see https://ibm.biz/BdXcBF. For more information about ObjectServer SQL commands, see https://ibm.biz/BdXcBX.
- Configure the Tivoli Netcool/OMNIbus Web GUI V8.1.0.4 as follows:
  - Install the **Netcool Operations Insight Extensions for IBM Tivoli Netcool/OMNIbus Web GUI** package. IBM Installation Manager installs this package separately from the Web GUI. It needs to be explicitly selected.
  - Check the `server.init` file to ensure that the **`scala*`** properties are set as follows:

    ```
    scala.app.keyword=OMNIbus_Keyword_Search
    scala.app.static.dashboard=OMNIbus_Static_Dashboard
    scala.datasource=omnibus
    scala.url=protocol://host:port
    scala.version=1.2.0.3
    ```

    This configuration needed for new environments and for environments that are upgraded from versions of Operations Analytics - Log Analysis that are earlier than 1.2.0.3.
  - Set up the Web GUI Administration API client, which is needed to install the event list tooling that launches Operations Analytics - Log Analysis. See http://www-01.ibm.com/support/knowledgecenter/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/webtop/wip/task/web_con_setwaapiuserandpw.html.
- Install and configure the Insight Packs as follows:
  1. Install the OMNIbusInsightPack_v1.3.1. If your environment is upgraded from a previous version of Netcool Operations Insight, upgrade to this version of the Insight Pack. See "Netcool/OMNIbus Insight Pack" on page 218.
  2. Create a data source.
  3. Obtain and install the Network Manager Insight Pack V1.3.0.0. See "Installing the Network Manager Insight Pack" on page 90.

Procedure

1. In $NCHOME/omnibus/extensions, run the **nco_sql** utility against the scala_itnm_configuration.sql file.

```
./nco_sql -user root -password myp4ss -server NCOMS
< /opt/IBM/tivoli/netcool/omnibus/extensions/scala/scala_itnm_configuration.sql
```

   Triggers are applied to the ObjectServer that delay the storage of events until the events are enriched by Network Manager IP Edition data from the NCIM database.
2. If the Gateway for Message Bus is not configured to forward event data to Operations Analytics - Log Analysis, perform the required configurations.[4]
3. Install the tools and menus to launch the custom apps of the Network Manager Insight Pack in the Operations Analytics - Log Analysis UI from the Web GUI.

   In $*WEBGUI_HOME*/extensions/LogAnalytics, run the **runwaapi** command against the scalaEventTopology.xml file.

```
    $WEBGUI_HOME/waapi/bin/runwaapi -user username -password password -file
scalaEventTopology.xml
```

   Where *username* and *password* are the credentials of the administrator user that are defined in the $*WEBGUI_HOME*/waapi/etc/waapi.init properties file that controls the WAAPI client.
4. On the host where the Network Manager GUI components are installed, install the tools and menus to launch the custom apps of the Network Manager Insight Pack in the Operations Analytics - Log Analysis GUI from the Network Views.
   a) In $NMGUI_HOME/profile/etc/tnm/topoviz.properties, set the **topoviz.unity.customappsui** property, which defines the connection to Operations Analytics - Log Analysis.
   For example:

```
# Defines the LogAnalytics custom App launcher URL
topoviz.unity.customappsui=https://server3:9987/Unity/CustomAppsUI
```

   b) In the $NMGUI_HOME/profile/etc/tnm/menus/ncp_topoviz_device_menu.xml file, define the **Event Search** menu item.

   Add the item <menu id="Event Search"/> in the file as shown:

```
<tool id="showConnectivityInformation"/>
                <separator/>
                <menu id="Event Search"/>
```

5. Start the Gateway for Message Bus in Operations Analytics - Log Analysis mode.
   For example:

```
$OMNIHOME/bin/nco_g_xml -propsfile $OMNIHOME/etc/G_SCALA.props
```

   The gateway begins sending events from Tivoli Netcool/OMNIbus to Operations Analytics - Log Analysis.

What to do next

- Configure single sign-on (SSO) between the products.

[4] At a high-level, this involves the following:

- Creating a gateway server in the Netcool/OMNIbus interfaces file
- Configuring the G_SCALA.props properties file, including specifying the xml1302.map
- Configuring the endpoint in the scalaTransformers.xml file
- Configuring the SSL connection, if required
- Configuring the transport properties in the scalaTransport.properties file

- Reconfigure your views in the Web GUI to display the **NmosObjInst** column. The tools that launch the custom apps of the Network Manager Insight Pack work only against events that have a value in this column. See http://www-01.ibm.com/support/knowledgecenter/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/webtop/wip/task/web_cust_settingupviews.html.

Related tasks

Installing the Network Manager Insight Pack
This topic explains how to install the Network Manager Insight Pack into the Operations Analytics - Log Analysis product and make the necessary configurations. The Network Manager Insight Pack is required only if you deploy the Networks for Operations Insight feature and want to use the topology search capability. For more information, see "Network Manager Insight Pack" on page 360. Operations Analytics - Log Analysis can be running while you install the Insight Pack.

Related reference

Supported products and components

Related information

Gateway for Message Bus documentation

Configuring single sign-on for the topology search capability

Configure single sign-on (SSO) between the Dashboard Application Services Hub that hosts the Network Manager IP Edition GUI components and Operations Analytics - Log Analysis so that users can switch between the two products without having to log in each time. First, create dedicated users in your LDAP directory, which must be used by both products for user authentication, and then configure the SSO connection.

Procedure

1. Create the dedicated users and groups in your LDAP directory.
   For example:
   a. Create a new Organization Unit (OU) named NetworkManagement.
   b. Under the NetworkManagement OU, create a new group named itnmldap.
   c. Under the NetworkManagement OU, create the following new users: itnm1, itnm2, itnm3, and itnm4.
   d. Add the new users to the itnmldap group.
2. In Dashboard Application Services Hub, assign the itnmldap group that you created in step "1" on page 364 to a Network Manager IP Edition user group that can access the Network Views.
   Network Manager IP Edition user roles are controlled by assignments to user groups. Possible user groups that can access the Network Views are Network_Manager_IP_Admin and Network_Manager_User.
3. Configure the SSO connection from the Operations Analytics - Log Analysis product to the Dashboard Application Services Hub instance in which Network Manager IP Edition is hosted.
   For more information about configuring SSO for Operations Analytics - Log Analysis, see the Operations Analytics - Log Analysis documentation.
   The following steps of the Operations Analytics - Log Analysis SSO configuration are important:
   - Assign Operations Analytics - Log Analysis roles to the users and groups that you created in step "1" on page 364.
   - In the $*SCALAHOME*/wlp/usr/servers/Unity/server.xml/server.xml file, ensure that the <webAppSecurity> element has a httpOnlyCookies="false" attribute. Add this line before the closing </server> element. For example:

```
    <webAppSecurity ssoDomainNames="hostname" httpOnlyCookies="false"/>
</server>
```

The httpOnlyCookies="false" attribute disables the httponly flag on the cookie that is generated by Operations Analytics - Log Analysis and is required to enable SSO with Network Manager IP Edition GUI.

Related tasks

Configuring SSO between Operations Analytics - Log Analysis V1.3.5 and Dashboard Application Services Hub

Configuring SSO between Operations Analytics - Log Analysis V1.3.3 and Dashboard Application Services Hub

Using Topology Search

After the topology search capability is configured, you can have Operations Analytics - Log Analysis show you the events that occurred within a specific time period on routes between two devices in the network topology. This capability is useful to pinpoint problems on the network, for example, in response to a denial of service attack on a PE device.

The custom apps of the Network Manager Insight Pack can be run from the Operations Analytics - Log Analysis and, depending on your configuration, from the Network Views in Network Manager IP Edition and the event lists in the Web GUI. The custom apps support searches on Layer 2 and Layer 3 of the topology. The custom apps use the network-enriched event data and the topology data from the Network Manager IP Edition NCIM database. They plot the lowest-cost routes across the network between two nodes (that is, network entities) and count the events that occurred on the nodes along the routes. You can specify different time periods for the route and events. The algorithm uses the speed of the interfaces along the routes to calculate the routes that are lowest-cost. That is, the fastest routes from start to end along which a packet can be sent. The network topology is based on the most recent discovery. Historical routes are not accounted for. If your network topology is changeable, the routes between the nodes can change over time. If the network is stable, the routes stay current.

Before you begin

- Knowledge of the events in your topology is required to obtain meaningful results from the topology search, for example, how devices are named in your environment, or with what information devices are enriched. Device names are usually indicative of their functions. This level of understanding helps you run searches in Operations Analytics - Log Analysis.
- Configure the products to enable the topology search capability. See “Configuring topology search” on page 362.
- To avoid reentering user credentials when launching between products, configure SSO. See “Configuring single sign-on for the topology search capability” on page 364.
- Create the network views that visualize the parts of the network that you are responsible for and want to search. See https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/admin/task/adm_crtnwview.html.
- Reconfigure your views in the Web GUI to display the **NmosObjInst** column. The tools that launch the custom apps of the Network Manager Insight Pack work only against events that have a value in this column. See http://www-01.ibm.com/support/knowledgecenter/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/webtop/wip/task/web_cust_settingupviews.html.

Procedure

The flow of this procedure is to select the two nodes, select the tool and a time period over which the tool searches the historical event data. Then, in the Operations Analytics - Log Analysis UI, select the route that you are interested in and view the events. You can run searches on the events to refine the results.

1. Run the topology search from one of the products, as follows:
   - Web GUI event lists:
     a. In an Event Viewer or AEL, select two rows that have a value in the **NmosObjInst** column.

b. Right click and click **Event Search** > **Find events between two nodes** > **Layer 2 Topology** or **Event Search** > **Find events between two nodes** > **Layer 3 Topology**, depending on which layer of the topology you want to search.

c. Click a time filter, or click **Custom** and select one.

- Network Manager IP Edition network views:

  a. Select two devices.

  b. Click **Event Search** > **Find Events Between 2 Nodes** > **Layer 2 Topology** or **Event Search** > **Find Events Between 2 Nodes** > **Layer 3 Topology** depending on which layer of the topology you want to search.

  c. Click a time filter, or click **Custom** and select one.

- Operations Analytics - Log Analysis UI. In the Operations Analytics - Log Analysis UI, the app requires search results before you can run it. In the search results, select the **NmosObjInst** column. The app finds the events between the two nodes on which each selected event originated.

  Important: Select the **NmosObjInst** cells only. Do not select the entire rows. If you select the entire rows, no results are found, or incorrect routes between the entities on the network are found.

  In the **Search Dashboards** section of the UI, click **NetworkManagerInsightPack** > **Find events between two nodes on layer 2 topology** or **Find events between two nodes on layer 3 topology**, depending which network layer you want to view.

See "Example" on page 367 for an example of how to run the apps from the Operations Analytics - Log Analysis UI.

The results of the search are displayed on the Operations Analytics - Log Analysis UI as follows:

Find alerts between two nodes on layer 2 topology
: This app shows the distribution of alerts on the least-cost routes between two network end points in a layer 2 topology. Charts show the alert distribution by severity and alert group for each route over the specified time period. The ObjectServer field for the alert group is AlertGroup. A list of the routes is displayed from which you can search the events that occurred on each route over the specified time period.

Find alerts between two nodes on layer 3 topology
: This app shows the distribution of alerts on the least-cost routes between two network end points in a layer 3 topology. Charts show the alert distribution by severity and alert group for each route over the specified time period. The ObjectServer field for the alert group is AlertGroup. A list of the routes is displayed from which you can search the events that occurred on each route over the specified time period.

The apps count the events that occurred over predefined periods of time, relative to the current time, or over a custom time period that you can specify. For the predefined time periods, the current time is calculated differently, depending on which product you run the apps from. Network Manager IP Edition uses the current time stamp. The Tivoli Netcool/OMNIbus Web GUI uses the time that is specified in the FirstOccurrence field of the events.

Restriction: The Web GUI and Operations Analytics - Log Analysis process time stamps differently. The Web GUI recognizes hours, minutes, and seconds but Operations Analytics - Log Analysis ignores seconds. This problem affects the **Show event dashboard by node** and **Search for events by node**. If the time stamp 8 January 2014 07:15:26 AM is passed, Operations Analytics - Log Analysis interprets this time stamp as 8 January 2014 07:15 AM. So, the results of subsequent searches might differ from the search that was originally run.

2. From the bar charts, identify the route that is of most interest. Then, on the right side of the UI, click the link that corresponds to that route.

   A search result is returned that shows all the events that occurred within the specified time frame on that network route.

3. Refine the search results.

You can use the patterns that are listed in **Search Patterns**. For example, to search the results for critical events, click **Search Patterns** > **Severity** > **Critical**. A search string is copied to the search field. Then, click **Search**.

4. Extend and refine the search as required.

   For more information about searches in Operations Analytics - Log Analysis, see one of the following links:

   - Operations Analytics - Log Analysis V1.3.5: https://www.ibm.com/support/knowledgecenter/SSPFMY_1.3.5/com.ibm.scala.doc/use/iwa_using_ovw.html
   - Operations Analytics - Log Analysis V1.3.3: https://www.ibm.com/support/knowledgecenter/SSPFMY_1.3.3/com.ibm.scala.doc/use/iwa_using_ovw.html

Example

An example of how to run the custom apps from the Operations Analytics - Log Analysis UI. This example searches between 2 IP addresses: 172.20.1.3 and 172.20.1.5.

1. To run a new search, click **Add search** and type `NodeAlias:"172.20.1.3" OR NodeAlias:"172.20.1.5"`. Operations Analytics - Log Analysis returns all events that have the NodeAlias 172.20.1.3, or the NodeAlias 172.20.1.5.
2. In the results display, switch to grid view. Scroll across until you see the **NmosObjInst** column. Identify 2 rows that have different **NmosObjInst** values.
3. For these rows, select the cells in the **NmosObjInst** column.
4. In the **Search Dashboards** section of the UI, click **NetworkManagerInsightPack** > **Find events between two nodes on layer 2 topology** or **Find events between two nodes on layer 3 topology**, depending which network layer you want to view.

Related concepts

Network Management tasks
Use this information to understand the tasks that users can perform using Network Management.

Chapter 12. IBM Networks for Operations Insight

Networks for Operations Insight adds network management capabilities to the Netcool Operations Insight solution. These capabilities provide network discovery, visualization, event correlation and root-cause analysis, and configuration and compliance management that provide service assurance in dynamic network infrastructures. It contributes to overall operational insight into application and network performance management.

For documentation that describes how to install Networks for Operations Insight, see Performing a fresh installation. For documentation that describes how to upgrade from an existing Networks for Operations Insight, or transition to Networks for Operations Insight, see “Upgrading on premises” on page 145.

Before you begin

The Networks for Operations Insight capability is provided through setting up the following products in Netcool Operations Insight:

- Network Manager IP Edition, see Network Manager Knowledge Center
- Netcool Configuration Manager, see http://www-01.ibm.com/support/knowledgecenter/SS7UH9/welcome

In addition, you can optionally add on performance management capability by setting up the Network Performance Insight product and integrating it with Netcool Operations Insight. Performance management capability includes the ability to display and drill into performance anomaly and flow data. For more information on Network Performance Insight, see https://www.ibm.com/support/knowledgecenter/SSCVHB.

About Networks for Operations Insight

Networks for Operations Insight provides dashboard functionality that enables network operators to monitor the network, and network planners and engineers to track and optimize network performance.

About Networks for Operations Insight dashboards

As a network operator you can monitor network performance at increasing levels of detail. As a network planner or engineer, you can display the top 10 interfaces based on network congestion, traffic utilization, and quality of service (QoS). You can also run reports to show historical traffic traffic utilization information over different periods of time up to the last 365 days.

If you are a network operator, then use the dashboards available in Networks for Operations Insight to monitor performance at the level of detail that you require:

- Use the **Network Health Dashboard** to monitor performance of all devices in a network view.
- Use the **Device Dashboard** to monitor performance at the device or interface level.
- Use the **Traffic Details dashboard** to monitor traffic flow details across a single interface.

If you are a network planner or engineer, then use “Network Performance Insight Dashboards” on page 408 to view top 10 information on interfaces across your network, including the following:

- Congestion
- Traffic utilization
- Quality of service

You can also run flow data reports for any device and interface over different periods of time up to the last 365 days.

Scenario: Monitoring bandwidth usage

If a user or application is using a lot of interface bandwidth this can cause performance degradation across the network. This scenario shows you how to set up the **Device Dashboard** to monitor bandwidth usage on selected interfaces, and how to navigate into **Traffic Details dashboard** to see exactly which user or application is using the most bandwidth on that interface.

The first steps involve setting up thresholds for bandwidth performance monitoring. Once this is done, you can monitor a device or interface at metric level, and navigate to more detailed information, such as network flow through an interface, to determine what is causing performance degradation.

The steps are described in the following table.

Table 47. Scenario for bandwidth performance monitoring

Action	More information
1. Ensure that the poll definitions snmpInBandwidth and snmpOutBandwidth exist and are set up to poll the interfaces for which you want to monitor bandwidth.	Network Manager documentation: • Creating poll policies: https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/poll/task/poll_crtpoll.html • Creating poll defintions: https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/poll/task/crtpolldef.html
2. Within these poll definitions, define anomaly threshold settings. Performance anomalies in the **Device Dashboard** and in the **Event Viewer** are generated based on these threshold settings.	"Defining anomaly thresholds" on page 400
3. Enable the collection of flow data on the devices and interfaces of interest.	"Defining traffic flow thresholds" on page 399
4. Once your settings in steps 2 and 3 have taken effect, launch the **Device Dashboard** and point it at the device of interest. Within the **Performance Insights** portlet, proceed as follows: 1. Select the **Interfaces** tab. 2. Click **Metrics** and select **snmpInBandwidth** or **snmpOutBandwidth** and find the interface of interest in the table. **Note:** If the interface of interest is showing a performance anomaly, this means that bandwidth thresholds are being exceeded on this interface. Perform steps 5 and 6 to determine who is using the bandwidth.	"Monitoring performance data" on page 393
5. Right-click the interface of interest and select **Show Traffic Details**. The **Traffic Details dashboard** opens in a separate tab and displays traffic flow through the selected interface.	"Displaying traffic data on an interface" on page 398
6. Use the controls in the **Traffic Details dashboard** to display the traffic flow view of interest. For example, to see which source device and application is using the most bandwidth, display the **Top Sources with Application** view.	"Traffic Details dashboard views" on page 402 "Monitoring NetFlow performance data from Traffic Details dashboard" on page 407

About the Network Health Dashboard

Use the **Network Health Dashboard** to monitor a selected network view, and display availability, performance, and event data, as well as configuration and event history for all devices in that network view.

Related concepts

Network Management tasks
Use this information to understand the tasks that users can perform using Network Management.

Monitoring the network using the Network Health Dashboard

Use this information to understand how to use the **Network Health Dashboard** to determine if there are any network issues, and how to navigate from the dashboard to other parts of the product for more detailed information.

The **Network Health Dashboard** monitors a selected network view, and displays device and interface availability within that network view. It also reports on performance by presenting graphs, tables, and traces of KPI data for monitored devices and interfaces. A dashboard timeline reports on device configuration changes and event counts, enabling you to correlate events with configuration changes. The dashboard includes the event viewer, for more detailed event information.

Monitoring the Network Health Dashboard

Monitor the **Network Health Dashboard** by selecting a network view within your area of responsibility, such as a geographical area, or a specific network service such as BGP or VPN, and reviewing the data that appears in the other widgets on the dashboard. If you have set up a default network view bookmark that contains the network views within your area of responsibility, then the network views in that bookmark will appear in the network view tree within the dashboard.

Before you begin

For more information about the network view tree in the **Network Health Dashboard**, see "Configuring the network view tree to display in the Network Health Dashboard" on page 378

About this task

Note: The minimum screen resolution for display of the **Network Health Dashboard** is 1536 x 864. If your screen is less than this minimum resolution, then you will see scroll bars on one or more of the widgets in the **Network Health Dashboard**.

Displaying device and interface availability in a network view

Using the **Unavailable Resources** widget you can monitor, within a selected network view, the number of device and interface availability alerts that have been open for more than a configurable amount of time. By default this widget charts the number of device and interface availability alerts that have been open for up to 10 minutes, for more than ten minutes but less than one hour, and for more than one hour.

About this task

To monitor the number of open device and interface availability alerts within a selected network view, proceed as follows:

Procedure

1. **Network Health Dashboard**
2. In the **Network Health Dashboard**, select a network view from the network view tree in the **Network Views** at the top left. The other widgets update to show information based on the network view that you selected.

   In particular, the **Unavailable Resources** widget updates to show device and interface availability in the selected network view.

A second tab, called "**Network View**", opens. This tab contains a dashboard comprised of the **Network Views** GUI, the **Event Viewer**, and the **Structure Browser**, and it displays the selected network view. You can use this second tab to explore the topology of the network view that you are displaying in the **Network Health Dashboard**.

For information about specifying which network view tree to display in the **Network Health Dashboard**, see "Configuring the network view tree to display in the Network Health Dashboard" on page 378.

3. In the **Unavailable Resources** widget, proceed as follows:

   To determine the number of unavailable devices and interface alerts, use the following sections of the chart and note the colors of the stacked bar segments and the number inside each segment.

   Restriction: By default, all of the bars described below are configured to display. However, you can configure the **Unavailable Resources** widget to display only specific bars. For example, if you configure the widget to display only the **Device Ping** and the **Interface Ping** bars, then only those bars will be displayed in the widget.

   Note: By default the data in the **Unavailable Resources** widget is updated every 20 seconds.

   SNMP Poll Fail
   Uses color-coded stacked bars to display the number of SNMP Poll Fail alerts within the specified timeframe.

   SNMP Link State
   Uses color-coded stacked bars to display the number of SNMP Link State alerts within the specified timeframe.

   Interface Ping
   Uses color-coded stacked bars to display the number of Interface Ping alerts within the specified timeframe.

   Device Ping
   Uses color-coded stacked bars to display the number of Device Ping alerts within the specified timeframe.

   Color coding of the stacked bars is as follows:

   *Table 48. Color coding in the **Unavailable Resources** widget*

Yellow	Number of alerts that have been open for up to 10 minutes.
Pink	Number of alerts that have been open for more than 10 minutes and up to one hour.
Blue	Number of alerts that have been open for more than one hour.

   Click any one of these bars to show the corresponding alerts for the devices and interfaces in the **Event Viewer** at the bottom of the **Network Health Dashboard**.

   Note: You can change the time thresholds that are displayed in this widget. The default threshold settings are 10 minutes and one hour. If your availability requirements are less stringent, then you could change this, for example, to 30 minutes and 3 hours. The change applies on a per-user basis.

   If none of the devices in the current network view is being polled by any one of these polls, then the corresponding stacked bar will always displays zero values. For example, If none of the devices in the current network view is being polled by the SNMP Poll Fail poll, then the **SNMP Poll Fail** bar will always displays zero values. If you are able to access the **Configure Poll Policies** panel in the **Network Polling GUI**, then you can use the **Device Membership** field on that table to see a list all of devices across all network views that are polled by the various poll policies.

Displaying overall network view availability

You can monitor overall availability of chassis devices within a selected network view using the **Percentage Availability** widget.

About this task

To display overall availability of chassis devices within a selected network view, proceed as follows:

Procedure

1. **Network Health Dashboard**
2. In the **Network Health Dashboard**, select a network view from the network view tree in the **Network Views** at the top left. The other widgets update to show information based on the network view that you selected.

   In particular, the **Percentage Availability** widget updates to show overall availability of chassis devices in network view. A second tab, called "**Network View**", opens. This tab contains a dashboard comprised of the **Network Views** GUI, the **Event Viewer**, and the **Structure Browser**, and it displays the selected network view. You can use this second tab to explore the topology of the network view that you are displaying in the **Network Health Dashboard**.

   For information about specifying which network view tree to display in the **Network Health Dashboard**, see “Configuring the network view tree to display in the Network Health Dashboard” on page 378.
3. In the **Percentage Availability** widget, proceed as follows:

   The **Percentage Availability** widget displays 24 individual hour bars. Each bar displays a value, which is an exponentially weighted moving average of ping results in the past hour; the bar only appears on the completion of the hour. The bar value represents a percentage availability rate rather than a total count within that hour. The color of the bar varies as follows:

   - Green: 80% or more.
   - Orange: Between 50% and 80%.
   - Red: Less than 50%.

Displaying highest and lowest performers in a network view

You can monitor highest and lowest poll data metrics across all devices and interfaces within a selected network view using the **Top Performers** widget.

About this task

To display highest and lowest poll data metrics across all devices and interfaces within a selected network view, proceed as follows:

Procedure

1. **Network Health Dashboard**
2. In the **Network Health Dashboard**, select a network view from the network view tree in the **Network Views** at the top left. The other widgets update to show information based on the network view that you selected.

   In particular, the **Top Performers** widget updates to show overall availability of chassis devices in network view. A second tab, called "**Network View**", opens. This tab contains a dashboard comprised of the **Network Views** GUI, the **Event Viewer**, and the **Structure Browser**, and it displays the selected network view. You can use this second tab to explore the topology of the network view that you are displaying in the **Network Health Dashboard**.

   For information about specifying which network view tree to display in the **Network Health Dashboard**, see “Configuring the network view tree to display in the Network Health Dashboard” on page 378.
3. In the **Top Performers** widget, proceed as follows:

Select from the following controls to display chart, table, or trace data in the **Top Performers** widget.

Metric
Click this drop-down list to display a selected set of poll data metrics. The metrics that are displayed in the drop-down list depend on which poll policies are enabled for the selected network view. Select one of these metrics to display associated data in the main part of the window.

Order
Click this drop-down list to display what statistic to apply to the selected poll data metric.

- Statistics available for all metrics, except the `SnmpLinkStatus` metric.

  From Top: Displays a bar chart or table that shows the 10 highest values for the selected metric. The devices or interfaces with these maximum values are listed in the bar chart or table.

  From Bottom: Displays a bar chart or table that shows the 10 lowest values for the selected metric. The devices or interfaces with these minimum values are listed in the bar chart or table.

- Statistics available for the `SnmpLinkStatus` metric. In each case, a bar chart or table displays and shows devices for the selected statistic.

  Unavailable: This statistic displays by default. Devices with this statistic are problematic.

  Admin Down Devices with this statistic are not problematic as Administrators change devices to this state.

  Available Devices with this statistic are not problematic.

Note: The widget lists devices or interfaces depending on which metric was selected:

- If the metric selected applies to a device, such as memoryUtilization, then the top 10 list contains devices.
- If the metric selected applies to an interface, such as ifInDiscards, then the top 10 list contains interfaces.

 **Show Chart**
Displays a bar chart with the 10 highest or lowest values. Show Chart is the display option when you first open the widget.

 **Show Table**
Displays a table of data associated with the 10 highest or lowest values.

 **Define Filter**

This button only appears if you are in **Show Table**  mode. Click here to define a filter to apply to the Top Performers table data.

The main part of the window contains the data in one of the following formats:

Chart
Bar chart with the 10 highest or lowest values. Click any bar in the chart to show a time trace for the corresponding device or interface.

Table
Table of data associated with the 10 highest or lowest values. The table contains the following columns:

- **Entity Name**: Name of the device or interface.
- **Show Trace**: Click a link in one of the rows to show a time trace for the corresponding device or interface.
- **Last Poll Time**: Last time this entity was polled.
- **Value**: Value of the metric the last time this entity was polled.

Trace
Time trace of the data for a single device or interface. Navigate within this trace by performing the following operations:

- Zoom into the trace by moving your mouse wheel forward.
- Zoom out of the trace by moving your mouse wheel backward.
- Double click to restore the normal zoom level.
- Click within the trace area for a movable vertical line that displays the exact value at any point in time.

Click one of the following buttons to specify which current or historical poll data to display in the main part of the window. This button updates the data regardless of which mode is currently being presented: bar chart, table, or time trace.

Restriction: If your administrator has opted not to store poll data for any of the poll data metrics in the **Metric** drop-down list, then historical poll data will not be available when you click any of the following buttons:

- **Last Day**
- **Last Week**
- **Last Month**
- **Last Year**

Current®
Click this button to display current raw poll data. When in time trace mode, depending on the frequency of polling of the associated poll policy, the time trace shows anything up to two hours of data.

Last Day
Click this button to show data based on a regularly calculated daily average.

- In bar chart or table mode, the top 10 highest or lowest values are shown based on a daily exponentially weighted moving average (EWMA).
- In time trace mode, a time trace of the last 24 hours is shown, based on the average values.

In the **Last Day** section of the widget EWMA values are calculated by default every 15 minutes and are based on the previous 15 minutes of raw poll data. The data presented in this section of the widget is then updated with the latest EWMA value every 15 minutes.

Last Week
Click this button to show data based on a regularly calculated weekly average.

- In bar chart or table mode, the top 10 highest or lowest values are shown based on a weekly exponentially weighted moving average (EWMA).
- In time trace mode, a time trace of the last 7 days is shown, based on the average values.

In the **Last Week** section of the widget EWMA values are calculated by default every 30 minutes and are based on the previous 30 minutes of raw poll data. The data presented in this section of the widget is then updated with the latest EWMA value every 30 minutes.

Last Month
Click this button to show data based on a regularly calculated monthly average.

- In bar chart or table mode, the top 10 highest or lowest values are shown based on a monthly exponentially weighted moving average (EWMA).
- In time trace mode, a time trace of the last 30 days is shown, based on the average values.

In the **Last Month** section of the widget EWMA values are calculated by default every two hours and are based on the previous two hours of raw poll data. The data presented in this section of the widget is then updated with the latest EWMA value every two hours.

Last Year
Click this button to show data based on a regularly calculated yearly average.

- In bar chart or table mode, the top 10 highest or lowest values are shown based on a yearly exponentially weighted moving average (EWMA).
- In time trace mode, a time trace of the last 365 days is shown, based on the average values.

In the **Last Year** section of the widget EWMA values are calculated by default every day and are based on the previous 24 hours of raw poll data. The data presented in this section of the widget is then updated with the latest EWMA value every day.

Displaying the Configuration and Event Timeline

You can display a timeline showing, for all devices in a selected network view, device configuration changes and network alert data over a time period of up to 24 hours using the **Configuration and Event Timeline** widget. Correlation between device configuration changes and network alerts on this timeline can help you identify where configuration changes might have led to network issues.

About this task

To display a timeline showing device configuration changes and network alert data for all devices in a selected network view, proceed as follows:

Procedure

1. **Network Health Dashboard**
2. In the **Network Health Dashboard**, select a network view from the network view tree in the **Network Views** at the top left. The other widgets update to show information based on the network view that you selected.

   In particular, the **Configuration and Event Timeline** updates to show configuration change and event data for the selected network view. A second tab, called "**Network View**", opens. This tab contains a dashboard comprised of the **Network Views** GUI, the **Event Viewer**, and the **Structure Browser**, and it displays the selected network view. You can use this second tab to explore the topology of the network view that you are displaying in the **Network Health Dashboard**.

   For information about specifying which network view tree to display in the **Network Health Dashboard**, see “Configuring the network view tree to display in the Network Health Dashboard” on page 378.
3. In the **Configuration and Event Timeline** widget, proceed as follows:

   Configuration changes displayed in the **Configuration and Event Timeline** can be any of the following. Move your mouse over the configuration change bars to view a tooltip listing the different types of configuration change made at any time on the timline.

   Note: If you do not have Netcool Configuration Manager installed, then no configuration data is displayed in the timeline.

   Changes managed by Netcool Configuration Manager

   These changes are made under full Netcool Configuration Manager control. The timeline differentiates between scheduled or policy-based changes, which can be successful (Applied) or unsuccessful (Not Applied), and one-time changes made using the IDT Audited terminal facility within Netcool Configuration Manager.

   Applied

   A successful scheduled or policy-based set of device configuration changes made under the control of Netcool Configuration Manager.

   Not Applied

   An unsuccessful scheduled or policy-based set of device configuration changes made under the control of Netcool Configuration Manager.

   IDT

   Device configuration changes made using the audited terminal facility within Netcool Configuration Manager that allows one-time command-line based configuration changes to devices.

Unmanaged changes

OOBC
Out-of-band-change. Manual configuration change made to device where that change is outside of the control of Netcool Configuration Manager.

Events are displayed in the timeline as stacked bars, where the color of each element in the stacked bar indicates the severity of the corresponding events. Move your mouse over the stacked bars to view a tooltip listing the number of events at each severity level. The X-axis granularity for both events and configuration changes varies depending on the time range that you select for the timeline.

*Table 49. X axis granularity in the **Configuration and Event Timeline***

If you select this time range	**Then the X axis granularity is**
6 hours	15 minutes
12 hours	30 minutes
24 hours	1 hour

For more detailed information on the different types of configuration change, see the Netcool Configuration Manager knowledge center at http://www-01.ibm.com/support/knowledgecenter/SS7UH9/welcome.

Select from the following controls to define what data to display in the **Configuration and Event Timeline**.

Time
Select the duration of the timeline:

- **6 Hours**: Click to set a timeline duration of 6 hours.
- **12 Hours**: Click to set a timeline duration of 12 hours.
- **24 Hours**: Click to set a timeline duration of 24 hours.

Events by Occurrence

- **First Occurrence**: Click to display events on the timeline based on the first occurrence time of the events.
- **Last Occurrence**: Click to display events on the timeline based on the last occurrence time of the events.

 **Show Table**
Displays the configuration change data in tabular form. The table contains the following columns.

Note: If you do not have Netcool Configuration Manager installed, then this button is not displayed.

- **Number**: Serial value indicating the row number.
- **Device**: Host name or IP address of the affected device.
- **Unit of Work (UoW)**: In the case of automated Netcool Configuration Manager configuration changes, the Netcool Configuration Manager unit of work under which this configuration change was processed.
- **Result**: Indicates whether the change was successful.
- **Start Time**: The time at which the configuration change began.
- **End Time**: The time at which the configuration change completed.
- **User**: The user who applied the change.
- **Description**: Textual description associated with this change.

 **Show Chart**
Click here to switch back to the default graph view.

Note: If you do not have Netcool Configuration Manager installed, then this button is not displayed.

Use the sliders under the timeline to zoom in and out of the timeline. The legend under the timeline shows the colors used in the timeline to display the following items:

- Event severity values.
- Configuration change types.

Note: If the integration with Netcool Configuration Manager has been set up but there is a problem with data retrieval from Netcool Configuration Manager, then the configuration change types shown in the legend are marked with the following icon: ☒

Configuring the Network Health Dashboard
As an end user, you can configure the **Network Health Dashboard** to display the data you want to see.

Configuring the network view tree to display in the Network Health Dashboard
As a user of the **Network Health Dashboard**, you can configure a default bookmark to ensure that you limit the data that is displayed in the **Network Health Dashboard** to the network views within your area of responsibility.

About this task
The network views tree in the **Network Health Dashboard** automatically displays the network views in your default network view bookmark. If there are no network views in your default bookmark, then a message is displayed with a link to the **Network Views** GUI, where you can add network views to your default bookmark. The network views that you add to your default bookmark will be displayed in the network tree within the **Network Health Dashboard**.

Complete the following steps to add network views to your default bookmark.

Procedure

1. Within the displayed message, click the link that is provided.
   The **Network Views** GUI opens in a second tab.
2. Follow the instructions in the following topic in the Network Manager Knowledge Center: https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/visualize/task/vis_addingnetworkviewstobookmark.html

Results
The network views tree in the **Network Health Dashboard** displays the network views in your newly configured default bookmark.

Configuring the Unavailable Resources widget
As a user of the **Network Health Dashboard**, you can configure which availability data is displayed in the **Network Health Dashboard** by the **Unavailable Resources** widget. For example you can configure the widget to display availability data based on ping polls only, and not based on SNMP polls. You can also configure the time duration thresholds to apply to availability data displayed in this widget. For example, by default the widget charts the number of device and interface availability alerts that have been open for up to 10 minutes, more than 10 minutes, and more than one hour. Yon can change these thresholds.

About this task

To configure which availability data is displayed by the **Unavailable Resources** widget, proceed as follows:

Procedure

1. **Network Health Dashboard**
2. In the **Unavailable Resources** widget, click **User Preferences** .
3. To configure the **Unavailable Resources** widget, use the following checkboxes and number steppers:

Device
Configure which device alerts to monitor in the **Unavailable Resources** widget in order to retrieve information on device availability. By default all of these boxes are checked.

Device Ping
Check the box to monitor Default Chassis Ping alerts. Selecting this option causes the **Unavailable Resources** widget to provide an indication of the number of open device ICMP (ping) polling alerts.

SNMP Poll Fail
Check the box to monitor SNMP Poll Fail alerts. Selecting this option causes the **Unavailable Resources** widget to provide an indication of the number of open SNMP Poll Fail alerts.

Interface
Configure which interface alerts to monitor in the **Unavailable Resources** widget in order to retrieve information on interface availability. By default all of these boxes are checked.

Interface Ping
Check the box to monitor Default Interface Ping alerts. Selecting this option causes the **Unavailable Resources** widget to provide an indication of the number of open interface ICMP (ping) polling alerts.

Link State
Check the box to monitor SNMP Link State alerts. Selecting this option causes the **Unavailable Resources** widget to provide an indication of the number of open SNMP Link State alerts.

Thresholds

Upper
Specify an upper threshold in hours and minutes. By default, the upper threshold is set to one hour. This threshold causes the chart in the **Unavailable Resources** widget to update as follows: when the amount of time that any availability alert in the selected network view remains open exceeds the one hour threshold, then the relevant bar in the **Unavailable Resources** chart updates to show this unavailability as a blue color-coded bar section.

Lower
Specify a lower threshold in hours and minutes. By default, the lower threshold is set to 10 minutes. This threshold causes the chart in the **Unavailable Resources** widget to update as follows: when the amount of time that any availability alert in the selected network view remains open exceeds the 10 minute threshold, then the relevant bar in the **Unavailable Resources** chart updates to show this unavailability as a as a pink color-coded bar section.

Configuring the Configuration and Event Timeline
You can configure which event severity values to display on the **Configuration and Event Timeline**.

About this task

To configure which event severity values to display on the **Configuration and Event Timeline**:

Procedure

1. **Network Health Dashboard**
2. In the **Configuration and Event Timeline** widget, click **User Preferences** .
3. To configure the **Configuration and Event Timeline**, use the following lists:

Available Severities
By default, lists all event severity values and these event severity values are all displayed in the **Configuration and Event Timeline**.

To remove an item from this list, select the item and click the right-pointing arrow. You can select and move multiple values at the same time.

Selected Severities
By default, no event severity values are displayed in this list. Move items from the **Available Severities** list to this list to show just those values in the **Configuration and Event Timeline**. For example, to show only Critical and Major in the **Configuration and Event Timeline**, move the Critical and Major items from the **Available Severities** list to the **Selected Severities** list.

To remove an item from this list, select the item and click the left-pointing arrow. You can select and move multiple values at the same time.

Administering the Network Health Dashboard

Perform these tasks to configure and maintain the **Network Health Dashboard** for users.

Before you begin
Device configuration change data can only be displayed in the **Configuration and Event Timeline** if the integration with Netcool Configuration Manager has been set up. For more information on the integration with Netcool Configuration Manager, see the following topic in the Network Manager Knowledge Center: https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/install/task/con_configintegrationwithncm.html

About this task

Note: The minimum screen resolution for display of the **Network Health Dashboard** is 1536 x 864. If your screen is less than this minimum resolution, then you will see scroll bars on one or more of the widgets in the **Network Health Dashboard**.

Configuring the Network Health Dashboard
As an administrator, you can configure how data is displayed, and which data is displayed in the **Network Health Dashboard**.

About this task

To fit the quantity of widgets onto a single screen, customers need a minimum resolution of 1536 x 864, or higher.

As an administrator, you can configure the **Network Health Dashboard** in a number of ways to meet the needs of your operators.

Changing the layout of the dashboard
You can change the layout of the dashboard. For example, you can reposition, or resize widgets. See the information about *Editing dashboard content and layout* on the Jazz for Service Management Knowledge Center: https://www.ibm.com/support/knowledgecenter/SSEKCU

Change the refresh period for all widgets on the Network Health Dashboard
The Network Manager widgets within the **Network Health Dashboard** update by default every 20 seconds. You can change this update frequency by performing the following steps.

Note: The **Event Viewer** widget updates every 60 seconds by default.

1. Edit the the following configuration file: `$NMGUI_HOME/profile/etc/tnm//nethealth.properties`.
2. Find the following line and update the refresh period to the desired value in seconds.

```
nethealth.refresh.period=60
```

3. Save the file.

4. Close and reopen the **Network Health Dashboard** tab to put the changes into effect.

Change the colors associated with event severity values used in the Configuration and Event Timeline

You can update the colors associated with event severity values used in the **Configuration and Event Timeline**, by performing the following steps:

1. Edit the following configuration file: `$NMGUI_HOME/profile/etc/tnm/status.properties.`
2. Find the properties `status.color.background.`*`severity_number`*, where *severity_number* corresponds to the severity number. For example 5 corresponds to Critical severity.
3. Change the RGB values for the severity values, as desired.
4. Save the file.

Disable launch of the Network View tab when selecting a network view in the Network Health Dashboard

When a user selects a network view in the **Network Health Dashboard**, by default a second tab is opened, called "**Network View**". This tab contains a dashboard comprised of the **Network Views** GUI, the **Event Viewer**, and the **Structure Browser**, and displaying the selected network view. If your network views are very large, then displaying this second tab can have an impact on system performance. To avoid this performance impact, you can disable the launch of the **Network View** tab by performing the following steps:

1. Edit the following configuration file: `$NMGUI_HOME/profile/etc/tnm/topoviz.properties`.
2. Find the following lines:

```
# Defines whether the dashboard network view tree fires a launchPage
event when the user clicks a view in the tree
topoviz.networkview.dashboardTree.launchpage.enabled=true
```

3. Set the property `topoviz.networkview.dashboardTree.launchpage.enabled` to `false`.
4. Save the file.

Troubleshooting the Network Health Dashboard

Use this information to troubleshoot the **Network Health Dashboard**.

Network Health Dashboard log files

Review the **Network Health Dashboard** log files to support troubleshooting activity.

The **Network Health Dashboard** log files can be found at the following locations:

*Table 50. Locations of **Network Health Dashboard** log files*

File	Location
Log file	$NMGUI_HOME/profile/logs/tnm/ncp_nethealth.0.log
Trace file	$NMGUI_HOME/profile/logs/tnm/ncp_nethealth.0.trace

Data sources for the Network Health Dashboard widgets

Use this information to understand from where the **Network Health Dashboard** widgets retrieve data. This information might be useful for troubleshooting data presentation issues in the **Network Health Dashboard**.

Configuration and Event Timeline widget

This widget is populated by the following integrations:

- Tivoli Netcool/OMNIbus integration that analyzes Tivoli Netcool/OMNIbus events and shows a count based on event severity in a specified period.
- Netcool Configuration Manager integration that that retrieves configuration change distribution.

Percentage Availability widget
The data source for this widget is the historical poll data table `pdEwmaForDay`. The widget displays data from the device poll PingResult from the `pdEwmaForDay` table, scoped as follows:

- Scope is the selected network view if called from the **Network Health Dashboard**
- Scope is the in-context devices or interfaces if called from a right-click command within a topology map.

Note: The widget is updated only at the end of the hour to which the data applies.

Top Performers widget
The data sources for this widgets are the various historical poll data tables:

- `pdEwmaForDay`
- `pdEwmaForWeek`
- `pdEwmaForMonth`
- `pdEwmaForYear`

The scope of the data is as follows:

- Scope is the selected network view if called from the **Network Health Dashboard**
- Scope is the in-context devices or interfaces if called from a right-click command within a topology map.

Unavailable Resources widget
This widget is populated by a Tivoli Netcool/OMNIbus integration that analyzes Tivoli Netcool/OMNIbus events and uses the event data to determine whether a device or interface is affected, and whether the issue is ICMP or SNMP-based.

Investigating data display issues in the Network Health Dashboard
If any of the widgets in the **Network Health Dashboard** are not displaying data, either there is no data to display, or there is an underlying problem that needs to be resolved As an administrator, you can configure poll policies and poll definition so that users are able to display the data that they need to see in the **Network Health Dashboard**. You can also explore other potential underlying issues, such as problems with the underlying systems that process and store historical poll data.

About this task
To configure poll policies, see the following topic in the Network Manager Knowledge Center: https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/poll/task/poll_creatingpollswithmultiplepolldefinitions.html.

When editing a poll policy editor, the following operations in the **Poll Policy Editor** are important for determining whether data from the poll policy will be available for display in the **Network Health Dashboard**:

Poll Enabled
Check this option to enable the poll policy.

Store?
Check this option to store historical data for the poll policy.

Note: Checking this option will activate the historical poll data storage system, which will store large amounts of data on your system. For more information on the historical poll data storage system, see the following topic in the Network Manager Knowledge Center: https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/poll/task/poll_administeringstorm.html

Network Views
In this tab, ensure that the poll policy is active in the network views that you want to monitor in the **Network Health Dashboard**.

Configure poll policies as follows in order to make data available in the various widgets of the **Network Health Dashboard**.

For information on the different Network Manager poll policies and poll definitions, see the following topic on the Network Manager Knowledge Center: https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/ref/reference/ref_pollingref.html

Unavailable Resources widget

Configure the following poll policies in order to make data available in the **Unavailable Resources** widget:

*Table 51. **Unavailable Resources** widget: which poll policies to configure*

If you want to show response data for	Then ensure that one or more of the following poll policies is enabled in the appropriate network views
Chassis devices based on ping polling	Any poll policy that uses one or more chassis ping poll definitions. An example of a poll policy of this type is the Default Chassis Ping poll policy
Interfaces based on ping polling	Any poll policy that uses one or more chassis ping poll definitions. An example of a poll policy of this type is the Default Interface Ping poll policy.
Chassis devices based on SNMP polling	Any poll policy that uses one or more SNMP poll definitions. An example of a poll policy of this type is the snmpInBandwidth poll policy.
Interfaces based on SNMP polling	SNMP Link State poll policy.

Percentage Availability widget

Enable the Default Chassis Ping poll policy in order to display overall chassis availability data in the **Percentage Availability** widget.

Top Performers widget

Metrics in the Top Performers widget

To show a specific metric in the **Top Performers** widget **Metric** drop-down list, you must enable the poll policy that contains a poll definition related to that metric. Alternatively, create a new poll definition and add it to an enabled poll policy.

Note: These must be poll definition that can be stored and that falls into one of the following types:

- Basic threshold
- Ping
- SnmpLinkState

For example, using the default poll policies and poll definitions provided with Network Manager, here are examples of poll policies to enable and the corresponding metric that will be made available in the **Metric** drop-down list:

*Table 52. **Top Performers** widget: examples of poll policies to configure*

To display this metric	Enable this poll policy	Which contains this poll definition
ifInDiscards	ifInDiscards	ifInDiscards
ifOutDiscards	ifOutDiscards	ifOutDiscards

Table 52. **Top Performers** widget: examples of poll policies to configure (continued)		
To display this metric	**Enable this poll policy**	**Which contains this poll definition**
snmpInBandwidth	snmpInBandwidth	snmpInBandwidth

Historical poll data in the Top Performers widget

You can display historical poll data for a metric in the **Top Performers** widget by clicking the **Last Day, Last Week**, **Last Month**, and **Last Year** buttons. To collect historical poll data to display in this way, you must select the option to store historical data for the related poll definition related to the metric. For example, using the default poll policies and poll definitions provided with Network Manager, here are examples of poll definitions to configure.

Note: Historical data will only be viewable in the **Top Performers** once it has been collected, processed, and stored in the NCPOLLDATA database. For example, if at the time of reading this you had selected the option to store poll data for a poll definition one month ago, then you will only see one month's worth of data in the **Last Year** option.

Table 53. **Top Performers** widget: examples of poll definitions to configure		
To display historical data for this metric	**Select the Store? option for this poll definition**	**Within this poll policy**
ifInDiscards	ifInDiscards	ifInDiscards
ifOutDiscards	ifOutDiscards	ifOutDiscards
snmpInBandwidth	snmpInBandwidth	snmpInBandwidth

Important: If you have correctly configured storage of historical poll data for the metrics you are interested in, but when you click any of the **Last Day**, **Last Week**, **Last Month**, and **Last Year** buttons you are not seeing any data, then there might be a problem with the underlying systems that process and store historical poll data. In particular, the Apache Storm system that processes historical poll data, might not be running, or Apache Storm might have lost connection to the NCPOLLDATA database, where historical poll data is stored. For more information, see the following topic in the Network Manager Knowledge Center: https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/poll/task/poll_administeringstorm.html

Configuration and Event Timeline

If a configuration with Netcool Configuration Manager was set up at installation time, but configuration change data does not appear in the **Configuration and Event Timeline** this might be due to integration issues. For more information on the integration with Netcool Configuration Manager, see the following topic in the Network Manager Knowledge Center: https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/install/task/con_configintegrationwithncm.html

Top Performers widget is unable to display values greater than 32 bit

The **Top Performers** widget is unable to display values greater than 32 bit. If no data is being displayed in the **Top Performers** widget for a selected metric, then this might be due to a number of factors. One possibility is that the value of data in that metric is greater than 32 bit.

If no data is being displayed in the **Top Performers** widget for a selected metric, then run the following SQL query to determine if there is an error, and if, what the error code is.

```
SELECT     errorcode, value, datalabel
FROM       ncpolldata.polldata pd
INNER JOIN ncpolldata.monitoredobject mo ON mo.monitoredobjectid
= pd.monitoredobjectid
WHERE datalabel =poll_of_interest
```

Error codes are listed in the ERROR_CODE values in KNP_POLL_DATA_COLLECTION Support document, at the following location: http://www.ibm.com/support/docview.wss?uid=swg21422092. The error code

112 indicates that metric contains a polled value that was greater than can be stored in a 32-bit integer field.

Percentage Availability widget takes a long time to refresh
If the **Percentage Availability** widget is taking a long time to refresh then one possible solution is to increase the number of threads available for this widget. This solution is most suitable for customers with large networks.

About this task

Increase the number of threads available by performing the following steps:

Procedure

1. Edit the the following configuration file: `$NMGUI_HOME/profile/etc/tnm//nethealth.properties`.
2. Find the following lines:

```
## Widget thread count for availability widget
nethealth.threads.availability=5
```

3. Increase the value of the `nethealth.threads.availability` property. The maximum possible value is 10.
4. Save the file.

Developing custom dashboards

You can create pages that act as "dashboards" for displaying information on the status of parts of your network or edit existing dashboards, such as the **Network Health Dashboard**. You can select from the widgets that are provided with Network Manager, Tivoli Netcool/OMNIbus Web GUI, and also from other products that are deployed in your Dashboard Application Services Hub environment.

About this task

For information on creating and editing pages in the Dashboard Application Services Hub, see the Jazz for Service Management information center at https://www.ibm.com/support/knowledgecenter/SSEKCU.

Before you begin

- – Determine which widgets you want on the page.
  - – If you want a custom Web GUI gauge on the page, develop the metric that will feed the gauge display.
  - – Decide which users, groups, or user roles you want to have access to the page and assign the roles accordingly.
  - – If you want the widgets to communicate in a custom wire, develop the wires that will control the communications between the widgets.

Related concepts

Network Management tasks
Use this information to understand the tasks that users can perform using Network Management.

Displaying an event-driven view of the network
You can configure a dashboard to contain a Dashboard Network Views widget and other widgets that are driven by network alert data. Under the configuration described here, when a user clicks a node in the network view tree in the Dashboard Network Views widget, the other widgets in the dashboard update to show data based on events for entities in the selected network view. This dashboard is useful for network operations centers that want to see the real-time situation on the network, as it provides a real-time view of network alerts and of the availability of devices and interfaces. No historical polling data is used by any of the widgets in this dashboard, so this provides an alternative to the **Network Health Dashboard** if polling data is not being stored.

Before you begin
Depending on the requirements you have of the page, perform some or all of the tasks described in "Developing custom dashboards" on page 385.

About this task

To develop a page that wires a Dashboard Network Views widget and other widgets that are driven by network alert data:

Procedure

1. Log in as a user that has the iscadmins role.
2. Create the page, assign it to a location in the navigation, and specify the roles that users need to view the page.

   The default location is **Default**, and in this task it is assumed that this default location is used. If you use a different location, then substitute your chosen location wherever you see the location **Default** used in this task.
3. Add the Dashboard Network Views widget to the page.
4. Add the **Configuration and Event Timeline** to the page.

   If you have configured the Network Manager integration with Netcool Configuration Manager then this widget displays a timeline up to a period of 24 hours, showing configuration change data and event data by first occurrence of the event. If you have not set up the integration, then the widget still displays event data on the timeline.
5. Add the **Unavailable Resources** widget to the page.

   This widget displays how many devices and interfaces within the selected network view are unavailable.
6. Add the **Event Viewer** to the page.
7. Click **Save and Exit** to save the page.

   Note: This page does not requires any wires. The Dashboard Network Views widget automatically broadcasts the NodeClickedOn event, and the other widgets automatically subscribe to this event and update their data accordingly.

Results
You can now click a network view in the network view tree in the Dashboard Network Views widget and the other widgets automatically update to show event data:

- The **Unavailable Resources** widget displays a bar chart showing how many devices and interfaces within the selected network view are unavailable. The exact data shown in this widget depends on whether the following poll policies are enabled:
  - Devices: Default Chassis Ping and SNMP Poll Fail poll policies must be enabled.
  - Interfaces: Default Interface Ping and SNMP Link State poll policies must be enabled.
- The **Configuration and Event Timeline** displays a timeline showing events by first occurrence, and, if the Netcool Configuration Manager integration is configured, configuration change data, for all entities in the network view.
- The **Event Viewer** shows events for all entities in the network view.

  Note: Clicking a bar in the **Unavailable Resources** widget further filters the **Event Viewer** to show only the availability events related to the devices or interfaces in that bar.

Displaying and comparing top performer data for entities in a network view
Create a dashboard containing multiple **Top Performers** widgets to enable you to compare historical poll data across multiple entities and metrics in a selected network view. This dashboard is particularly useful for background investigation and analysis in order to determine how devices and interfaces are performing over time and whether there are any underlying issues.

Before you begin
Depending on the requirements you have of the page, perform some or all of the tasks described in "Developing custom dashboards" on page 385.

About this task

To develop a page that wires a **Network Views** widget and multiple **Top Performers** widgets to enable you to compare historical poll data across multiple entities and metrics in a selected network view:

Procedure

1. Log in as a user that has the iscadmins role.
2. Create the page, assign it to a location in the navigation, and specify the roles that users need to view the page.

   The default location is **Default**, and in this task it is assumed that this default location is used. If you use a different location, then substitute your chosen location wherever you see the location **Default** used in this task.
3. Add the **Network Views** widget to the page.
4. Add two **Top Performers** widgets to the page.

   Note: Adding two **Top Performers** widgets enables you to perform basic comparisons, such as displaying metric traces on the same device or interface over two different time periods. You can add more than two **Top Performers** widgets, and this will provide the ability to perform comparisons across a wider range of data; for example, adding four **Top Performers** widgets enables you to display metric traces on the same device or interface over four different time periods.
5. Click **Save and Exit** to save the page.

   Note: This page does not requires any wires. The **Network Views** widget automatically broadcasts the NodeClickedOn event, and the other widgets automatically subscribe to this event and update their data accordingly.

What to do next
You can use this dashboard to compare metric traces or charts.

Example: comparing metric traces on the same device or interface over different time periods
Use the custom dashboard that contains the two **Top Performers** widgets to compare metric traces on the same device or interface over different time periods; for example, you might see a spike in the current raw data trace for a metric such as snmpInBandwidth on a specific interface. To determine if this is just an isolated spike or a more serious ongoing issue, you can, on the same dashboard, also display a trace for the same snmpInBandwidth metric on the same interface over a longer time period, such as the last day or last week, and then visually determine if there have been continual incidences of high snmpInBandwidth on this interface over the last day or week.

About this task

To use the dashboard to compare metric traces on the same device or interface over different time periods, proceed as follows:

Procedure

1. In the **Network Views** widget, select a network view.

   The two **Top Performers** widgets update to show data for the selected network view.
2. From each of the **Top Performers** widgets, click the **Metric** drop-down list and select a metric of interest; for example snmpInBandwidth. Select the same metric on both **Top Performers** widgets; this ensures that the top entity in both chart is always the same.
3. In one of the **Top Performers** widgets, click the top bar to show the trace for the entity with the top value in that chart.

This displays a time-based trace of current raw data for the snmpInBandwidth metric.

4. In the other **Top Performers** widget, click the top bar to show the trace for the entity with the top value in that chart.

   You are now showing the identical time trace in both widgets.

5. In the second **Top Performers** widget, change the timeframe; for example, click **Last Day**.

   You are now showing a current raw data trace of snmpInBandwidth data in the first widget, and a trace of the last day's worth of snmpInBandwidth data for the same interface in the second widget, and you can compare the more transient raw data in the first widget with data averages over last day.

Example: comparing different metric traces on the same device or interface

Use the custom dashboard that contains the two Top Performers widgets to compare different metric traces on the same device or interface; for example, you might see a number of incidences of high snmpInBandwidth on a specific interface over the last day. To determine if the high incoming SNMP bandwidth usage on this interface is affecting the outgoing SNMP bandwidth usage on that same interface, you can, on the same dashboard, also display a trace for the snmpOutBandwidth metric on the same interface and also over the last day, and then visually compare the two traces.

About this task

To use the dashboard to compare different metric traces on the same device or interface, proceed as follows:

Procedure

1. In the **Network Views** widget, select a network view.
2. From each of the **Top Performers** widgets, click the **Metric** drop-down list and select a metric of interest; for example snmpInBandwidth. Select the same metric on both **Top Performers** widgets; this ensures that the top entity in both chart is always the same.
3. In one of the **Top Performers** widgets, click the top bar to show the trace for the entity with the top value in that chart.

   This displays a time-based trace of current raw data for the snmpInBandwidth metric.

4. In the other **Top Performers** widget, click the top bar to show the trace for the entity with the top value in that chart.

   You are now showing the identical time trace in both widgets.

5. In the second **Top Performers** widget, click the **Metric** drop-down list and select snmpOutBandwidth.

   You are now displaying a trace of incoming SNMP bandwidth usage on the interface with the highest incoming SNMP bandwidth usage in the network view on one widget, and a trace of outgoing SNMP bandwidth usage on that same interface. You can now visually compare the two traces to see if there is any correlation.

Example: comparing different Top 10 metric charts

Use the custom dashboard that contains the two Top Performers widgets to compare different Top 10 metric charts. This enables you to see the potential impact of one metric on another across the devices that are showing the highest performance degradation on the first metric. For example, you might want to compare the chart showing those devices showing the highest ten incoming SNMP bandwidth usage values, with the chart showing those devices showing the highest ten outgoing SNMP bandwidth usage values.

About this task

To use the dashboard to compare different Top 10 metric charts, proceed as follows:

Procedure

1. In the **Network Views** widget, select a network view.

2. In one of the **Top Performers** widgets, click the **Metric** drop-down list and select a metric of interest; for example snmpInBandwidth.

   The **Top Performers** widget updates to show a bar chart of the ten interfaces in the network view with the highest incoming SNMP bandwidth usage values.
3. In the other **Top Performers** widget, click the **Metric** drop-down list and select a second metric of interest; for example snmpOutBandwidth.

   The **Top Performers** widget updates to show a bar chart of the ten interfaces in the network view with the highest outgoing SNMP bandwidth usage values.

Results
You can now compare the two charts to see if there is any correlation between the data.

Displaying network view event data in a gauge group
You can use wires to configure a **Network Views** widget and a Dashboard Application Services Hub gauge group to pass data between each other. Under the configuration described here, when a user clicks a node in the network view tree in the **Network Views** widget, the gauge group updates to show a number of status gauges: you can configure as many status gauges as desired. In the example described here, three gauges are configured: Severity 3 (minor), Severity 4 (major), and Severity 5 (critical), together with a number within each status gauge indicating how many events at that severity are currently present on the devices in that network view. The instructions in this topic describe a possible option for wiring the two widgets.

Before you begin
Depending on the requirements you have of the page, perform some or all of the tasks described in "Developing custom dashboards" on page 385.

About this task

To develop a page that wires a **Network Views** widget and a Dashboard Application Services Hub gauge group:

Procedure

1. Log in as a user that has the iscadmins role.
2. Create the page, assign it to a location in the navigation, and specify the roles that users need to view the page.

   The default location is **Default**, and in this task it is assumed that this default location is used. If you use a different location, then substitute your chosen location wherever you see the location **Default** used in this task.
3. Add the **Network Views** widget to the page.
4. Add the Dashboard Application Services Hub gauge group to the page.
5. Edit the gauge group widget.
6. Select a dataset for the gauge group. In the **Gauge Group: Select a Dataset** window, search for the **Netcool/OMNIbus WebGUI** > **All data** > **Filter Summary** dataset.

   One way to do this is as follows:

   a) In the search textbox at the top left of the **Gauge Group: Select a Dataset** window, type `filter`.

   b) Click **Search**. This search retrieves two **Filter Summary** datasets.

   c) Select the dataset that has a provider title labeled **Provider: Netcool/OMNIbus WebGUI** > **Datasource: All data**.
7. Configure how you want the gauge group to be displayed. In the **Gauge Group: Visualization Settings** window, add three value status gauges by performing the following steps:.

   a) Click **Choose Widget** and select **ValueStatus Gauge** from the drop-down list. Then click **Add**

   b) Add two more **ValueStatus Gauge** widgets, following the instruction in the previous step.

There should now be three **ValueStatus Gauge** widgets listed in the **Selected Widgets** list.

8. Configure the three value status gauges to show the following:
   - First value status gauge will display the number of Severity 3 (minor) events, within the Severity 3 (minor) symbol, .
   - Second value status gauge will display the number of Severity 4 (major) events, within the Severity 4 (major) symbol, .
   - Third value status gauge will display the number of Severity 5 (critical) events, within the Severity 5 (critical) symbol, .

   Perform the following steps to configure the Severity 3 (minor) value status gauge:
   a) Select the first value status gauge item in the **Selected Widgets** list.
   b) Click **Required Settings**.
   c) Click the **Value** drop-down list and select **Severity 3 Event Count** from the drop-down list.
   d) Click **Optional Settings**.
   e) Click the **Label above Gauge** drop-down list and select **Severity 3 Event Count Name** from the drop-down list.
   f) In the **Minor** spinner set a threshold value of 0 by typing 0.
      This threshold value causes any number of Severity 3 (minor) events to generate a Severity 3 value status gauge.

   Perform the following steps to configure the Severity 4 (major) value status gauge:
   a) Select the first value status gauge item in the **Selected Widgets** list.
   b) Click **Required Settings**.
   c) Click the **Value** drop-down list and select **Severity 4 Event Count** from the drop-down list.
   d) Click **Optional Settings**.
   e) Click the **Label above Gauge** drop-down list and select **Severity 4 Event Count Name** from the drop-down list.
   f) In the **Major** spinner set a threshold value of 0 by typing 0.
      This threshold value causes any number of Severity 4 (major) events to generate a Severity 4 value status gauge.

   Perform the following steps to configure the Severity 5 (critical) value status gauge:
   a) Select the first value status gauge item in the **Selected Widgets** list.
   b) Click **Required Settings**.
   c) Click the **Value** drop-down list and select **Severity 5 Event Count** from the drop-down list.
   d) Click **Optional Settings**.
   e) Click the **Label above Gauge** drop-down list and select **Severity 5 Event Count Name** from the drop-down list.
   f) In the **Critical** spinner set a threshold value of 0 by typing 0.
      This threshold value causes any number of Severity 5 (critical) events to generate a Severity 5 value status gauge.
9. Click **Save and Exit** to save the page.
10. From the page action list, select **Edit Page**.
11. Click **Show Wires**  and then, in the **Summary of wires** section of the window. click **New Wire**.
12. Specify the wires that connect the **Network Views** widget to the Dashboard Application Services Hub gauge group.

- In the **Select Source Event for New Wire** window, click **Network Views** > **NodeClickedOn**, and then click **OK**.
- In the **Select Target for New Wire** window, click **Default** > **This page *name_of_page*** > **Event Viewer** , where *name_of_page* is the name of the page that you created in step 2.
- In the **Transformation** window, select **Show Gauge Events**, and then click **OK**

13. Close the **Summary of wires** section of the window by clicking the X symbol at the top right corner.
14. Click **Save and Exit** to save the page.

Results
You can now click a network view in the network view tree in the **Network Views** widget and have the gauge group update to show three status values: Severity 3 (minor), Severity 4 (major), and Severity 5 (critical), together with a number within each status gauge indicating how many events at that severity are currently present on the devices in the selected network view.

Event information for Network Health Dashboard widgets
Refer to this table to get information about the publish events and subscribe events for **Network Health Dashboard** widgets. Use this event information when you create a new custom widget and you want to wire your custom widget with an existing **Network Health Dashboard** widget.

*Table 54. Event information for **Network Health Dashboard** widgets*

Widget name	Event type	Event name	Event description
Configuration and Event Timeline	Subscribe event	`NodeClickedOn`	Subscribes to a `NodeClickedOn` event and displays data based on the events *ViewId* and the *datasource*.
Percentage Availability	Subscribe event	`NodeClickedOn`	Subscribes to a `NodeClickedOn` event and displays data based on the events *ViewId* and the *datasource*.
Network Manager Polling Chart	Subscribe event	`NodeClickedOn`	Subscribes to a `NodeClickedOn` event and displays data based on the events *ViewId* and the *datasource*.
Unavailable Resources	Publish event	`showEvents`	Click a bar in the displayed graph and the widget publishes a `showEvents` event that contains the name of the transient filter.
	Subscribe event	`NodeClickedOn`	Subscribes to a `NodeClickedOn` event and displays data based on the events *ViewId* and the *datasource*.

Device Dashboard

Use the **Device Dashboard** to troubleshoot network issues by navigating the network topology and seeing performance anomalies and trends on any device, link, or interface.

The content of the **Device Dashboard** varies depending on whether you installed Network Performance Insight.

- If Network Performance Insight is installed, then the **Device Dashboard** includes the **Performance Insights** widget. The **Performance Insights** widget shows performance measure values, anomalies, and trends for the selected entity. For more information, see "Monitoring performance data" on page 393.
- If Network Performance Insight is not installed, then the **Device Dashboard** does not include the **Performance Insights** widget. Instead, it displays the **Top Performers** widget, showing performance measure values only for the selected entity. For more information, see "Displaying highest and lowest performers in a network view" on page 373.

Related tasks

Installing the Device Dashboard
Install the Device Dashboard to view event and performance data for a selected device and its interfaces on a single dashboard.

Troubleshooting network issues using the Device Dashboard

Use the **Device Dashboard** to troubleshoot any network issue on a device, link, or interface.

Starting the Device Dashboard

You can start the **Device Dashboard** from an event in the **Event Viewer**, from a device in the Network Views, Network Hop View, or Path Views. You can also start the **Device Dashboard** from the **Top Performers** widget within the **Network Health Dashboard**.

About this task

Open the **Device Dashboard** using one of the following options:

- Open the **Device Dashboard** from the **Top Performers** widget in the **Network Health Dashboard**, by completing one of the following tasks.
  - Press `Ctrl` and click a bar in the Chart view.
  - Click an entity name in the Table view.

  The **Device Dashboard** opens in a new Dashboard Application Services Hub tab. The **Device Dashboard** opens in the context of the device or interface that is associated with the element that you clicked.
- Open the **Device Dashboard** from any of the topology GUIs. Right-click a device, link, or interface and click **Performance Insight** > **Show Device Dashboard**.

  The topology GUIs include the **Network Hop View**, **Network Views**, **Path Views GUI**, and **Structure Browser**.

  The **Device Dashboard** opens in a new tab, in the context of the device, link, or interface associated with the element that you clicked.
- To open the **Device Dashboard** from the **Event Viewer**, right-click an event and click **Performance Insight** > **Show Device Dashboard**.

  The **Device Dashboard** opens in a new tab, in the context of the device or interface that is associated with the element that you clicked.

Changing Device Dashboard focus

You can change the **Device Dashboard** focus to a different device or link.

Procedure

1. To switch focus to a device, click the device in the **Device** tab in the **Topology** widget. If the device is not shown in the **Topology** widget, complete the following steps.
   a) Search for the device by using the **Network Hop View** search feature. For more information, see https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/visualize/task/vis_searchfordevices.html.
      The **Topology** widget is updated to center on the device that is selected in the search results.
   b) In the **Topology** widget, click the device of interest.
2. To switch focus to a link, click the link in the **Device** tab in the **Topology** widget. If the link is not shown in the **Topology** widget, complete the following steps.
   a) Search for a device at one end of the link by using the **Network Hop View** search feature. For more information, see https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/visualize/task/vis_searchfordevices.html.
      The **Topology** widget is updated to center on the device that is selected in the search results.
   b) In the **Topology** widget, click the link of interest.

Results

All the widgets on the **Device Dashboard** update to show data for the device, interface, or link.

Monitoring performance data

You can monitor performance metrics for a device, link, or interface using the **Performance Insights** widget within the **Device Dashboard**.

Procedure

1. Launch the **Device Dashboard** as described in "Starting the Device Dashboard" on page 392.
2. Ensure that the device, link, or interface of interest is selected. To change the dashboard context, see "Changing Device Dashboard focus" on page 393.
3. In the **Performance Insights** widget, proceed as follows:

   Select from the following controls to display performance metric values, anomalies, and trends in the **Performance Insights** widget.

   Note: Performance anomalies are determined based on the application of static thresholds to performance data. Anomaly detection is an early warning system that indicates that a device, interface, or link, needs attention.

   Severity
   This takes the form of a square in a color that indicates the anomaly severity for this performance metric.

   - Red: indicates a higher severity anomaly.
   - Orange: indicates a lower severity anomaly.
   - Green: no anomaly.

   Links
   For each link displays the NCIM topology database entity identifiers of the devices at either end of the link. Roll over this field for more details of the devices, such as hostname and IP address.

   Connections
   Within each link, lists the connections that make up the link. Each connection is identified using the NCIM topology database entity identifiers of the interfaces at either end of the link. Roll over this field for more details of the interfaces, such as interface name, and interface description.

Devices

This tab contains performance data for the devices that you selected in the **Topology** portlet. If you selected interfaces, it shows the devices that contain them. If you selected links, it shows the devices at either end.

If there are performance anomalies associated with any of the metrics, then the number of the worst severity metrics is shown in a colored square in the tab header.

- A number in a red square indicates that there are higher severity metrics, and indicates the number of higher severity metrics. If a red square is showing, there might also be lower severity metrics.
- A number in an orange square indicates that there are lower severity metrics, but no higher severity metrics, and indicates the number of lower severity metrics.

Filter

Type any string to filter the rows displayed in the table.

Device

Lists the device or devices selected in the **Topology** widget. Expand the device node to see the metrics associated with the device.

Metric

Lists the performance metrics available for the associated device.

Last 30 Minutes

Displays a sparkline showing the trend of the performance metric over the last 30 minutes. The current value is shown using a colored dot at the end of the sparkline. The color of the dot indicates whether there is an anomaly associated with this performance metric.

- Red: indicates a higher severity anomaly.
- Orange: indicates a lower severity anomaly.
- Black: no anomaly.

Severity

Bullet graph, showing a central bar in a color that indicates the anomaly severity. The thresholds that have been applied to the performance metric are shown in different levels of gray shading in the wider bar. The color of the central bar specifies the anomaly severity.

- Red: indicates a higher severity anomaly.
- Orange: indicates a lower severity anomaly.
- Black: no anomaly.

Value

Indicates the value of the performance metric based on the most recent poll action.

Links

This tab contains performance metric anomaly and trend data for one or more links. In particular, it displays data for each connection on each of the selected links.

Note: The **Links** tab only appears if you selected one or more links in the **Topology** widget.

If there are performance anomalies associated with any of the metrics, then the number of the worst severity metrics is shown in a colored square in the tab header.

- A number in a red square indicates that there are higher severity metrics, and indicates the number of higher severity metrics. If a red square is showing, there might also be lower severity metrics.
- A number in an orange square indicates that there are lower severity metrics, but no higher severity metrics, and indicates the number of lower severity metrics.

Filter

Type any string to filter the rows displayed in the table.

Links
For each link displays the NCIM topology database entity identifiers of the devices at either end of the link. Roll over this field for more details of the devices, such as hostname and IP address.

Connections
Within each link, lists the connections that make up the link. Each connection is identified using the NCIM topology database entity identifiers of the interfaces at either end of the link. Roll over this field for more details of the interfaces, such as interface name, and interface description.

Metric
Lists the performance metrics available for the associated connection.

Last 30 Minutes
Displays a sparkline showing the trend of the performance metric over the last 30 minutes. The current value is shown using a colored dot at the end of the sparkline. The color of the dot indicates whether there is an anomaly associated with this performance metric.

- Red: indicates a higher severity anomaly.
- Orange: indicates a lower severity anomaly.
- Black: no anomaly.

Severity
Bullet graph, showing a central bar in a color that indicates the anomaly severity. The thresholds that have been applied to the performance metric are shown in different levels of gray shading in the wider bar. The color of the central bar specifies the anomaly severity.

- Red: indicates a higher severity anomaly.
- Orange: indicates a lower severity anomaly.
- Black: no anomaly.

Value
Indicates the value of the performance metric based on the most recent poll action.

Interfaces
Depending on the selection in the **Topology** widget, this tab contains performance data for the following interfaces:

If you selected in the Topology widget	Then this tab displays performance metric anomaly and trend data for	For more information, see
One or more devices	All the interfaces on these devices.	Interfaces tab: Content when a device or interface is selected in the Topology widget
One or more interfaces	The selected interfaces.	
One or more links	The interfaces at either end of the connections that make up these links.	Interfaces tab: Content when a link is selected in the Topology widget

If there are performance anomalies associated with any of the metrics, then the number of the worst severity metrics is shown in a colored square in the tab header.

- A number in a red square indicates that there are higher severity metrics, and indicates the number of higher severity metrics. If a red square is showing, there might also be lower severity metrics.
- A number in an orange square indicates that there are lower severity metrics, but no higher severity metrics, and indicates the number of lower severity metrics.

If there is performance data associated with the interface, you can view it by opening the **Traffic Details** widget in a new tab. Right-click an interface and select **Show Traffic Details**.

Interfaces tab: Content when a device or interface is selected in the Topology widget

Filter
Type any string to filter the rows displayed in the table.

Metric
Drop-down list that lists all of the metrics available on the selected interfaces.

Device
Lists the device or devices that contain the interfaces selected in the **Topology** widget. Expand the device node to see the all the interfaces in that device and associated status and metric data.

Interface
Interface name of the selected interfaces or of the interfaces on the selected devices.

Metric
Performance metric selected from the **Metric** drop-down list.

Last 30 Minutes
Displays a sparkline showing the trend of the performance metric over the last 30 minutes. The current value is shown using a colored dot at the end of the sparkline. The color of the dot indicates whether there is an anomaly associated with this performance metric.

- Red: indicates a higher severity anomaly.
- Orange: indicates a lower severity anomaly.
- Black: no anomaly.

Severity
Bullet graph, showing a central bar in a color that indicates the anomaly severity. The thresholds that have been applied to the performance metric are shown in different levels of gray shading in the wider bar. The color of the central bar specifies the anomaly severity.

- Red: indicates a higher severity anomaly.
- Orange: indicates a lower severity anomaly.
- Black: no anomaly.

Value
Indicates the value of the performance metric based on the most recent poll action.

Interfaces tab: Content when a link is selected in the Topology widget

Filter
Type any string to filter the rows displayed in the table.

Links
For each link displays the NCIM topology database entity identifiers of the devices at either end of the link. Roll over this field for more details of the devices, such as hostname and IP address.

Connections
Within each link, lists the connections that make up the link. Each connection is identified using the NCIM topology database entity identifiers of the interfaces at either end of the link. Roll over this field for more details of the interfaces, such as interface name, and interface description.

Device
Specifies the device at each end of the connection.

Interface
Specifies the interface at each end of the connection.

Metric
Performance metric selected from the **Metric** drop-down list.

Last 30 Minutes
Displays a sparkline showing the trend of the performance metric over the last 30 minutes. The current value is shown using a colored dot at the end of the sparkline. The color of the dot indicates whether there is an anomaly associated with this performance metric.

- Red: indicates a higher severity anomaly.
- Orange: indicates a lower severity anomaly.
- Black: no anomaly.

Severity
Bullet graph, showing a central bar in a color that indicates the anomaly severity. The thresholds that have been applied to the performance metric are shown in different levels of gray shading in the wider bar. The color of the central bar specifies the anomaly severity.

- Red: indicates a higher severity anomaly.
- Orange: indicates a lower severity anomaly.
- Black: no anomaly.

Value
Indicates the value of the performance metric based on the most recent poll action.

Displaying performance timelines

You can view performance metrics for a device, link, or interface in the **Performance Timeline** portlet in the **Device Dashboard**.

About this task

To display performance timelines, complete the following steps.

Procedure

1. Start the **Device Dashboard** as described in "Starting the Device Dashboard" on page 392.
2. Ensure that the device, link, or interface of interest is selected. To change the dashboard context, see "Changing Device Dashboard focus" on page 393.
3. In the **Performance Insights** portlet, click the sparkline that is associated with the metric for which you want to display a timeline.

   You might need to change to a different tab, for example, to the **Interfaces** tab, to find the metric that you want.

   The **Performance Timeline** portlet updates and displays a timeline of the performance metric that you clicked.

   The last 30 minutes of data are displayed on the upper timeline. By default, the last 12 hours of data are displayed on the lower timeline, which is called the viewfinder. You can configure the amount of data that is displayed in the viewfinder by updating the portlet configuration settings. The window in the viewfinder indicates which section of the longer timeline is being displayed in the upper timeline.
4. Optional: Investigate the performance data by performing the following actions.
   - To display a different time window, click in the viewfinder.The timeline updates to display the time period that you clicked.
   - To display a larger or smaller time window, drag the handles of the viewfinder window inwards or outwards. By default, the time period is 30 minutes.
   - To view extra information about a particular point, hover over the point in the timeline.

Displaying traffic data on an interface

You can display data about the traffic that is flowing through an interface by opening the **Traffic Details** portlet from an interface in the **Performance Insights** portlet.

About this task

To display traffic data for an interface, complete the following steps.

Procedure

1. Start the **Device Dashboard** as described in "Starting the Device Dashboard" on page 392.
2. Click the device of interest in the **Device** tab in the **Topology** portlet.
3. Click the **Interfaces** tab in the **Performance Insights** portlet.
4. Right-click an interface of interest and select **Show Traffic Details**.

Results

The **TrafficDetails** portlet in Network Performance Insight opens, and shows traffic information for the selected interface.

Configuring the Performance Insights widget

Both administrators and end users can configure the **Performance Insights** widget within the **Device Dashboard**.

About this task

To configure the **Performance Insights** widget, perform the following steps.

Procedure

1. In the **Performance Insights** widget, click **Edit Preferences**  .
2. To configure communication between the **Performance Insights** widget and Network Performance Insight, use the following controls:

   NPI Hostname
   Specify the hostname of Network Performance Insight server that provides performance data for this widget.

   Important: You must only ever change this setting based on instructions from a system administrator.

   Port
   Specify the port on the Network Performance Insight host. This is usually either 8080 or 9443.
3. To configure presentation and refresh of data on the **Performance Insights** widget, use the following controls:

   Time Period
   Specify how many minutes of sparkline data to display for each performance metric in the **Performance Insights** widget. Default value is 30 minutes.

   Refresh Rate
   Specify how often to refresh the data in the **Performance Insights** widget. Default value is 60 seconds.
4. Click **OK**.

Configuring thresholds

You can configure thresholds that determine when anomalies or events are raised.

About this task

You can configure traffic flow thresholds, which raise Tivoli Netcool/OMNIbus events when the thresholds are breached. Traffic flow thresholds use data from Network Performance Insight.

You can configure anomaly thresholds for performance measures. Anomaly thresholds use data from Network Manager.

Defining traffic flow thresholds

You can view details about the traffic flowing across an interface in the **Traffic Details** portlet. This feature is available only for interfaces on which traffic flow is enabled.

About this task

You can view details about the traffic flowing across an interface by opening the **Traffic Details** portlet from the **Interfaces** tab in the **Performance Insights** portlet. Right-click an interface that has traffic flow enabled and select **Show Traffic Details**.

To view the flow data, you must first enable one or more devices to capture flow data, and configure one or more interfaces on those devices to capture flow data. Optionally, configure traffic flow thresholds to raise Netcool/OMNIbus events when the thresholds are breached.

Flow protocols vary between router providers. For example, Cisco uses the NetFlow protocol. Flow is usually enabled on key interfaces, to avoid using large amounts of resources. Flow data is stored by Network Performance Insight. The **Traffic Details** portlet is a Network Performance Insight portlet.

To enable traffic flow on interfaces, complete the following steps.

Procedure

1. Enable the collection of flow data on the devices of interest, using the appropriate protocol.
   For more information, see the following topic in the Network Performance Insight Knowledge Center:
   - Network Performance Insight: Configuring flow devices
2. Configure the collection of flow data on the interfaces of interest.
   For more information, see the following topic in the Network Performance Insight Knowledge Center:
   - Network Performance Insight: Configuring flow interfaces
3. Optional: Configure a traffic flow threshold on one or more interfaces.
   For more information, see the following topic in the Network Performance Insight Knowledge Center:
   - Network Performance Insight: Configuring flow thresholds

   Note: It is not necessary to set a flow threshold in order to view flow data.

Defining performance thresholds for anomaly detection

If your Netcool Operations Insight solution is integrated with Network Performance Insight, then you can define static thresholds for anomaly detection. You define a static threshold for a given KPI within the poll definition that polls for that KPI.

About this task

For full information on how to configure poll definitions, see the following topic in the Network Manager Knowledge Center: https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/poll/task/crtpolldef.html.

About poll definitions and static thresholds

You specify static thresholds on performance measures on devices and interfaces, by defining the thresholds within the relevant poll definition. If these static thresholds are violated for any given performance measure on a device or interface, Netcool/OMNIbus events will be generated at an appropriate severity level, and if the context of the **Device Dashboard** is switched to the relevant device or interface, then the anomaly value and any associated performance measure trends will be shown on the dashboard.

Related tasks

Network Manager V4.2 documentation: Poll definition parametersUse this information to understand the parameters of a Network Manager poll definition.

Network Manager V4.2 documentation: Threshold polling During threshold polling, predefined formulas are applied to selected MIB variables, and if the threshold is exceeded by the MIB variable, then an event is generated. A Clear event is generated when the value of the MIB variable either falls below the threshold value, or falls below a different clear value.

Anomaly detection

By defining static thresholds for a performance measure in the Network Manager **Poll Definition Editor**, you can optionally specify definitive threshold values. If these threshold values are overriden for the specified number of consecutive occurrences, then a performance anomaly with the appropriate severity level is generated.

Defining anomaly thresholds

Define anomaly thresholds to detect anomalies in performance measures on devices, links, and interfaces, and display these anomalies in the **Performance Insights**.

Before you begin

Anomaly thresholds are defined within poll definitions.

Restriction:

You can define anomaly thresholds only for Basic threshold, Chassis Ping, and Interface Ping poll definitions. The poll definition must be set to store historical data and specified within an active poll policy.

About this task

To define an anomaly threshold, complete the following steps:

Procedure

1. Create a poll definition, or modify an existing poll definition.
   - To apply anomaly thresholds to performance measures such as CPU and memory utilization on devices, and incoming and outgoing bandwidth utilization on interfaces, create or modify a basic threshold poll definition. Creating a basic threshold poll definition is described in https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/poll/task/poll_crtbasictholddef.html.
   - To apply anomaly thresholds to performance measures associated with ping operations against devices and interfaces, such as packet loss and ping time, create or modify a chassis or interface ping poll definition. Creating a chassis or interface ping poll definition is described in https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/poll/task/poll_creatingchassisandifpingpolldefs.html.

   When you define a new poll definition, all of the tabs in the **Poll Definition Editor** are editable, except for the **NPI Anomaly Threshold** tab, which is blank. Complete the other tabs but do not complete the **NPI Anomaly Threshold**. Complete this tab later in the procedure.
2. Ensure that the poll definition is specified in an active poll policy.
   a) If necessary, create a poll policy. For more information about creating poll policies, see https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/poll/task/poll_crtpoll.html.

b) Set the poll definition to store historical data, by selecting the **Poll Policy Properties** tab in the **Poll Policy Editor** and clicking the **Store?** check box next to the name of the poll definition.

c) Enable the poll policy. For more information about enabling poll policies, see https://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/poll/task/poll_enablepoll.html.

3. Open the poll definition and define anomaly thresholds by following the instructions in the next step.
4. Click the **NPI Anomaly Threshold** tab and specify static thresholds for triggering performance measurement anomalies in the Device Dashboard, and associated Tivoli Netcool/OMNIbus events.

   a) In the **Upper Limit** and **Lower Limit** fields, specify the upper and lower threshold limits.

   b) In the **Consecutive Occurrences** spinner, specify the number of consecutive threshold violations that must occur before an anomaly and an event is generated.

   c) In the **Type** drop-down list, specify the threshold type. There are three types of threshold. In each threshold type, the current value is compared to the upper and lower threshold limits.

   Upper type threshold

   Use this type of threshold when the desired value of the performance measure is lower than the threshold limits.

   - If the current value exceeds the lower limit but is less than the upper limit for the specified number of consecutive occurrences, then a lower severity (orange) anomaly is generated on the Device Dashboard, and a Major Tivoli Netcool/OMNIbus event is generated.
   - If the current value exceeds the upper limit for the specified number of consecutive occurrences, a higher severity (red) anomaly is generated on the Device Dashboard, and a Critical Tivoli Netcool/OMNIbus event is generated.

   Lower type threshold

   Use this type of threshold when the desired value of the performance measure is higher than the threshold limits.

   - If the current value is less than the upper limit but higher than the lower limit for the specified number of consecutive occurrences, then a lower severity (orange) anomaly is generated on the Device Dashboard, and a Major Tivoli Netcool/OMNIbus event is generated.
   - If the current value is less than the lower limit, for the specified number of consecutive occurrences, then a higher severity (red) anomaly is generated on the Device Dashboard, and a Critical Tivoli Netcool/OMNIbus event is generated.

   Band type threshold

   Use this type of threshold when the desired value of the performance measure is within a band specified by the lower and upper threshold limits. If the current value falls outside of the band (either above or below the band) for the specified number of consecutive occurrences, a higher severity anomaly is generated on the Device Dashboard, and a Critical Tivoli Netcool/OMNIbus event is generated.

5. Click **Save**.

Administering the Device Dashboard

Perform these tasks to administer the **Device Dashboard** for users.

Changing the Network Performance Insight server for all users

If you need to change the Network Performance Insight server that provides data to the **Device Dashboard**, then you must perform the following steps,

Procedure

1. Edit the `npi.properties` file.

   This file is located by default at `$NMGUI_HOME/profile/etc/tnm/npi.properties`, where $NMGUI_HOME is by default `/opt/IBM/netcool/gui/precision_gui`.

2. In the `npi.properties` file change the following properties to point to the new Network Performance Insight server:
   - `npi.server.name`
   - `npi.host.name`
3. Instruct each end user to go to the **Device Dashboard** and change the portlet preferences there to point to the new Network Performance Insight server. End users must do this because the portlet preferences override the properties defined in the `npi.properties` file.
4. Stop and then restart the WebSphere Application Server.

Traffic Details dashboard

Use the Traffic Details dashboard to monitor the network performance details of a particular interface. Network Performance Insight provides built-in and interactive dashboards that cover the entire traffic data representation.

The Flow data that is collected by Network Performance Insight is shown from Traffic Details dashboard. It displays the traffic details at interface level.

You can launch the Traffic Details portlet that is available as a widget from the following dashboards:

- From Network Health Dashboard

  The traffic details for an interface are populated in the **Network View** page, from a selected network view.
- From Device Dashboard

  Right-click an interface of interest and select **Show Traffic** from the **Interfaces** tab in the **Performance Insights** portlet.
- From Event Viewer or AEL

  Right-click a Flow event from the Event Viewer or the AEL and select **Flow Dashboard**, to view the traffic details for the event.

Traffic Details dashboard views

Use this information to see the list of views that are available in the Traffic Details dashboard.

Traffic Details dashboard populates the performance metrics based on the collected IP network traffic information as the packet enters or exits an interface of a device. This resource provides a view of the applications responsible for traffic volume, either inbound or outbound, through the interface over the selected time period. This data provides granular traffic details at interface level.

You can view the Flow data that is collected by Network Performance Insight from Traffic Details dashboard. All the widgets in the dashboard display top 10 metric values for the resources.

The Traffic Details dashboard views are as follows:

- Top Applications
- Top Applications with ToS
- Top Autonomous System Conversations
- Top Destination Autonomous Systems
- Top Source Autonomous Systems
- Top Conversations
- Top Conversations with Application
- Top Conversations with ToS
- Top Destinations
- Top Destinations with Application

- Top IP Group Conversations
- Top IP Group Conversations with Application
- Top IP Group Conversations with Protocol
- Top IP Group Conversations with ToS
- Top Destination IP Groups
- Top Destination IP Groups with Application
- Top Destination IP Groups with Protocol
- Top Destination IP Groups with ToS
- Top Source IP Groups
- Top Source IP Groups with Application
- Top Source IP Groups with Protocol
- Top Source IP Groups with ToS
- Top Protocols
- Top Protocols with Application
- Top Protocols with Conversation
- Top Protocols with Destination
- Top Protocols with Source
- Top QoS Hierarchies with Queue ID
- Top Sources
- Top Sources with Application
- Top ToS

Note: All the data on the Traffic Detail dashboard views is populated based on your local timezone.

To view all the Top Talker resource views from Traffic Details dashboard, enable the Flow aggregations. For more information, see Configuring Flow aggregations.

Related information
Built-in aggregation definitions

Displaying NetFlow performance data from Network Health Dashboard

IBM Networks for Operations Insight is an optional feature that can be added to a deployment of the base Netcool Operations Insight solution to provide service assurance in dynamic network infrastructures. The capabilities of Networks for Operations Insight include network discovery, visualization, event correlation and root-cause analysis, and configuration and Compliance Management that provide service assurance in dynamic network infrastructures.

The Network Health Dashboard monitors a selected network view, and displays device and interface availability within that network view. It also reports on performance by presenting graphs, tables, and traces of KPI data for monitored devices and interfaces. Traffic Details portlet can be launched for the interfaces that are available in the Structured Browser.

The Network Health Dashboard includes the Event Viewer for detailed event information. You can start the Traffic Details portlet to display the NetFlow traffic details for the interface that violated a set threshold value.

Adding the Network Performance Insight widget in Network Health dashboard

Dashboards or pages, are an arrangement of one or more widgets in the work area and contain the widgets that are needed to complete tasks. Users whose roles have Editor or Manager access to a dashboard can edit a dashboard's layout and content. You can add multiple widgets in a screen. When you are adding widgets, you can also rearrange the widgets as needed.

About this task

By default, the **Network Health Dashboard** page displays the **Network View**, **Structure Browser** , and **Event Viewer** widgets.

To view the traffic details from Network Health Dashboard for the first time, you need to add the Network Performance Insight widget.

Procedure

1. Log in to Jazz for Service Management server.
2. Click the **Incident** icon () and select **Network Health Dashboard** under **Network Availability**.
3. In the **Network Health Dashboard**, select a network view.

   A second tab, called **Network View**, opens.
4. In the tab bar, click **Page Actions** icon ( ) and select **Edit Page**.

   The dashboard is changed to show the widget palette and a series of buttons in the tab bar. The menu that is associated with the **Edit** options icon for each widget is updated so that you can edit its layout and content.
5. Click the *NPI* folder from the widget palette.

   The *NPI* folder name is based on the launch-in-context tool that is created for Network Performance Insight.
6. Click and drag the **Traffic Details** widget from the palette.

   To assist you in positioning the widget, use the background layout grid. You can change the size of the layout grid and have widgets snap to the layout guide lines through the **Layout** button in the tab bar.
7. Click **Save and Exit** to exit the dashboard from the edit mode after you complete the editing.

Related information

Editing dashboard content and layout

Traffic Details from Network Health Dashboard

Network Health Dashboard displays the traffic details of a particular network.

The Traffic Details widget data is populated from NCIM view that is part of Network Performance Insight database structure, where it joins multiple tables into a single virtual table.

NCIM view represents the subset of data that is discovered from the Tivoli Network Manager and Network Performance Insight Flow tables.

The discovered data by Tivoli Network Manager is mapped with Flow records by using the following fields:

Table 55. Mappings table

Flow table	**NCIM view**
exporter ip	*device ip address*
if index	*interface index*

The following SNMP fields in the Traffic Details widget are populated from the NCIM view:

Note: If the discovered interfaces are not mapped with Network Performance Insight flow data, you can't see the SNMP fields in the table on Traffic Details dashboard.

Table 56. SNMP fields in Traffic Details dashboard		
Traffic Details fields	**Description**	**Mapping**
Device	The device name.	The *exporter ip* from Flow table is mapped to the *device name* in the NCIM view. This mapping results to the device name populated in the traffic details widget.
Index	A unique number that is associated with the physical or logical interface.	The *if index* from Flow table is mapped to the *interface index* in the NCIM view. This mapping results to the index value populated in the traffic details widget.
Description	The interface name.	The interfaces that are discovered by Tivoli Network Manager are mapped with the collected Network Performance Insight flow data as *interface name*. This mapping results to the description of the device that is populated in the traffic details widget.
Speed	The value of the traffic flow through network interfaces, which measures the speed of the data transferred.	The *speed* from Flow table is mapped to the *interface speed* from NCIM view. This mapping results to the speed value populated in the traffic details widget.

Launching Traffic Details dashboard from Network Health dashboard
The Traffic Details widget displays the details of an interface from a selected network device.

About this task
Browse from Network Health dashboard to view the specific traffic details of an interface in the **Traffic Details** page.

By default, the **Traffic Details** page displays the details for Top Ingress interfaces at Ingress level, Top Egress interfaces at Egress level. Whereas, Top Interfaces and all Top Networks display the traffic details as **Both**.

Procedure

1. Log in to Jazz for Service Management server.
2. Click the **Incident** icon () and select **Network Health Dashboard** under **Network Availability**.

   The dashboard page populates the configured network devices.
3. Select a view from the **Network Views** bookmark that you configured from the **Network Health Dashboard**.

   The other widgets update to show information based on the network view that you selected.

   The **Network View** dashboard opens in another tab. This dashboard contains Network Views GUI, the **Event Viewer**, the **Structure Browser** , and the **Traffic Details**, and it displays the selected network view.
4. Double-click a network from the **Network View**.

   For example, double-click **All Routers**.
5. Click an entity or device from the **Network View**.

The selected entity details are displayed on the **Structure Browser** .

6. Click the **Show Interfaces** icon () from the **Structure Browser** .

   List of interfaces for the entity is displayed.
7. Click an interface.

   The traffic details data with the interface details is displayed on the **Traffic Details**.
8. Select an entity from **View** list.

   For information on monitoring traffic details, see "Monitoring NetFlow performance data from Traffic Details dashboard" on page 407.
9. Optional: Click the **Maximize** icon () from the upper right of the **Traffic Details** widget.

   The traffic details dashboard is displayed in full screen mode.

Displaying NetFlow performance data from Event Viewer

You can monitor and manage network performance from events that are generated by Tivoli Netcool/ OMNIbus on Web GUI.

About this task

You can access the events from the following widgets:

- Managing events in the Event Viewer

  Use the JavaScript Event Viewer to monitor and manage events. You can access Event Viewers in any page in Dashboard Application Services Hub that hosts an Event Viewer widget.
- Monitoring events in Active Event Lists

  The Active Event List (AEL) is an interactive Java applet for displaying alert data from the ObjectServer. Communication between the ObjectServer and the AEL is bidirectional. The AEL presents alert information from the `alerts.status` table in the ObjectServer to operators. Operators can perform actions against alerts such as changing the results from the alert properties in the `alerts.status` table from the AEL.

Procedure

1. Log in to Jazz for Service Management server.
2. Click **Incident (**  **)** > **Events** > **Event Viewer** in the navigation bar.
3. Right-click an event that is labeled as **NPI/Flow** under the **Manager** column from the **Event Viewer** and select one of the following commands. Both of these commands launch the **Traffic Details** dashboard.
   - **Flow Dashboard**
   - **Performance Insight** > **Show Traffic**

   The **Traffic Details** dashboard that is associated with the selected event is displayed in another window.

   Note: You can launch the Traffic Details dashboard from Flow events only. Flow events are marked in the **Event Viewer** using the value **NPI/Flow** in the **Manager** column.
4. Optional: Right-click a Flow event from **Incident (**  **)** > **Events** > **Active Event List (AEL)** and select **Flow Dashboard**.

   The Traffic Details dashboard that is associated with the selected event is displayed in another window.

Related information

Configuring launch-in-context integration with Network Performance Insight
Event severity levels

Monitoring NetFlow performance data from Traffic Details dashboard

After launching the Traffic Details dashboard, you can display the different views of NetFlow data for a selected interface.

About this task

By default, the **Traffic Details** page displays the data about the top ingress, top egress, or both flows for an interface.

Procedure

1. Select **Traffic Details** for **Ingress**, **Egress**, or **Both** from the list.
2. Select any dashboard view from the **View** list.

   For more information about dashboard views, see "Traffic Details dashboard views" on page 402.
3. Click  to select a start date from the **Start** field.
4. Click  to select start time.
5. Click  to select an end date from the **End** field.
6. Click  to select end time.
7. Click **Update** to update the details for selected date and time.

   Two red lines are displayed in the graph that notify the Upper Threshold and Lower Threshold limits are crossed. When you hover over the area charts, a tooltip is flashed giving you the details for that particular source.

   The graphical presentation of data is updated for the selected date and time.
8. Click  to refresh the page.
9. Select one or more interfaces from the legend on the right.

   It displays the top 10 interface details.
10. Clear the check box for a particular interface.

    The details for that interface are hidden.
11. Check the **Remaining** check box.

    It displays the details for the remaining traffic on the interfaces.

    Note: If the upper limit and lower limit of a threshold is crossed, two extra check boxes for Upper Threshold and Lower Threshold are seen in the legend.

    The table at the bottom displays top 10 interface details. The table displays the following details for an interface:

    Rank
    Display the number in ascending order.

    Grouping
    Displays the interface for which you are viewing the data. For example, if you select Top Protocols from the **View** list, the grouping that is displayed is for Protocol.

    The columns change depending on the view selected. The main grouping keys for the Traffic Details dashboard can be defined as:

Table 57. Grouping	
Grouping Key	**Description**
Application	Applications are mapped based on port, protocol, and IP address or network.
Destination	Destination can be a host computer to which the network flow comes from a source computer.
Destination AS	Autonomous systems that a specific route passes through to reach one destination.
Destination IP Group	A group of destinations that you want to control by specifying the IP addresses. A grouping of endpoints for traffic accounting, billing purpose.
Protocol	It is a standard that is defined on how a network conversation is established. It delivers the packets from the source host to the destination host.
Source	Source is the IP address from which traffic is originated.
Source AS	Autonomous systems that a specific route passes through from a source to reach one destination.
Source IP Group	A group of sources that you want to control by specifying the IP addresses. A grouping of endpoints for traffic accounting, billing purpose.
Source ToS	The class of service, which examines the priority of the traffic.
QoS Hierarchy	QoS behavior at multiple policy levels, which provides the visibility of how the defined traffic classes are performing.
QoS Queue ID	QoS Queues provide bandwidth reservation and prioritization of traffic as it enters or leaves a network device.

Octets
Displays the amount of data that is used in KB and Bytes.

Percentage
Percentage of traffic on the grouping that occupies the traffic.

Note: If the top 10 interface table is not shown, reduce the zoom percentage of your current browser.

Related information
Built-in aggregation definitions

Network Performance Insight Dashboards

After the system is configured as per your requirements, Network Performance Insight can start collecting, aggregating, and storing the network performance data. The data is rendered on various ready-to-use dashboards that it offers.

Categories of Network Performance Insight Dashboards:

- Network Performance Overview dashboard
- Network Performance Overview by Deviation dashboard

- “WiFi Overview dashboard” on page 435
- “IP Links Performance Overview dashboard” on page 440
- “Load Balancer dashboards” on page 444
- NetFlow dashboards
- On Demand Filtering dashboards

Network Performance Insight Dashboards can be grouped functionally as follows:

- Operational dashboards
- Analytical dashboards
- Strategic dashboards
- On-demand dashboards

Important: When Network Performance Insight is used to collect, aggregate, and render the NetFlow data alone, the following dashboards are applicable.

- “NetFlow dashboards” on page 466
- On Demand Filtering - Flow dashboard

The following dashboards that display performance data are not applicable in this scenario:

- Network Performance Overview dashboard
- Network Performance Overview by Deviation dashboard
- “WiFi Overview dashboard” on page 435
- On Demand Filtering - IPSLA dashboard
- On Demand Filtering - Device Health dashboard
- On Demand Filtering - HTTP Operations dashboard
- On Demand Filtering - Timeseries Data dashboard

Getting started with Network Performance Insight Dashboards

This information provides instructions and general information on how to use the Network Performance Insight Dashboards that render network performance data from Network Performance Insight.

The Network Performance Insight Dashboards report the network performance data that is gathered and stored by the Network Performance Insight and its components.

Network Performance Insight Dashboards are federated on Dashboard Application Services Hub portal and you can derive the following information from them:

- Summary level views of the network and see how your network resources are performing.
- Detailed views from the listener and drill-down widgets. You can switch between different metric views to analyze and monitor your network.
- Monitor Top Talker resource views that help to understand network insights at a more granular level.
- Share this information with other stakeholders by generating and sending the report information in PDF, CSV, or XLS format.

Accessing the Network Performance Insight Dashboards

Access the Network Performance Insight Dashboards from Dashboard Application Services Hub portal.

Procedure

1. Log in to Dashboard Application Services Hub portal with `npiadmin` and `netcool` credentials.

2. Click **Console Integrations** icon () in the navigation bar and select **Dashboards** under **Performance**.

   The page loads with menu bar to navigate to different Network Performance Insight Dashboards.

   a) Click **Home** to see the two types of Network Performance Overview dashboards from the menu bar.

   - **Network Performance Overview**

     Network Performance Overview: Top 10 dashboard offers summary level current view of the managed network.

   - **Network Performance Overview by Deviation**

     Network Performance Overview by Deviation: Top 10 Deviations dashboard offers summary level deviation view of the managed network.

   - **WiFi Overview**

     WiFi Overview dashboard offers summary level current view of WiFi network

   b) Click **NetFlow** to see Network Traffic Overview, Applications Response Overview dashboards and all the NetFlow built-in Top N resource views.

   You can see the following dashboards:

Dashboard group	Available views
Network Traffic Overview	Network Traffic Overview: Top 10 Traffic dashboard that contains various widgets.
Applications Response Overview	Applications Response Overview: Top 10 dashboard that contains various widgets. • **Applications Response Time**
Applications	• **Top Applications** • **Top Applications with ToS**
Conversations	• **Top Conversations** • **Top Conversations with Application** • **Top Conversations with ToS**
Sources	• **Top Sources** • **Top Sources with Application**
Protocols	• **Top Protocols** • **Top Protocols with Application** • **Top Protocols with Conversation** • **Top Protocols with Destination** • **Top Protocols with Source**
Destinations	• **Top Destinations** • **Top Destinations with Application**
ToS	• **Top ToS**

Dashboard group	Available views
Autonomous Systems	• **Top Autonomous System Conversations** • **Top Destination Autonomous Systems** • **Top Source Autonomous Systems**
QoS Queue	• **QoS Queue Drops**
IP Address Grouping	• **Top IP Group Conversations** • **Top IP Group Conversations with Application** • **Top IP Group Conversations with Protocol** • **Top IP Group Conversations with ToS** • **Top Destination IP Groups** • **Top Destination IP Groups with Application** • **Top Destination IP Groups with Protocol** • **Top Destination IP Groups with ToS** • **Top Source IP Groups** • **Top Source IP Groups with Application** • **Top Source IP Groups with Protocol** • **Top Source IP Groups with ToS**

c) Click **On Demand Filtering** to see the following views:

- **Device Health**
- **IPSLA**
- **Flow**
- **HTTP Operations**
- **Timeseries Data**

Accessing the Network Performance Insight Dashboards

Access the Network Performance Insight Dashboards from Dashboard Application Services Hub portal.

Procedure

1. Log in to Dashboard Application Services Hub portal with `npiadmin` and `netcool` credentials.
2. Click **Console Integrations** icon () in the navigation bar and select **Dashboards** under **Performance**.

   The page loads with menu bar to navigate to different Network Performance Insight Dashboards.

   a) Click **Home** to see the two types of Network Performance Overview dashboards from the menu bar.

   - **Network Performance Overview**

     Network Performance Overview: Top 10 dashboard offers summary level current view of the managed network.
   - **Network Performance Overview by Deviation**

     Network Performance Overview by Deviation: Top 10 Deviations dashboard offers summary level deviation view of the managed network.
   - **WiFi Overview**

WiFi Overview dashboard offers summary level current view of WiFi network

- **IP Links Performance Overview**

  The IP links performance is monitored through a heat map, which provides an overall view of outbound one way link performance for speech applications between a source and destination server present at various geographical locations.

- **Load Balancer**

  Under this category, the following dashboards are available:

  - **Load Balancer Overview**
  - **GTM Details**
  - **LTM Details**
  - **Virtual Server Details**
  - **Pool Details**
  - **Pool Member Details**

b) Click **NetFlow** to see Network Traffic Overview, Applications Response Overview dashboards and all the NetFlow built-in Top N resource views.

You can see the following dashboards:

Dashboard group	Available views
Network Traffic Overview	Network Traffic Overview: Top 10 Traffic dashboard that contains various widgets.
Applications Response Overview	Applications Response Overview: Top 10 dashboard that contains various widgets. • **Applications Response Time**
Applications	• **Top Applications** • **Top Applications with ToS**
Conversations	• **Top Conversations** • **Top Conversations with Application** • **Top Conversations with ToS**
Source	• **Top Sources** • **Top Sources with Application**
Protocols	• **Top Protocols** • **Top Protocols with Application** • **Top Protocols with Conversation** • **Top Protocols with Destination** • **Top Protocols with Source**
Destinations	• **Top Destinations** • **Top Destinations with Application**
ToS	• **Top ToS**

Dashboard group	Available views
Autonomous Systems	• **Top Autonomous System Conversations** • **Top Destination Autonomous Systems** • **Top Source Autonomous Systems**
QoS Queue	• **QoS Queue Drops**
IP Address Grouping	• **Top IP Group Conversations** • **Top IP Group Conversations with Application** • **Top IP Group Conversations with Protocol** • **Top IP Group Conversations with ToS** • **Top Destination IP Groups** • **Top Destination IP Groups with Application** • **Top Destination IP Groups with Protocol** • **Top Destination IP Groups with ToS** • **Top Source IP Groups** • **Top Source IP Groups with Application** • **Top Source IP Groups with Protocol** • **Top Source IP Groups with ToS**

c) Click **On Demand Filtering** to see the following views:

- **Device Health**
- **IPSLA**
- **Flow**
- **HTTP Operations**
- **Timeseries Data**

Generic functions of Network Performance Insight Dashboards

Use this information to understand the generic interactivity and filtering functions that are available on Network Performance Insight Dashboards.

Procedure

- Generic interactivity that is applicable for all Network Performance Insight Dashboards:

  a) Click **Auto Refresh** () icon to enable or disable auto refresh.

  Note: By default, the data is auto refreshed every one minute

  If auto refresh option is enabled, the dashboard metrics are refreshed with latest values.

  b) Click **Email PDF** () icon to email the dashboard view as a PDF.
  Complete the following steps to email a PDF:

  a. In the **Email the PDF** file window that is displayed, enter the following details:

  1) **Subject**

  2) **Email To**

  Note: Ensure that the email address is valid and in correct format.

3) **Content**

b. Click **Send**

c) Click **Save As** ( ) icon and select **PDF**, **CSV**, or **XLS** to save and export the dashboard to the selected file format.

Note:

– In a PDF file format, the complete data is populated in the next page on a tabular format. You need to click the grid to view the complete data.

– **PDF** option is not available for On Demand Filtering dashboards.

d) To maximize the widget display, click .

e) To minimize the widget display, click .

- Generic interactivity that is applicable for only the WiFi Overview, NetFlow and On Demand Filtering dashboards and widgets:

a) To change to a different chart type, click and select a different chart type from the widget.

The widget renders according to the selected chart type.

b) To hide the widget, click .

c) To display the widget, click .

d) From the **Datasource** icon ( ), click to list the options to choose other performance metric.

Note: The Datasource icon is available on Network Performance Overview, Network Performance Overview by Deviation, and NetFlow dashboards widgets.

e) Click the column header from any grid widgets to sort in ascending or descending order.

You can sort the data in numerical or alphabetical, depending on what type of data is populated in the grid widget.

f) The following icon on a widget indicates that it is a drill-down chart. Click one of the elements from the chart widget.

For example, in a bar chart, click one of the bars to further drill down to one hierarchy level down which correspond to the bar that you selected.

- Widget-level filters for WiFi Dashboards

Widget-level filters are available for the Grid chart type. In WiFi dashboards, the two widgets, Access Points Details and WiFi Network Details are represented in the table format.

Filter details:

1. Click the icon that is available besides the column name.

   A pop-up window that has a list, a field for entering filter inputs, and two buttons **Apply Filter** and **Clear Filter** appear.
2. From the first list field, select any one of the 6 values:

   – **Equals**: Select this value if you want to include any column value in the filter results.

   – **Not Equal**: Select this value if you want to exclude any column value in the filter results.

   – **Starts With**: Select this value if you want to filter the column values that start with a particular letter.

- **Ends With**: Select this value if you want to filter the column values that end with a particular letter
- **Contains**: Select this value if you want to filter the column value that is based on a particular letter or a sequence of letters it contains.
- **Not Contains**: Select this value if you want to filter the column value based on a particular letter or a sequence of letters it does not contain.

3. In the field, enter the value for which you want to apply the filter.
4. After you enter the filter value in the field above, two grouping options and another list and filter field appear. You can group the filter results of the two list and field values. The grouping options are:
   - **AND**: Conditions from both the first and the second list and field pair must match with the values of column.
   - **OR**: Conditions from any one of the two list and field pairs must match with the values of column.
5. Click **Apply Filter**.
6. Click **Clear Filter** and then **Apply Filter** to see the original table.

- TimeZoom

  TimeZoom is a feature enhancement for Line charts to make them more readable. When the data points are too many at a particular time stamp that it gets difficult to identify and distinguish the data points from one another, TimeZoom feature increases the time granularity and zooms into the chart to sparsely place the data points. For instance, if the time stamp records data at an hourly rate, applying the TimeZoom displays the data points that are recorded at a shorter interval say half hour and place the data points at a comprehensible distance. TimeZoom is represented in the form of a horizontal scroll bar, which can zoom out and in as you move it. It is placed at the bottom of the widget.

  Following widgets are enhanced with Timezoom feature:

  - Wifi Client Count (drill down from Wifi Overview)
  - Wifi Interference and Noise Performance (drill down from Wifi Overview)
  - Device Health History
  - IPSLA History
  - HTTP Response Time
  - Timeseries Data
  - Application Response Time (drill down from Application Response Overview)

- Filtering options that are applicable for Network Performance Insight Dashboards.

  The filter options and list differ based on the different types of Network Performance Insight Dashboards.

Table 58. Filter options

Filter name	Filter description	Network Performance Insight Dashboards
Aggregation Type	The list contains the available NetFlow dashboards.	On Demand Filtering Flow
AP Radio Channel	The list contains available AP radio Channels in the WiFi Network.	WiFi Overview Dashboard
Application	The list contains the available applications.	Application Response Time

Table 58. Filter options (continued)

Filter name	Filter description	Network Performance Insight Dashboards
Business Hour	The list contains the peak or a non-peak business hours option. You can select either **Business Hour** or **Non-Business Hour** or both as **ALL**.	Network Performance Overview Network Performance Overview by Deviation Network Traffic Overview All Load Balancer dashboards IP Links Performance Overview
Controller	The list contains the controllers that are connected to the WiFi network.	WiFi Overview Dashboard
Destination	The list contains the destination IP addresses for the specific source IP addresses.	Source and Destination Details
KPI	The list contains the available performance metrics.	– On Demand Filtering Device Health – On Demand Filtering HTTP Operations – On Demand Filtering IPSLA
LTM	The list contains the available LTM devices.	– Pool Details – Pool Member Details – Virtual Server Details
Site	The list contains the locations to view the traffic levels.	Network Performance Overview Network Performance Overview by Deviation Network Traffic Overview All Load Balancer dashboards IP Links Performance Overview
SLA Test	The IPSLA operations (SLA Test) list contains options from the following functional areas: – Availability monitoring – Network monitoring – Application monitoring – Voice monitoring – Video monitoring	On Demand Filtering IPSLA
Source	The list contains the configured or discovered devices.	On Demand Filtering HTTP Operations On Demand Filtering IPSLA IP Links Performance Overview

Table 58. Filter options (continued)

Filter name	Filter description	Network Performance Insight Dashboards
Sort By	Select from the list to sort the result by top or bottom.	On Demand Filtering Device Health On Demand Filtering HTTP Operations On Demand Filtering IPSLA
Target	The list contains the available target servers.	Application Response Time
Time Period	The dashboard data is populated based on the Network Performance Insight server time zone. – **Last Hour**, filters the last 1 hour of the current time. – **Last 6 Hours**, filters the last 6 hours of the current time. – **Last 12 Hours**, filters the last 12 hours of the current time. – **Last 24 Hours**, filters the last 24 hours from the current time and day. – **Last 7 Days**, filters the last 7 days from the current time and day. – **Last 30 Days**, filters the last 30 days from the current time and day. – **Last 365 Days**, filters the last 365 days from the current time and day. – **Custom**, from the **Time Period Selection**, you can filter based on a specific date and time range. View the start and end time on the dashboard title bar. The start and end time is displayed according to the filter you select. For example, if the current time is *3.00 PM* on *11/13/17* and you select the filter for *Last 24 Hours*. The dashboard displays the time period as: 11/12/17, 3:00 PM - 11/13/17, 2:59 PM *<Timezone>*.	Network Performance Overview Network Performance Overview by Deviation WiFi Overview Network Traffic Overview All Load Balancer dashboards IP Links Performance Overview Source and Destination Details Applications Response Overview Applications Response Time NetFlow top talker dashboards On Demand Filtering Device Health On Demand Filtering Flow On Demand Filtering HTTP Operations On Demand Filtering IPSLA On Demand Filtering Timeseries Data **Note:** – For Network Performance Overview by Deviation and Network Traffic Overview dashboards, only **Last Hour** and **Last 24 Hours** time period filter options are made available. – For Network Performance Overview dashboard, only **Last Hour**, **Last 24 Hours**, **Last 7 Days**, **Last 30 Days** and **Custom** time period filter options are made available.

Table 58. Filter options (continued)		
Filter name	**Filter description**	**Network Performance Insight Dashboards**
Top N	From the list, you can choose *N* number of Top Performers to compare historical poll data across multiple entities and metrics in a selected network view.	On Demand Filtering Flow All Load Balancer dashboards
Resource Type	The list contains the available types of resources.	On Demand Filtering Timeseries Data
Resource Name	The list contains the available resource names.	On Demand Filtering Timeseries Data

What to do next

Click **Apply Filter** to apply the filter selection.

The dashboard reloads the data according to the filter selections.

Site and Business hour filter condition

Conditions that are applicable for the **Site** and **Business Hour** filters. These filter options are available on Network Performance Overview, Network Performance Overview by Deviation, Network Traffic Overview, and Load Balancer dashboards.

The table describes the conditions that can be applied to the filters and the time periods that are based on it.

Table 59. Site and Business Hour filter conditions

Dashboard	Filter options			Condition
	Site	**Business Hour**	**Time Period**	
Network Performance Overview	ALL	ALL	Last Hour **Note:** Last Hour is available only when you select All for both Site and Business Hour. Last 24 Hours Last 7 Days Last 30 Days Custom	When you select **ALL** from the **Site** filter option, both business hour and non-business hour data is queried and displayed.
	<site_name>	ALL Business Hour Non-Business Hour	Last 24 Hours Last 7 Days Last 30 Days Custom **Note:** Use Custom filter to get more granularity.	When you select a site from the **Site** filter option, the **Time Period** filter option list starts from last 24 hours and more.
Network Performance Overview by Deviation	ALL	ALL	Last Hour Last 24 Hours	When you select **ALL** from the **Site** filter option, both business hour and non-business hour data is queried and displayed.
Network Traffic Overview	*<site_name>*	• ALL • Business Hour • Non-Business Hour	Last 24 Hours	When you select a site from the Site filter option, the Time Period filter option list displays only last 24 hours.

Table 59. Site and Business Hour filter conditions (continued)				
Dashboard	**Filter options**			**Condition**
	Site	**Business Hour**	**Time Period**	
Load Balancer dashboards	ALL	ALL	• Last Hour • Last 6 Hours • Last 12 Hours • Last 24 Hours • Last 7 Days • Last 30 Days • Last 365 Days • Custom **Note:** Use Custom filter to get more granularity.	When you select a site from the **Site** filter option, the **Time Period** filter option list displays only last 24 hours.
	<site_name>	• ALL • Business Hour • Non-Business Hour	• Last 24 Hours • Last 7 Days • Last 30 Days • Last 365 Days • Custom **Note:** Use Custom filter to get more granularity.	When you select a site from the **Site** filter option, the **Time Period** filter option list starts from last 24 hours and more.

Related tasks

Configuring site grouping

Network Performance Overview dashboards

Network Performance Overview summarizes a high-level view of your network, application, and device performance data in a single location. It gives you full control over how you investigate and analyze the measurements. You can navigate to the detailed views from here that focus on specific diagnostics.

Network Performance Overview dashboards provide advanced monitoring of your network from a single pane of glass and you can drill down to specific details from here. Two types of overview dashboards are available that display data by current values and by deviation.

All the widgets in the overview dashboards display the metrics by current values for top 10 NetFlow enabled resources in your network.

You can access network traffic congestion, network traffic utilization, application performance, quality of service, and device performance easily and quickly from the widgets in this dashboard for the following areas:

Congestion

Network congestion can occur due to high data volumes and not enough capacity to support the applications that are involved. It can lead to delays and packet drops and cause brownouts. Some of the reasons for congestion might be as follows:

- Network data in the route is more than the network capacity.

- Poor network design
- Inefficient internet routing. Border Gateway Protocol (BGP) that sends all traffic through shortest logical path is not congestion aware and transit paths can become overloaded.
- Incorrect QoS configuration or no QoS implementation

Typical effects include queuing delays, packet loss, or refusing new connections to the resource. From these widgets, you can understand which interfaces are experiencing congestion by observing the inbound and outbound packet discard deviations.

This data can be correlated with QoS queue drops and total application response time.

- When a router receives NetFlow data at a rate faster than it can transmit, the device buffers the extra traffic until bandwidth is available. This buffering process is called queuing. You can use QoS to define the priorities to transmit the records from the queue.
- The application response time is represented as total delay, which is the sum of max client network delay, max server network delay and application delay.

Quality of Service

When real-time data is passed in a network (for example, video, audio, or VoIP), implementing the Quality of Service is crucial for a good user experience. These dashboards help to understand the round-trip time for the successful echo operations.

You can use this dashboard to troubleshoot issues with business-critical applications by determining the round-trip delay times, testing connectivity to the devices, and probe loss.

Voice performance is monitored with the help of the widgets in the Quality of Service section.

Traffic

To ensure adequate network bandwidth for users to use the business-critical applications, you can correlate bandwidth utilization on different interfaces and the applications that consume the bandwidth. You can then see the applications that are impacted on the interfaces with high-bandwidth utilization.

Device load

Faulty devices can create a bottleneck to your network. High CPU utilization and memory loads can adversely affect the performance of the device and network. To understand the critical health of the monitored devices, the Device load widgets are useful.

Network Performance Overview dashboard

All the widgets in this dashboard display the metrics by current values for top 10 resources in your network.

Available widgets and their interactivity

The master-listener and drill-down widget interactions can be seen in the following ways:

- For any widgets with the ( ) drill-down icon, from the chart, click any bar that represents the interface or metric to open the drill-down dashboard page, which contains the associated widgets.
- For the other widgets that are shown in the diagram to be having a master-listener interaction, click any of the bars that represent the interface to change the listener widget that displays the related data for the selected interface.

The diagram shows the master-listener, and drill-down interactions between the available widgets:

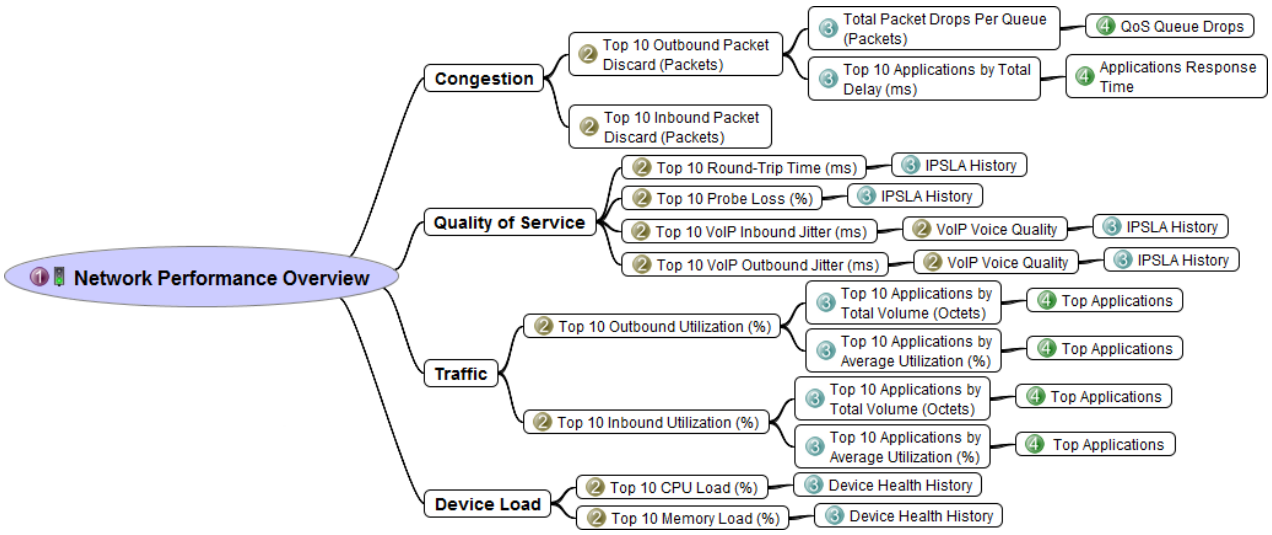

Network Performance Overview

1. Click **Home** > **Network Performance Overview**.

   The **Network Performance Overview: Top 10** dashboard loads.

   This dashboard displays the metrics values for top 10 resources.
2. You can filter data based on the following filters:
   - **Site**, the list contains the locations to view the traffic levels.
   - **Business Hour** , the list contains of **Business Hour**, **Non-Business Hour**, and both as **ALL**.
   - **Time Period**, the list contains of **Last Hour**, **Last 24 Hours**, **Last 7 Days**, **Last 30 Days** and **Custom**.

   Note:

   - You can categorize your enterprise network based on different geographical areas and sites. You can configure the sites by specifying the IP address ranges and the time period for business hours and non-business hours. To configure these properties, see Configuring site grouping.
   - Refer to the conditions applied when you're using the **Site** and **Business Hour** filter options.

Table 60. Widget interactions

Area monitored	Master widgets	Listener widgets	Drill-down dashboard
Congestion	Top 10 Outbound Packet Discard (Packets) Click any bar that represents the interface outbound packet discard to refresh the listener widgets to display the related data for the selected interface.	Total Packet Drops Per Queue (Packets) Click **Find out more details about Packet Drops per Queue. Click here** link to open the QoS Queue Drops dashboard in a new tab.	QoS Queue Drops
		Top 10 Applications by Total Delay (ms) Click any bar that represent an application in Top 10 Applications by Total Delay (ms), the **Application Response Time** page opens in a new tab.	Applications Response Time
	Top 10 Inbound Packet Discard (Packets)	N/A	N/A

Table 60. Widget interactions (continued)

Area monitored	Master widgets	Listener widgets	Drill-down dashboard
Quality of Service	Top 10 Round-Trip Time (ms) Click any bar that represents the round-trip time between a source and destination device IP addresses to open the IPSLA History dashboard in a new tab.	N/A	IPSLA History
	Top 10 Probe Loss (%) Click any bar that represents the probe loss between a source and destination device IP addresses to open the IPSLA History dashboard in a new tab.	N/A	
	Top 10 VoIP Inbound Jitter (ms) or Top 10 VoIP Outbound Jitter (ms) Click any bar that represents the VoIP Inbound or Outbound Jitter between a source and destination device IP addresses to refresh the listener widgets. **Note:** Click the **Data Source** icon () to toggle between Top 10 VoIP Inbound Jitter (ms) and Top 10 VoIP Outbound Jitter (ms) widgets.	VoIP Voice Quality The widget displays the following metrics for the same devices: • MOS • IcPIF Click **Source - Destination: *<IP_Address>* - *<IP_Address>*** to open the IPSLA History dashboard in a new tab.	

Table 60. Widget interactions (continued)

Area monitored	Master widgets	Listener widgets	Drill-down dashboard
Traffic	Top 10 Outbound Utilization (%) Click any bar that represents the interface outbound utilization to refresh the listener widgets to display the related data for the selected interface.	Top 10 Applications by Total Volume (Octets) Top 10 Applications by Average Utilization (%) Click **Find out more details about the Application. Click here** link to open the **Top Applications: Traffic Volume Details for Interface** dashboard a new tab.	Top Applications
	Top 10 Inbound Utilization (%) Click any bar that represents the interface inbound utilization to refresh the listener widgets to display the related data for the selected interface.		
Device load	Top 10 CPU Load (%) Click any bar that represents the CPU Load on a resource, the Device Health History dashboard for the same resource opens in a new tab.	N/A	Device Health History
	Top 10 Memory Load (%) Click any bar that represents the Memory Load on a resource, the Device Health History dashboard for the same resource opens in a new tab.	N/A	

Table 61. Available widgets

Widget name	Chart type	Typical uses
Top 10 Outbound Packet Discard (Packets)	Bar	Measures the number of outbound packets that are discarded even though no errors are detected to free the buffer space. It might be due to resource limitations on the outbound interface.
Total Packet Drops Per Queue (Packets)	Bar	Measures the number of packet drops per traffic class queue per interface.

Table 61. Available widgets (continued)

Widget name	Chart type	Typical uses
Top 10 Inbound Packet Discard (Packets)	Bar	Measures the number of inbound packets that are discarded even though no errors are detected to free the buffer space. It might be due to resource limitations on the inbound interface.
Top 10 Applications by Total Delay (ms)	Bar	Measures maximum total time that it takes an application to respond to user requests. The total delay is the sum of max client network delay, max server network delay, and application delay.
Top 10 Round-Trip Time (ms)	Bar	Measures the time that is taken between sending a UDP echo request message from a source device to the destination device and receiving a UDP echo reply from the destination device. It is useful in troubleshooting issues with business-critical applications by determining the round-trip delay times and testing connectivity between devices.
Top 10 Probe Loss (%)	Bar	Measures the percentage of probes lost. it shows the completion status of an RTT operation.
Top 10 VoIP Inbound Jitter (ms)	Bar	Measures the variance in latency over time in incoming direction between two devices in your network.
Top 10 VoIP Outbound Jitter (ms)	Bar	Measures the variance in latency over time in outgoing direction between two devices in your network.

Table 61. Available widgets (continued)

Widget name	Chart type	Typical uses
VoIP Voice Quality	Badge	Measures voice performance that is important for delivering voice quality services. This widget shows two metrics: • Mean Opinion Scores (MOS). A common benchmark to determine voice quality is MOS. With MOS, a wide range of listeners can judge the quality of voice samples on a scale of 1 (bad quality) to 5 (excellent quality). The Cisco IOS implementation of MOS takes RTT delay and packet loss into the MOS calculation. However, jitter is not included. The following colors on the widget denote the severity of degradation of voice quality: – Green, when the MOS value is greater than 4 and less than or equal to 5 – Amber, when the MOS value is greater than 2 and less than or equal to 4 – Red, when the MOS value is greater than 1 and less than or equal to 2 • Impairment/Calculated Planning Impairment Factor (IcPIF) IcPIF values are expressed in a typical range of 5 (low impairment) to 55 (high impairment). Typically, IcPIF values that are numerically less than 20 are considered adequate. The following colors on the widget denote the severity of degradation of voice quality: – Green, when the IcPIF value is greater than 0 and less than or equal to 14 – Amber, when the IcPIF value is greater than 14 and less than or equal to 34 – Red, when the IcPIF value is i greater than 34 and less than or equal to 93
Top 10 Outbound Utilization (%)	Bar	Measures outbound bandwidth utilization for the highest Outbound Packet Discards on the interfaces.
Top 10 Applications by Total Volume (Octets)	Bar	Measures the corresponding applications with highest bandwidth utilization by volume.
Top 10 Inbound Utilization (%)	Bar	Measures the bandwidth utilization for incoming traffic for the highest Inbound Packet Discards on the interfaces.

Table 61. Available widgets (continued)		
Widget name	**Chart type**	**Typical uses**
Top 10 Applications by Average Utilization (%)	Bar	Measures the corresponding applications with highest bandwidth utilization by percentage.
Top 10 CPU Load (%)	Bar	Measure the total CPU utilization in percentage across all vendor devices in your network. Currently, Network Performance Insight supports the following vendors: • Cisco • Huawei • Juniper
Top 10 Memory Load (%)	Bar	Measure the total memory usage in percentage across all vendor devices in your network. Currently, Network Performance Insight supports the following vendors: • Cisco • Huawei • Juniper

Network Performance Overview by Deviation dashboard
For all the widgets that display the metrics by deviation, the values are calculated by computing deviation for the current data that is compared against an average value on the same day of week over the last four weeks.

Available widgets and their interactivity

The master-listener and drill-down widget interactions can be seen in the following ways:

- For any widgets with the ( ) drill-down icon, from the chart, click any bar that represents the interface or metric to open the drill-down dashboard page, which contains the associated widgets.
- For the other widgets that are shown in the diagram to be having a master-listener interaction, click any of the bars that represent the interface to change the listener widget that displays the related data for the selected interface.

The diagram shows the master-listener, and drill-down interactions between the available widgets:

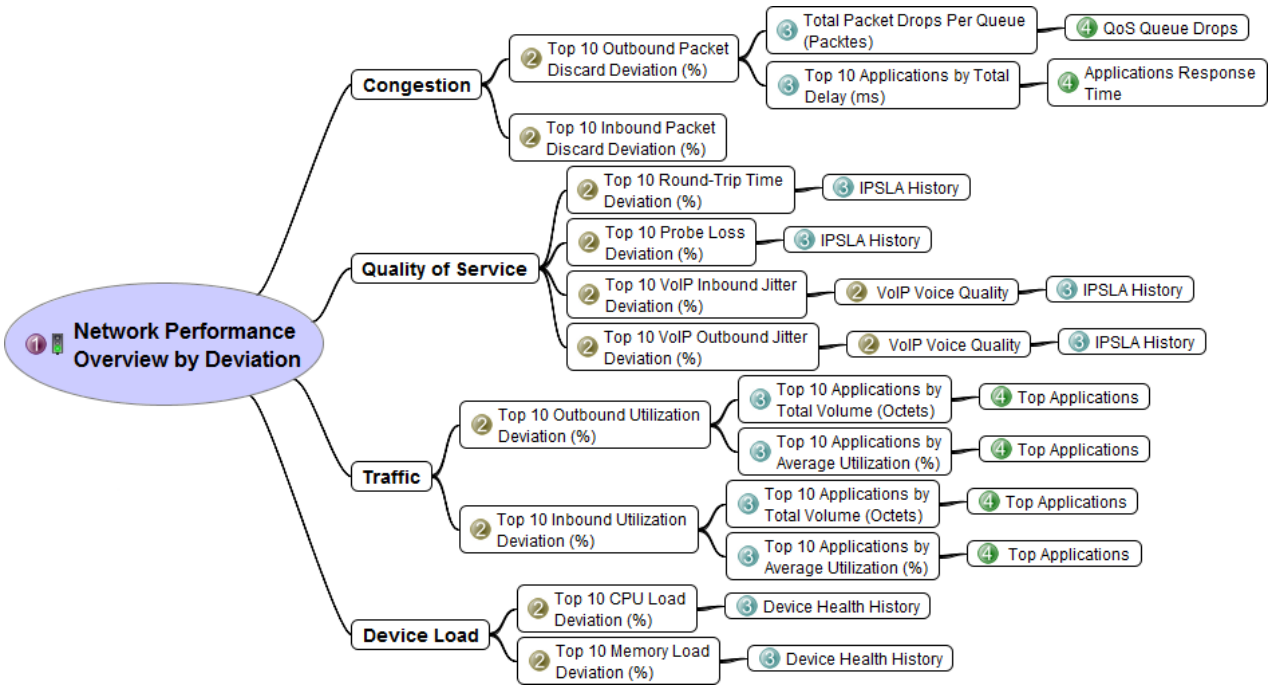

Network Performance Overview

1. Click **Home** > **Network Performance Overview**.

   The **Network Performance Overview: Top 10** dashboard loads.

   This dashboard displays the metrics values for top 10 resources.

2. You can filter data based on the following filters:

   - **Site**, the list contains the locations to view the traffic levels.
   - **Business Hour** , the list contains of **Business Hour**, **Non-Business Hour**, and both as **ALL**.
   - **Time Period**, the list contains of **Last Hour**, and **Last 24 Hours**.

   Note:

   - You can categorize your enterprise network based on different geographical areas and sites. You can configure the sites by specifying the IP address ranges and the time period for business hours and non-business hours. To configure these properties, see Configuring site grouping.
   - Refer to the conditions applied when you're using the **Site** and **Business Hour** filter options.

Table 62. Widget interactions

Area monitored	Master widgets	Listener widgets	Drill-down dashboard
Congestion	Top 10 Outbound Packet Discard Deviation (%) Click any bar that represents the interface outbound packet discard deviation to refresh the listener widgets to display the related data for the selected interface.	Total Packet Drops Per Queue (Packets) Click **Find out more details about Packet Drops per Queue. Click here** link to open the QoS Queue Drops dashboard in a new tab.	QoS Queue Drops
		Top 10 Applications by Total Delay (ms) Click any bar that represent an application in Top 10 Applications by Total Delay (ms), the **Application Response Time** page opens in a new tab.	Applications Response Time
	Top 10 Inbound Packet Discard Deviation (%)	N/A	N/A

Table 62. Widget interactions (continued)

Area monitored	Master widgets	Listener widgets	Drill-down dashboard
Quality of Service	Top 10 Round-Trip Time Deviation (%) Click any bar that represents the round-trip time between a source and destination device IP addresses to open the IPSLA History dashboard in a new tab.	N/A	IPSLA History
	Top 10 Probe Loss Deviation (%) Click any bar that represents the probe loss between a source and destination device IP addresses to open the IPSLA History dashboard in a new tab.	N/A	
	Top 10 VoIP Inbound Jitter Deviation (%) or Top 10 VoIP Outbound Jitter Deviation (%) Click any bar that represents the VoIP Inbound or Outbound Jitter Deviation between a source and destination device IP addresses to refresh the listener widgets. **Note:** Click the **Data Source** icon () to toggle between Top 10 VoIP Outbound Jitter Deviation (%) and Top 10 VoIP Inbound Jitter Deviation (%) widgets.	VoIP Voice Quality The widget displays the following metrics for the same devices: • MOS • IcPIF Click **Source - Destination: *<IP_Address>* - *<IP_Address>*** to open the IPSLA History dashboard in a new tab.	

Table 62. Widget interactions (continued)

Area monitored	Master widgets	Listener widgets	Drill-down dashboard
Traffic	Top 10 Outbound Utilization Deviation (%) Click any bar that represents the interface outbound utilization deviation to refresh the listener widgets to display the related data for the selected interface.	Top 10 Applications by Total Volume (Octets) Top 10 Applications by Average Utilization (%) Click **Find out more details about the Application. Click here** link to open the **Top Applications: Traffic Volume Details for Interface** dashboard a new tab.	Top Applications
	Top 10 Inbound Utilization Deviation (%) Click any bar that represents the interface inbound utilization deviation to refresh the listener widgets to display the related data for the selected interface.		
Device load	Top 10 CPU Load Deviation (%) Click any bar that represents the CPU Load Deviation on a resource, the Device Health History dashboard for the same resource opens in a new tab.	N/A	Device Health History
	Top 10 Memory Load Deviation (%) Click any bar that represents the Memory Load Deviation on a resource, the Device Health History dashboard for the same resource opens in a new tab.	N/A	

Table 63. Available Widgets

Widget name	Chart type	Typical uses
Top 10 Outbound Packet Discard Deviation (%)	Bar	Measures the number of outbound packets that are discarded even though no errors are detected to free the buffer space. It might be due to resource limitations on the outbound interface.

Table 63. Available Widgets (continued)

Widget name	Chart type	Typical uses
Total Packet Drops Per Queue (Packets)	Bar	Measures the number of packet drops per traffic class queue per interface.
Top 10 Inbound Packet Discard Deviation (%)	Bar	Measures the number of inbound packets that are discarded even though no errors are detected to free the buffer space. It might be due to resource limitations on the inbound interface.
Top 10 Applications by Total Delay (ms)	Bar	Measures maximum total time that it takes an application to respond to user requests. The total delay is the sum of max client network delay, max server network delay, and application delay.
Top 10 Round-Trip Time Deviation (%)	Bar	Measures the time that is taken between sending a UDP echo request message from a source device to the destination device and receiving a UDP echo reply from the destination device. It is useful in troubleshooting issues with business-critical applications by determining the round-trip delay times and testing connectivity between devices.
Top 10 Probe Loss Deviation (%)	Bar	Measures the percentage of probes lost. it shows the completion status of an RTT operation.
Top 10 VoIP Inbound Jitter Deviation (%)	Bar	Measures the variance in latency (jitter) deviation over time in incoming direction between two devices in your network.
Top 10 VoIP Outbound Jitter (ms)	Bar	Measures the variance in latency (jitter) deviation over time in outgoing direction between two devices in your network.

Table 63. Available Widgets (continued)

Widget name	Chart type	Typical uses
VoIP Voice Quality	Badge	Measures voice performance that is important for delivering voice quality services. This widget shows two metrics: • Mean Opinion Scores (MOS). A common benchmark to determine voice quality is MOS. With MOS, a wide range of listeners can judge the quality of voice samples on a scale of 1 (bad quality) to 5 (excellent quality). The Cisco IOS implementation of MOS takes RTT delay and packet loss into the MOS calculation. However, jitter is not included. The following colors on the widget denote the severity of degradation of voice quality: – Green, when the MOS value is between 4 - 5 – Amber, when the MOS value is between 2 - 3 – Red, when the MOS value is between 0 - 1 • Impairment/Calculated Planning Impairment Factor (IcPIF) IcPIF values are expressed in a typical range of 5 (low impairment) to 55 (high impairment). Typically, IcPIF values numerically less than 20 are considered adequate. The following colors on the widget denote the severity of degradation of voice quality: – Green, when the IcPIF value is between 0 - 13 – Amber, when the IcPIF value is between 14 - 33 – Red, when the IcPIF value is between 34 - 93
Top 10 Outbound Utilization Deviation (%)	Bar	Measures outbound bandwidth utilization deviation for the highest Outbound Packet Discards on the interfaces.
Top 10 Applications by Total Volume (Octets)	Bar	Measures the corresponding applications with highest bandwidth utilization by volume.
Top 10 Inbound Utilization Deviation (%)	Bar	Measures the bandwidth utilization deviation for incoming traffic for the highest Inbound Packet Discard Deviation on the interfaces.
Top 10 Applications by Average Utilization (%)	Bar	Measures the corresponding applications with highest bandwidth utilization by percentage.

Table 63. Available Widgets (continued)		
Widget name	**Chart type**	**Typical uses**
Top 10 CPU Load Deviation (%)	Bar	Measure the total CPU utilization deviation in percentage across all vendor devices in your network. Currently, Network Performance Insight supports the following vendors: • Cisco • Huawei • Juniper
Top 10 Memory Load Deviation (%)	Bar	Measure the total memory usage deviation in percentage across all vendor devices in your network. Currently, Network Performance Insight supports the following vendors: • Cisco • Huawei • Juniper

WiFi Overview dashboard

WiFi overview dashboards are instrumental for enterprises in monitoring the health and performance of the WiFi network. This dashboard represents key performance indicators (KPI), in the form of widgets, of a WiFi network to monitor the network. You can navigate further from these widgets to analyze specific diagnostics.

The WiFi Overview is based on the following three key components:

- Client Count: The number of clients that are connected to any particular WiFi network (SSID) or Access point (AP) informs the network administrator about the amount of load on that SSID or Access point. At any given point in time, if the client count increases, network administrator can take a decision on moving clients between Access Points or SSIDs and reduce the load on network.
- Signal Characteristics: WiFi overview dashboards, display signal quality details by using two characteristics as follows.
  - Received Signal Strength Indicator (RSSI): RSSI is an estimated measure of power of frequency of a signal that a client receives from the access points. It varies with the distance between clients and access points, hence indicates the network administrator to adjust the client-access point locations and maintain the overall network performance.
  - Signal-to-Noise ratio (SNR): Signal-to-Noise ratio, is the comparison between the power of signal and the presence of noise in the network. This ratio is an indicator of the quality of the signal. The more the noise the less healthy a signal is. If a client is experiencing low signal-to-noise ratio, a network administrator adjusts the placement of AP to avoid physical obstacles or adjust the client count of the network.
- Worst Performing channels: The worst performing channels widget in the dashboard, gives the details of the channels that are facing high intensity of interference and noise. It helps to identify the channels that are acting as loopholes in the network, and gives the network administrator a lead to analyze the network to improve the performance.

Refer the widget details to understand the fundamentals of WiFi overview dashboards.

Available widgets and their interactivity

The drill-down widget interactions can be seen in the following ways:

- For any widgets with the ( ) drill-down icon, from the chart, click any bar that represents the interface or metric to open the drill-down dashboard page, which contains the associated widgets.

The diagram shows the interactions between widgets and their drill-down widgets:

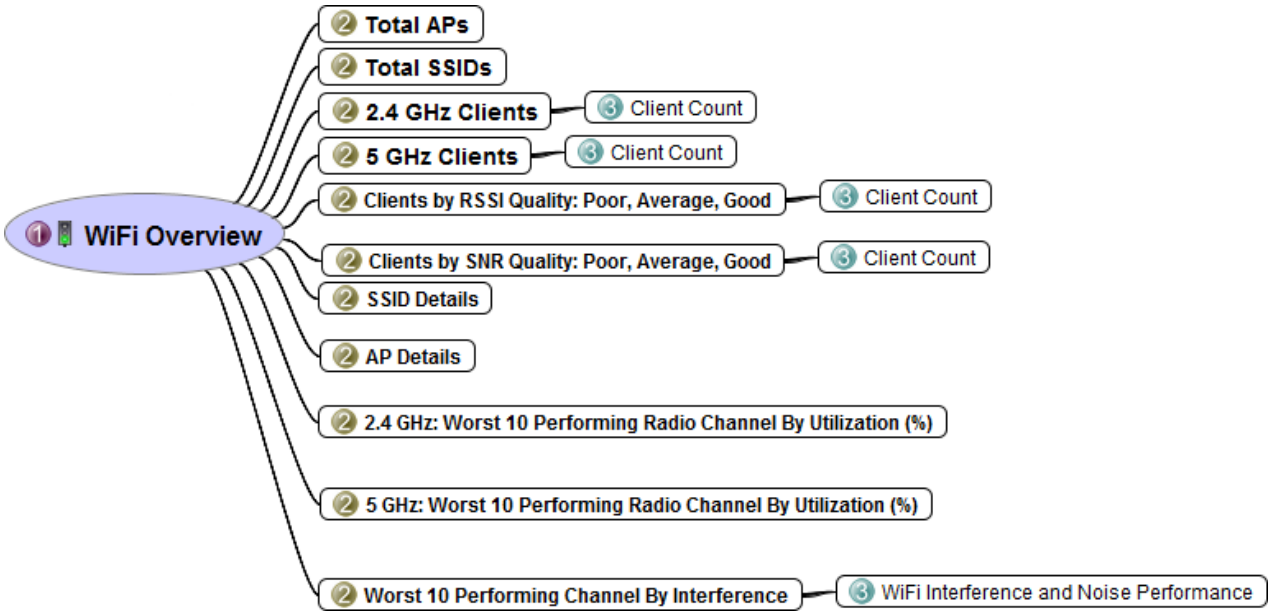

Network Performance Overview

1. Click **Home** > **WiFi Overview**.

   The **WiFi Overview** dashboard loads.

   This dashboard displays the metrics values for KPIs.
2. You can filter data based on the following filters:
   - **Controller**, contains the IP address of a controller that manages the Access Points of the network.
   - **Time Period**, the list contains of **Last Hour**, **Last 6 Hours**, **Last 12 Hours**, **Last 24 Hours**, **Last 7 Days**, **Last 30 Days**, **Last 365 Days** and **Custom**.

Table 64. Widget interactions

Master widgets	Drill-down dashboard
Total APs	N/A
Total SSIDs	N/A
2.4 GHz Clients	WiFi Client Count
5 GHz Clients	
Clients by RSSI Quality: Poor	WiFi Client Count
Clients by RSSI Quality: Average	
Clients by RSSI Quality: Good	
Clients by SNR Quality: Poor	WiFi Client Count
Clients by SNR Quality: Average	
Clients by SNR Quality: Good	
WiFi Network Details	N/A

Table 64. Widget interactions (continued)	
Master widgets	**Drill-down dashboard**
Access Points Details	N/A
2.4 GHz: Worst 10 Performing Radio By Channel Utilization (%)	N/A
5 GHz: Worst 10 Performing Radio By Channel Utilization (%)	N/A
Worst 10 Performing Channel By Interference	WiFi Interference and Noise Performance

Table 65. Available Widgets		
Widget name	**Chart type**	**Typical uses**
Total APs	Badge	Displays the number of Access Points that are connected to the selected controller.
Total SSIDs	Badge	Displays the number of SSID (WiFi Network) configured to the selected controller.
2.4 GHz Clients	Badge	Displays the client count to 2.4 GHz frequency band of signal.
5 GHz Clients	Badge	Displays the client count to 5 GHz frequency band of signal.
Clients by RSSI Quality: Poor	Badge	Displays the number of clients who are receiving a poor signal strength. The standard RSSI range is less than or equal to -70 dBm.
Clients by RSSI Quality: Average	Badge	Displays the number of clients who are receiving an average signal strength. The standard RSSI range is -46 to -69 dBm
Clients by RSSI Quality: Good	Badge	Displays the number of clients who are receiving a good signal strength. The standard RSSI range is greater than or equal to -45 dBm.
Clients by SNR Quality: Poor	Badge	Displays the number of clients experiencing the lowest signal-to-noise ratio. The standard SNR range is 0 to 12 dBm. If more number of clients are in the Poor range, network performance is hampered.
Clients by SNR Quality: Average	Badge	Displays the number of clients experiencing moderate level of signal-to-noise ratio. The standard SNR range is 13 to 19 dBm.

Table 65. Available Widgets (continued)		
Widget name	**Chart type**	**Typical uses**
Clients by SNR Quality: Good	Badge	Displays the number of clients experiencing the highest signal-to-noise ratio. The standard SNR range is equal or greater than 20 dBm. The more number of clients in this range the better for network performance.
Wi-Fi Network Details	Grid	Displays the list of attributes of SSID. 1. Status of SSID as Enabled (1) or Disabled (0) 2. QoS Profile (0: Bronze, 1: Silver, 2: Gold, 3: Platinum) 3. Number of clients connected to the SSID
Access Points Details	Grid	Displays the list of attributes of Access Points (AP) 1. Location of AP 2. IP address of AP 3. Status as Enabled (1) or Disabled (2) 4. Number of clients connected to the AP
2.4 GHz: Worst 10 Performing Radio By Channel Utilization (%)	Bar	Displays the channels that are utilizing most of the bandwidth of 2.4 GHz frequency.
5 GHz: Worst 10 Performing Radio By Channel Utilization (%)	Bar	Displays the channels that are utilizing most of the bandwidth of 5 GHz frequency.
Worst 10 Performing Channel By Interference	Bar	Displays channels that are experiencing highest intensity of interference.

WiFi Client Count

WiFi Client Count *Dashboard* gives you the number of clients that are attached to 2.4 GHz and 5 GHz frequencies and the client count that is affected by RSSI and SNR quality trend of WiFi signal. This client count is distributed based on a timestamp. A time series representation of the client count helps you to identify a definite pattern of client experience in terms of frequency and quality of signal.

Access WiFi Client Count

1. Click **Home** > **WiFi Overview**

   Clients by RSSI Quality

   - Click **Poor**, **Average**, or **Good** from the badge widget.

   Clients by SNR Quality

   - Click **Poor**, **Average**, or **Good** from the badge widget.

Client Count widgets

You have the options to filter the existing Client Count metrics by using the following filters.

1. From the filter options, choose the **Controller** and **Time Period**, and click **Apply Filter**.

   The dashboard refreshes the data according to the filter attribute values.

Table 66. Client Count Widgets

Widget Name	Chart Type	Description
2.4GHz:Client Count Trend	Timeseries	Displays the number of clients that are connected to 2.4 GHz frequency based on timestamp
5GHz:Client Count Trend	Timeseries	Displays the number of clients that are connected to 5 GHz frequency based on timestamp
Client Count by RSSI Quality Trend	Timeseries	Displays the RSSI quality (signal strength) experienced by a number of clients based on timestamp. The client count is divided into three categories of signal strength: Poor, Average, Good.
Client Count by SNR Quality Trend	Timeseries	Displays the SNR quality (Signal-to-Noise ratio) experienced by a number of clients based on timestamp. The client count is divided into three categories of signal-to-noise ratio: Poor, Average, Good.

WiFi Interference and Noise Performance

A drill-down widget of Top 10 worst performing Channels by Interference, this widget displays the characteristics of a signal. You can identify the interference and noise ratio for channels. The timeseries chart helps notice a pattern of signal health for particular controller and channel.

Access WiFi Interference and Noise Performance

1. Click **Home** > **WiFi Overview**

   Click any bar on the **Top 10 Worst Performing Channels by Interference.**

WiFi Interference and Noise Performance widgets

You have the options to filter the existing WiFi Interference and Noise Performance metrics with the following filters.

1. From the filter options, choose the **Controller**, **AP Radio Channel** and **Time period** and click **Apply Filter**.

   The dashboard refreshes the data according to the filter attribute values.

Table 67. WiFi Interference and Noise Performance widgets		
Widget Name	**Chart Type**	**Description**
Interference score (%) Trend	Timeseries	Displays the interference score for a selected controller and AP radio channel. For the selected controller, you can change the AP radio channel to identify a channel with less interference score.
Interference Power (dBm) Trend	Timeseries	Displays the intensity of interference for the selected channel and AP controller. Interference is measured in decibels (dBm). For the selected controller, you can change the AP radio channel to identify a channel with less interference power.
Noise (dBm) Trend	Timeseries	Displays the intensity of noise present in the selected channel of the selected controller. For the selected controller, you can change the AP radio channel to identify a channel with less noise in its signal.

IP Links Performance Overview dashboard

The fundamental usage of IP Links Performance overview dashboard is for the network administrator to reduce the mean time taken to identify an event in a network, its effects on the network and its root cause to efficiently manage the network performance.

The IP links performance is monitored through a heatmap which provides an overall view of outbound one way link performance for speech applications between a source and destination server present at various geographical locations. The performance of IP links is affected by factors such as Latency, Jitter and Packet Loss. These three factors are defined as follow:

- Latency: Latency is a measurement of delay of a message from source to destination. It is measured is milliseconds.
- Jitter: Jitter is defined as the variation in arrival time of messages from source to destination. In other words, Jitter is a variation in latency. It is measured is milliseconds.
- Packet Loss: Packet loss is a percentage of packets lost with respect to the packets sent.

This dashboard displays the Latency between the source and destination servers in a particular geographical area. Based on the Latency value, network administrator can drill down to links having highest Latency and can get more information about Jitter and Packet Loss between the servers.

Recommended Average Outbound One Way Latency for Speech Application

A new chart type called Heat map is introduced in this dashboard to represent latency. A heat map is a graphical representation of data where the individual values contained in a matrix are represented as colors.

The heat map looks as follows:

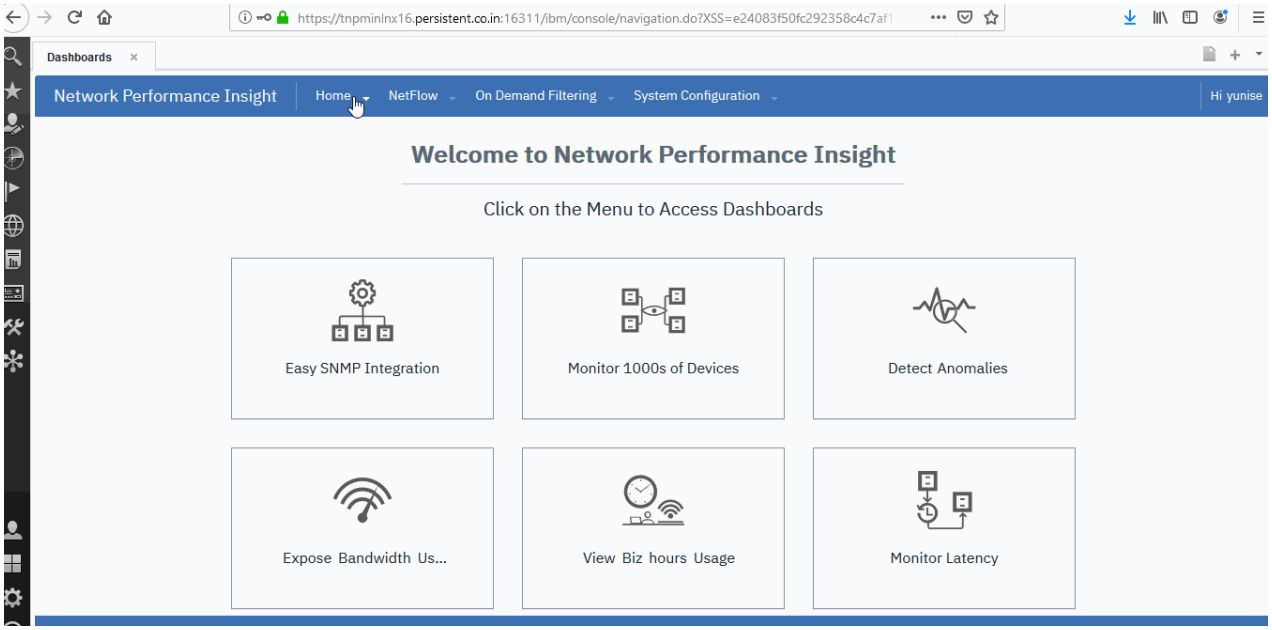

The parameters of Heat map are:

- Rows: The rows of the map represent Source IP address
- Columns: The columns of the map represent Destination IP address.
- Values: The color coded values of the map represent latency between the source and destination IP links. Click any value in the heat map to drill down to the Source and Destination Details dashboard.

  The standard ranges for Latency and the color associated to each range are given at the bottom of the heat map. These ranges are taken from ITU-T G.114 standards for one way latency for speech applications.

  The standard ranges are as follows:

Table 68. Recommended standard ranges for Latency in speech applications

Color Code	Description
Excellent (0-150 ms)	This color depicts that the Latency between source and destination IP links is between 0-150 ms. The IP links are performing excellent with this range of Latency.
Very Good (151-237 ms)	This color depicts that the Latency between source and destination IP links is between 151-237 ms. The IP links are performing very well with this range of Latency.
Good (238-337 ms)	This color depicts that the Latency between source and destination IP links is between 237-337 ms. The IP links are performing well with this range of Latency.

Table 68. Recommended standard ranges for Latency in speech applications (continued)	
Color Code	**Description**
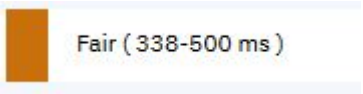	This color depicts that the Latency between source and destination IP links is between 338-500 ms. The IP links are performing fair with this range of Latency.
	This color depicts that the Latency between source and destination IP links is between 501-700 ms. The IP links are performing poor with this range of Latency.
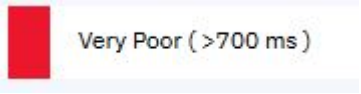	This color depicts that the Latency between source and destination IP links is greater than 700 ms. The IP links are performing very poor with this range of Latency.

The values in the heat map which are blank and do not come under any of the color codes depict that there is no data flow and communication between the source and destination IP links.

IP Links Performance Overview

1. Click **Home** > **IP Links Performance Overview**
   This dashboard displays IP Links performance in the form of heat map.
2. You can filter data based on the following filters:
   - **Site**, the list contains the locations of different networks containing various source and destination servers.
   - **Business Hour** , the list contains of **Business Hour**, **Non-Business Hour**, and both as **ALL**.
   - **Time Period**, the list contains the following options:
     - **Last Hour**
     - **Last 6 Hours**
     - **Last 12 Hours**
     - **Last 24 Hours**
     - **Last 7 Days**
     - **Last 30 Days**

   Note:

   - You can categorize your enterprise network based on different geographical areas and sites. You can configure the sites by specifying the IP address ranges and the time period for business hours and non-business hours. To configure these properties, see Configuring site grouping.
   - Refer to the conditions applied when you're using the **Site** and **Business Hour** filter options.

Source and Destination Details dashboard

This is a drill-down dashboard of IP Links Performance Overview dashboard. In the IP Link Performance Overview dashboard, when you click on the colored boxes of the heat map, it drills down to the Source and Destination Details dashboard.

The latency, jitter and packet loss are interdependent in a way that when packet loss is more, user will experience more Latency and in turn the jitter value will get affected, as there will be fluctuations in the latency.

Source and Destination Details

1. Click **Home** > **IP Links Performance Overview**
   Click any colored value on the **Recommended Average Outbound One Way Latency for Speech Application**.
2. You can filter the data based on following filters:
   - **Source**, the list contains IP addresses of Source servers.
   - **Destination**, the list contains IP addresses of Destination servers.
   - **Time Period**, the list contains the following options:
     - **Last Hour**
     - **Last 6 Hours**
     - **Last 12 Hours**
     - **Last 24 Hours**
     - **Last 7 Days**
     - **Last 30 Days**
3. Click **Apply filter**

Table 69. Available Widgets

Widget name	Chart type	Typical uses
Source-Destination	Badge	Displays the IP addresses of the selected Source and Destination.
Maximum Latency • Outbound (ms) • Inbound (ms)	Badge	Displays the Outbound and Inbound Latency in milliseconds for source and destination IP links. Standard range for Outbound and Inbound Latency: • Green: 0-337 ms • Yellow: 338-700 ms • Red: Greater than 700 ms
Maximum Jitter • Outbound (ms) • Inbound (ms)	Badge	Displays the Outbound and Inbound Jitter in milliseconds for source and destination IP links. Standard range for Outbound and Inbound Jitter: • Green: 0-30 ms • Yellow: 31-40 ms • Red: Greater than 40 ms

Table 69. Available Widgets (continued)		
Widget name	**Chart type**	**Typical uses**
Maximum Packet Loss • Outbound (%) • Inbound (%)	Badge	Displays the Outbound and Inbound Packet Loss in percentage for source and destination IP links. Standard range for Outbound and Inbound Packet Loss: • Green: 0%-1% • Yellow: 1.01%-5% • Red: Greater than 5%
Latency, Jitter and Packet Loss Trend • Outbound • Inbound	Timeseries	Displays the Latency in milliseconds, Jitter in milliseconds and Packet Loss in percentage for source and destination IP links at a specific timestamps.

Load Balancer dashboards

The Load Balancer dashboards monitor network load balancing with F5 BIG IP technology. The Load Balancer dashboards provide an insight into the distribution of network traffic across server resources in multiple geographies. These servers can be on premises or hosted on cloud. These dashboards provide visualizations on the performance and availability of your global applications.

Data centers that are spread geographically can slow down user requests and network traffic needs that must be managed instantly and load balance during peak demands and downtime. The number of application connection requests and utilization from users can exceed the capacity of a server that hosts the application. The load balancing mechanism helps to distribute the inbound requests and processing load of responses across a group of servers that run the same application effectively.

Load balancer components

Typically, the load balancer services that are based on F5 technology consists of the following components:

- Global Traffic Manager (GTM)

  GTM is also called as BIG-IP DNS that improves the performance and availability of your global applications by sending users to the closest or best-performing physical, virtual, or cloud environment. It also hyperscales and secures your DNS infrastructure from DDoS attacks.

  BIG-IP's module built to monitor the availability and performance of global resources and use that information to manage network traffic patterns.

- Local Traffic Manager (LTM)

  BIG-IP's module that manages and optimizes traffic for network applications and clients. BIG-IP LTM treats all network traffic as stateful connection flows. Even connectionless protocols such as User Datagram Protocol (UDP) and Internet Control Message Protocol (ICMP) are tracked as flows.

- Virtual servers

  A virtual server allows BIG-IP systems to send, receive, process, and relay network traffic. A virtual server is a proxy of the actual server (physical, virtual, or container). Combined with a virtual IP address, which is the application endpoint that is presented to the outside world.

- Pools

A configuration object that groups pool members together to receive and process network traffic that is determined by a specified load balancing algorithm. Collections of similar services available on any number of hosts.

- Pool members

  A pool member is a node and service port to which BIG-IP LTM can load balance the network traffic. Nodes can be members of multiple pools.

  The pool member includes the definition of the application port and the IP address of the physical or virtual server. Refer to it as the service. Unique load balancing and health monitoring based on the services instead of the host.

- Nodes

  A configuration object represented by the IP address of a device on the network.

- Wide IPs

  The Fully Qualified Domain Name (FQDN) of a service.

Load balancing methods

Static load balancing methods do not use any traffic metrics from the node to distribute traffic. Dynamic load balancing methods use traffic metrics from the node to distribute traffic.

Health monitors keep a close eye on the health of a resource to deem it available or unavailable. They are independent of load balancing methods.

Performance monitors measure the hosts performance and dynamically send traffic to hosts in the pool. They work with corresponding dynamic load balancing methods. Health monitors can be applied at the node level or at the pool level, but performance monitors can be applied at the node level only.

For more information about the different load balancing methods and algorithms that can be used, see Understanding F5 Load Balancing Methods.

Load Balancer Overview dashboard

With the help of F5 Load Balancer from F5 Networks Inc, network traffic is diverted from servers that are overloaded to the other servers that can handle the load. The F5 load balancer service consists of many components. All the components and the key performance metrics that are collected by them are displayed in the overview dashboard.

Available widgets and their interactivity

The Load Balancer Overview dashboard and its widgets:

The master-listener and drill-down widget interactions can be seen in the following ways:

- For any widgets with the ( ) drill-down icon, from the chart, click any bar that represents the interface or metric to open the drill-down dashboard page, which contains the associated widgets.
- For the other widgets that are shown in the diagram that have a master-listener interaction, click any of the values to change the listener widget that displays the related data for the selected value.

The diagram shows the master-listener, and drill-down interactions between the available widgets:

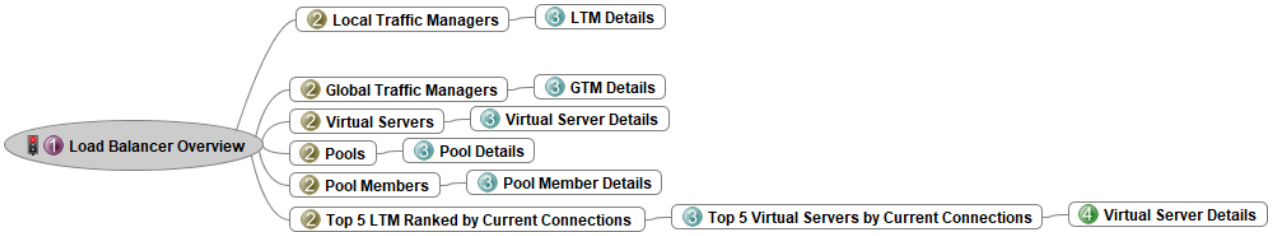

Load Balancer Overview

1. Click **Home** > **Load Balancer** > **Load Balancer Overview**.

   This dashboard displays the F5 load balancer components.
2. You can filter data based on the following filters:
   - **Site**, the list contains the locations to view the traffic levels.
   - **Business Hour** , the list contains of **Business Hour**, **Non-Business Hour**, and both as **ALL**.
   - **Time Period**, the list contains the following options:
     - **Last Hour**
     - **Last 6 Hours**
     - **Last 12 Hours**
     - **Last 24 Hours**
     - **Last 7 Days**
     - **Last 30 Days**
     - **Last 365 Days**
     - **Custom**

Note:

- You can categorize your enterprise network based on different geographical areas and sites. You can configure the sites by specifying the IP address ranges and the time period for business hours and non-business hours. To configure these properties, see Configuring site grouping.
- Refer to the conditions applied when you're using the **Site** and **Business Hour** filter options.

3. Click **Apply Filter**.

Table 70. Widget interactions

Master widgets	Listener widgets	Drill-down widgets
Local Traffic Managers	N/A	LTM Details
Global Traffic Managers	N/A	GTM Details
Virtual Servers	N/A	Virtual Server Details
Pools	N/A	Pool Details
Pool Members	N/A	Pool Member Details
Top 5 LTM - Ranked by Current Connections	Top 5 Virtual Servers by Current Connections *<LTM_Device>*	N/A
Top 5 Virtual Servers by Current Connections	N/A	Virtual Server Details

Table 71. Available widgets

Widget name	Chart type	Description
Local Traffic Managers	Badge	Displays the number of Local Traffic Managers that are available in your load-balancing service. It also displays the state of the LTMs. The value and color code indicate the number of LTMs that are unavailable, partially available, and available. Explanation for the color codes is as follows: • Green: Available • Orange: Partially Available • Red: Unavailable

Table 71. Available widgets (continued)

Widget name	Chart type	Description
Virtual Servers	Badge	Displays the number of Virtual Servers that are available in your load-balancing service. It also displays the state of the Virtual Servers. The value and color code indicate the number of Virtual Servers that are unavailable, partially available, and available. Explanation for the color codes is as follows: • Green: Available • Orange: Partially Available • Red: Unavailable
Global Traffic Managers	Badge	Displays the number of Global Traffic Managers that are available in your load-balancing service. It also displays the state of the GTMs. The value and color code indicate the number of GTMs that are unavailable, partially available, and available. Explanation for the color codes is as follows: • Green: Available • Orange: Partially Available • Red: Unavailable
Pools	Badge	Displays the number of Pools that are available in your load-balancing service. It also displays the state of the Pools. The value and color code indicate the number of Pools that are unavailable, partially available, and available. Explanation for the color codes is as follows: • Green: Available • Orange: Partially Available • Red: Unavailable

Table 71. Available widgets (continued)

Widget name	Chart type	Description
Pool Members	Badge	Displays the number of Pool Members that are available in your load-balancing service. It also displays the state of the Pool Members. The value and color code indicate the number of Pools that are unavailable, partially available, and available. Explanation for the color codes is as follows: • Green: Available • Orange: Partially Available • Red: Unavailable
Top 5 LTM - Ranked by Current Connections	Grid	Displays the top five LTMs by current connections. Current or active connections are the connections that are available to that LTM at a particular time that is based on the selected time period. The metrics that are displayed in the Grid widget are as follows: • `LTM` • `Current Connections` • `Connections per Second` • `Inbound Throughput (bps)` • `Outbound Throughput (bps)` • `CPU Utilization (%)` • `Memory Utilization (%)`
Top 5 Virtual Servers by Current Connections	Line	Displays the top five Virtual Servers by current connections for the selected LTM. Current or active connections are the connections that are available to that LTM at a particular time that is based on the selected time period. The Time Zoom feature is available for this widget.

GTM Details dashboard

F5® BIG-IP® Global Traffic Manager™ (GTM) distributes DNS and user application requests based on business policies, data center and cloud service conditions, user location, and application performance. BIG-IP GTM delivers flexible global application management in virtual and cloud environments.

Available widgets and their interactivity

The GTM Details and its widgets:

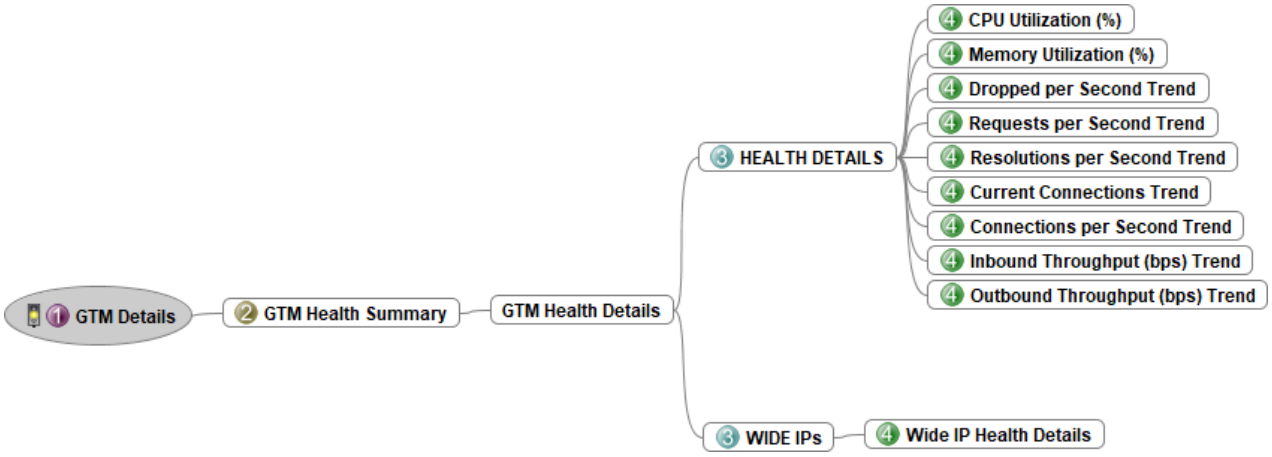

GTM Details

1. Click **Home** > **Load Balancer** > **GTM Details**.

   This dashboard displays the F5 BIG-IP LTM information in your load-balancing environment.
2. You can filter data based on the following filters:
   - **Site**, the list contains the locations to view the traffic levels.
   - **Business Hour** , the list contains **Business Hour**, **Non-Business Hour**, and both as **ALL**.
   - **Top N**, the list contains **10**, **20**, **50**, and **ALL** to display the top N performers.
   - **Time Period**, the list contains the following options:
     - **Last Hour**
     - **Last 6 Hours**
     - **Last 12 Hours**
     - **Last 24 Hours**
     - **Last 7 Days**
     - **Last 30 Days**
     - **Last 365 Days**
     - **Custom**

   Note:

   - You can categorize your enterprise network based on different geographical areas and sites. You can configure the sites by specifying the IP address ranges and the time period for business hours and non-business hours. To configure these properties, see Configuring site grouping.
   - Refer to the conditions applied when you're using the **Site** and **Business Hour** filter options.
3. Click **Apply Filter**.

Table 72. Widget interactions		
Master widgets	**Listener widgets**	**Drill-down widgets**
GTM Health Summary	GTM Health Details	N/A
	Health Details • CPU Utilization (%) • Memory Utilization (%) • Dropped per Second Trend • Requests per Second Trend • Resolutions per Second Trend • Current Connections Trend • Connections per Second Trend • Inbound Throughput (bps) Trend • Outbound Throughput (bps) Trend	N/A
	WIDE IPS • Wide IP Health Details	N/A

Table 73. Available widgets

Widget name	Chart type	Description
GTM Health Summary	Grid	Displays the list of Global Traffic Managers that are available in your load-balancing service. This widget displays the metrics to monitor the health of the GTM. Click any GTM from the table to see the details. The following metrics are monitored: • IP Address • Requests per Second • Max Requests per Second • Resolutions per Second • Max Resolutions per Second • Dropped per Second • Max Dropped per Second • Current Connections • Max Current Connections • Connections per Second • Max Connections per Second • Inbound Throughput (bps) • Max Inbound Throughput (bps) • Outbound Throughput (bps) • Max Outbound Throughput (bps)
GTM Health Details	Badge	Displays the selected GTM identifies and the status of the GTMs that are available and their status.

Table 73. Available widgets (continued)

Widget name	Chart type	Description
Status	Badge	It displays the state of the GTM. The state of a GTM is derived from the status of Wide IPs. • If all Wide IPs are available, status is 1, then the color is Green. GTMs are all available. • If Wide IPs are partially available with a mix status of 1,2,3,4 then the color is Orange. GTMs are partially available. • If all Wide IPs are not available with a status 3,4, then the color is Red. GTMs are unavailable. Explanation for the color codes is as follows: • Green: Available • Orange: Partially Available • Red: Unavailable
CPU Utilization (%)	Gauge	CPU availability that is used across all available cores in the selected GTM.
Memory Utilization (%)	Gauge	Memory availability that is used across all available cores in the selected GTM.
Dropped per Second Trend	Line	The number of requests that are dropped per second. **Note:** The TimeZoom feature is available for all the Line charts in this dashboard.
Requests per Second Trend	Line	The number of requests that are satisfied per second.
Resolutions per Second Trend	Line	The number of transaction requests that are completed per second.
Current Connections Trend	Line	Current or active connections that are available to that GTM at a particular time that is based on the selected time period. It is a measure of the number of client/server requests that can be handled.
Connections per Second Trend	Line	Active GTM connections that are available per second.

Table 73. Available widgets (continued)		
Widget name	**Chart type**	**Description**
Inbound Throughput (bps) Trend	Line	Measured in bits per second. Typically, it is the amount of data that is sent to and from the GTM to the client.
Outbound Throughput (bps) Trend	Line	Measured in bits per second. Typically, it is the amount of data that is sent to and from the client to the GTM.
Wide IP Health Details	Grid	The main configuration element in a GTM is called a Wide IP or WIP. A Wide IP equates to the common URL that you are load balancing. For example, www.yourcompany.com. A pool or pools are usually attached to a WIP, which contain the IP addresses it's intelligently resolving. BIG-IP® DNS selects pools based on the order in which they are listed in a wide IP. The Wide IP contains one or more pools, which in turn contain one or more virtual servers. • Status • Wide IP Name • Requests per Second • Max Requests per Second • Resolutions per Second • Max Resolutions per Second • Dropped per Second • Max Dropped per Second

LTM Details dashboard

F5® BIG-IP® Local Traffic Manager™ (LTM) is an intelligent application traffic management tool. It monitors and controls the network traffic to make sure that the applications are always available, secure, and fast. BIG-IP LTM includes static and dynamic load balancing to eliminate single points of failure.

Available widgets and their interactivity

The LTM Details and its widgets:

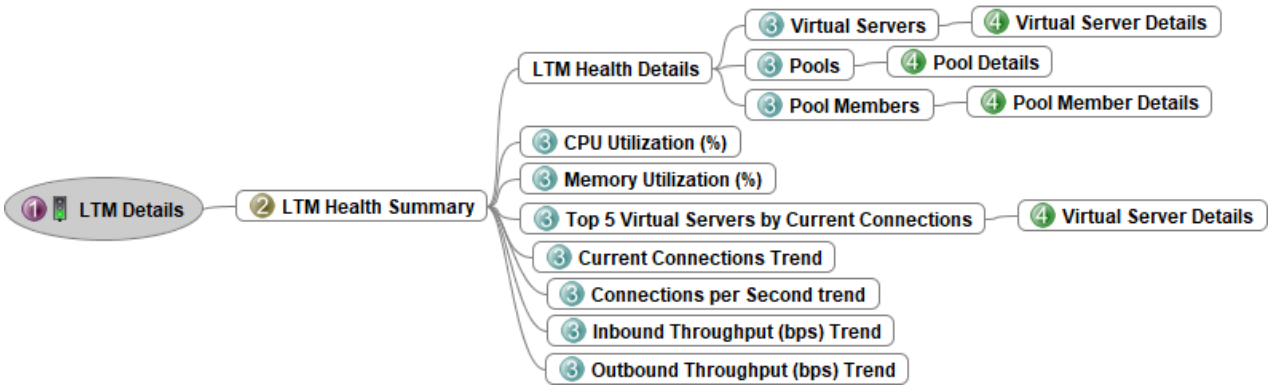

LTM Details

1. Click **Home** > **Load Balancer** > **LTM Details**.

   This dashboard displays the F5 BIG-IP LTM information in your load-balancing environment.
2. You can filter data based on the following filters:
   - **Site**, the list contains the locations to view the traffic levels.
   - **Business Hour** , the list contains of **Business Hour**, **Non-Business Hour**, and both as **ALL**.
   - **Top N**, the list contains **10**, **20**, **50**, and **ALL** to display the top N performers.
   - **Time Period**, the list contains the following options:
     - **Last Hour**
     - **Last 6 Hours**
     - **Last 12 Hours**
     - **Last 24 Hours**
     - **Last 7 Days**
     - **Last 30 Days**
     - **Last 365 Days**
     - **Custom**

   Note:

   - You can categorize your enterprise network based on different geographical areas and sites. You can configure the sites by specifying the IP address ranges and the time period for business hours and non-business hours. To configure these properties, see Configuring site grouping.
   - Refer to the conditions applied when you're using the **Site** and **Business Hour** filter options.
3. Click **Apply Filter**.

Table 74. Widget interactions

Master widgets	Listener widgets	Drill-down widgets
LTM Health Summary	LTM Health Details	N/A
	Virtual Servers	Virtual Servers Details
	Pools	Pool Details
	Pool Members	Pool Member Details
	CPU Utilization (%)	N/A
	Memory Utilization (%)	N/A
	Top 5 Virtual Servers by Current Connections	Virtual Servers Details
	Current Connections Trend	N/A
	Connections per Second Trend	N/A
	Inbound Throughput (bps) Trend	N/A
	Outbound Throughput (bps) Trend	N/A

Table 75. Available widgets

Widget name	Chart type	Description
LTM Health Summary	Grid	Displays the list of Local Traffic Managers that are available in your load-balancing service. This widget displays the metrics to monitor the health of the LTMs. Click any LTM from the table to see the details. The following metrics are monitored: • LTM • IP Address • Current Connections • Max Current Connections • Connections per Second • Max Connections per Second • Inbound Throughput (bps) • Max Inbound Throughput (bps) • Outbound Throughput (bps) • Max Outbound Throughput (bps) • CPU Utilization (%) • Max CPU Utilization (%) • Memory Utilization (%) • Max Memory Utilization (%)

Table 75. Available widgets (continued)

Widget name	Chart type	Description
LTM Health Details • Status • Virtual Servers • Pools • Pool Members	Badge	Status displays the state of the LTMs. The state of a LTM is derived from the status of Virtual Servers. • If all Virtual Servers are available, status is 1, then the color is Green. LTM is available. • If Virtual Servers are partially available with a mixed status of 1,2, 3, 4, then the color is Orange. LTM is partially available. • If all Virtual Servers are not available with a status of 4, then the color is Red. LTM is unavailable. Explanation for the color codes is as follows: • Green: Available • Orange: Partially Available • Red: Unavailable Displays the total number of Virtual Servers, Pools, and Pool Members that are associated with a selected LTM. It also displays the state of the Virtual Servers, Pools, and Pool Members for the LTM. The value and color code indicate the number of devices that are unavailable, partially available, and available. Explanation for the color codes is as follows: • Green: Available • Orange: Partially Available • Red: Unavailable
CPU Utilization (%)	Gauge	CPU availability that is used across all available cores in the selected LTM.
Memory Utilization (%)	Gauge	Memory availability that is used across all available cores in the selected LTM.

Table 75. Available widgets (continued)		
Widget name	**Chart type**	**Description**
Top 5 Virtual Servers by Current Connections	Line	Top five Virtual Servers that are available by the current connections in the selected LTM. It drills down to the specific Virtual Servers Details dashboard. **Note:** The TimeZoom feature is available for all the Line charts in this dashboard.
Current Connections Trend	Line	Current or active connections that are available to that LTM at a particular time that is based on the selected time period. It is a measure of the number of client/ server requests that can be handled.
Connections per Second Trend	Line	Active LTM connections that are available per second.
Inbound Throughput (bps) Trend	Line	Measured in bits per second. Typically, it is the amount of data that is sent to and from the LTM to the client.
Outbound Throughput (bps) Trend	Line	Measured in bits per second. Typically, it is the amount of data that is sent to and from the client to the LTM.

Related concepts

Properties of the Load Balancer dashboards

Virtual Server Details dashboard

A virtual server is one of the most important components of any BIG-IP system. The F5 Virtual Server is a traffic management object on your F5 BIG-IP device. It is the representation of multiple servers to the user as a single server. The F5 Virtual Server is a virtual IP that serves user requests. It transmits the requests to the pool that you configure. It is represented by a virtual IP address and a service, such as *<IP_address>*:80. The primary purpose of a virtual server is to distribute traffic.

Available widgets and their interactivity

The Virtual Server Details and its widgets:

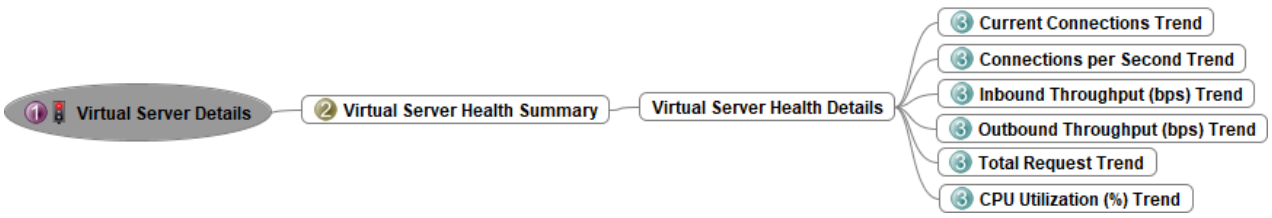

Virtual Server Details

1. Click **Home** > **Load Balancer** > **Virtual Server Details**.

This dashboard displays the F5 BIG-IP Virtual Servers information in your load-balancing environment.

2. You can filter data based on the following filters:
   - **Site**, the list contains the locations to view the traffic levels.
   - **LTM**, the list contains the available LTM devices in your load-balancing environment.
   - **Business Hour** , the list contains of **Business Hour**, **Non-Business Hour**, and both as **ALL**.
   - **Top N**, the list contains **10**, **20**, **50**, and **ALL** to display the top N performers.
   - **Time Period**, the list contains the following options:
     - **Last Hour**
     - **Last 6 Hours**
     - **Last 12 Hours**
     - **Last 24 Hours**
     - **Last 7 Days**
     - **Last 30 Days**
     - **Last 365 Days**
     - **Custom**

   Note:

   - You can categorize your enterprise network based on different geographical areas and sites. You can configure the sites by specifying the IP address ranges and the time period for business hours and non-business hours. To configure these properties, see Configuring site grouping.
   - Refer to the conditions applied when you're using the **Site** and **Business Hour** filter options.
3. Click **Apply Filter**.

Table 76.

Master widgets	Listener widgets	Drill-down widgets
Virtual Server Health Summary	Virtual Server Health Details	N/A
	Current Connections Trend	N/A
	Connections per Second Trend	N/A
	Inbound Throughput (bps) Trend	N/A
	Outbound Throughput (bps) Trend	N/A
	Total Request Trend	N/A
	CPU Utilization (%) Trend	N/A

Widget name	Chart type	Description
Virtual Server Health Summary	Grid	Displays the list of Virtual Servers that are in the load balancer environment and their associated metrics in the columns: • Status Status of the virtual server is represented as follows: – Green: Available

Widget name	Chart type	Description
		– Orange: Partially Available – Red: Unavailable • Virtual Server • IP Address • Connection Limit • Current Connections • Max Current Connections • Connections per Second • Max Connections per Second • Total Request • Max Total Request • Inbound Throughput (bps) • Max Inbound Throughput (bps) • Outbound Throughput (bps) • Max Outbound Throughput (bps) • CPU Utilization (%) • Max CPU Utilization (%)
Virtual Server Health Details	Badge	The selected Virtual Server health details. The Virtual Server is selected from Virtual Server Health Summary table.
Current Connections Trend	Line	Current or active connections that are available to that Virtual Server at a particular time that is based on the selected time period. It is a measure of the number of client/server requests that can be handled. It also displays the Connection Limit metric. **Note:** The TimeZoom feature is available for all the Line charts in this dashboard.
Connections per Second Trend	Line	Active Virtual Server connections that are available per second.
Inbound Throughput (bps) Trend	Line	Measured in bits per second. Typically, it is the amount of data that is sent to and from the Virtual Server to the client.
Outbound Throughput (bps) Trend	Line	Measured in bits per second. Typically, it is the amount of data that is sent to and from the client to the Virtual Server.

Widget name	Chart type	Description
Total Request Trend	Line	Total number of requests that are handled by the Virtual Server at a particular time period.
CPU Utilization (%) Trend	Line	CPU availability is used across all available cores in the selected Virtual Server.

Related information

Virtual Servers

Pool Details dashboard

A load-balancing pool consists of a set of devices such as web servers that can be logically grouped. These pools can receive and process traffic. Instead of sending client traffic to the destination IP address specified in the client request, Local Traffic Manager sends the request to any of the servers that are members of that pool. It helps to efficiently distribute the load on your server resources.

Available widgets and their interactivity

The Pool Details and its widgets:

Pool Details

1. Click **Home** > **Load Balancer** > **Pool Details**.

   This dashboard displays the F5 BIG-IP Pool information in your load-balancing environment.
2. You can filter data based on the following filters:
   - **Site**, the list contains the locations to view the traffic levels.
   - **LTM**, the list contains the available LTM devices in your load-balancing environment.
   - **Business Hour** , the list contains of **Business Hour**, **Non-Business Hour**, and both as **ALL**.
   - **Top N**, the list contains **10**, **20**, **50**, and **ALL** to display the top N performers.
   - **Time Period**, the list contains the following options:
     - **Last Hour**
     - **Last 6 Hours**
     - **Last 12 Hours**
     - **Last 24 Hours**
     - **Last 7 Days**
     - **Last 30 Days**
     - **Last 365 Days**
     - **Custom**

   Note:

   - You can categorize your enterprise network based on different geographical areas and sites. You can configure the sites by specifying the IP address ranges and the time period for business hours and non-business hours. To configure these properties, see Configuring site grouping.

- Refer to the conditions applied when you're using the **Site** and **Business Hour** filter options.

3. Click **Apply Filter**.

Table 77. Widget interactions

Master widgets	Listener widgets	Drill-down widgets
Pool Health Summary	Pool Health Details	N/A
	Current Connections Trend	N/A
	Connections per Second Trend	N/A
	Inbound Throughput (bps) Trend	N/A
	Outbound Throughput (bps) Trend	N/A

Table 78. Available widgets

Widget name	Chart type	Description
Pool Health Summary	Grid	Displays the Pools that are in the load balancer environment and their associated metrics in the columns: • Status Status of the Pool member is represented as follows: – Green: Available – Orange: Partially Available – Red: Unavailable • Pool • Load Balancing Algorithm • Current Connections • Max Current Connections • Connections per Second • Max Connections per Second • Inbound Throughput (bps) • Max Inbound Throughput (bps) • Outbound Throughput (bps) • Max Outbound Throughput (bps)
Pool Health Details	Badge	The selected pool health details. Pool identifier is selected from the Pool Health Summary table.

Table 78. Available widgets (continued)		
Widget name	**Chart type**	**Description**
Current Connections Trend	Line	Current or active connections that are available to that pool at a particular time that is based on the selected time period. It is a measure of the number of client/ server requests that can be handled. **Note:** The TimeZoom feature is available for all the Line charts in this dashboard.
Connections per Second Trend	Line	Active pool connections that are available per second.
Inbound Throughput (bps) Trend	Line	Measured in bits per second. Typically, it is the amount of data that is sent to and from the pool to the client.
Outbound Throughput (bps) Trend	Line	Measured in bits per second. Typically, it is the amount of data that is sent to and from the client to the pool.

Related information

Pools

Pool Member Details dashboard

A pool consists of pool members. A pool member is a logical object that represents a physical node on the network. A pool is assigned to a virtual server. The BIG-IP system directs traffic that comes into the virtual server to a member of that pool. An individual pool member can belong to one or multiple pools, depending on your network traffic configuration.

Available widgets and their interactivity

The Pool Member Details and its widgets:

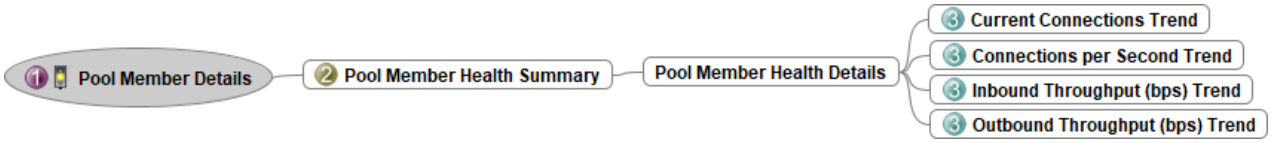

Pool Member Details

1. Click **Home** > **Load Balancer** > **Pool Member Details**.

   This dashboard displays the F5 BIG-IP Pool Member details information in your load-balancing environment. Typically, it shows the pool members that are configured on a specific device. You can track and monitor the number of pool members, and the virtual server and IP address of the device on which the pool members are configured.
2. You can filter data based on the following filters:
   - **Site**, the list contains the locations to view the traffic levels.
   - **LTM**, the list contains the available LTM devices in your load-balancing environment.
   - **Business Hour** , the list contains of **Business Hour**, **Non-Business Hour**, and both as **ALL**.

- **Top N**, the list contains **10**, **20**, **50**, and **ALL** to display the top N performers.
- **Time Period**, the list contains the following options:
  - **Last Hour**
  - **Last 6 Hours**
  - **Last 12 Hours**
  - **Last 24 Hours**
  - **Last 7 Days**
  - **Last 30 Days**
  - **Last 365 Days**
  - **Custom**

Note:

- You can categorize your enterprise network based on different geographical areas and sites. You can configure the sites by specifying the IP address ranges and the time period for business hours and non-business hours. To configure these properties, see Configuring site grouping.
- Refer to the conditions applied when you're using the **Site** and **Business Hour** filter options.

3. Click **Apply Filter**.

Table 79. Widget interactions

Master widgets	Listener widgets	Drill-down widgets
Pool Member Health Summary	Pool Member Health Details	N/A
	Current Connections Trend	N/A
	Connections per Second Trend	N/A
	Inbound Throughput (bps) Trend	N/A
	Outbound Throughput (bps) Trend	N/A

Table 80. Available widgets

Widget name	Chart type	Description
Pool Member Health Summary	Grid	Displays the list of all the Pool members that are in the load balancer environment and their associated metrics in the columns: • Status Status of the Pool member is represented as follows: – Green: Available – Orange: Partially Available – Red: Unavailable • Pool Member • IP addresses of the Pool members • Connection Limit • Current Connections • Max Current Connections • Connections per Second • Max Connections per Second • Inbound Throughput (bps) • Max Inbound Throughput (bps) • Outbound Throughput (bps) • Max Outbound Throughput (bps)
Pool Member Health Details	Badge	The selected pool member health details. IP address of the pool member that is selected from the Pool Member Health Summary table.
Current Connections Trend	Line	Current or active connections that are available to that pool member at a particular time that is based on the selected time period. It is a measure of the number of client/server requests that can be handled. It also displays the Connection Limit metric. **Note:** The TimeZoom feature is available for all the Line charts in this dashboard.
Connections per Second Trend	Line	Active connections that are available per second.

Table 80. Available widgets (continued)		
Widget name	**Chart type**	**Description**
Inbound Throughput (bps) Trend	Line	Measured in bits per second. Typically, it is the amount of data that is sent to and from the pool member to the client.
Outbound Throughput (bps) Trend	Line	Measured in bits per second. Typically, it is the amount of data that is sent to and from the client to the pool member.

NetFlow dashboards

Network flow monitoring is often used to resolve network performance issues and ensure Quality of Service (QoS) for key applications and services.

Network flow monitoring gives visibility to an effective network and infrastructure management. Network Performance Insight Dashboards track the flow of applications and key services over all areas of the network, such as devices, servers, and more and offer insights into network bandwidth utilization.

Network Performance Insight Dashboards populates the performance metrics based on the collected IP network traffic information as the packet enters or exits an interface of a device.

When you select an interface for a device, this resource provides a view of the applications responsible for traffic volume, either inbound or outbound, through the interface over the selected time period. This data provides granular details about network traffic that passes through an interface.

Ensure that you've configured your Network Performance Insight correctly before using the NetFlow dashboards.

Drill-down dashboard views

The diagram shows the flow dashboard groups and the available views.

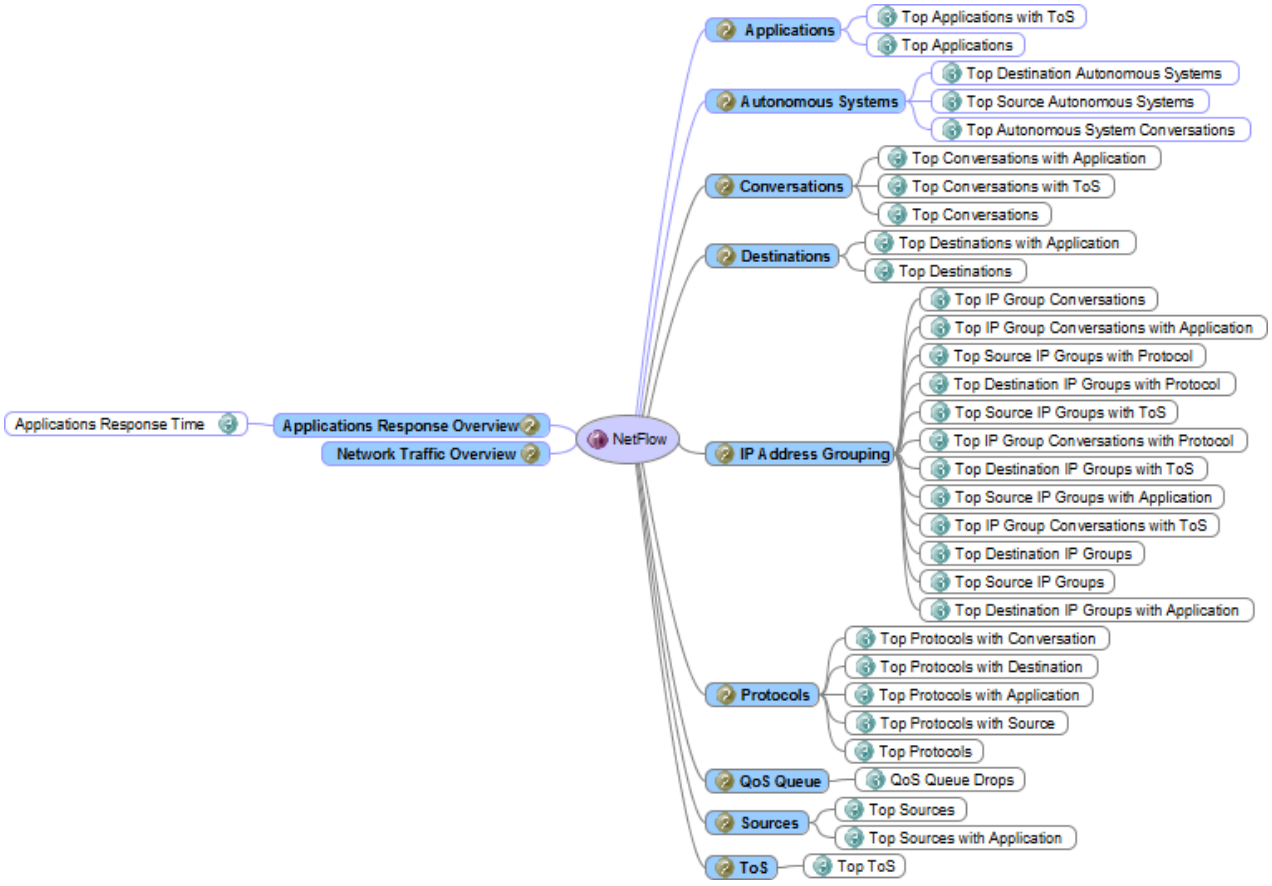

Network Traffic Overview dashboard

Network traffic overview dashboard provides comprehensive bandwidth analysis and performance metrics monitoring capabilities by monitoring the worst performing interfaces. This helps you to understand further the interface utilization trend and also the IP traffic composition that is related to that interface based on the flow data collected.

Use Network Traffic Overview dashboard to monitor network performance details of a particular interface, the traffic trends, and the usage patterns of your network. It identifies the network Top Talkers on applications that use the most network bandwidth.

Available widgets and their interactivity

The master-listener and drill-down widget interactions can be seen in the following ways:

- For any widgets with the ( ) drill-down icon, from the chart, click any bar that represents the interface or metric to open the drill-down dashboard page, which contains the associated widgets.
- For the other widgets that are shown in the diagram to be in a master-listener or drill-down interaction, click any of the bars that represent the interface to change the listener widget that displays the related data for the selected interface.

The diagram shows the master-listener, and drill-down interactions between the available widgets:

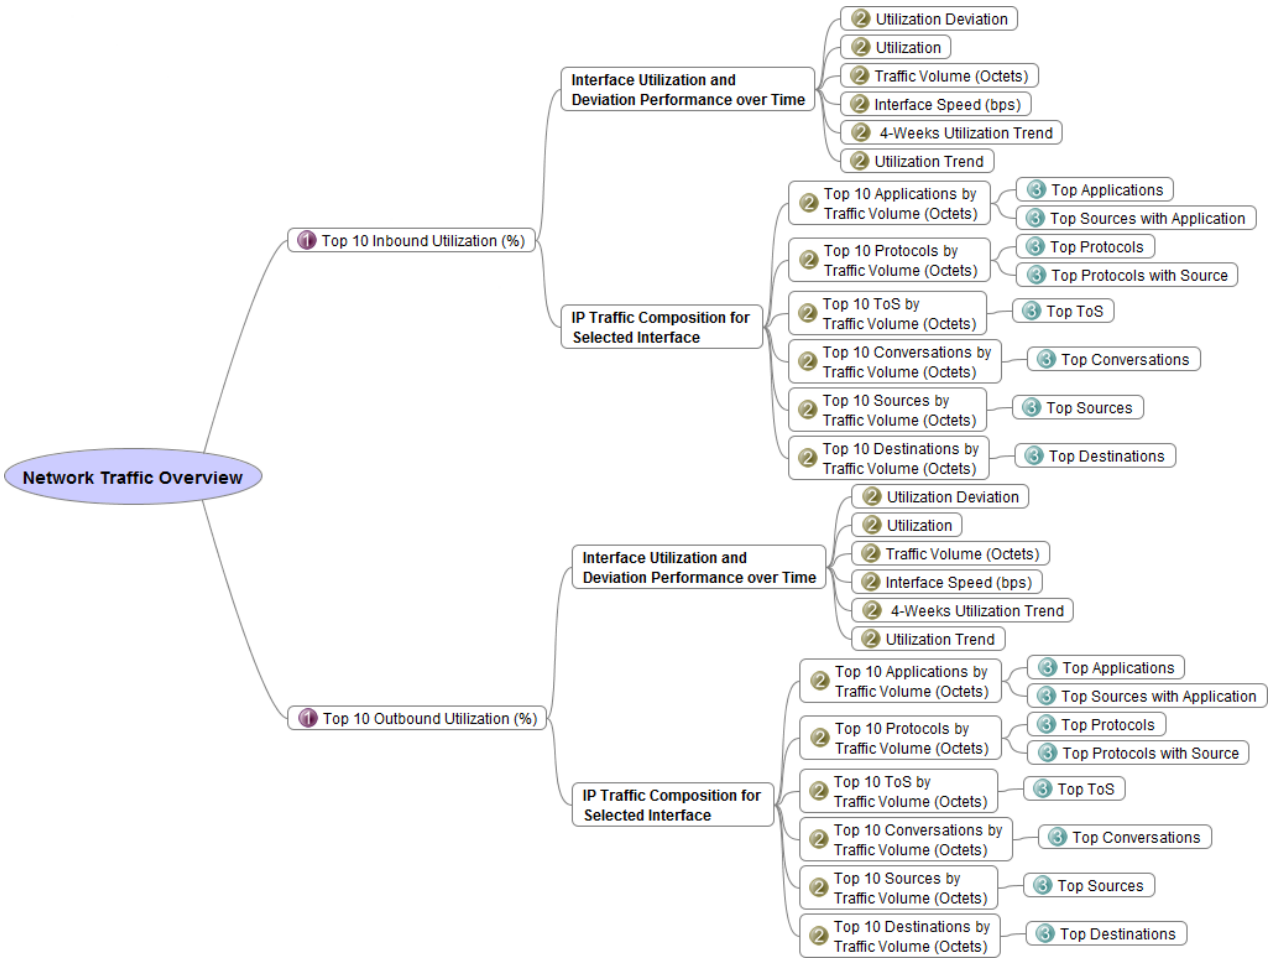

Network Traffic Overview

1. Click **NetFlow** > **Network Traffic Overview**.

   The **Network Traffic Overview: Top 10 Traffic** dashboard loads.

   This dashboard provides a view of the top 10 traffic details that are ranked in order by traffic volume.
2. You can filter data based on the following filters:
   - **Site**, the list contains the locations to view the traffic levels.
   - **Business Hour** , the list contains of **Business Hour**, **Non-Business Hour**, and both as **ALL**.
   - **Time Period**, the list contains of **Last Hour**, and **Last 24 Hours**.

   Note:

   - You can categorize your enterprise network based on different geographical areas and sites. You can configure the sites by specifying the IP address ranges and the time period for business hours and non-business hours. To configure these properties, see Configuring site grouping.
   - Refer to the conditions applied when you're using the **Site** and **Business Hour** filter options.

Table 81. Widget interactions

Master widgets	Listener widgets	Drill-down dashboard
Top 10 Inbound Utilization (%) or Top 10 Outbound Utilization (%) Click any bar that represents the interface inbound or outbound utilization to refresh the listener widgets to display the related data for the selected interface.	The listener widgets from the Interface Utilization and Deviation Performance over Time pane are: • Utilization Deviation • Utilization • Traffic Volume (Octets) • Interface Speed (bps) • 4-Weeks Utilization Trend • Utilization Trend	N/A
	The listener widgets from the IP Traffic Composition for Selected Interface pane are: • Top 10 Applications by Traffic Volume (Octets). Click **Find out more details about top users with application. Click here** to open the Top Sources with Application dashboard. • Top 10 Protocols by Traffic Volume (Octets). Click **Find out more details about top users with protocol. Click here** to open the Top Protocols with Source dashboard. • Top 10 ToS by Traffic Volume (Octets) • Top 10 Conversations by Traffic Volume (Octets) • Top 10 Sources by Traffic Volume (Octets) • Top 10 Destinations by Traffic Volume (Octets) Click the listener widgets title bar to drill down to the NetFlow dashboards for a more detailed information on the flow data.	• Top Applications • Top Sources with Application • Top Protocols • Top Protocols with Source • Top ToS • Top Conversation • Top Sources • Top Destinations For more information, see “NetFlow dashboards” on page 466.

Table 82. Available Widgets		
Widget name	**Chart type**	**Typical uses**
Top 10 Inbound Utilization (%)	Bar	Display the top 10 worst performing inbound or outbound interfaces, which are based on the interface utilization metric that is calculated from the flow data collected. The data is displayed based on the selected time period.
Top 10 Outbound Utilization (%)	Bar	
Utilization Deviation	Badge	Display the interface utilization deviation percentage. The utilization deviation metric is calculated by comparing the interface utilization values of the previous four weeks over the same time period.
Utilization	Badge	Display the interface utilization percentage.
Traffic Volume (Octets)	Badge	Display the traffic volume in octets for the selected interface.
Interface Speed (bps)	Badge	Display the interface speed in bits per second (bps).
4-Weeks Utilization Trend	Bar	Displays the interface utilization trend of the selected interface. **Note:** The deviation is calculated based on the 30-minutes aggregation table.
Utilization Trend	Timeseries	Displays the time-series utilization performance of the selected interface and time period.
Top 10 Applications by Traffic Volume (Octets)	Donut	Displays the top 10 applications, which consumed the most network bandwidth for the selected interface.
Top 10 Protocols by Traffic Volume (Octets)	Donut	Displays the top 10 protocols, which consumed the most network bandwidth for the selected interface.
Top 10 ToS by Traffic Volume (Octets)	Donut	Displays the top 10 Type of Services (ToS) which consumed the most network bandwidth for the selected interface.
Top 10 Conversations by Traffic Volume (Octets)	Donut	Displays the top 10 conversation (the source - destination IP addresses) which consumed the most network bandwidth for the selected interface.
Top 10 Sources by Traffic Volume (Octets)	Donut	Displays the top 10 sources IP, which consumed the most network bandwidth for the selected interface.

Table 82. Available Widgets (continued)		
Widget name	**Chart type**	**Typical uses**
Top 10 Destinations by Traffic Volume (Octets)	Donut	Displays the top 10 destinations IP, which consumed the most network bandwidth for the selected interface.

Network Traffic Overview dashboard for Interim Fix2

The enhanced version of Network Traffic Overview consists of visually advanced Sankey diagrams which display both Traffic flow and volume. Sankey diagrams are a specific type of flow diagram, in which the width of the arrows is shown proportionally to the flow quantity.

The six widgets of IP Traffic composition for selected interface are replaced with two new widgets which are represented by Sankey diagrams.
The new widgets are as follows:

- Top 10 Sources with Applications and Destinations
- Top 10 Destinations with Applications and Sources

These two widgets show the traffic flow and traffic volume between Source and Destination servers and applications.

Available Widgets and their interactions

Table 83. Widget interactions		
Master widgets	**Listener widgets**	**Drill-down dashboard**
Top 10 Inbound Utilization (%) or Top 10 Outbound Utilization (%) Click any bar that represents the interface inbound or outbound utilization to refresh the listener widgets to display the related data for the selected interface.	The listener widgets from the Interface Utilization and Deviation Performance over Time pane are: • Utilization Deviation • Utilization • Traffic Volume (Octets) • Interface Speed (bps) • 4-Weeks Utilization Trend • Utilization Trend	N/A

Table 84. Available Widgets		
Widget name	**Chart type**	**Typical uses**
Top 10 Inbound Utilization (%)	Bar	Display the top 10 worst performing inbound or outbound interfaces, which are based on the interface utilization metric that is calculated from the flow data collected. The data is displayed based on the selected time period.
Top 10 Outbound Utilization (%)	Bar	
Utilization Deviation	Badge	Display the interface utilization deviation percentage. The utilization deviation metric is calculated by comparing the interface utilization values of the previous four weeks over the same time period.

Table 84. Available Widgets (continued)		
Widget name	**Chart type**	**Typical uses**
Utilization	Badge	Display the interface utilization percentage.
Traffic Volume (Octets)	Badge	Display the traffic volume in octets for the selected interface.
Interface Speed (bps)	Badge	Display the interface speed in bits per second (bps).
4-Weeks Utilization Trend	Bar	Displays the interface utilization trend of the selected interface. **Note:** The deviation is calculated based on the 30-minutes aggregation table.
Utilization Trend	Timeseries	Displays the time-series utilization performance of the selected interface and time period.
Top 10 Sources with Applications and Destination	Sankey	Displays the top 10 sources IP, which consumed the most network bandwidth when interacting with Applications. The bandwidth consumed by Applications when interacting with Destinations is also displayed. The thicker the width of the arrows of Sankey diagram the more is the bandwidth.
Top 10 Destinations with Applications and Sources	Sankey	Displays the top 10 destinations IP, which consumed the most network bandwidth when interacting with Applications. The bandwidth consumed by Applications when interacting with Sources is also displayed. The thicker the width of the arrows of Sankey diagram the more is the bandwidth.

Applications Response Overview dashboard

Applications Response Overview dashboard displays the overview of the worst performing application-target server pairings based on the Applications Response Time metrics that are monitored.

Available widgets and their interactivity

The master-listener and drill-down widget interactions can be seen in the following ways:

- For any widgets with the ( ) drill-down icon, from the chart, click any bar that represents the interface or metric to open the drill-down dashboard page, which contains the associated widgets.
- For the other widgets that are shown in the diagram to be in a master-listener or drill-down interaction, click any of the bars that represent the interface to change the listener widget that displays the related data for the selected interface.

The diagram shows the master-listener, and drill-down interactions between the available widgets:

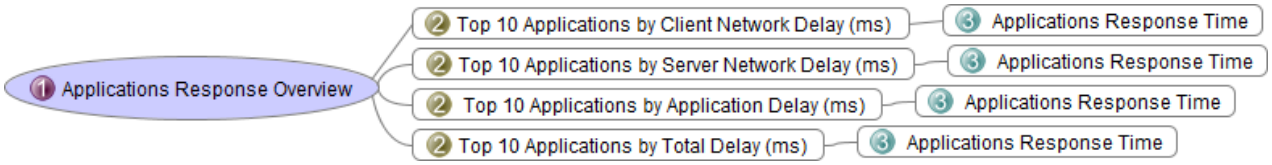

Applications Response Overview

1. Click **NetFlow** > **Applications Response Overview**.

   The **Applications Response Overview: Top 10** dashboard loads.

   This dashboard provides a view of the top 10 worst performing application-target server pairings based on the Applications Response Time metrics that are monitored.

2. From the filter options, choose the **Device**, **Interface**, and **Time period** and click **Apply Filter**.

   Note: You can select the filter values or time ranges that you want to display in the dashboard to which the filter is assigned.

   The dashboard refreshes the data according to the filter attribute values.

Table 85. Widget interactions

Master widgets	Listener widgets	Drill-down dashboard
Top 10 Applications by Client Network Delay (ms) Click any bar that represents the application to drill down to Applications Response Time page.	N/A	Applications Response Time You can drill down to the Applications Response Time page from these widgets to view applications response time issues in a network.
Top 10 Applications by Server Network Delay (ms) Click any bar that represents the application to drill down to Applications Response Time page.		
Top 10 Applications by Application Delay (ms) Click any bar that represents the application to drill down to Applications Response Time page.		
Top 10 Applications by Total Delay (ms) Click any bar that represents the application to drill down to Applications Response Time page.		

Table 86. Available Widgets

Widget name	Chart type	Typical uses
Top 10 Applications by Client Network Delay (ms)	Bar	Display the top 10 applications delay time of a client based on the Applications Response Time metrics that are monitored. The application client delay is shown in milliseconds based on the selected filter attribute values.

Table 86. Available Widgets (continued)		
Widget name	**Chart type**	**Typical uses**
Top 10 Applications by Server Network Delay (ms)	Bar	Display the top 10 applications delay time of a server based on the Applications Response Time metrics that are monitored. The application server delay is shown in milliseconds based on the selected filter attribute values.
Top 10 Applications by Application Delay (ms)	Bar	Display the top 10 applications delay time of an application based on the Applications Response Time metrics that are monitored. The application response delay time is shown in milliseconds based on the selected filter attribute values.
Top 10 Applications by Total Delay (ms)	Bar	Display the top 10 applications total delay time based on the Applications Response Time metrics that are monitored. The total application delay time is shown in milliseconds based on the selected filter attribute values.

Applications dashboards
This topic gives you an overview of the Applications dashboards usage.

The Applications dashboards is divided into two parts. To access the Applications dashboards, click **NetFlow** > **Applications** and select any one of the dashboards:

- **Top Applications**
- **Top Applications with ToS**

Top Applications

This dashboard provides a view of the top 10 applications responsible for monitored traffic on your network, ranked in order of traffic volume for an interface.

Table 87. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Applications	Donut	Displays the top 10 applications name with its total traffic volume in octets. The percentage of the application is based on the selected application that is shown by the widget. The individual application in the legend adds up to 100%. This percentage can be absolute or relative.

Table 87. Available Widgets (continued)		
Widget Name	**Chart Type**	**Description**
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the application uses the interface bandwidth in octets, bps, or percentage.
Top 10 Applications	Grid	Displays the Top 10 traffic data (in octets and packets) of device applications through the selected interface. The columns that are displayed depend on the flow direction set, either Inbound or Outbound for the selected time period.

Top Applications with ToS

This dashboard provides a top 10 view of network traffic segmented by Type of Service (ToS) method for an interface.

Note: Dashboard filters allow you to choose different views of data to be displayed on the active dashboard tab. You can select the filter values or time ranges that you want to display in dashboards to which the filter is assigned.

Table 88. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Applications with ToS	Donut	Displays the top 10 applications name by Type of Service (ToS) method for an interface. The percentage of the Applications with ToS is based on the selected application with ToS that is shown by the widget. The individual application in the legend adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time. You can choose to display the Utilization trend in percentage or the Throughput Trend in bit per second (bps). It describes how the application uses the interface bandwidth in octets, bps, or percentage.

Table 88. Available Widgets (continued)		
Widget Name	**Chart Type**	**Description**
Top 10 Applications with ToS	Grid	Displays the Top 10 traffic data (in octets and packets) of device applications with ToS through the selected interface. The columns that are displayed depend on the flow direction set, either Inbound or Outbound for the selected time period.

Conversations dashboards

This topic gives you an overview of the Conversations dashboards usage.

To access the Conversations dashboards, click **NetFlow** > **Conversations** and select any one of the dashboards:

- **Top Conversations**
- **Top Conversations with ToS**
- **Top Conversations with Application**

Top Conversations

This dashboard provides a view of the top 10 most bandwidth consuming conversations, which are conducted over your monitored network. Conversations are listed with the amount of data that is transferred in the conversation, in both octets and packets.

Table 89. Available Widgets		
Widget Name	**Chart Type**	**Typical uses**
Top 10 Conversations	Donut	Displays the top 10 conversations name with its total traffic volume in octets. The percentage of the conversations is based on the selected conversation that is shown by the widget. The individual conversation in the legend adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput Trend in bit per second (bps). It describes how the conversation traffic uses the interface bandwidth in octets, bps, or percentage.
Top 10 Conversations	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected conversations through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top Conversations with ToS

This dashboard provides a view of the top 10 most bandwidth consuming conversations with ToS conducted over your monitored network. Conversations with ToS are listed with the amount of data that is transferred in the conversation, in both octets and packets.

Table 90. Available Widgets

Widget Name	Chart Type	Description
Top 10 Conversations with ToS	Donut	Displays the top 10 conversations with ToS name with its total traffic volume in octets. The percentage of the conversations with ToS is based on the selection on the widget. The individual conversation in the legend adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the conversation traffic uses the interface bandwidth in octets, bps, or percentage.
Top 10 Conversations with ToS	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected conversations with ToS through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top Conversations with Application

This dashboard provides a view of the top 10 most bandwidth consuming conversations with applications that are conducted over your monitored network. Conversations with application are listed with the amount of data that is transferred in the conversation, in both octets and packets.

Note:

You can select the filter values or time ranges that you want to display in dashboards to which the filter is assigned.

Table 91. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Conversations with Application	Donut	Displays the top 10 conversations with application name with its total traffic volume in octets. The percentage of the conversations with application is based on the selection on the widget. The individual conversation with application in the legend adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the conversation traffic uses the interface bandwidth in octets, bps, or percentage.
Top 10 Conversations with Application	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected conversations with application through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Sources dashboards

This topic gives you an overview of the Sources dashboards usage.

To access the Sources dashboards, click **NetFlow** > **Sources** and select any one of the dashboards:

- **Top Sources**
- **Top Sources with Application**

Top Sources

This dashboard provides a view of the top 10 sources of traffic used most for traffic on your monitored network.

Note:

You can select the filter values or time ranges that you want to display in dashboards to which the filter is assigned.

Table 92. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Sources	Donut	Displays the top 10 sources name with its total traffic volume in octets. The percentage of the sources is based on the selected sources that are shown by the widget. The individual source adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the sources traffic uses the interface bandwidth in octets, bps, or percentage.
Top 10 Sources	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected sources through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top Sources with Application

This dashboard provides a view of the top 10 sources of traffic with application used most for traffic on your monitored network.

Table 93. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Sources with Application	Donut	Displays the top 10 sources with application name with its total traffic volume in octets. The percentage of the sources with applications is based on the selection on the widget. The individual sources with application adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the sources with applications traffic utilize the interface bandwidth in octets, bps, or percentage.

Table 93. Available Widgets (continued)		
Widget Name	**Chart Type**	**Description**
Top 10 Sources with Applications	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected sources with applications through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Protocols dashboards

This topic gives you an overview of the Protocols dashboards usage.

To access the Protocols dashboards, click **NetFlow** > **Protocols** and select any one of the dashboards:

- **Top Protocols**
- **Top Protocols with Application**
- **Top Protocols with Source**
- **Top Protocols with Destination**
- **Top Protocols with Conversation**

Top Protocols

This dashboard provides a view of the top 10 of the protocols used most for traffic on your monitored network.

Table 94. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Protocols	Donut	Displays the top 10 protocols name with its total traffic volume in octets. The percentage of the protocols is based on the selected destinations that are shown by the widget. The individual protocol adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the protocols traffic uses the interface bandwidth in octets, bps, or percentage.

Table 94. Available Widgets (continued)		
Widget Name	**Chart Type**	**Description**
Top 10 Protocols	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected protocols through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top Protocols with Application

This dashboard provides a view of the top 10 of the protocols with application used most for traffic on your monitored network.

Table 95. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Protocols with Application	Donut	Displays the top 10 protocols with application name and its total traffic volume in octets. The percentage of the protocols with application is based on the selection on the widget. The individual protocol with application adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the protocols with application traffic utilize the interface bandwidth in octets, bps, or percentage.
Top 10 Protocols with Application	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected protocols with application through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top Protocols with Source

This dashboard provides a view of the top 10 of the protocols with source used most for traffic on your monitored network.

Table 96. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Protocols with Source	Donut	Displays the top 10 protocols with source name and its total traffic volume in octets. The percentage of the protocols with source is based on the selection on the widget. The individual protocol with source adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the protocols with source traffic utilize the interface bandwidth in octets, bps, or percentage.
Top 10 Protocols with Source	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected protocols with source through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top Protocols with Destination

This dashboard provides a view of the top 10 of the protocols with destination used most for traffic on your monitored network.

Table 97. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Protocols with Destination	Donut	Displays the top 10 protocols with destination name and its total traffic volume in octets. The percentage of the protocols with destinations is based on the selection on the widget. The individual protocol with destination adds up to 100%. This percentage can be absolute or relative.

Table 97. Available Widgets (continued)		
Widget Name	**Chart Type**	**Description**
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the protocols with destination traffic utilize the interface bandwidth in octets, bps, or percentage.
Top 10 Protocols with Destination	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected protocols with destination through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top Protocols with Conversation

This dashboard provides a view of the top 10 of the protocols with conversation used most for traffic on your monitored network.

Note:

You can select the filter values or time ranges that you want to display in dashboards to which the filter is assigned.

Table 98. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Protocols with Conversation	Donut	Displays the top 10 protocols with conversation name and its total traffic volume in octets. The percentage of the protocols with conversation is based on the selection on the widget. The individual protocol with conversation adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the protocols with conversations traffic utilize the interface bandwidth in octets, bps, or percentage.

Table 98. Available Widgets (continued)		
Widget Name	**Chart Type**	**Description**
Top 10 Protocols with Conversation	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected protocols with conversation through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Destinations dashboards

This topic gives you an overview of the Destinations dashboards usage.

To access the Destinations dashboards, click **NetFlow** > **Destinations** and select any one of the dashboards:

- **Top Destinations**
- **Top Destinations with Application**

Top Destinations

This dashboard provides a view of the top 10 domains that serve as destinations of traffic on the network, ranked by percentage of the total traffic over the specified time period.

Table 99. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Destinations	Donut	Displays the top 10 destinations name with its total traffic volume in octets. The percentage of the destinations is based on the selected destinations that are shown by the widget. The individual destination adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the destinations traffic uses the interface bandwidth in octets, bps, or percentage.
Top 10 Destinations	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected destinations through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top Destinations with Application

This dashboard provides a view of the top 10 destinations with application responsible for monitored traffic on your network, ranked in order of traffic volume.

Note:

You can select the filter values or time ranges that you want to display in dashboards to which the filter is assigned.

Table 100. Available Widgets

Widget Name	Chart Type	Description
Top 10 Destinations with Application	Donut	Displays the top 10 destinations with application name with its total traffic volume in octets. The percentage of the destinations with application is based on the selection on the widget. The individual destination with application adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the destinations traffic uses the interface bandwidth in octets, bps, or percentage.
Top 10 Destinations with Application	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected destination with application through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

ToS dashboards

This topic gives you an overview of the ToS dashboard usage.

Top ToS

1. Click **NetFlow** > **ToS** > **Top ToS**.

   The **Top ToS: Traffic Volume Details for Interface** dashboard loads.

   This dashboard provides a view of the top 10 most bandwidth consuming Type of Service (ToS) for an interface.

   Note:

   You can select the filter values or time ranges that you want to display in dashboards to which the filter is assigned.

Table 101. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 ToS	Donut	Displays the top 10 ToS name with its total traffic volume in octets. The percentage of the ToS is based on the selected ToS shown by the widget. The individual source adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the ToS traffic uses the interface bandwidth in octets, bps, or percentage.
Top 10 ToS	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected ToS through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Autonomous Systems dashboards

This topic gives you an overview of the Autonomous Systems dashboards usage.

To access the Autonomous Systems dashboards, click **NetFlow** > **Autonomous Systems dashboards** and select any one of the dashboards:

- **Top Autonomous System Conversations**
- **Top Source Autonomous Systems**
- **Top Source Autonomous Systems**

Top Autonomous System Conversations

This dashboard provides a view of the top 10 list of Autonomous System conversations with highest bandwidth consumption. Autonomous System Conversations are listed with the amount of data that is transferred, in both octets and packets, and the percentage of traffic utilization generated by the autonomous system over the specified time period.

Table 102. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Autonomous System Conversations	Donut	Displays the top 10 Autonomous System conversations name with its total traffic volume in octets. The percentage of the conversation is based on the selected conversation that is shown by the widget. The individual conversation in the legend adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput Trend in bit per second (bps). It describes how the conversation of Autonomous Systems traffic uses the interface bandwidth in octets, bps, or percentage.
Top 10 Autonomous System Conversations	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected autonomous systems conversation through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top Source Autonomous Systems

This dashboard provides a view of the top 10 list of source Autonomous Systems with highest bandwidth consumption. Source Autonomous Systems are listed with the amount of data that is transferred, in both octets and packets, and the percentage of traffic utilization generated by the autonomous system over the specified time period.

Table 103. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Source Autonomous Systems	Donut	Displays the top 10 source Autonomous Systems name with its total traffic volume in octets. The percentage of the source is based on the selected source that is shown by the widget. The individual application in the legend adds up to 100%. This percentage can be absolute or relative.

Table 103. Available Widgets (continued)		
Widget Name	**Chart Type**	**Description**
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time. You can choose to display the Utilization trend in percentage or the Throughput Trend in bit per second (bps). It describes how the conversation of Autonomous Systems traffic uses the interface bandwidth in octets, bps, or percentage.
Top 10 Source Autonomous Systems	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected autonomous systems sources through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top Destination Autonomous Systems

This dashboard provides a view of the top 10 list of destination Autonomous Systems with highest bandwidth consumption. Destination Autonomous Systems are listed with the amount of data that is transferred, in both octets and packets, and the percentage of traffic utilization generated by the autonomous system over the specified time period.

Note:

You can select the filter values or time ranges that you want to display in dashboards to which the filter is assigned.

Table 104. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Destination Autonomous Systems	Donut	Displays the top 10 destination Autonomous Systems name with its total traffic volume in octets. The percentage of the destination is based on the selected destination that is shown by the widget. The individual destination in the legend adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the destination of Autonomous Systems traffic uses the interface bandwidth in octets, bps, or percentage.

Table 104. Available Widgets (continued)		
Widget Name	**Chart Type**	**Description**
Top 10 Destination Autonomous Systems	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected autonomous systems destination through the selected interface of a device. Displays the Top 10 traffic data (in octets and packets) of device destination of Autonomous Systems through the selected interface. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

QoS Queue dashboards
This topic gives you an overview of the QoS Queue dashboard usage.

QoS Queue Drops

1. Click **NetFlow** > **QoS Queue** > **QoS Queue Drops**.

   The **QoS Queue Drops: Outbound Traffic Details for Interface** dashboard loads.

   This dashboard provides a view of the top 10 most QoS Queue Drops for traffic on your monitored network for an outbound interface of a device.

   It provides the visibility of how the defined traffic classes are performing in terms of packet drops.

   Note:

   You can select the filter values or time ranges that you want to display in dashboards to which the filter is assigned.

Table 105. Available Widgets		
Widget Name	**Chart Type**	**Description**
QoS Queue Drops	Donut	Displays the top 10 QoS Queue Drops ID with its total drops in packets. The percentage of the QoS Queue Drops is based on the selection on the widget. The individual QoS Queue Drop adds up to 100%. This percentage can be absolute or relative.
Packet Drop Trend	Timeseries	This timeseries widget displays the Packet Drop trend in packets across a selected time period.
QoS Queue Drops	Grid	Displays the Top 10 QoS Queue Drops data in packets that flows in the selected queue for an outbound interface of a device.

IP Address Grouping dashboards
This topic gives you an overview of the IP Address Grouping dashboards usage.

To access the IP Address Grouping dashboards, click **NetFlow** > **IP Address Grouping** and select any one of the dashboards:

- **Top Source IP Groups with Application**
- **Top Source IP Groups with ToS**

- **Top Source IP Groups with Protocol**
- **Top Destination IP groups with Application**
- **Top Destination IP Groups with ToS**
- **Top Destination IP Groups with Protocol**
- **Top IP Group Conversations with Application**
- **Top IP Group Conversations with ToS**
- **Top IP Group Conversations with Protocol**
- **Top Source IP Groups**
- **Top Destination IP Groups**
- **Top IP Group Conversations**

Top Source IP Groups with Application

This dashboard provides a view of the top 10 source hosts that contribute to traffic on the network with applications that are associated with IP address group, which is responsible for the most traffic on your network.

Table 106. Available Widgets

Widget Name	Chart Type	Description
Top 10 Source IP Groups with Application	Donut	Displays the top 10 source IP Groups with application name and its total traffic volume in octets. The percentage of the source IP Groups with application is based on the selection on the widget. The individual source IP Groups with application adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the source IP Groups with applications traffic utilize the interface bandwidth in octets, bps, or percentage.
Top 10 Source IP Groups with Application	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected source IP Groups with application through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top Source IP Groups with ToS

This dashboard provides a view of the top 10 source hosts that contribute to traffic on the network with ToS associated with IP address groups, which is responsible for the most traffic on your network.

Table 107. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Source IP Groups with ToS	Donut	Displays the top 10 source IP Groups with ToS name and its total traffic volume in octets. The percentage of the source IP Groups with ToS is based on the selection on the widget. The individual source IP Groups with ToS adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the source IP Groups with ToS traffic utilize the interface bandwidth in octets, bps, or percentage.
Top 10 Source IP Groups with ToS	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected source IP Groups with ToS through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top Source IP Groups with Protocol

This dashboard provides a view of the top 10 source hosts that contribute to traffic on the network with protocols that are associated with IP address group, which is responsible for the most traffic on your network.

Table 108. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Source IP Groups with Protocol	Donut	Displays the top 10 source IP Groups with protocol name and its total traffic volume in octets. The percentage of the source IP Groups with protocol is based on the selection on the widget. The individual source IP Groups with protocol adds up to 100%. This percentage can be absolute or relative.

Table 108. Available Widgets (continued)		
Widget Name	**Chart Type**	**Description**
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the source IP Groups with protocols traffic utilize the interface bandwidth in octets, bps, or percentage.
Top 10 Source IP Groups with Protocol	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected source IP Groups with protocol through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top Destination IP Groups with Application

This dashboard provides a view of the top 10 domains that serve as destinations of traffic on the network with applications that are associated with an IP address group, which is responsible for the most traffic on your network.

Table 109. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Destination IP Groups with Application	Donut	Displays the top 10 destination IP Groups with application name and its total traffic volume in octets. The percentage of the destination IP Groups with application based on the selection on the widget. The individual destination IP Groups with application adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the destination IP Groups with applications traffic utilize the interface bandwidth in octets, bps, or percentage.

Table 109. Available Widgets (continued)		
Widget Name	**Chart Type**	**Description**
Top 10 Destination IP Groups with Application	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected destination IP Groups with application through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top Destination IP Groups with ToS

This dashboard provides a view of the top 10 domains that serve as destinations of traffic on the network with ToS that are associated with an IP address group, which is responsible for the most traffic on your network.

Table 110. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Destination IP Groups with ToS	Donut	Displays the top 10 destination IP Groups with ToS name and its total traffic volume in octets. The percentage of the destination IP Groups with ToS is based on the selection on the widget. The individual destination IP Groups with ToS adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the destination IP Groups with ToS traffic utilize the interface bandwidth in octets, bps, or percentage.
Top 10 Destination IP Groups with ToS	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected destination IP Groups with ToS through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top Destination IP Groups with Protocol

This dashboard provides a view of the top 10 domains that serve as destinations of traffic on the network with protocol that are associated with an IP address group, which is responsible for the most traffic on your network.

Table 111. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Destination IP Groups with Protocol	Donut	Displays the top 10 destination IP Groups with protocol name and its total traffic volume in octets. The percentage of the destination IP Groups with protocols based on the selection on the widget. The individual destination IP Groups with protocol adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the destination IP Groups with protocols traffic utilize the interface bandwidth in octets, bps, or percentage.
Top 10 Destination IP Groups with Protocol	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected destination IP Groups with protocol through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top IP Group Conversations with Application

This dashboard provides a view of the top 10 bandwidth consuming conversations with application that is associated with an IP address group, which is responsible for the most traffic on your network.

Table 112. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 IP Group Conversations with Application	Donut	Displays the top 10 IP Group conversations with application name and its total traffic volume in octets. The percentage of the IP Group conversations with application is based on the selection on the widget. The individual IP Group conversations with application adds up to 100%. This percentage can be absolute or relative.

Table 112. Available Widgets (continued)		
Widget Name	**Chart Type**	**Description**
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the IP Group conversations with application traffic utilize the interface bandwidth in octets, bps, or percentage.
Top 10 IP Group Conversations with Application	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected IP Group conversations with application through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top IP Group Conversations with ToS

This dashboard provides a view of the top 10 bandwidth consuming conversations with ToS associated with an IP address group, which is responsible for the most traffic on your network.

Table 113. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Conversation IP Groups with ToS	Donut	Displays the top 10 conversation IP Groups with ToS name and its total traffic volume in octets. The percentage of the conversation IP Groups with ToS based on the selection on the widget. The individual conversation IP Groups with ToS adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the conversation IP Groups with ToS traffic utilize the interface bandwidth in octets, bps, or percentage.

Table 113. Available Widgets (continued)		
Widget Name	**Chart Type**	**Description**
Top 10 Conversation IP Groups with ToS	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected conversation IP Groups with ToS through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top IP Group Conversations with Protocol

This dashboard provides a view of the top 10 bandwidth consuming conversations with protocols that are associated with an IP address group, which is responsible for the most traffic on your network.

Table 114. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Conversation IP Groups with Protocol	Donut	Displays the top 10 conversation IP Groups with protocol name and its total traffic volume in octets. The percentage of the conversation IP Groups with protocol based on the selection on the widget. The individual conversation IP Groups with protocol adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the conversation IP Groups with protocols traffic utilize the interface bandwidth in octets, bps, or percentage.
Top 10 Conversation IP Groups with Protocol	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected conversation IP Groups with protocol through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top Source IP Groups

This dashboard provides a view of the top 10 source hosts that contribute to traffic on the network that are associated with IP address group, which is responsible for the most traffic on your network.

Table 115. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Source IP Groups	Donut	Displays the top 10 source IP Groups and its total traffic volume in octets. The percentage of the source IP Groups based on the selection on the widget. The individual source IP Groups adds up to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the source IP Groups with applications traffic utilize the interface bandwidth in octets, bps, or percentage.
Top 10 Source IP Groups	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected source IP Groups through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top Destination IP Groups

This dashboard provides a view of the top 10 domains that serve as destinations of traffic on the network that are associated with an IP address group, which is responsible for the most traffic on your network.

Table 116. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 Destination IP Groups	Donut	Displays the top 10 destination IP Groups and its total traffic volume in octets. The percentage of the destination IP Groups based on the selection on the widget. The individual destination IP Groups adds up to 100%. This percentage can be absolute or relative.

Table 116. Available Widgets (continued)		
Widget Name	**Chart Type**	**Description**
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the destination IP Groups with applications traffic utilize the interface bandwidth in octets, bps, or percentage.
Top 10 Destination IP Groups	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected destination IP Groups through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Top IP Group Conversations

This dashboard provides a view of the top 10 bandwidth consuming conversations that is associated with an IP address group, which is responsible for the most traffic on your network.

Note:

You can select the filter values or time ranges that you want to display in dashboards to which the filter is assigned.

Table 117. Available Widgets		
Widget Name	**Chart Type**	**Description**
Top 10 IP Group Conversations	Donut	Displays the top 10 IP Group conversations and its total traffic volume in octets. The percentage of the IP Group conversations is based on the selection on the widget. The individual IP Group conversations add ups to 100%. This percentage can be absolute or relative.
Traffic Volume Trend	Timeseries	This timeseries widget displays the Traffic Volume trend in octets across a selected time period. You can choose to display the Utilization trend in percentage or the Throughput trend in bit per second (bps). It describes how the IP Group conversations traffic utilize the interface bandwidth in octets, bps, or percentage.

Table 117. Available Widgets (continued)		
Widget Name	**Chart Type**	**Description**
Top 10 IP Group Conversations	Grid	Displays the Top 10 traffic data (in octets and packets) that flows in the selected conversation IP Groups through the selected interface of a device. The columns display depend on the flow direction, either Inbound or Outbound for the selected time period.

Applications Response Time dashboard
This topic gives you an overview of the Applications Response Time dashboard usage.

Access Flow Applications Response Time view

The Applications Response Time page allows you to view network response time issues for a particular application or category.

You can drill down to the Applications Response Time page from the following dashboards:

1. Click **Home** > **Network Performance Overview**.

   Network Performance Overview: Top 10 page loads.

   - From the Congestion pane, click any bar for an interface from the **Top 10 Applications by Total Delay (ms)** bar widget.
2. Click **Home** > **Network Performance Overview by Deviation**.

   Network Performance Overview by Deviation: Top 10 Deviation page loads.

   - From the Congestion pane, click any bar for an interface from the **Top 10 Applications by Total Delay (ms)** bar widget.
3. Click **NetFlow** > **Applications Response Overview**.

   Applications Response Overview: Top 10 dashboard loads.

   You can drill down to the **Applications Response Time: Response Time for Interface** page from these widgets by clicking any of the bars that represent the application to view the network response time issues:

   - **Top 10 Applications by Client Network Delay (ms)**
   - **Top 10 Applications by Server Network Delay (ms)**
   - **Top 10 Applications by Application Delay (ms)**
   - **Top 10 Applications by Total Delay (ms)**

Applications Response Time widgets

You have the options to filter the existing Applications Response Time metrics using the following filters.

1. From the filter options, choose the **Device**, **Interface**, **Application**, **Target**, and **Time period** and click **Apply Filter**.

   Note: You can select the filter values or time ranges that you want to display in the dashboard to which the filter is assigned.

   The dashboard refreshes the data according to the filter attribute values.

Table 118. Applications Response Time Widgets		
Widget Name	**Chart Type**	**Description**
Applications Response Time	Badge	It displays the following Applications Response Time metrics in milliseconds. • Max Client Network Delay (ms) • Max Server Network Delay (ms) • Max Application Delay (ms) • Max Total Delay (ms)
Response Time Trend	Timeseries	It shows the Applications Response Time metrics that display the application response time delay that is contributed from the client, server and application server side. The following are the application response time metrics that is shown in milliseconds (ms): • Max Client Network delay • Max Server Network delay • Max Application delay • Max Total delay
Response Time in milliseconds (ms)	Grid	

Navigate to different NetFlow dashboards

This section describes the navigation from an existing NetFlow dashboard to another NetFlow Top 10 dashboard view.

You can choose to navigate to a different NetFlow dashboard, by clicking the aggregation type to view options available.

For example,

1. Click **NetFlow** > **Applications** > **Top Applications**

   The **Top Applications: Traffic Volume Details for Interface** dashboard loads.
2. From the **Select aggregation type to view** pane, you can click any of the options to view the NetFlow dashboard.

   The selected NetFlow Top 10 dashboard loads in a new tab.

Aggregation type to view

The views are categorized into the following types:

- Aggregation
- IP Grouping aggregation

The table explains which NetFlow dashboard is loaded based on your aggregation type selection.

Table 119. Aggregation type to view	
Aggregation	**NetFlow Dashboards**
>> Applications	Top Applications
>> Applications with ToS	Top Applications with ToS
>> Sources	Top Sources

Table 119. Aggregation type to view (continued)	
Aggregation	**NetFlow Dashboards**
>> Sources with Application	Top Sources with Application
>> Destinations	Top Destinations
>> Destinations with Application	Top Destinations with Application
>> Conversations	Top Conversations
>> Conversations with Application	Top Conversations with Application
>> Conversations with ToS	Top Conversations with ToS
>> Protocols	Top Protocols
>> Protocols with Application	Top Protocols with Application
>> Protocols with Conversation	Top Protocols with Conversation
>> Protocols with Destination	Top Protocols with Destination
>> Protocols with Source	Top Protocols with Source
>> AS Conversations	Top Autonomous System Conversations
>> Destination AS	Top Destination Autonomous Systems
>> Source AS	Top Source Autonomous Systems
>> ToS	Top ToS

Table 120. IP Grouping aggregation type to view	
IP Grouping aggregation	**NetFlow Dashboards**
>> Source IP Groups	Top Source IP Groups
>> Source IP Groups with Application	Top Source IP Groups with Application
>> Source IP Groups with ToS	Top Source IP Groups with ToS
>> Source IP Groups with Protocol	Top Source IP Groups with Protocol
>> Destination IP Groups	Top Destination IP Groups
>> Destination IP Groups with Application	Top Destination IP Groups with Application
>> Destination IP Groups with ToS	Top Destination IP Groups with ToS
>> Destination IP Groups with Protocol	Top Destination IP Groups with Protocol
>> IP Group Conversations	Top IP Group Conversations

Table 120. IP Grouping aggregation type to view (continued)	
IP Grouping aggregation	**NetFlow Dashboards**
>> IP Groups Conversations with Application	Top IP Group Conversations with Application
>> IP Groups Conversations with ToS	Top IP Group Conversations with ToS
>> IP Groups Conversations with Protocol	Top IP Group Conversations with Protocol

On Demand Filtering dashboards

You can view the network issues for troubleshooting with On Demand Filtering dashboard, which displays information about your network.

On Demand Filtering helps you to investigate in detail the performance of a specific interface over a period for a set of KPIs. You can identify the latency or jitter in your network for further troubleshooting.

On Demand Filtering consists of Device Health, Flow, HTTP Operations, IPSLA, and Timeseries Data dashboards. It provides real-time and historical KPI that actively monitors the network status and network quality experience. Probes are one of the most effective ways to gain insights to facilitate root-cause analysis across network interfaces. You can also view the historical trend of loss and latency at the network interface level for troubleshooting network connectivity issues.

The Device Health dashboard shows the network health-related KPIs such as CPU, memory, interface traffic, and utilization.

IPSLA dashboard supports Cisco, Huawei, and Juniper IPSLA KPIs on jitter and latency quality for sensitive network services while the Flow dashboard displays the total octet for all the different aggregation views.

The HTTP Operations dashboard shows the HTTP server response time (RTT). The measurements consist of three types:

- DNS lookup, the RTT taken to perform domain name lookup.
- TCP Connect, the RTT taken to perform a TCP connection to the HTTP server.
- HTTP transaction time, the RTT taken to send a request and get a response from the HTTP server.

Timeseries Data dashboard populates the SNMP performance metrics that are available from the timeseries database.

Rapid SNMP device onboarding is a new feature to bring new devices such as routers, switches, and servers into a network environment smoothly and seamlessly. You can create your own custom discovery formulas, collection formulas, and metrics and deploy them in the Network Performance Insight system. Network Performance Insight can start discovery and polling from the new devices and their resources with in a day. The collected metrics are stored in the timeseries database and the metrics are populated on the Timeseries Data dashboard.

Categories of On Demand Filtering dashboards:

- Device Health
- Flow
- HTTP Operations
- IPSLA
- Timeseries Data

Device Health dashboard

On Demand filtering Device Health dashboard helps to identify any anomalies from just visualizing the network entities trend charts. It can be further drill down to check the historical data for further troubleshooting.

Device Health

1. Click **On Demand Filtering** > **Device Health**.

   The **Device Health** dashboard loads with device health monitoring details for the selected KPI over the selected time period.

2. From the filter options, choose the **Device**, **KPI**, **Sort By**, and **Time period** and click **Apply Filter**.

   Note: You can select the filter values or time ranges that you want to display in dashboards to which the filter is assigned.

   The dashboard refreshes the data according to the filter attribute values.

   The Device Health dashboard shows the network health-related KPIs such as CPU, memory, interface traffic, and utilization. You can detect the anomalies from monitoring the spikes, dips, or irregular trend in the performance data from these network-related health entities dynamic charts. Each page displays a maximum of four dynamic charts.

   Note: You cannot see any plotted data in the dynamic chart when your query result has less than four data points. The dynamic chart populates the data only if the query returns more than four data points.

Table 121. Total network-related health entities dynamic charts

Widget name	Description	Drill-down dashboard
Trend For *<Network related health Entities>* **Note:** *<Network related health Entities>* denotes either Total Interfaces, CPU, Memory, or Temperature.	Displays the selected KPI trend over the selected time period. It's a dynamic chart that allows drill down to the **Device Health History** dashboard for the selected KPI. Each of the dynamic chart dynamically plots the total interfaces, CPU, Memory, or Temperature based on the filter attribute values.	It displays the **Device Health History** dashboard for the selected KPI and with other filter attribute values. The Device Health History dashboard consist of two sets of information: a. A timeseries chart that displays the total network related health entities for the selected time period. b. Predefined time range of historical data for the selected KPI for: • Last 24 Hours • Last 7 Days • Last 30 Days • Last 365 Days

3. Click the identified data point from the dynamic chart to drill down.

   The **Device Health History** dashboard for the selected KPI page loads in a new tab. It displays the detailed Device Health data and the historical data based on the filter attribute values from Device Health On Demand Filtering dashboard.

Table 122. Available widgets

Widget name	Chart type	Description
<Time Period> <KPI> >> <Device IP> >> <Resource Name> For example: **Last Hour Inbound Utilization (%) >> 10.55.239.100 >> Fa0/1**	Timeseries	It displays the selected KPI trend over the selected time period. By default, the information is shown based on the filter attribute values from the Device Health On Demand Filtering dashboard. **Note:** The TimeZoom feature is available for all the Timeseries charts in this dashboard.
Last 24 Hours *<KPI> >> <Device IP> >> <Resource Name>* For example: **Last 24 Hours Inbound Utilization (%) >> 10.55.239.100 >> Fa0/1**	Timeseries	It displays the entity history data of last 24 hours from the current time of the selected device. The historical data is displayed in accordance to the filter attribute values.
Last 7 Days *<KPI> >> <Device IP> >> <Resource Name>* For example: **Last 7 Days Inbound Utilization (%) >> 10.55.239.100 >> Fa0/1**	Timeseries	It displays the entity history data of last 7 days from the current time of the selected device. The historical data is displayed in accordance to the filter attribute values.
Last 30 Days *<KPI> >> <Device IP> >> <Resource Name>* For example: **Last 30 Days Inbound Utilization (%) >> 10.55.239.100 >> Fa0/1**	Timeseries	It displays the entity history data of last 30 days from the current time of the selected device. The historical data is displayed in accordance to the filter attribute values.
Last 365 Days *<KPI> >> <Device IP> >> <Resource Name>* For example: **Last 365 Days Inbound Utilization (%) >> 10.55.239.100 >> Fa0/1**	Timeseries	It displays the entity history data of last 365 days from the current time of the selected device. The historical data is displayed in accordance to the filter attribute values.
• Min • Avg • Max	Badge	**Note:** These badges are available for all the widgets in the dashboard. • Min indicates the minimum value of the KPI within the time period. • Avg indicates the average value of the KPI within the time period. • Max indicates the maximum value of the KPI within the time period.

4. You can select the filter values or time ranges that you want to display in the charts from the **Device Health History** dashboard filter options.

   From the filter options, choose the **Device**, **KPI**, **Resource**, and **Time period** and click **Apply Filter**.

   The dashboard refreshes the data according to the filter attribute values.

IPSLA dashboard

On Demand Filtering IPSLA dashboard helps you to identify the probes between the source and destination IP addresses. You can identify the latency or jitter in your network over a period for a set of KPIs for further troubleshooting.

IPSLA

1. Click **On Demand Filtering** > **IPSLA**.

   The **IPSLA** dashboard loads.

   IPSLA dashboard can be used to assess the network quality that uses network parameters such as delay, loss, jitter, and more to identify network issues that cause service quality degradation.

2. From the filter options, choose the **SLA Test**, **Source**, **Sort By**, **KPI**, and **Time period** and click **Apply Filter**.

   Note:

   - You can select the filter values or time ranges that you want to display in dashboards to which the filter is assigned.
   - The KPI list is populated based on the selected Service Level Agreement (SLA) Test list option.

   The dashboard refreshes the data according to the filter attribute values.

   This dashboard provides the IPSLA metric monitoring details for the selected KPI over the selected time period.

   You can detect the anomalies from monitoring the spikes, dips, or irregular trend in the performance data from the **Total probes** dynamic charts. Each page displays a maximum of four dynamic charts.

   Note: You cannot see any plotted data in the dynamic chart when your query result has less than four data points. The dynamic chart populates the data only if the query returns more than four data points.

   Table 123. Total Probes dynamic charts

Widget name	Description	Drill-down dashboard
Trend For *<Source IP> - <Destination IP>*	It's a dynamic chart that allows drill down to the **IPSLA History** dashboard for the source IP. Each of the dynamic chart dynamically plots the probes that help to identify the network connectivity issues on end-to-end between the selected router and devices by using IP addresses. The data is displayed in accordance to the filter attribute values.	It displays the IPSLA History dashboard for the selected source and with other filter attribute values. The **IPSLA History** dashboard consist of two sets of information: a. A timeseries chart that displays the probe counts for the selected time period. b. Predefined time range of historical data for the selected KPI for: • Last 24 Hours • Last 7 Days • Last 30 Days • Last 365 Days

3. Click the identified data point from the dynamic chart to drill down.

   The **IPSLA History** dashboard for the selected source and destination IP page loads in a new tab. It displays the detailed IPSLA data and the historical data based on the filter attribute values from IPSLA On Demand Filtering dashboard.

Table 124. Available widgets		
Widget name	**Chart type**	**Description**
<Time Period> Probe Count (count) >> Source *<IP>* >> Destination *<IP>* For example: **Last Hour Probe Count (count) >> Source 10.25.238.209 >> Destination 4.10.110.50**	Timeseries	It displays the probe count of the selected IPSLA KPI trend. By default, the information is shown based on the filter attribute values from the IPSLA On Demand Filtering dashboard. **Note:** The TimeZoom feature is available for all the Timeseries charts in this dashboard.
Probe History Data		
Last 24 Hours Probe Count (count) >> Source *<IP>>>* Destination *<IP>* For example: **Last 24 Hours Probe Count (count) >> Source 10.25.238.209 >> Destination 4.10.110.50**	Timeseries	Displays the probe history data of last 24 hours from the current time of the selected source and destination IP. The historical data is displayed in accordance to the filter attribute values.
Last 7 Days Probe Count (count) >> Source *<IP>>>* Destination *<IP>* For example: **Last 7 Days Probe Count (count) >> Source 10.25.238.209 >> Destination 4.10.110.50**	Timeseries	Display the probe history data of last 7 days from the current time of the selected source and destination IP. The historical data is displayed in accordance to the filter attribute values.
Last 30 Days Probe Count (count) >> Source *<IP>>>* Destination *<IP>* For example: **Last 30 Days Probe Count (count) >> Source 10.25.238.209 >> Destination 4.10.110.50**	Timeseries	Display the probe history data of last 30 days from the current time of the selected source and destination IP. The historical data is displayed in accordance to the filter attribute values.
Last 365 Days Probe Count (count) >> Source *<IP>>>* Destination *<IP>* For example: **Last 365 Days Probe Count (count) >> Source 10.25.238.209 >> Destination 4.10.110.50**	Timeseries	Display the probe history data of last 365 days from the current time of the selected source and destination IP. The historical data is displayed in accordance to the filter attribute values.

Table 124. Available widgets (continued)		
Widget name	**Chart type**	**Description**
• Min • Avg • Max	Badge	**Note:** These badges are available for all the widgets in the dashboard. • Min indicates the minimum value of the KPI within the time period. • Avg indicates the average value of the KPI within the time period. • Max indicates the maximum value of the KPI within the time period.

4. You can select the filter values or time ranges that you want to display in the charts from the **IPSLA History** dashboard filter options.

   From the filter options, choose the **SLA Test**, **Source**, **Destination**, **KPI**, and **Time period** and click **Apply Filter**.

   The dashboard refreshes the data according to the filter attribute values.

Flow dashboard

On demand filtering Flow dashboard helps to identify any anomalies from just visualizing the network interfaces traffic volume trend charts. It can be further drill down to check the historical data for further troubleshooting.

NetFlow

1. Click **On Demand Filtering** > **Flow**.

   The **Flow** dashboard loads.

   Flow dashboard displays the total flow interfaces for either inbound or outbound traffic. For example, it displays the total flow interface for the selected device based on the aggregation type and time range.

   By default, the dashboard displays the total flow interface for Application aggregation type.

2. From the filter options, choose the **Device**, **Direction**, **Aggregation Type**, **Top N**, and **Time period** and click **Apply Filter**.

   Note: You can select the filter values or time ranges that you want to display in dashboards to which the filter is assigned.

   The dashboard refreshes the data according to the filter attribute values.

   You can detect the anomalies from monitoring the spikes, dips, or irregular trend in the performance data from the **Total Flow Interface(s): <N>** dynamic charts. Each page displays a maximum of two dynamic charts. This dynamic chart dynamically charts the interfaces.

   Note:

   - *<N>* denotes the number of total flow enabled interfaces.
   - You will not see any data plotted in the dynamic chart when your query result has less than four data points. The dynamic chart only populates the data if the query returns more than four data points.

Table 125. Total Flow Interface(s): <N> dynamic charts		
Widget name	**Description**	**Drill-down dashboard**
Trend For *<Interface Name>*	It's a dynamic chart that allows drill down to the **Flow History** dashboard for the selected aggregation type view. Each of the dynamic charts shows either inbound or outbound flow traffic volume trend of an interface per aggregation view for the selected device. The data is displayed in accordance to the filter attribute values.	It displays the Flow History dashboard for the selected aggregation type view with the filter attribute values. The **Flow History** dashboard consist of two sets of information: a. A timeseries chart that displays the traffic volume in octets for the selected time period. b. Predefined time range of historical data for the selected interface of an aggregation type view for: • Last 24 Hours • Last 7 Days • Last 30 Days • Last 365 Days

3. Click the identified data point from the dynamic chart to drill down.

   The **Flow History** dashboard for the selected aggregation type view page loads in a new tab. It displays the detailed flow data and the historical data for the selected aggregation type view with the filter attribute values.

Table 126. Available widgets		
Widget name	**Chart type**	**Description**
<Time Period> *<Device IP>* >> *<Interface Name>* >> *<Aggregation Type>* For example: **Last Hour 10.55.239.219 >> 10.55.239.219-Fa0/1 >> Applications**	Timeseries	It displays the flow traffic volume in octets. By default, the information is shown based on the filter attribute values from the Flow On Demand Filtering dashboard.
Flow History Data		
Last 24 Hours *<Device IP>* >> *<Interface Name>* >> *<Aggregation Type>* For example: **Last 24 Hours 10.55.239.219 >> 10.55.239.219-Fa0/1 >> Applications**	Timeseries	Displays flow traffic volume in octets for the selected interface for the last 24 hours from the current time. The historical data is displayed in accordance to the filter attribute values.

Table 126. Available widgets (continued)		
Widget name	**Chart type**	**Description**
Last 7 Days *<Device IP>* >> *<Interface Name>* >> *<Aggregation Type>* For example: **Last 7 Days 10.55.239.219 >> 10.55.239.219-Fa0/1 >> Applications**	Timeseries	Displays flow traffic volume in octets for the selected interface for the last 7 days from the current time. The historical data is displayed in accordance to the filter attribute values.
Last 30 Days *<Device IP>* >> *<Interface Name>* >> *<Aggregation Type>* For example: **Last 30 Days 10.55.239.219 >> 10.55.239.219-Fa0/1 >> Applications**	Timeseries	Displays flow traffic volume in octets for the selected interface for the last 30 days from the current time. The historical data is displayed in accordance to the filter attribute values.
Last 365 Days *<Device IP>* >> *<Interface Name>* >> *<Aggregation Type>* For example: **Last 365 Days 10.55.239.219 >> 10.55.239.219-Fa0/1 >> Applications**	Timeseries	Displays flow traffic volume in octets for the selected interface for the last 365 days from the current time. The historical data is displayed in accordance to the filter attribute values.

4. You can select the filter values or time ranges that you want to display in the charts from the **Flow History** dashboard filter options.

   From the filter options, choose the **Device**, **Interface**, **Direction**, **Aggregation Type**, **Top N**, and **Time period** and click **Apply Filter**.

   The dashboard refreshes the data according to the filter attribute values.

HTTP Operations dashboard

On demand filtering HTTP Operations dashboard helps to display the round-trip time (RTT) between a Cisco device and an HTTP server to retrieve a web page. It can be further drill down to view the response time trend over a period for a set of KPIs for further troubleshooting.

HTTP Operations

1. Click **On Demand Filtering** > **HTTP Operations**.

   The **HTTP Operations** dashboard loads with the round-trip time (RTT) to help to analyze how an HTTP server is performing.

2. From the filter options, choose the **Source**, **Sort By**, **KPI**, and **Time period** and click **Apply Filter**.

   Note:

   - You can select the filter values or time ranges that you want to display in dashboards to which the filter is assigned.
   - The KPI list is populated based on the IP Service Level Agreement (SLA) HTTP Operation.

   This dashboard provides the HTTP Operations metric monitoring details for the selected KPI over the selected time period.

   You can detect the anomalies from monitoring the spikes, dips, or irregular trend in the performance data from the **Total probes** dynamic charts. Each page displays a maximum of four dynamic charts.

Note: You will not see any data plotted in the dynamic chart when your query result has less than four data points. The dynamic chart only populates the data if the query returns more than four data points.

Table 127. Total Probes dynamic charts

Widget name	Description	Drill-down dashboard
Trend For *<Source IP> - <Destination IP>*	Its a dynamic chart that allows drill down to the **HTTP Response Time: IPSLA HTTP Operations** page for the source IP. Each of the dynamic chart dynamically plot the probes, which helps to identify round-trip time (RTT) between a Cisco device and an HTTP server to retrieve a web page. The data is displayed in accordance to the filter attribute values.	It displays the **HTTP Response Time: IPSLA HTTP Operations** page for the selected source and with other filter attribute values. The **HTTP Response Time: IPSLA HTTP Operations** dashboard consist of two sets of information: a. A timeseries chart that displays the response time trend for the selected *<Source IP> - <Destination IP>*. b. A grid chart that displays the response time in milliseconds.

3. Click the identified data point from the dynamic chart to drill down.

   The **HTTP Response Time: IPSLA HTTP Operations** for the selected source and destination IP page loads in a new tab.

Table 128. Available widgets

Widget name	Chart type	Description
Response Time Trend	Timeseries	The widget displays the HTTP response time for the HTTP operation metrics. By default, the information is shown based on the filter attribute values from the HTTP Operations On Demand Filtering dashboard.
Response Time in milliseconds (ms)	Grid	

4. You can select the filter values or time ranges that you want to display in the charts from the **HTTP Response Time: IPSLA HTTP Operations** page filter options.

   From the filter options, choose the **Source**, **Destination**, and **Time period** and click **Apply Filter**.

   The dashboard refreshes the data according to the filter attribute values.

Timeseries Data dashboard

On demand filtering Timeseries Data dashboard helps to display SNMP performance metrics. It also allows you to monitor for new device types and vendors.

Timeseries Data dashboard populates the SNMP performance metrics that are available from the timeseries database.

It can also be used for Rapid SNMP device onboarding scenario where you can view the new SNMP metrics that are created and deployed in the Network Performance Insight system.

Note: To use Timeseries Data dashboard in Rapid SNMP device onboarding scenario, perform the following tasks:

- Create the custom formulas and metrics for the newly on-boarded devices.

- Deploy them in Network Performance Insight system to start using the content to discover and poll the devices and resources.

Timeseries Data

1. Click **On Demand Filtering** > **Timeseries Data**.

   The **Timeseries Data** dashboard loads.

   The Timeseries data dashboard displays the collected performance metrics that are stored in the Network Performance Insight Dashboards timeseries database.

2. From the filter options, choose the **Device**, **Resource Type**, **Resource Name**, and **Time period** and click **Apply Filter**.

   Note: You can select the filter values or time ranges that you want to display in dashboards to which the filter is assigned.

   The **List of KPI** widget populates the available KPIs for the selected device and resource type.

3. From the **List of KPI** widget check-boxes, select the required KPIs to be populated on the **KPI(s) Trend** chart.

   a. Select the KPIs check boxes.

   b. Right click the selection and click **KPI(s) Trend**.

   The **KPI(s) Trend** timeseries chart refreshes according to the selected KPI list.

Table 129. Widget interactions

Widget name	Description	Drill-down dashboard
List of KPI	List the available KPIs for the selected device and resource type. You can view one or more KPI trend on the timeseries chart by selecting the KPI check boxes. You can set the KPI list filter condition: • Click the define filter icon ( ). The **Filter** pop-up window loads. • Set the conditions from the **Filter** pop-up window and click **Filter** The KPI list refreshes according to the filter conditions set.	There's no drill-down from this widget. The **KPI(s) Trend**- *<'KPI Name'>* timeseries chart refreshes according to the selected KPI list.

Table 129. Widget interactions (continued)		
Widget name	**Description**	**Drill-down dashboard**
KPI(s) Trend	A dynamic timeseries chart that displays the KPI trend for the selected KPIs. It's a dynamic chart that allows drill down to the **Timeseries Data History** page for the KPI and filter attribute values. **Note:** You will not see any data plotted in the dynamic chart when your query result has less than four data points. The dynamic chart only populates the data if the query returns more than four data points.	The **Timeseries Data History** page displays the predefined time range of history data for the selected KPI and resource type of: • Last 24 Hours • Last 7 Days • Last 30 Days • Last 365 Days

4. Click the identified data point from the **KPI(s) Trend** dynamic timeseries chart to drill down.

   The **Timeseries Data History** page loads in a new tab. It displays the selected KPI's data for the predefined time range of historical data.

Table 130. Historical Data widgets		
Widget name	**Chart type**	**Description**
Last 24 Hours >> Resource Type *<resource_type>* >> Resource Name *<resource_name>* For example: **Last 24 Hours >> Resource Type interface >> Resource Name ge-1/0/0**	Timeseries	Displays the KPI's timeseries data for the selected device and resource type for the last 24 hours from the current time. The historical data is displayed in accordance to the filter attribute values.
Last 7 Days >> Resource Type *<resource_type>* >> Resource Name *<resource_name>* For example: **Last 7 Days >> Resource Type interface >> Resource Name ge-1/0/0**	Timeseries	Displays the KPI's timeseries data for the selected device and resource type for the last 7 days from the current time. The historical data is displayed in accordance to the filter attribute values.
Last 30 Days >> Resource Type *<resource_type>* >> Resource Name *<resource_name>* For example: **Last 30 Days >> Resource Type interface >> Resource Name ge-1/0/0**	Timeseries	Displays flow traffic volume in octets for the selected interface for the last 30 days from the current time. The historical data is displayed in accordance to the filter attribute values.
Last 365 Days >> Resource Type *<resource_type>* >> Resource Name *<resource_name>* For example: **Last 365 Days >> Resource Type interface >> Resource Name ge-1/0/0**	Timeseries	Displays the KPI's timeseries data for the selected device and resource type for the last 365 days from the current time. The historical data is displayed in accordance to the filter attribute values.

5. You can select the filter values or time ranges that you want to display in the charts from the **Timeseries Data History** dashboard filter options.

   From the filter options, choose the **Device**, and **Resource Name**, and click **Apply Filter**.

   The dashboard refreshes the data according to the filter attribute values.

Related information

Rapid SNMP device onboarding

Chapter 13. Configuring integration to IBM Connections

IBM Tivoli Netcool/Impact provides integration to IBM Connections by using a Netcool/Impact `IBMConnections` action function. The `IBMConnections` action function allows users to query forums and topics lists, create a new forum, create a new topic, and update existing topics. The `IBMConnections` action function package is available in the directory `$IMPACT_HOME/integrations/IBMConnections`.

About this task

Complete the following steps to configure Netcool/Impact to integrate with the IBM Connections Server.

Procedure

1. Go to the directory `$IMPACT_HOME/add-ons/IBMConnections`, this directory is the IBM Connections integration package. The package includes the following subdirectory:

   importData
   A project that includes policies, data source, data type, and a service. The project serves an example to show how to connect, create, update, and topics in IBM Connections
2. Import the project `$IMPACT_HOME/bin/nci_ import <ServerName> __Extraction_Directory__/importData`.
3. Within the `$IMPACT_HOME/etc/<NCI>_server.props` file, add the following parameters:

   `impact.ibmconnections.forum.title.maxsize= number`. Default value is 0. Any string size can be used.
   `impact.ibmconnections.forum.content.maxsize= number`. Default value is 0. Any string size can be used
   `impact.ibmconnections.topic.title.maxsize= number`. Default value is 255, for a topic and a reply.
   `impact.ibmconnections.topic.content.maxsize = number`. Default value is 0. Any string size can be used
4. Restart Netcool/Impact servers.

Related tasks

Installing Netcool/OMNIbus and Netcool/Impact

IBM Connections Overview

IBM Connections is a leading social software platform that can help your organization to engage the right people, accelerate innovation, and deliver results.

This integrated, security-rich platform helps people engage with networks of experts in the context of critical business processes. Now everyone can act with confidence and anticipate and respond to emerging opportunities

For information about IBM Connections, refer to

http://www-03.ibm.com/software/products/en/conn

Parameters for the `IBMConnections` function

The integration between Netcool/Impact and IBM Connections uses a new policy action function that is called the `IBMConnections` function. The `IBMConnections` function can be used within any policy to

connect to the IBM Connections Server and to perform an action on the IBM Connections Server. The function accepts two input parameters, the **Action Option** parameter and the **Impact Object** parameter.

Action Option Parameter

The **Action Option** parameter accepts one of the following action entries, which are case insensitive. Some action entries require property information that is case-sensitive.

In the content of an IBM Connections forum, topic, or reply, you can use HTML formatting tags br, b, and a. For more information about the supported HTML tags, see http://www-03.ibm.com/software/products/en/conn.

CREATEFORUM

Creates a forum.

Enter the following property information that is case-sensitive. The tags must be created before they pass to a variable name.

`props.ForumTitle=`*`title`*`;`
`props.ForumContent=`*`full text of the body`*`;`
`props.ForumTags=List_Of_Tags;` Is optional, the object must be a Netcool/Impact object.
`Tags=NewObject();`
`Tags.Tag1=`*`some tag`*`;`
`Tags.Tag`*`2`*`=`*`some tag2`*`;` Is optional if want more than one tag.

CREATETOPIC

Creates a topic.

Enter the following property information that is case-sensitive:

`props.TopicTitle=`*`title`*`;`
`props.TopicContent=`*`full text of the body`*`;`
`props.ForumId=`*`forum id`*: Where the *forum id* is an ID and not a forum name.

DELETEFORUM

Deletes the forum name that was created by the logged in user and any topic or reply belonging to it.

Enter the following property information that is case-sensitive:

`props.ForumId=`*`forumId`*`;` Or `props.ForumId=`*`forumTitle`*`;`
`props.FirstMatchOnly=`*`true`*`;` Or `props.FirstMatchOnly=`*`false`*`;` The `props.FirstMatchOnly` property deletes the first matching forum or matching topic that it finds, or else it deletes any matching forum or matching topic and its default value is `true`.

DELETEPUBLICFORUM

Deletes the given public forum name and any topic or reply belonging to it.

Enter the following property information that is case-sensitive:

`props.ForumId=`*`forumId`*`;` Or `props.ForumId=`*`forumTitle`*`;`
`props.FirstMatchOnly=`*`true`*`;` Or `props.FirstMatchOnly=`*`false`*`;`The `props.FirstMatchOnly` property deletes the first matching forum or matching topic that it finds, or else it deletes any matching forum or matching topic and its default value is `true`.

DELETEREPLY

Deletes a topic reply in the topic in the forum id.

Enter the following property information that is case-sensitive:

`props.ForumId=`*`forumId`*`;` Where the *forumId* is an ID and not a title.
`props.TopicTitle=`*`title`*`;` Or `props.TopicTitle=`*`id`*`;`
`props.ReplyTitle=`*`replytitle`*`;` Or `props.ReplyTitle=`*`replyid`*`;`

`props.FirstMatchOnly=true`; Or `props.FirstMatchOnly=false`;The `props.FirstMatchOnly` property deletes the first matching forum or matching topic that it finds, or else it deletes any matching forum or matching topic and its default value is `true`.

DELETETOPIC

Deletes a topic in the forum id.

Enter the following property information that is case-sensitive:

`props.ForumId=forumId`; Where the *forumId* is an ID and not a title.
`props.TopicTitle=title`; Or `props.TopicTitle=id`;
`props.FirstMatchOnly=true`; Or `props.FirstMatchOnly=false`;The `props.FirstMatchOnly` property deletes the first matching forum or matching topic that it finds, or else it deletes any matching forum or matching topic and its default value is `true`.

GETCOMMUNITYFORUMID

Gets the ID of the community that is created by the logged in user

Enter the following property information that is case-sensitive:

`props.CommunityName=Community name`;
`props.ForumName=forum name`;

GETFORUMTOPICS

Gets list of topics for the forum id

Enter the following property information that is case-sensitive:

`props.ForumId=forum id`;

GETMYCOMMUNITYID

Gets the ID of the community for the logged in user

Enter the following property information that is case-sensitive:

`props.CommunityName=Community name`;

GETMYFORUMID

Gets the ID of the forum that is created by the logged in user.

Enter the following property information that is case-sensitive:

`props.ForumName=actual forum name that is created by the logged in user`

GETMYFORUMS

Gets all the forums for the logged in user

GETPUBLICCOMMUNITYID

Gets the ID of the given public community ID

Enter the following property information that is case-sensitive:

`props.CommunityName=Community name`;

GETPUBLICFORUMID

Gets the ID of the given public forum name

Enter the following property information that is case-sensitive:

`props.ForumName=actual public forum name`

GETPUBLICFORUMS

Gets list of all public forums

GETTOPICREPLIES

Gets list of replies for a topic

Enter the following property information that is case-sensitive:

`props.TopicId=topic id`; Where *topic id* must be the topic id not the topic name.

REPLYTOTOPIC

Creates a reply to an existing topic

Enter the following property information that is case-sensitive:

`props.TopicId=`*`topic id`*`;` Where *topic id* is a topic ID not a topic name
`props.ReplyTitle=`*`title`*`;`
`props.ReplyContent=`*`full text of the body`*`;`

`Impact Object` Parameter

The **`Impact Object`** parameter accepts the following property information. The authentication, and connection property information is mandatory.

```
props = NewObject();
props.Protocol=https;
props.Host=IBM Connections Server Host/IP;
props.Port=_PORT_;
props.Username=userName;
```

`props.Password=`*`password`*`;` The password can be encrypted by using either the Netcool/Impact `nci_crypt` tool or the policy function `Encrypt()`. If the password is encrypted, you must use the property `props.DecryptPassword=true;`

IBMConnections Project and artifacts

The imported `IBMConnections` project includes artifacts that are categorized into four categories.

Data sources

- `IBMConnectionsObjectServerDSA`
- `Internal`

Data types

- `TopicCreationTracker`
  - An internal data type that is used to track the topic creation to avoid duplicate names.
- `InternetOutageEvents`
  - ObjectServer data type that can be used to view the critical events in the UI Data Provider widgets. It populates the severity as status data type to show colorful images.

Policies

- `IBMConnectionsUtils`
  - Includes a utility function to extract value from a data item.
- `IBMConnectionsUtilsCaller`
  - Shows how to call the utility function in `IBMConnectionsUtils`.
- `IBMConnectionsUtilsJS`
  - For JavaScript policies.
- `IBMConnectionsUtilsCallerJS`
  - For JavaScript policies.
- `NetworkMonitorExample`
  - Example policy that is run by an event reader to create and update topics.

- NetworkMonitorForOpView
  - Example policy that is run by the operator view from the ObjectServer Event List tool.
- Opview_IBMConnectionsOpView
  - Is run by the AEL tool.

Services

- NetworkMonitorExample
  - Connects to the object server data source and uses a default filter of Node in ('US','France','UK')and Identifier Like 'Monitoring Network for'. You can change the default filter at any time. It runs the policy NetworkMonitorExample.

Automatic topic management

The IBM Connections integration package includes NetworkMonitorExample event reader service that connects to the Netcool/OMNIbus Object Server and filters for specific events. When there is a match, the service runs the policy that either updates the topic by sending a reply to the topic, or creates a topic if a topic does not exist.

The forum that is used in this example is the same name as the AlertKey in the Object Server event, forumName = @IBMConnections_Forum

You can use the sql scripts in the $IMPACT_HOME/integrations/IBMConnections/db directory to create extra fields in the Object Server or you can use existing fields.

The topic title is created as a combination of a hardcoded string and the node field from the event:

```
topicTitleVar="Network Monitor is down on node: @Node@" ;
IBMConnectionsUtils.extractParametersAndSubstitute
(topicTitleVar,EventContainer,result);
topicTitle = result;
```

The policy checks if the topic was created by querying the internal data type TopicCreationTracker. If the topic exists, the policy sends a reply instead of creating a new one.

Automatic topic management with event management tools

The operator view policy is updated to run the NetworkMonitorForOpView policy that automatically updates an existing topic or creates a new topic.

Procedure

1. Create the event management tool, refer to the Netcool/OMNIbus documentation.
2. In the event management tool, select the executable box tab and enter the following text: start "" "https://<impactgui_server>:<port>/opview/displays/NCICLUSTER-IBMConnectionsOpView.html?Node=@Node&Serial=@Serial&Severity=@Severity&Acknowledged=@Acknowledged&AlertKey=@AlertKey&AlertGroup=@AlertGroup&Summary=@Summary"

   NCICLUSTER Is the default cluster but if you are using a different cluster name then update the URL with your cluster name.

   The above text is an example of text to enter for an Object Server on Windows.
3. Add the new event management tool to the AlertsStatus tools menu, refer to the Netcool/OMNIbus documentation.

4. When you right-click on an event, click the tool to start the URL and run the policy. The operator view is started and gives a notification that the topic is created or updated, along with a link to the topic URL to use.

Related information

Creating event management tools

Enabling historical events

Create a connection to the historical database in Impact to view historical events in the Event Viewer.

Procedure

1. Log in to Dashboard Application Services Hub and select the **Data Model** tab.
2. Click the **New Data Source** icon.
3. Point to **Database SQL** and select your database type. For example **Db2**.
4. In the **Data Source Name** field enter: **`historicalEventsDatasource`**.
5. Enter your Username and Password in the fields provided and click the **Save** icon.
6. In the left-hand navigation pane, right-click **ImpactHistoricalEventData** and select **New Data Type**.
7. In the **Data Type Name** field enter: **`historicalEventData`**.
8. Click **Refresh**.

Chapter 14. Troubleshooting Netcool Operations Insight

Consult these troubleshooting notes to help determine the cause of the problem and what to do about it.

Troubleshooting Netcool Operations Insight on premises

Use this information to troubleshoot problems that may occur when installing or using Netcool Operations Insight on premises.

Troubleshooting Event Analytics (on premises)

Use the following troubleshooting information to resolve problems with Event Analytics.

If your problem is not listed in this topic, then refer to the Release notes for additional issues.

Analytics data

Use the following troubleshooting information to resolve problems with Event Analytics analytics data.

Two or more returned seasonal events appear to be identical

It is possible for events to have the same **Node**, **Summary**, and **Alert Group** but a different **Identifier**. In this scenario, the event details of two (or more) events can appear to be identical because the **Identifier** is not displayed in the details.

Error displaying Seasonal Event Graphs in Microsoft Internet Explorer browser

The Seasonal Event Graphs do not display in a Microsoft Internet Explorer browser.

This problem happens because Microsoft Internet Explorer requires the Microsoft Silverlight plug-in to display the Seasonal Event Graphs.

To resolve this problem, install the Microsoft Silverlight plug-in.

Within the Event Viewer, you are unable to view seasonal events and error ATKRST103E is logged

When you complete the following type of steps, then within the event viewer the seasonal events are not viewable and error ATKRST103E is logged.

1. Open the event viewer and select to edit the widget from the widget menu.
2. From the list on the edit screen, select the `Impact Cluster` data provider.
3. Select to view either the seasonality report and the report name.
4. Save the configuration.

To resolve the problem, view seasonal events by using the provided seasonal events pages and view related events parent to child relationships by using the Tivoli Netcool/OMNIbus data provider.

Event relationships display in the Event Viewer, only if the parent and child events match the filter

The Event Viewer is only able to show relationships between events if the parent and the child events are all events that match the filter. There are some use cases for related events where parent or child events might not match the filter.

Background

Netcool/OMNIbus Web GUI is able to show the relationships between events in the Event Viewer, if the Event Viewer view in use has an associated Web GUI relationship. This relationship defines which field in an event contains the identifier of the event's parent, and which field contains the identifier for

the current event. For more information about defining event relationships, see http://www-01.ibm.com/support/knowledgecenter/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/webtop/wip/task/web_cust_jsel_evtrelationshipmanage.html.

The relationship function works from the set of events that are included in the event list, and the event list displays the events that match the relevant Web GUI filter. See the following example. If you have a filter that is called `Critical` to show all critical events, the filter clause is `Severity = 5`, then relationships between these events are shown provided the parent and child events in the relationships all have `Severity = 5`. If you have a parent event that matches the filter `Severity = 5` but has relationships to child events that have a major severity `Severity = 4`, these child relations are not seen in the event list because the child events do not match the filter. Furthermore, these child relations are not included in the set of events that are returned to the Event Viewer by the server.

Resolution

To resolve this problem, you must define your filter with appropriate filter conditions that ensures that related events are included in the data that is returned to the Event Viewer by the server. The following example builds on the example that is used in the *Background* section.

1. Make a copy of the `Critical` filter and name the copy `CriticalAndRelated`. You now have two filters. Use the original filter when you want to see only critical events. You use the new filter to see related events, even if events are not critical.
2. Manually modify the filter condition of the `CriticalAndRelated` filter to include the related events. To manually modify this filter condition, use the advanced mode of the Web GUI filter builder. The following example conditions are based on the current example.

   The main filter condition is `Severity = 5`.
   In an event, the field that denotes the identifier of the parent event is called `ParentIdentifier`.
   The value of the `ParentIdentifier` field, where populated, is the `Identifier` of an event.
   If `ParentIdentifier` is `0`, this value is a default value and does not reference another event.

   - Including related child events. To include events that are the immediate child events of events that match the main filter, set this filter condition.

     ```
     Severity = 5
     OR
     ParentIdentifier IN (SELECT Identifier FROM alerts.status WHERE Severity = 5)
     ```

   - Including related parent events. To include events that are the immediate parent of events that match the main filter, set this filter condition.

     ```
     Severity = 5
     OR
     Identifier IN (SELECT ParentIdentifier from alerts.status WHERE Severity = 5)
     ```

   - Including related sibling events. To include events that are the other child events of the immediate parents of the event that matches the main filter (the siblings of the events that match the main filter), set this filter condition.

     ```
     Severity = 5
     OR
     ParentIdentifier IN (SELECT ParentIdentifier from alerts.status WHERE
     Severity = 5 AND ParentIdentifier > 0)
     ```

   - Including related parents, children, and siblings together. Combine the previous types of filter conditions so that the new `CriticalAndRelated` filter retrieves critical events, and the immediate children of the critical events, and the immediate parents of the critical events, and the immediate children of those parent events (the siblings). You must have this filter condition.

     ```
     Severity = 5
     OR
     ParentIdentifier IN (SELECT Identifier FROM alerts.status WHERE Severity = 5)
     OR
     ```

```
Identifier IN (SELECT ParentIdentifier from alerts.status WHERE Severity = 5)
OR
ParentIdentifier IN (SELECT ParentIdentifier from alerts.status WHERE
Severity = 5 AND ParentIdentifier > 0)
```

- Including related events that are more than one generation away. In the previous examples, the new filter conditions go up to only one level, up or down, from the initial set of critical events. However, you can add more filter conditions to retrieve events that are more than one generation away from the events that match the main filter. If you want to retrieve grandchildren of the critical events (that is, two levels down from the events that match the main filter condition) and immediate children, set this filter condition.

```
-- The initial set of Critical events
Severity = 5
OR
-- Children of the Critical events
ParentIdentifier IN (SELECT Identifier FROM alerts.status WHERE Severity =
5)
-- Children of the previous "child events"
OR
ParentIdentifier IN (SELECT Identifier FROM alerts.status WHERE
    ParentIdentifier IN (SELECT Identifier FROM alerts.status WHERE Severity = 5) )
```

Use a similar principal to retrieve parent events that are two levels up, and siblings of the parent events. To pull this scenario together, set this filter condition.

```
-- The initial set of Critical events
Severity = 5

OR

-- Children of the Critical events
ParentIdentifier IN (SELECT Identifier FROM alerts.status WHERE
Severity = 5)

OR

-- Children of the previous "child events"
ParentIdentifier IN (SELECT Identifier FROM alerts.status WHERE
    ParentIdentifier IN (SELECT Identifier FROM alerts.status WHERE
Severity = 5) )

OR

-- Parents of the Critical events
Identifier IN (SELECT ParentIdentifier from alerts.status WHERE Severity = 5)

OR

-- Parents of the previous "parent events"
Identifier IN (SELECT ParentIdentifier from alerts.status WHERE
    Identifier IN (SELECT ParentIdentifier from alerts.status WHERE Severity = 5) )
OR

-- Other children of the Critical events' parents
ParentIdentifier IN (SELECT ParentIdentifier from alerts.status WHERE
Severity = 5 AND ParentIdentifier > 0)

OR

-- Other children of the Critical events' grandparents
ParentIdentifier IN (SELECT ParentIdentifier from alerts.status WHERE
    Identifier IN (SELECT ParentIdentifier from alerts.status WHERE
Severity = 5 AND ParentIdentifier > 0) AND ParentIdentifier > 0)
```

You can continue this principal to go beyond two levels in the hierarchy. However, with each additional clause the performance of the query degrades due to the embedded subquerying. Therefore, there might be a practical limit to how far away the related events can be.

Configurations

Use the following troubleshooting information to resolve problems with Event Analytics configurations.

Aggregate fields missing from History database.

If existing aggregate (or rollup) fields no longer exist in the History database the Event Analytics wizard displays a warning.

To resolve this problem, Complete the following steps:

1. Complete one of the following options:
   - Add the missing field to the History database, and re-open the wizard.
   - Use a database view instead of the Event History database table to add a dummy field for the missing aggregate field. For more information about creating a database view for the Event Analytics wizard, see "Mapping customized field names" on page 279.
   - Open the Event Analytics wizard and edit the aggregate field to use a different field. Save the configuration.
   -
2. Open the Event Analytics wizard and delete the aggregate field.
3. Save the configuration.

Seasonal Event configuration stops running before completion.

The Seasonal Event configuration does not complete running. No errors are displayed. The report progress does not increase.

This problem occurs if a Seasonal Event Report is running when the Netcool/Impact back-end server goes offline while the Impact UI server is still available. No errors are displayed in the Impact UI and no data is displayed in the widgets/dashboards.

To resolve this problem, ensure that the Netcool/Impact servers are running. Edit and rerun the Seasonal Event Report.

Incomplete, stopped, and uninitiated configurations

Configuration run operations do not complete, are stalled on the Configure Analytics window, or fail to start.

These problems occur if the services are not started after Event Analytics is installed, or the Netcool/Impact server is restarted.

To resolve these problems, complete the following steps.

1. In the Netcool/Impact UI, select the **Impact Services** tab.
2. Ensure that each of the following services is started. To start a service, right-click the service and select **Start**.

```
LoadRelatedEventPatterns
ProcessClosedPatternInstances
ProcessPatternGroupsAllocation
ProcessRelatedEventConfig
ProcessRelatedEventPatterns
ProcessRelatedEventTypes
ProcessRelatedEvents
ProcessSeasonalityAfterAction
ProcessSeasonalityConfig
ProcessSeasonalityEvents
ProcessSeasonalityNonOccurrence
```

```
UpdateSeasonalityExpiredRules
```

Event Analytics: Configurations fail to run due to event count queries that take too long.

Configurations fail to run due to large or unoptimized datasets that cause the Netcool/Impact server to timeout and reports fails to complete.

To resolve this issue, increase the Netcool/Impact server timeout value to ensure that the Netcool/Impact server processes these events before it times out. As a result of increasing this server timeout value, the Netcool/Impact server waits for the events to be counted, thus ensuring that the reports complete and display in the appropriate portlet.

Edit the Netcool/Impact `impact.server.timeout` value, at

```
$IMPACT_HOME/etc/ServerName_server.props
```

By default, the `impact.server.timeout` property is set to 120000 milliseconds, which is equal to 2 minutes. The recommendation is to specify a server timeout value of at least 5 minutes. If the issue continues, increase the server timeout value until the reports successfully complete and display in the appropriate portlet.

Running a Seasonal Event configuration displays an error message `Error creating report. Seasonality configuration is invalid`

The Seasonal Event configuration does not run. An error message is displayed.

```
Error creating report.
Seasonality configuration is invalid. Verify settings and retry.
```

This problem occurs when Event Analytics is not correctly configured before you run a Seasonal Event Report.

To resolve this problem, review the Event Analytics installation and configuration guides to ensure that all of the prerequisites and configuration steps are complete. Also, if you use a table name that is not the standard REPORTER_STATUS, you must verify the settings that are documented in the following configuration topics.

"Configuring Db2 database connection within Netcool/Impact" on page 259
"Configuring Oracle database connection within Netcool/Impact" on page 261
"Configuring MS SQL database connection within Netcool/Impact" on page 263

Seasonality and Related Event configuration runs time out when you use large data sets

Before the seasonality policy starts to process a report, the seasonality policy issues a database query to find out how many rows of data need to be processed. This database query has a timeout when the database contains many rows and the database is not tuned to process the query. Within the *<impact install>*`/logs/impact_server.log` file, the following message is displayed.

```
02 Sep 2014 13:00:28,485 ERROR [JDBCVirtualConnectionWithFailOver] JDBC Connection
Pool recieved
error trying to connect to data source at: jdbc:db2://localhost:50000/database
02 Sep 2014 13:02:28,500 ERROR [JDBCVirtualStatement] JDBC execute failed twice.
com.micromuse.common.util.NetcoolTimeoutException: TransBlock [Executing SQL query:
select count(*)
as COUNT from Db2INST1.PRU_REPORTER where ((Severity >= 4) AND ( FIRSTOCCURRENCE >
'2007-
09-02 00:00:00.000' )) AND ( FIRSTOCCURRENCE < '2014-09-02 00:00:00.000')] timed
out after
120000ms.
```

Check that you have indexes for the FIRSTOCCURRENCE field and any additional filter fields that you specified, for example, `Severity`. Use a database tuning utility, or refresh the database statistics, or contact your database administrator for help. Increase the `impact.server` timeout to a value greater than the default of `120s`, see http://www-01.ibm.com/support/docview.wss?uid=swg21621488.

Seasonal or related events configurations hang, with error ATKRST132E

When you start cluster members, replication starts and the Netcool/Impact database goes down. Any running seasonality reports or related events configurations hang and this error message is logged in the Netcool/Impact server log.

```
ATKRST132E An error occurred while transferring a request to the following remote
provider: 'Impact_NCICLUSTER.server.company.com'. Error Message is
'Cannot access data provider - Impact_NCICLUSTER.server.company.com'.
```

To resolve this problem, do a manual restart or a scheduled restart of the affected reports or configurations.

Event Analytics configuration `Finished with Warnings`

The seasonality report or related events configuration completes with a status of `Finished with Warnings`. This message indicates that a potential problem was detected but it is not of a critical nature. You should review the log file for more information (`$NCHOME/logs/impactserver.log`). The following is an example of a warning found in `impactserver.log`:

```
11:12:38,366 WARN  [NOIProcessRelatedEvents] WARNING: suggested pattern :
RE-sqa122-last36months-Sev3-Default_Suggestion4 includes too many types,
could be due to configuration of types/patterns.
The size of the data exeeceded the column limit.
The pattern will be dropped as invalid.
```

Event Analytics configuration `Finished with Errors`

One reason for an Event Analytics configuration to complete with a status of `Finished with Errors` is because the suggested patterns numbering is not sequential. This can be because, for example, the pattern type found is invalid or the string is too long to be managed by the Derby database. You should review the log file for more information (`$NCHOME/logs/impactserver.log`).

Export

Use the following troubleshooting information to resolve problems with Event Analytics export operations.

Export of Event Analytics reports causes log out of DASH

If Netcool/Impact and DASH are installed on the same server, a user might be logged out of DASH when exporting Event Analytics reports from DASH. The problem occurs when the **Download export result** link is clicked in DASH. A new browser tab is opened and the DASH user is logged out from DASH.

To avoid this issue, configure SSO between DASH and Netcool/Impact. For more information, see https://www.ibm.com/support/knowledgecenter/SSSHYH_7.1.0.16/com.ibm.netcoolimpact.doc/admin/imag_configure_single_signon.html.

Failover

Use the following troubleshooting information to resolve failover problems with Event Analytics.

Configuring Netcool/Impact for ObjectServer failover

Netcool/Impact does not process new events for Event Analytics after ObjectServer failover. Seasonal event rule actions are not applied if the Netcool/Impact server is not configured correctly for ObjectServer failover as new events are processed. For example, if a seasonal event rule creates a synthetic event, the synthetic event does not appear in the event list, or if a seasonal event rule changes the column value for an event, the value is unchanged.

This problem occurs when Netcool/Impact is incorrectly configured for ObjectServer failover.

To resolve this problem, extra Netcool/Impact configuration is required for ObjectServer failover. To correctly configure Netcool/Impact, complete the steps in the *Managing the OMNIbusEventReader with an ObjectServer pair for New Events or Inserts* topic in the Netcool/Impact documentation: https://

www.ibm.com/support/knowledgecenter/SSSHYH_7.1.0.16/com.ibm.netcoolimpact.doc/common/dita/ts_serial_value_omnibus_eventreader_failover_failback.html

When configured, Netcool/Impact uses the failover ObjectServer to process the event.

Patterns
Use the following troubleshooting information to resolve problems with Event Analytics patterns.

The pattern displays 0 groups and 0 events

The events pattern that is created and displayed in the **Group Sources** table in the View Related Events portlet displays 0 groups and 0 events

The pattern displays 0 groups and 0 events for one of the following reasons.

- The pattern creation process is not finished. The pattern creation process can take a long time to complete due to large datasets and high numbers of suggested patterns.
- The pattern creation process was stopped before it completed.

To confirm the reason that the pattern displays 0 groups and 0 events, complete the following steps.

1. To confirm that the process is running,
   a. Append the policy name to the policy logger file from the **Services** tab, **Policy Logger** service. For more information about configuring the Policy logger, see https://www.ibm.com/support/knowledgecenter/SSSHYH_7.1.0.16/com.ibm.netcoolimpact.doc/user/policy_logger_service_window.html.
   b. Check the following log file.

      ```
      $IMPACT_HOME/logs/<serverName>_policylogger_PG_ALLOCATE_PATTERNS_GROUPS.log
      ```

   If the log file shows that the process is running, wait for the process to complete. If the log file shows that the process stopped without completing, proceed to step 2.
2. To force reallocation for all configurations and patterns run the PG_ALLOCATE_PATTERNS_GROUPS_FORCE from Global projects policy with no parameters from the UI.
3. Monitor the $IMPACT_HOME/logs/<serverName>_policylogger_PG_ALLOCATE_PATTERNS_GROUPS_FORCE.log log file to track the completion of the process.

Event pattern with the same criteria already exists (error message)
An error message is displayed if you create a pattern that has a duplicate pattern criterion selected. Check the following log file to determine which pattern is the duplicate:

```
$IMPACT_HOME/logs/<serverName>_policylogger_PG_SAVEPATTERN.log
```

Performance
Use the following troubleshooting information to resolve performance problems with Event Analytics.

Improving Event Analytics performance due to large search results

If you are performing an upgrade of Event Analytics from an earlier version, the upgrade repopulates the existing data from the previous version and aligns this data with the new schema, tables, and views. It is possible that you might see degradation in the performance of Event Analytics operations. Examples of degradation in performance include but are not limited to:

- Reports can hang.
- Reports complete, but no data is displaying for seasonal events.

To improve any degradation in the performance of Event Analytics operations due to the upgrade to 1.3.1 or later releases, run the SE_CLEANUPDATA policy as follows:

1. Log in to the server where IBM Tivoli Netcool/Impact is stored and running. You must log in as the administrator (that is, you must be assigned the `ncw_analytics_admin` role).
2. Navigate to the policies tab and search for the SE_CLEANUPDATA policy.
3. Open this policy by double-clicking it.
4. Select to run the policy by using the run button on the policy screen toolbar.

The SE_CLEANUPDATA policy cleans up the data. Specifically, the SE_CLEANUPDATA policy:

- Does not remove or delete any data from the results tables. The results tables hold all the original information about the analysis.
- Provides some additional views and tables on top of the original tables to enhance performance.
- Combines some information from related events, seasonal events, rules, and statistics.
- Cleans up only the additional tables and views.

Related Event Details page is slow to load
To avoid this problem, create an index on the Event History Database for the SERVERSERIAL and SERVERNAME columns.

```
create index myServerIndex on Db2INST1.REPORTER_STATUS (SERVERSERIAL , SERVERNAME )
```

It is the responsibility of the database administrator to construct (and maintain) appropriate indexes on the REPORTER history database. The database administrator should review the filter fields for the reports as a basis for an index, and should also review if an index is required for Identity fields.

Export of large Related Event configuration fails

The export a configuration with more then 2000 Related Event groups fails. An error message is displayed.

```
Export failed.
An invalid response was received from the server.
```

To resolve this issue, increase the Java Virtual Machine memory heap size settings from the default values. For Netcool/Impact the default value of the Xmx is 2400 MB. In JVM, Xmx sets the maximum memory heap size. To improve performance, make the heap size larger than the default setting of 2400 MB. For details about increasing the JVM memory heap size, see https://www.ibm.com/support/knowledgecenter/SSSHYH_7.1.0.16/com.ibm.netcoolimpact.doc/admin/imag_monitor_java_memory_status_c.html.

Troubleshooting Operations Management on IBM Cloud Private

Use this information to support troubleshooting of errors that might occur when deploying or using Operations Management on IBM Cloud Private.

Viewing installation logs

View installation logs for information on the success of the installation process.

Viewing Kubernetes logs
To view information on the success of the installation process for Netcool Operations Insight on IBM Cloud Private, run the `kubectl logs` command from the command line.

About this task
There are three levels of detail at which you can report on the progress of pod and container installation:

Displaying pod status
Run the following command to see overall status for each pod.

```
kubectl get pods
```

Run the following command if the namespace is non-default.

```
kubectl get pods --namespace
```

Where *namespace* is the name of the non-default namespace.

Displaying a log file

Run the following command to display log files for a specific pod or container within that pod.

```
kubectl logs name_of_pod [-c name_of_container]
```

The following section lists the relevant commands for the different Operations Management on IBM Cloud Private pods and containers.

Primary ObjectServer

```
kubectl logs name_of_objserv-primary-pod
```

Backup ObjectServer

```
kubectl logs name_of_objserv-backup-pod -c ncobackup-agg-b
```

Failover gateway

```
kubectl logs name_of_objserv-backup-pod -c ncobackup-agg-gate
```

WebGUI

```
kubectl logs name_of_webgui-pod -c webgui
```

Log Analysis

```
kubectl logs name_of_log-analysis-pod -c unity
```

XML gateway

```
kubectl logs name_of_log-analysis-pod -c gateway
```

Primary Impact Server

```
kubectl logs name_of_impactcore-primary-pod -c nciserver
```

Backup Impact Server

```
kubectl logs name_of_impactcore-backup-pod -c nciserver
```

Impact GUI Server

```
kubectl logs name_of_impactgui-pod -c impactgui
```

Db2

```
kubectl logs name_of_db2ese-pod -c db2ese
```

Proxy

```
kubectl logs name_of_proxy-pod
```

OpenLDAP

```
kubectl logs name_of_openLDAP-pod
```

Cloud Native Analytics

```
kubectl logs name_of_cassandra-pod
```

```
kubectl logs name_of_couchdb-pod
```

```
kubectl logs name_of_ea-noi-layer-eanoigateway-pod
```

```
kubectl logs name_of_ea-noi-layer-eanoiactionservice-pod
```

```
kubectl logs name_of_ibm-hdm-analytics-dev-inferenceservice-pod
```

```
kubectl logs name_of_ibm-hdm-analytics-dev-eventsqueryservice-pod
```

```
kubectl logs name_of_ibm-hdm-analytics-dev-archivingservice-pod
```

```
kubectl logs name_of_ibm-hdm-analytics-dev-servicemonitorservice-pod
```

```
kubectl logs name_of_ibm-hdm-analytics-dev-policyregistryservice-pod
```

```
kubectl logs name_of_ibm-hdm-analytics-dev-ingestionservice-pod
```

```
kubectl logs name_of_ibm-hdm-analytics-dev-trainer-pod
```

```
kubectl logs name_of_ibm-hdm-analytics-dev-collater-aggregationservice-pod
```

```
kubectl logs name_of_ibm-hdm-analytics-dev-normalizer-aggregationservice-pod
```

```
kubectl logs name_of_ibm-hdm-analytics-dev-dedup-aggregationservice-pod
```

```
kubectl logs name_of_spark-master-pod
```

```
kubectl logs name_of_spark-slave-pod
```

```
kubectl logs name_of_ea-ui-api-graphql-pod
```

```
kubectl logs name_of_ibm-hdm-common-ui-uiserver-pod
```

```
kubectl logs name_of_kafka-pod
```

```
kubectl logs name_of_zookeeper-pod
```

```
kubectl logs name_of_redis-sentinel-pod
```

```
kubectl logs name_of_redis-server-pod
```

Following a log file

Run the following command to stream a log file for a specific pod or container within that pod.

```
kubectl logs -f name_of_pod [-c name_of_container]
```

The following section lists the relevant commands for the different Operations Management on IBM Cloud Private pods and containers.

Primary ObjectServer

```
kubectl logs -f --tail=1 name_of_objserv-primary-pod
```

Backup ObjectServer

```
kubectl logs -f --tail=1 name_of_objserv-backup-pod -c ncobackup-agg-b
```

Failover gateway

```
kubectl logs -f --tail=1 name_of_objserv-backup-pod -c ncobackup-agg-gate
```

WebGUI

```
kubectl logs -f --tail=1 name_of_webgui-pod -c webgui
```

Log Analysis

```
kubectl logs -f --tail=1 name_of_log-analysis-pod -c unity
```

XML gateway

```
kubectl logs -f --tail=1 name_of_log-analysis-pod -c gateway
```

Primary Impact Server

```
kubectl logs -f --tail=1 name_of_impactcore-primary-pod -c nciserver
```

Backup Impact Server

```
kubectl logs -f --tail=1 name_of_impactcore-backup-pod -c nciserver
```

Impact GUI Server

```
kubectl logs -f --tail=1 name_of_impactgui-pod -c impactgui
```

Db2

```
kubectl logs -f --tail=1 name_of_db2ese-pod -c db2ese
```

Proxy

```
kubectl logs -f --tail=1 name_of_proxy-pod
```

OpenLDAP

```
kubectl logs -f --tail=1 name_of_openLDAP-pod
```

Cloud Native Analytics

```
kubectl logs -f --tail=1 name_of_cassandra-pod
```

```
kubectl logs -f --tail=1 name_of_couchdb-pod
```

```
kubectl logs -f --tail=1 name_of_ea-noi-layer-eanoigateway-pod
```

```
kubectl logs -f --tail=1 name_of_ea-noi-layer-eanoiactionservice-pod
```

```
kubectl logs -f --tail=1 name_of_ibm-hdm-analytics-dev-inferenceservice-pod
```

```
kubectl logs -f --tail=1 name_of_ibm-hdm-analytics-dev-eventsqueryservice-pod
```

```
kubectl logs -f --tail=1 name_of_ibm-hdm-analytics-dev-archivingservice-pod
```

```
kubectl logs -f --tail=1 name_of_ibm-hdm-analytics-dev-servicemonitorservice-pod
```

```
kubectl logs -f --tail=1 name_of_ibm-hdm-analytics-dev-policyregistryservice-pod
```

```
kubectl logs -f --tail=1 name_of_ibm-hdm-analytics-dev-ingestionservice-pod
```

```
kubectl logs -f --tail=1 name_of_ibm-hdm-analytics-dev-trainer-pod
```

```
kubectl logs -f --tail=1 name_of_ibm-hdm-analytics-dev-collater-aggregationservice-pod
```

```
kubectl logs -f --tail=1 name_of_ibm-hdm-analytics-dev-normalizer-aggregationservice-pod
```

```
kubectl logs -f --tail=1 name_of_ibm-hdm-analytics-dev-dedup-aggregationservice-pod
```

```
kubectl logs -f --tail=1 name_of_spark-master-pod
```

```
kubectl logs -f --tail=1 name_of_spark-slave-pod
```

```
kubectl logs -f --tail=1 name_of_ea-ui-api-graphql-pod
```

```
kubectl logs -f --tail=1 name_of_ibm-hdm-common-ui-uiserver-pod
```

```
kubectl logs -f --tail=1 name_of_ea-ui-api-graphql-pod
```

```
kubectl logs -f --tail=1 name_of_kafka-pod
```

```
kubectl logs -f --tail=1 name_of_zookeeper-pod
```

```
kubectl logs -f --tail=1 name_of_redis-sentinel-pod
```

```
kubectl logs -f --tail=1 name_of_redis-server-pod
```

Related information

Kubernetes documentation: kubectl command referenceThis documentation lists all Kubernetes commands and provides examples.

View not displayed

Panels are not displayed for incompatible views in the Operations Management on IBM Cloud Private UI.

Problem

Panels are hidden if the selected view does not contain specific columns. The Investigate Incident, Temporal Group, and Seasonal Event pages require specific columns to be added to a view for the UI to function correctly.

Resolution

Add the following columns to your view:

- **CEACorrelationDetails**
- **CEASeasonalDetails**

ImageRepository field is empty

The **ImageRepository** field in the IBM Cloud Private UI is not pre-populated.

Problem

If you are installing on IBM Cloud Private with OpenShift, the **ImageRepository** field might be empty.

Resolution

Use the following command to help you find the value that you need to supply.

```
oc get route docker-registry
```

Use the value of host that is returned by this command, with the namespace of your Operations Management on IBM Cloud Private deployment appended as the value for **ImageRepository**.

Install timeout error

How to debug the cause of an installation timeout.

About this task

When you attempt an installation of Operations Management on IBM Cloud Private and receive a timeout error, use the following procedure to help you to debug the cause of the error.

```
helm install . -f ~/values.yaml --tls --name my-noi-release
Error: timed out waiting for the condition
```

Procedure

1. Find the tiller pod name, and search for entries about the failed install in the tiller pod's logs, as in the following example where the failed deployment is called *my-noi-release*.

   ```
   kubectl get pod -n kube-system | grep tiller
   tiller-deploy-5d8494fb8-ggrzl                                    1/1       Running
   1          91d
   kubectl logs tiller-deploy-5d8494fb8-ggrzl -n kube-system | grep my-noi-release

   [tiller] 2019/06/06 09:14:43 warning: Release my-noi-release pre-install ibm-netcool-prod/
   templates/preinstallhook.yaml could not complete: timed out waiting for the condition
   [tiller] 2019/06/06 09:14:43 deleting pre-install hook my-noi-release-verifysecrets for
   release my-noi-release due to "hook-failed" policy
   ```

   The pre-install hook job *my-noi-release-verifysecrets* failed, and was deleted.
2. Attempt the installation again, as in the following example:

   ```
   helm install . -f ~/values.yaml --tls  --name my-noi-release
   ```

   While this command is running, use another shell to monitor the *my-noi-release-verifysecrets* job for problems:

   ```
   kubectl describe job my-noi-release-verifysecrets
   Warning  FailedCreate  8s (x2 over 19s)  job-controller  Error creating: pods "my-noi-
   release-verifysecrets-" is forbidden: error looking up service account testnew/noi-service-
   account: serviceaccount "noi-service-account" not found
   ```

   The warning message shows that the expected service account, *noi-service-account*, was not found.
3. To see where this service account is created, search the charts in the Operations Management on IBM Cloud Private deployment directory for ServiceAccount, as in the following example:

   ```
    find . -type f | xargs grep "^kind: ServiceAccount"
   ./templates/serviceaccount.yaml:kind: ServiceAccount
   ```

   and then view that file:

   ```
   view ./templates/serviceaccount.yaml
   {{ if .Values.global.rbac.create }}
   {{- if (not (eq .Values.global.rbac.serviceAccountName "default" )) }}
   {{- include "sch.config.init" (list . "sch.chart.config.values") -}}
   ```

   From this chart, it can be seen that the service account is created in `./templates/serviceaccount.yaml` only if *Values.global.rbac.create* is true. To get the installation to complete

successfully, the service account must either be manually created before installation, or *Create required RBAC RoleBindings* (*global.rbac.create* in the ibm.netcool `values.yaml` file) must be set to true. For more information, see #unique_146/unique_146_Connect_42_create_rbac_id.

Data fetch error

When you fetch seasonality data in the **Incident Viewer**, an error is displayed.

Problem

The following error occurs when fetching seasonality data in the **Incident Viewer**.

```
An error occurred while fetching data from the server. Please make sure you have an active
internet connection.
Code: FETCH_ERROR_NETWORK
```

This error happens when there is a cookie conflict with the IBM WebSphere Application Server cookie, LTPAToken2.

Resolution

To work around this issue, clear your browser cookies.

New view not available

When you add a view in the **Event Viewer**, the view should be available in the drop-down list.

Problem

After you add a new view and click investigate on a parent event in the **Event Viewer**, the new view is not displayed in the view drop-down list in the **Incident Viewer**.

Resolution

This is a known issue. To display a new view in the list, it must contain all of the following fields as display columns, sort columns, or relationships:

- ParentIdentifier
- Identifier
- CEACorrelationDetails
- CEASeasonalDetails

NoHostAvailable error

When you restart all cassandra pods with the **`kubectl delete pod`** command, they should be available with no errors.

Problem

After you restart all cassandra pods, log in to cassandra, and run a query, the `NoHostAvailable` error is displayed.

Resolution

List the cassandra pods and restart one of them, as in the following example:

```
kubectl get pod |grep cass
noi-cassandra-0                                                   1/1     Running
0          75m
noi-cassandra-1                                                   1/1     Running
0          2m6s
noi-cassandra-2                                                   1/1     Running
```

```
0          76m
kubectl delete pod noi-cassandra-1
```

Datasources are not persistent

When Netcool Operations Insight is running on IBM Cloud Private, and you create a new datasource, on restart of the Dashboard Application Services Hub container (`webgui`), the datasource is no longer present.

About this task

After you create a new datasource by clicking **Administration** > **Datasources** in the Dashboard Application Services Hub GUI, when you restart the Dashboard Application Services Hub container (`webgui`), the datasource is no longer present.

Procedure

If you want to visualize the contents of an on-premises ObjectServer in a Web GUI instance running under IBM Cloud Private, then set up a gateway between your on-premises ObjectServer and the ObjectServer running under IBM Cloud Private.

Customizations to the default Netcool/Impact are not persisted

When Netcool Operations Insight is running on IBM Cloud Private, any customizations to the default Netcool/Impact connection are not persisted. The default connection is recreated every time the `webgui` container is restarted.

Procedure

To connect with a customized Netcool/Impact connection, you must create a new connection. Any additional Netcool/Impact connection that you create is preserved.

Note: This workaround requires you to connect to an external on-premises Netcool/Impact connection and not to the Netcool/Impact container within the IBM Cloud Private cluster.

Communication between the proxy server and the IBM Cloud Private object server drops

The proxy server and the IBM Cloud Private object server are disconnected after 15 minutes.

Problem

The proxy server and IBM Cloud Private object server are disconnected.

Resolution

The default timeout value for communication to the IBM Cloud Private object server is 15 minutes. You can modify the default value through the proxy server configmap. Set the **`connectionTimeoutMs`** value in milliseconds.

Restart of all Cassandra pods causes errors for connecting services

When all the Cassandra pods go down simultaneously, the following error is displayed by the Cloud Native Analytics user interface when the pods come back up:

```
An error occurred while fetching data from the server. The response from the server was '500'.
Please try again later.
```

`kubectl get events` also outputs a warning:

```
Warning FailedToUpdateEndpoint Endpoints Failed to update endpoint
```

Resolution
Use the following procedure to resolve this problem.

1. Scale Cassandra down to 0 instances with this command:

   ```
   kubectl scale --replicas=0 StatefulSet/helm_release_name-cassandra
   ```

   Where *helm_release_name* is the name of your Operations Management on IBM Cloud Private Helm release.
2. Use `kubectl get pods | grep cass` to verify that there are no Cassandra pods running.
3. Scale Cassandra back up to one instance.

   ```
   kubectl scale --replicas=1 StatefulSet/helm_release_name-cassandra
   ```

   Where *helm_release_name* is the name of your Operations Management on IBM Cloud Private Helm release.
4. Use `kubectl get pods | grep cass` to verify that there is one Cassandra pod running.
5. Repeat step 3, incrementing *replicas* each time until the required number of Cassandra pods are running. Wait for each Cassandra pod to come up before incrementing the replica count to start another.
6. Verify that the cluster is running with this command:

   ```
   kubectl exec -ti helm_release_name-cassandra-0 bash
   [cassandra@m86-cassandra-0 /]$ nodetool status
   ```

   where *helm_release_name* is the name of your Operations Management on IBM Cloud Private Helm release.
   Expect to see UN for all nodes in cluster, as in this example:

   ```
   Datacenter: datacenter1

   Status=Up/Down
   |/ State=Normal/Leaving/Joining/Moving
   -- Address Load Tokens Owns (effective) Host ID Rack
   UN 10.1.26.101 598.28 KiB 256 100.0% bbd34cab-9e91-45c1-bfcb-1fe59855d9b3 rack1
   UN 10.1.150.13 654.78 KiB 256 100.0% 555f00c8-c43d-4962-a8a0-72eed028d306 rack1
   UN 10.1.228.111 560.28 KiB 256 100.0% 8741a69b-acdb-4736-bc74-905d18ebdafa rack1
   ```
7. Restart the pods that connect to Cassandra with the following commands

   ```
   kubectl delete helm_release_name-ibm-hdm-analytics-dev-policyregistryservice
   kubectl delete helm_release_name-ibm-hdm-analytics-dev-eventsqueryservice
   kubectl delete helm_release_name-ibm-hdm-analytics-dev-archivingservice
   ```

   Where *helm_release_name* is the name of your Operations Management on IBM Cloud Private Helm release.
8. Relogin to the UI.

Load helm charts fails

Problem

When the helm charts are loaded, the following error is displayed:

```
You have not accepted the terms of the license agreement. You must review and accept the terms
by setting license=accept.
```

Resolution

Ensure that *license=accept* is used when you run the `load-helm-chart` command, as in the following example:

```
cloudctl catalog load-helm-chart --archive charts/ibm-netcool-prod-2.1.1.tgz license=accept --
registry cluster-CA-domain:8500/namespace
```

where *cluster-CA-domain* is the name of the certificate authority domain that was used in the `config.yaml` file during IBM Cloud Private installation.

StatefulSet pods with local storage are stuck in Pending state

Problem

If you are using local storage and a node in the cluster goes down, then StatefulSet pods remain bound to that node. These pods are unable to restart on another node and are stuck in `Pending` state because they have a Persistent Volume on the node that is down.

Resolution

The persistent volumes and persistent volume claims for these pods must be removed to allow the pods to be reassigned to other nodes. There is a script in the ITOM Developer Center that can be used to do this: http://developer.ibm.com/itom/2019/10/31/cleanup_local_storage

Run the script with the following command:

```
./cleanPVCandPV.sh
```

Unable to add new groups using WebSphere Application Server

About this task

The following error message is displayed if you try to add a new group using WebSphere Application Server console:

```
CWWIM4520E The 'javax.naming.directory.SchemaViolationException: [LDAP: error code 65 - object
class 'groupOfNames' requires attribute 'member']
```

As a workaround, create the group in LDAP instead using the following procedure.

Procedure

1. Log in to the LDAP Proxy Server pod.

   ```
   kubectl exec -it helm_release_name-openldap-0 /bin/bash
   ```

   Where *helm_release_name* is the name of your Operations Management on IBM Cloud Private Helm release.
2. Create the new group
   a) Create an LDAP Data Interchange Format file to define the new group.

      For example:

      ```
      vi test-group.ldif
      ```

   b) Define the contents of the LDIF file that you created by using a format similar to this example:

      ```
      dn: cn=newgroup,ou=groups,dc=mycluster,dc=icp
      cn: newgroup
      owner: uid=newgroup,ou=users,dc=mycluster,dc=icp
      description: newgroup test
      ```

```
objectClass: groupOfNames
member: uid=icpadmin,ou=users,dc=mycluster,dc=icp
```

Where:

- the value of `uid` and `cn` are the name of the new group
- the value of *dc* is the domain components that were specified for the suffix and baseDN. By default the value of this parameter is `dc=mycluster,dc=icp`.

c) Run the following command to create the new group

```
ldapadd -h localhost  -p  389 -D "cn=admin,dc=mycluster,dc=icp" -w password -f ./test-
group.ldif
```

New user does not inherit roles from assigned group

About this task

When a new user is created and added to a group with WebSphere Application Server Console, the user is not assigned the roles for that group as they should be. You must add the roles that are required for that user with Web GUI.

Procedure

1. Select **Console Settings->User Roles** and search for your user in **Available Users.**
2. Select the new user with the missing roles from the displayed results, and then select the required roles for your user.

Cannot launch WebSphere Application Server from DASH on RHOCP environment

About this task

If your deployment is on an OpenShift environment and you attempt to access **Console Settings ->Websphere Application Server Console**, then the following error is returned:

```
"502 Bad Gateway The server returned an invalid or incomplete response."
```

Use the following procedure as a workaround.

Procedure

Access WebSphere Application Server directly with a URL in this format:

```
https://was.helm_release_name.master_node_name:<port>/ibm/console
```

Where

- *helm_release_name* is the helm release name of the Operations Management on IBM Cloud Private installation.
- *master_node_name* is the hostname of the master node.
- *port* is the value that you specified for *ingress_https_port* in your configuration yaml file when you installed ICP.

Troubleshooting Event Search

Use this information to support troubleshooting of errors that might occur when deploying or using Event Search.

Troubleshooting event search

How to resolve problems with your event search configuration.

You must log in each time you switch between interfaces

The problem occurs if single sign-on (SSO) is not configured. If the Web GUI and Operations Analytics - Log Analysis are on the same host computer, you must log in each time you switch between the interfaces in your browser.

This problem happens because each instance of WebSphere Application Server uses the same default name for the LTPA token cookie: `LtpaToken2`. When you switch between the interfaces, one WebSphere Application Server instance overwrites the cookie of the other and your initial session is ended.

The ways of resolving this problem are as follows:

- Customize the domain name in the Web GUI SSO configuration:
  1. In the administrative console of the WebSphere Application Server that hosts the Web GUI, click **Security** > **Global security**. Then, click **Authentication** > **Web security** and click **Single sign-on (SSO).**.
  2. Enter an appropriate domain name for your organization, for example, `abc.com`. By default, the domain name field is empty and the cookie's domain is the host name. If you also customize the domain name in the Operations Analytics - Log Analysis WebSphere Application Server, to avoid any conflict ensure that the two domain names are different.
  3. Restart the Dashboard Application Services Hub server.
- Use the fully qualified domain name for accessing one instance of WebSphere Application Server and the IP address for accessing the other. For example, always access the Web GUI by the fully qualified domain name and always access Operations Analytics - Log Analysis by the IP address. To configure the Web GUI to access Operations Analytics - Log Analysis by the IP address:
  1. In the *$WEBGUI_HOME*`/etc/server.init` file, change the value of the **`scala.url`** property to the IP address of the host, For example:

     ```
     https://3.127.46.125:9987/Unity
     ```

  2. Restart the Dashboard Application Services Hub server. See http://www-01.ibm.com/support/knowledgecenter/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/webtop/wip/task/web_adm_server_restart.html.

Operations Analytics - Log Analysis session times out after 2 hours

This problem occurs if SSO is not configured. The first time that you start the Operations Analytics - Log Analysis product from an event list in the Web GUI, you are prompted to log in to Operations Analytics - Log Analysis. You are automatically logged out after 2 hours and must reenter your login credentials every 2 hours. This problem occurs because the default expiration time of the LTPA token is 2 hours.

To resolve this problem, change the session timeout in the Operations Analytics - Log Analysis product as follows:

1. In the $*SCALA_HOME*/wlp/usr/servers/Unity/server.xml file, increase the value of the <ltpa expiration="120m"/> attribute to the required value, in minutes. For example, to change the session timeout to 540 minutes:

```
</oauthProvider>
  <ltpa expiration="540"/>
  <webAppSecurity ssoDomainNames="hostname" httpOnlyCookies="false"/>
</server>
```

2. Restart the Operations Analytics - Log Analysis WebSphere Liberty Profile.

Launch to Operations Analytics - Log Analysis fails on Firefox in non-English locales

This problem is a known issue when you launch from the Active Event List (AEL) into the Firefox browser.

If your browser is set to a language other than US English (en_us) or English (en), you might not be able to launch into Operations Analytics - Log Analysis from the Web GUI AEL.

This problem happens because Operations Analytics - Log Analysis does not support all the languages that are supported by Firefox.

To work around this problem, try setting your browser language to an alternative language version. For example, if the problem arises when the browser language is French[fr], set the language to French[fr-fr]. If the problem arises when the browser language is German[de-de], set the language to German[de].

Right-click tool fail to start Operations Analytics - Log Analysis from event lists

The following error is displayed when you start the tools from the right-click menu of an event list:

```
CTGA0026E: The APP name in the query is invalid or it does not exist
```

This error occurs because the custom app that is defined in the $*WEBGUI_HOME*/etc/server.init file does not match the file names in the Tivoli Netcool/OMNIbus Insight Pack.

To resolve this problem, set the **scala.app.keyword** and **scala.app.static.dashboard** properties in the server.init file accordingly.

- If the properties are set as follows, the version of the Insight Pack needs to be V1.1.0.2:

```
scala.app.keyword=OMNIbus_Keyword_Search
scala.app.static.dashboard=OMNIbus_Static_Dashboard
```

- If the properties are set as follows, the version of the Insight Pack needs to V1.1.0.1 or V1.1.0.0:

```
scala.app.keyword= OMNIbus_SetSearchFilter
scala.app.static.dashbaord=OMNIbus_Event_Distribution
```

If you need to change the values of these properties, restart the Dashboard Application Services Hub server afterwards. See http://www-01.ibm.com/support/knowledgecenter/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/webtop/wip/task/web_adm_server_restart.html.

Error message displayed when dynamic dashboard is run

The following error is displayed when you run a dynamic dashboard from the Operations Analytics - Log Analysis product:

```
undefined not found in results data
```

This error is a known defect in the Operations Analytics - Log Analysis product. To resolve, it close and then reopen the dynamic dashboard.

Error message displayed on "Show event dashboard by node" tool from event lists

An error message is displayed when you start the **Show event dashboard by node** tool from an event list.

This error is caused by incompatibility between the version of the Insight Pack and the and the version of the Operations Analytics - Log Analysis product.

Ensure that the versions are compatible. See "Required products and components" on page 217. For more information about checking which version of the Insight Pack is installed, see "Checking the version of the Insight Pack" on page 165.

Chart display in Operations Analytics - Log Analysis changes without warning

The sequence in which charts are displayed on the Operations Analytics - Log Analysis GUI can changes intermittently. This problem is a known defect in the Operations Analytics - Log Analysis product and has no workaround or solution.

addIndex script error: Unexpected ant version

When you are running the `addIndex` script to create or update a data source type, the script fails with an error that contains text similar to the following: `installation does not contain expected ant version`.

```
./addIndex.sh -i
Prompt: installation does not contain expected ant version
```

This error could be caused by any of the following issues:

- Operations Analytics - Log Analysis is not installed on the machine on which you are running the `addIndex` script.
- Operations Analytics - Log Analysis is installed on the machine on which you are running the `addIndex` script, but the script does not support the version of Operations Analytics - Log Analysis.

To resolve this run the `addIndex` script on a machine where Operations Analytics - Log Analysis is installed. Ensure that Operations Analytics - Log Analysis V1.3.5 is installed.

addIndex script error: Duplicate fields in template file

When you are running the `addIndex` script to create or update a data source type, the script fails with an error that contains text similar to the following: `filesets.json does not exist`.

```
./addIndex.sh -i
....
generatebasedsv:
      [exec] Duplicate field name: field_name
....
BUILD FAILED
filepath/addIndex.xml:68:  Replace: source file filepath/
data_source_type_nameInsightPack_v1.3.1.0/metadata/filesets.json
does not exist
```

Where:

- *field_name* is the name of the duplicate field name.
- *filepath* is the system-dependent path to the script.
- *data_source_type_name* is the name of your custom data source type.

The `omnibus1100_template.properties` template file is used to define the fields in the data source type. This error is caused by definition of duplicate fields in the `omnibus1100_template.properties` file.

To resolve this, edit the `omnibus1100_template.properties` file and remove any duplicate fields. Then rerun the `addIndex` script.

Related tasks

Configuring single sign-on for the event search capability

Configure single sign-on (SSO) between Web GUI and Operations Analytics - Log Analysis so that users can switch between the two products without having to log in each time.

Searching for events

Checking the version of the Insight Pack

To ensure compatibility between the versions of the Tivoli Netcool/OMNIbus Insight Pack, the Web GUI and the Operations Analytics - Log Analysis product, run the **pkg_mgmt** command to check which version of the Insight Pack is installed.

Restarting the Dashboard Application Services server

WebSphere Application Server Liberty: Starting and stopping a server from the command prompt

Chapter 15. Reference

Reference information for Netcool Operations Insight.

Accessibility features for Netcool Operations Insight

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

The following Netcool Operations Insight components have accessibility features.

- Jazz for Service Management
- Netcool/OMNIbus Web GUI
- Netcool/Impact
- IBM Tivoli Network Manager

See the relevant product Knowledge Centers for more detailed information on accessibility for each component.

Related information

Jazz for Service Knowledge CenterClick here and search for "accessibility" to retrieve more details on accessibility in Jazz for Service Management.

Netcool/OMNIbus WebGUI Knowledge Center: accessibilityClick here for more details on accessibility in Netcool/OMNIbus Web GUI.

Network Manager Knowledge Center: accessibilityClick here for more details on accessibility in Network Manager.

Netcool/Impact Knowledge CenterClick here and search for "accessibility" to retrieve more details on accessibility in Netcool/Impact.

Scope-based grouping

Events in a scope-based group are grouped together because they share a common attribute, such as a resource. For more information about scope-based grouping, see https://www.ibm.com/support/knowledgecenter/SSSHTQ_8.1.0/com.ibm.netcool_OMNIbus.doc_8.1.0/use/omn_con_concept_extsbgeventswithea.html.

Configmap reference

This section lists the pods that have configmaps and explains which parameters you can configure in each configmap.

Related tasks

Customizing applications using config maps

Customize the one or more applications by editing the relevant config map and restarting the associated pod.

Primary Netcool/OMNIbus ObjectServer configmap

This topic explains the structure of the configmap for the primary IBM Tivoli Netcool/OMNIbus ObjectServer pod, `ncoprimary`, and lists the data elements that can be configured with this configmap.

Edit this configmap to customize, and add custom automations and triggers to, the primary Netcool/OMNIbus ObjectServer.

Contents

The following table lists the data elements that are contained in the primary Netcool/OMNIbus ObjectServer configmap:

Table 131. Data elements in the primary Netcool/OMNIbus ObjectServer configmap

Data elements	Description	More information
`agg-p-props-append`	Properties that are specified in this data element are appended to the end of the Netcool/OMNIbus ObjectServer properties file on pod restart.	Netcool/OMNIbus V8.1 documentation: Using the ObjectServer properties and command-line options
`agg-p-sql-extensions`	Use this element to add a new SQL extension, such as a trigger or an automation, to the Netcool/OMNIbus ObjectServer on pod restart.	Netcool/OMNIbus V8.1 documentation: ObjectServer SQL

Examples of each of the data elements in this configmap are provided.

Data element: `agg-p-props-append`

The following data element appends a `MessageLevel: 'debug'` property to the `.props` file of the Primary ObjectServer.

```
agg-p-props-append: |
  MessageLevel: 'debug'
```

Data element: `agg-p-sql-extensions`

The following data element adds a custom database that contains a single table to the primary Netcool/OMNIbus ObjectServer.

```
agg-p-sql-extensions: |
  -- create a custom db
  create database mydb;
  go

  -- create a custom table
  create table mydb.mytable persistent
  (
  col1 incr primary key,
  col2 varchar(255)
  );
  go
```

Backup Netcool/OMNIbus ObjectServer configmap

Learn about the structure of the configmap for the backup IBM Tivoli Netcool/OMNIbus ObjectServer pod, `ncobackup-agg-b`. Edit this configmap to customize, and add custom automations and triggers to, the backup Netcool/OMNIbus ObjectServer, and to customize the operation of the bidirectional gateway that connects the primary and backup Netcool/OMNIbus ObjectServers.

Contents

The following table lists the data elements that are contained in the backup Netcool/OMNIbus ObjectServer configmap:

Table 132. Data elements in the backup Netcool/OMNIbus ObjectServer configmap		
Data elements	**Description**	**More information**
`agg-b-props-append`	Use this element to append a new property to the end of the Netcool/OMNIbus ObjectServer properties file on pod restart.	Netcool/OMNIbus V8.1 documentation: Using the ObjectServer properties and command-line options
`agg-b-sql-extensions`	Use this element to add a new SQL extension, such as a trigger or an automation, to the Netcool/OMNIbus ObjectServer on pod restart.	Netcool/OMNIbus V8.1 documentation: ObjectServer SQL
`agg-gate-map-replace`	The gateway map definition in this element replaces the definition in the pod (AGG_GATE.map) on pod restart.	Netcool/OMNIbus V8.1 documentation: Failover configuration
`agg-gate-props-append`	Properties that are listed in this element are appended to the gateway properties file prior on pod restart.	Netcool/OMNIbus V8.1 documentation: Common gateway properties and command-line options
`agg-gate-startup-cmd-replace`	The gateway startup command definition in this element replaces the definition in the pod (AGG_GATE.startup.cmd) on pod restart.	Netcool/OMNIbus V8.1 documentation: Startup command file
`agg-gate-tblrep-def-replace`	The gateway table replication definition in this element replaces the definition in the pod (AGG_GATE.tblrep.def) on pod restart.	Netcool/OMNIbus V8.1 documentation: Table replication definition file

Examples of each of the data elements in this configmap are provided.

Data element: `agg-b-props-append`

The following data element appends a `MessageLevel: 'debug'` property to the `.props` file of the backup Netcool/OMNIbus ObjectServer.

```
agg-b-props-append: |
  MessageLevel: 'debug'
```

Data element: `agg-b-sql-extensions`

The following data element adds a custom database containing a single table to the backup Netcool/OMNIbus ObjectServer.

```
agg-b-sql-extensions: |
  -- create a custom db
  create database mydb;
  go

  -- create a custom table
  create table mydb.mytable persistent
  (
```

```
col1 incr primary key,
col2 varchar(255)
);
go
```

Data element: agg-gate-map-replace

The following data element replaces the definition in the pod (AGG_GATE.map).

```
agg-gate-map-replace: |
  # My test map
  CREATE MAPPING StatusMap
  (
  'Identifier' = '@Identifier' ON INSERT ONLY,
  'Node'       = '@Node' ON INSERT ONLY
  );
```

Data element: agg-gate-props-append

The following data element is appended to the gateway properties file before startup.

```
agg-gate-props-append: |
  MessageLevel: 'debug'
agg-gate-startup-cmd-replace: |
  SHOW PROPS;
```

Data element: agg-gate-startup-cmd-replace

The gateway startup command definition in this data element replaces the definition in the pod (AGG_GATE.startup.cmd).

```
agg-gate-startup-cmd-replace: |
  SHOW PROPS;
```

Data element: agg-gate-tblrep-def-replace

The Gateway table replication definition in this data element replaces the definition in the pod (AGG_GATE.tblrep.def).

```
agg-gate-tblrep-def-replace: |
  # My test table replication
  REPLICATE ALL FROM TABLE 'alerts.status'
    USING MAP 'StatusMap';
```

Netcool/Impact core server configmap

Learn about the structure of the configmap for the Netcool/Impact core server pod, `nciserver`. Edit this configmap to customize the properties, logging features, and memory status monitoring for the primary and backup Netcool/Impact core server pods. You can also customize the Derby database that runs inside the Netcool/Impact server pod.

Contents

The following table lists the data elements that are contained in the Primary Netcool/Impact core server configmap:

Table 133. Data elements in the Primary Netcool/Impact core server configmap		
Data elements	**Description**	**More information**
`impactcore-server-props-update`	Properties that are listed in this data element are used to update the `NCI_*_server.props` file on pod restart.	To find information on the parameters that can be configured by using the `NCI_server.props` file, go to the Netcool/Impact documentation Welcome page and search for "NCI_server.props". Scroll down to the properties that are of interest to you.
`impactcore-log4j-props-update`	Properties that are listed in this data element are used to update the `impactserver-log4j.properties` file on pod restart.	Netcool/Impact documentation: Log4j properties files
`impactcore-jvm-options-replace`	Properties that are listed in this data element are used to replace the properties in the `jvm.options` file on pod restart.	Netcool/Impact documentation: Memory status monitoring
`impactcore-derby-sql-extension`	The SQL recorded in this data element is applied to the Netcool/Impact Derby database on pod restart.	Netcool/Impact documentation: Managing the database server

Examples of each of the data elements in this configmap are provided.

Data element: `impactcore-server-props-update`

Properties that are listed in this data element are used to update the `NCI_*_server.props` file on pod restart.

```
  impactcore-server-props-update: |
    impact.server.timeout=123456
    impact.servicemanager.storelogs=false
    # the following should just be appended
    fred=banana
```

Data element: `impactcore-log4j-props-update`

Properties that are listed in this data element are used to update the `impactserver-log4j.properties` file on pod restart.

```
  impactcore-log4j-props-update: |
    log4j.rootCategory=DEBUG

    log4j.appender.NETCOOL=org.apache.log4j.RollingFileAppender
    log4j.appender.NETCOOL.threshold=DEBUG
    log4j.appender.NETCOOL.layout=org.apache.log4j.PatternLayout
    log4j.appender.NETCOOL.layout.ConversionPattern=%d{DATE} %-5pC3PO [%c{1}] %m%n
    log4j.appender.NETCOOL.append=false
    log4j.appender.NETCOOL.file=/opt/IBM/tivoli/impact/logs/impactserver.log
    log4j.appender.NETCOOL.bufferedIO=true
    log4j.appender.NETCOOL.maxBackupIndex=4
    log4j.appender.NETCOOL.maxFileSize=21MB
```

Data element: `impactcore-jvm-options-replace`

Properties that are listed in this data element are used to replace the properties in the `jvm.options` file on pod restart.

```
  impactcore-jvm-options-replace: |
    -Xms512M
    -Xmx4096Mdd
    -Dclient.encoding.override=UTF-8
    -Dhttps.protocols=SSL_TLSv2
    #-Xgc:classUnloadingKickoffThreshold=100
    -Dcom.ibm.jsse2.overrideDefaultTLS=true
```

Data element: `impactcore-derby-sql-extension`

The SQL recorded in this data element is applied to the Netcool/Impact Derby database on pod restart.

```
impactcore-derby-sql-extensions: |
    CREATE SCHEMA MYSCHEMA;
    SET SCHEMA MYSCHEMA;

    CREATE TABLE MYTABLE (
      keyvalue character varying (256),
      value character varying (256)
    );
    INSERT INTO MYTABLE VALUES ('mykey1', 'myvalue1');
```

Netcool/Impact GUI server configmap

Learn about the structure of the configmap for the IBM Tivoli Netcool/Impact GUI server pod, `impactgui`. Edit this configmap to customize the logging features of the Netcool/Impact GUI server, and to customize the Derby database that runs inside the Netcool/Impact server.

Contents

The following table lists the data elements that are contained in the Netcool/Impact GUI server configmap:

Table 134. Data elements in the Netcool/Impact GUI server configmap

Data elements	Description	More information
`server-props-update`	Properties that are listed in this data element are used to update the `server.props` file on pod restart.	To find information on the parameters that can be configured by using the `server.props` file, go to the Netcool/Impact documentation Welcome page and search for "server.props". Scroll down to the properties that are of interest to you.
`impactcore-log4j-props-update`	Properties that are listed in this data element are used to update the `impactserver-log4j.properties` file on pod restart.	Netcool/Impact documentation: Log4j properties files

Examples of each of the data elements in this configmap are provided.

Data element: `server-props-update`

Properties that are listed in this data element are used to update the `server.props` file on pod restart.

```
        server-props-update: |
          impact.cluster.network.call.timeout=60
```

Data element: `impactcore-log4j-props-update`

Properties that are listed in this data element are used to update the `impactserver-log4j.properties` file on pod restart.

```
  impactcore-log4j-props-update: |
    log4j.rootCategory=DEBUG
```

Proxy configmap

Edit this configmap to configure parameters such as connection timeouts and to enable or disable transport layer security (TLS) encryption.

Contents

The proxy is configured with the *{{ .Release.Name }}*`-proxy-config` configmap, where *{{ .Release.Name }}* is the unique name that is assigned to the deployment, for example, **noi**. The configmap is used to configure parameters such as connection timeouts and to enable or disable TLS encryption, as in the following example:

```
  connectionTimeoutMs: "900000"
  externalHostOrIP: mycluster.icp
   revision: "1"
  routes: |
    [{"Port":6001, "Service": "{{ .Release.Name }}-objserv-agg-primary:4100"},
     {"Port":6002, "Service": "{{ .Release.Name }}-objserv-agg-backup:4100"}
    ]
  tlsEnabled: "true"
  tlsHandshakeTimeoutMs: "2000"
```

Where *mycluster.icp* is the public domain name of the cluster.

The following table lists the data elements that are contained in the proxy configmap:

Table 135. Data elements in the proxy configmap

Data elements	Description
`connectionTimeoutMs`	Use this element to specify the connection timeout in milliseconds.
`externalHostOrIP`	Use this element to specify the external Host or IP address. After deployment, do not edit this value.
`revision`	Version number of configmap. Must be incremented when a change is made to the configmap, or the change will not be active. Must be an integer greater than 0.
`routes`	Use this element to specify the gateway routes. After deployment, do not edit this value.
`tlsEnabled`	Use this element to enable or disable TLS encryption.

Table 135. Data elements in the proxy configmap (continued)	
Data elements	**Description**
`tlsHandshakeTimeoutMs`	Use this element to specify the TLS handshake timeout in milliseconds.

After you update the configmap the changes are automatically applied to the existing proxy pod.

LDAP Proxy configmap

ldap_proxy_configmap is the configmap for the LDAP proxy pod, `openldap`. Edit this configmap to configure connections to your own LDAP server when you have *LDAPmode* set to `proxy` rather than `standalone`. This configmap is not used when *LDAPmode* is set to `standalone`. For more information, see "Configuring Installation Parameters for Operations Management on IBM Cloud Private" on page 125.

Contents

The following table lists the data elements that are contained in the `openldap` configmap:

Table 136. Data elements in the `openldap` configmap		
Data elements	**Description**	**More information**
`ldap-proxy-slapd-replace:`	Replaces the contents of the `slapd.conf` file, which configures the connection to your LDAP server.	"Configuring Installation Parameters for Operations Management on IBM Cloud Private" on page 125

Dashboard Application Services Hub configmap

Learn about the structure of the configmap for the Dashboard Application Services Hub pod, `webgui`. Edit this configmap to customize the properties of the Web GUI Event Viewer.

Contents

The following table lists the data elements that are contained in the Dashboard Application Services Hub configmap :

Table 137. Data elements in the Dashboard Application Services Hub configmap		
Data elements	**Description**	**More information**
`server-init-update`	Properties that are listed in this data element are used to overwrite the environmental and server session properties of the Web GUI server that are stored in the `server.init` initialization file.	Netcool/OMNIbus documentation: server.init properties

Examples of each of the data elements in this configmap are provided.

Data element: `server-init-update`

Properties that are listed in this data element are used to overwrite the environmental and server session properties of the Web GUI server that are stored in the `server.init` initialization file.

```
  server-init-update: |
    eventviewer.pagesize.max:20000

columngrouping.allowedcolumns=Acknowledged,AlertGroup,Class,Customer,Location,Node,NodeAlias,Nmo
```

```
sCauseType,NmosManagedStatus,Severity,Service
    columngrouping.maximum.columns:3

alerts.status.sort.displayvalue=Acknowledged,Class,ExpireTime,Flash,NmosCauseType,NmosManagedSta
tus,OwnerGID,OwnerUID,SupressEscl,TaskList,Type,X733EventType,X733ProbableCause
    dashboard.edit.render.mode:applet
    dashboard.render.mode:applet
    webtop.keepalive.interval:3
    datasource.failback.delay:120
    users.global.filter.mode:1
    users.global.view.mode:1
    users.group.filter.mode:1
```

Gateway for Message Bus configmap

Learn about the structure of the configmap for the Gateway for Message Bus. This configmap is associated with the `scala` pod. Edit this configmap to customize the properties of the Gateway for Message Bus, which defines which data is transferred from the Netcool/OMNIbusObjectServer to the Operations Analytics - Log Analysis to support the Event Search capability.

Contents

The following table lists the data elements that are contained in the Gateway for Message Bus configmap :

Table 138. Data elements in the Gateway for Message Bus configmap

Data elements	Description	More information
`xml-gate-props-append`	Properties that are listed in this data element are appended to XML gateway properties file on pod startup.	Netcool/OMNIbus documentation: Gateway for Message Bus properties file
`xml-gate-map-replace`	Properties that are listed in this data element replace the gateway map definition in the `LA_GATE.map` file on pod startup.	Netcool/OMNIbus documentation: Gateway for Message Bus map definition file
`xml-gate-tblrep-def-replace`	Properties that are listed in this data element replace the table replication in the `LA_GATE.tblrep.def` file on pod startup.	Netcool/OMNIbus documentation: Gateway for Message Bus table replication definition file
`xml-gate-startup-cmd-replace`	Properties that are listed in this data element replace the startup command definition in the `LA_GATE.startup.cmd` file on pod startup.	Netcool/OMNIbus documentation: Gateway for Message Bus startup command file

Examples of each of the data elements in this configmap are provided.

Data element: `xml-gate-props-append`

Properties that are listed in this data element are appended to XML gateway properties file on pod startup.

```
xml-gate-props-append: |
 MessageLevel: 'debug'
```

Data element: `xml-gate-map-replace`

Properties that are listed in this data element replace the gateway map definition in the `LA_GATE.map` file on pod startup.

```
xml-gate-map-replace: |CREATE LOOKUP SeverityLkTable
(
    { 0, 'Clear'                                                    },
    { 1, 'Indeterminate'                                            },
    { 2, 'Warning'                                                  },
    { 3, 'Minor'                                                    },
    { 4, 'Major'                                                    },
    { 5, 'Critical'                                                 }
) DEFAULT = TO_STRING('@Severity');

CREATE LOOKUP TypeLkTable
(
    { 0, 'Type Not Set'                                             },
    { 1, 'Problem'                                                  },
    { 2, 'Resolution'                                               },
    { 3, 'Visionary Problem'                                        },
    { 4, 'Visionary Resolution'                                     },
    { 7, 'ISM New Alarm'                                            },
    { 8, 'ISM Old Alarm'                                            },
    { 11, 'More Severe'                                             },
    { 12, 'Less Severe'                                             },
    { 13, 'Information'                                             }
) DEFAULT = TO_STRING('@Type');

CREATE LOOKUP ClassLkTable
(
    { 0, 'Default Class'                                            },
    { 95, 'Fujitsu FLEXR+'                                          }
) DEFAULT = TO_STRING('@Class');

    # My test map
    CREATE MAPPING StatusMap
    (
        'LastOccurrence'  =     '@LastOccurrence',
        'Summary'          =    '@Summary',
        'NmosObjInst'      =    '@NmosObjInst',
        'Node'             =    '@Node',
        'NodeAlias'        =    '@NodeAlias' NOTNULL '@Node',
        'LastOccurrence'   =    '@LastOccurrence',
        'Severity'         =    Lookup('@Severity', 'SeverityLkTable'),
        'AlertGroup'       =    '@AlertGroup',
        'AlertKey'         =    '@AlertKey',
        'Identifier'       =    '@Identifier',
        'Location'         =    '@Location',
        'Type'             =    Lookup('@Type', 'TypeLkTable'),
        'Tally'            =    '@Tally',
        'Class'            =    Lookup('@Class', 'ClassLkTable'),
        'OmniText'         =    '@Manager' + ' ' + '@Agent',
        'ActionCode'       =    ACTION_CODE,
        'ServerName'       =    '@ServerName',
        'ServerSerial'     =    '@ServerSerial'
    );
```

Data element: `xml-gate-tblrep-def-replace`

Properties that are listed in this data element replace the table replication in the `LA_GATE.tblrep.def` file on pod startup.

```
  xml-gate-tblrep-def-replace: |
    # My test table replication definition
    REPLICATE FT_INSERT,FT_UPDATE FROM TABLE 'alerts.status'
        USING MAP 'StatusMap';
```

Data element: `xml-gate-startup-cmd-replace`

Properties that are listed in this data element replace the startup command definition in the `LA_GATE.startup.cmd` file on pod startup.

```
xml-gate-startup-cmd-replace: |
  SHOW PROPS;
```

Configuration share configmap

The configmap for the configuration share configures pod to pod file sharing, and should not be edited.

Cassandra configmap

The `cassandra` configmap is optionally used by Cassandra to ensure that nodes are started consecutively. If node startup fails, you can edit the configmap to start another node.

Contents

The following table lists the data elements that are contained in the `cassandra` configmap:

Table 139. Data elements in the `cassandra` configmap

Data elements	Description	More information
`bootstrapping.node`	The name of the node that is bootstrapping (starting), or the last node to start.	https://github.ibm.com/hdm/common-cassandra/wiki/seed-options

ASM-UI configmap

This configmap has the configuration for the IBM Agile Service Manager user interface. It should not be edited.

Cloud Native Analytics gateway configmap

The `ea-noi-layer-eanoigateway` configmap is created during the helm install and should not be edited.

CouchDB configmap

The `couchdb` configmap is created during the helm install and should not be edited.

Kafka configmap

The `kafka` configmap is created during the helm install and should not be edited.

Zookeeper configmap

The `zookeeper` configmap is created during the helm install and should not be edited.

Insight packs

Insight packs are used together with IBM Operations Analytics - Log Analysis to provide Event search and Topology search capabilities.

You can find more information on the insight packs in the following sections of this documentation:

- Installing the insight packs
  - Installing the Tivoli Netcool/OMNIbus Insight Pack
  - “Installing the Network Manager Insight Pack” on page 90
- Configuring the insight packs

– Configuring the Tivoli Netcool/OMNIbus Insight Pack
– Configuring the Network Manager Insight Pack

All of the insight pack documentation can be found in the following PDF documents:

Document title	Link to document	Description
IBM Operations Analytics - Log Analysis: Netcool/OMNIbus Insight Pack README	Click here to download.	Documents the installation, operation, and customization options for the Insight Pack that enables the Event Search integration between IBM Operations Analytics - Log Analysis and Netcool/OMNIbus.
IBM Operations Analytics - Log Analysis: Network Manager Insight Pack README	Click here to download.	Documents the installation, operation, and customization options for the Insight Pack that enables the topology search integration between IBM Operations Analytics - Log Analysis and Network Manager.

Netcool Operations Insight audit log files

Records of user activity and report history are contained in the Netcool Operations Insight audit log files.

The log files can be found at the following locations:

Audit log

`$IMPACT_HOME/logs/NCI_NOI_Audit.log`

The audit log is a record of all user interactions with Netcool Operations Insight.

Report history log

`$IMPACT_HOME/logs/NCI_NOI_Report_History.log`

The report history log records the following data for executed Event Analytics reports:

- `Run ID`
- `Date`
- `Report Name`
- `Report Type`
- `Status`
- `Start Date`
- `End Date`
- `Duration`
- `Number of Events`
- `Seasonal Events`
- `Seasonality Related Events Count`
- `Related Events`
- `Related Events Groups`
- `Related Events Group Size`
- `Suggested Patterns`
- `Filter`
- `Additional Related Events Filter`

Sample report history log:

```
"RunId_1519402261","2018-02-23
16:37:43.000","Seasonal_and_Related_events","Combined","COMPLETED","2017-07-01
00:00:00.0","2017-07-04
23:59:59.0","00:08:15","459418","3141","193731","244","3569","93","96","38","10","7","0","",
"((AlertGroup != 'Synthetic Event - Parent') OR AlertGroup IS NULL)"
"RunId_1519405376","2018-02-23
17:30:53.000","Related_Events_1","RelatedEvents","COMPLETED","2017-07-01
00:00:00.0","2017-07-04
23:59:59.0","","459418","0","0","244","3569","93","96","38","10","7","0","","((AlertGroup !
= 'Synthetic Event - Parent') OR AlertGroup IS NULL)"
"RunId_1523626088","2018-04-13
09:30:56.000","Seasonallity_only","Seasonality","COMPLETED","2015-02-01
00:00:00.0","2015-03-01
23:59:59.0","00:02:09","116961","3540","0","0","0","0","0","0","0","0","0","",""
```

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive Armonk, NY
10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
958/NH04
IBM Centre, St Leonards
601 Pacific Hwy
St Leonards, NSW, 2069
Australia

IBM Corporation
896471/H128B
76 Upper Ground
London SE1 9PZ
United Kingdom

IBM Corporation
JBF1/SOM1
294 Route 100
Somers, NY, 10589-0100
United States of America

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

AIX, Db2, IBM, the IBM logo, ibm.com®, Informix, iSeries, Netcool, OS/390®, Passport Advantage, pSeries, Service Request Manage, System p, System z, Tivoli, the Tivoli logo, Tivoli Enterprise Console®, TotalStorage, WebSphere, xSeries, z/OS, and zSeries are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Adobe, Acrobat, Portable Document Format (PDF), PostScript, and all Adobe-based trademarks are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Appendix A. Release notes

IBM Netcool Operations Insight V1.6.0.1 is available. Compatibility, installation, and other getting-started issues are addressed in these release notes.

Contents

Description

Netcool Operations Insight combines real-time event consolidation and correlation capabilities of Netcool Operations Insight with Event Search and Event Analytics. It further delivers seasonality analysis to assist in detecting regularly occurring issues. Netcool Operations Insight also enables real-time enrichment and correlation to enable agile responses to alerts raised across disparate systems including application topology.

Compatibility

IBM Netcool Operations Insight includes the product and component versions listed in the following topic: "V1.6.0.1 Product and component version matrix" on page 13. This topic also includes information on the eAssemblies and fix packs required to download and install.

For more information about the Netcool Operations Insight products and components, see "Products and components on premises" on page 9

Release history

Full release history is given on the following web page.

https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Netcool%20OMNIbus/page/Release%20history

System requirements

For information about hardware and software compatibility of each component, and detailed system requirements, see the IBM Software Product Compatibility Reports website:

http://www-969.ibm.com/software/reports/compatibility/clarity/index.html

Tip: When you create a report, search for Netcool Operations Insight and select your version (for example, V1.4). In the report, additional useful information is available through hover help and additional links.
For example, to check the compatibility with an operating system for each component, go to the **Operating Systems** tab, find the row for your operating system, and hover over the icon in the **Components** column. For more detailed information about restrictions, click the **View** link in the **Details** column.

New product features and functions in V1.6.0.1

1.6.0.1

Updated product versions in V1.6.0.1

The Netcool Operations Insight V1.6.0 solution includes features delivered by the fix packs and fix pack extensions of the products and versions listed in the following topic:

"V1.6.0.1 Product and component version matrix" on page 13

The products are available for download from Passport Advantage and Fix Central.

More information: "Products and components on premises" on page 9

New features in V1.6.0.1

The following features and functions are available in the Netcool Operations Insight V1.6.0.1 product:

Operations Management on IBM Cloud Private

The Operations Management on IBM Cloud Private solution provides the following new features and functions.

Cloud Native Analytics self-monitoring

A self-monitoring policy can be enabled to provide assurance that Cloud Native Analytics is processing events. For more information, see "Cloud Native Analytics Service Monitoring" on page 354.

Topology Analytics

View topological context for your Cloud Native Analytics events, where there is an associated resource. For more information, see "Topology Analytics on IBM Cloud Private" on page 356.

Connecting an on-premises Object Server to Operations Management on IBM Cloud Private

After you successfully deploy Operations Management on IBM Cloud Private, you can connect to an existing on-premises installation to create an event feed between your on-premises and cloud installations. For more information, see "Connecting an on-premises IBM Tivoli Netcool/OMNIbus ObjectServer to Operations Management on IBM Cloud Private" on page 196.

Upgrade

You can upgrade from V1.6.0 to V1.6.0.1. For more information, see "Upgrading Operations Management on IBM Cloud Private from 1.6.0 to 1.6.0.1" on page 153. For information about rolling back the upgrade, see "Rolling back Operations Management on IBM Cloud Private from V1.6.0.1 to V1.6.0" on page 161.

New product features and functions in V1.6.0

Updated product versions in V1.6.0

The Netcool Operations Insight V1.6.0 solution includes features delivered by the fix packs and fix pack extensions of the products and versions listed in the following topic:

"V1.6.0 Product and component version matrix" on page 17

The products are available for download from Passport Advantage and Fix Central.

More information: "Products and components on premises" on page 9

New features in V1.6.0

The following features and functions are available in the Netcool Operations Insight V1.6.0 product.

Operations Management on IBM Cloud Private

The Operations Management on IBM Cloud Private solution provides the following new features and functions.

Cloud Native Analytics

Cloud Native Analytics formulates event grouping policies by identifying seasonal and temporal patterns of events within a monitored environment. You can automatically deploy these policies, or you can manually select which policies are deployed. Events are consolidated into incidents, from which you can drill down into event details and time-lines,

and see the seasonal, temporal and scope-based groups to which events belongs. For more information, see “Cloud Native Analytics on IBM Cloud Private” on page 342.

More configurable deployment options
Increased configuration options allow you to configure passwords, secrets and pod access, or to allow Operations Management on IBM Cloud Private to configure these for you during installation. For more information, see “Preparing for installation on IBM Cloud Private” on page 98.

Red Hat OpenShift support
Added support for Red Hat OpenShift. The Operations Management on IBM Cloud Private solution can be deployed on IBM Cloud Private with OpenShift. When you install IBM Cloud Private with OpenShift, IBM Cloud Private provides the IBM Cloud Private experience, management, and operations for applications and uses OpenShift's Kubernetes and Docker registry that is already installed by Red Hat. For more information, see IBM Cloud Private documentation: IBM Cloud Private with OpenShift.

Increased Security Assurance
Installation of Operations Management on IBM Cloud Private must now be performed with user-defined passwords and secrets or with randomly generated passwords and secrets, instead of with default passwords. For more information, see “Configuring passwords and secrets” on page 111.

Netcool/Impact scalability
The number of Netcool/Impact core server pods can be increased or decreased while Operations Management on IBM Cloud Private is running. For more information, see “Scaling the Netcool/Impact service” on page 141.

Certification Levels
Operations Management on IBM Cloud Private now has RedHat image certification, and IBM CloudPak level 2 certification, with improved resiliency, scaling and security. For more information, see https://www.ibm.com/support/knowledgecenter/en/SSBS6K_3.2.0/app_center/cloud_paks_over.html

Operations Management for Operations Insight
The base Netcool Operations Insight solution provides the following new features and functions.

Event source integration
IBM Cloud Event Management can be configured to forward public or private cloud events to Netcool Operations Insight , where they will appear in the Event Viewer. Netcool Operations Insight events can also be configured to display in IBM Cloud Event Management. For more information, see “Connecting event sources to your IBM Netcool Operations Insight on premises deployment” on page 189

Known problems at eGA
The following problems with IBM Netcool Operations Insight were known at the time of eGA.

- “Solution-level problems” on page 562
- Operations Management on IBM Cloud Private
- “Device Dashboard” on page 566
- Event Analytics
- “Event Search” on page 572
- “Globalization” on page 574
- “Network Health Dashboard” on page 575
- “Operations Analytics - Log Analysis” on page 575

Solution-level problems

Save and close button in View Builder not working

When the View Builder is opened from the Event Viewer the **Save and close** button doesn't close the window. Selecting the button in the pop-up View Builder displays a close button, which will then work to close the window.

Unable to create new users in LDAP using WebSphere Application Server

When a new user is created using the WebSphere Application Server, the `UniqueName` attribute references the defaultFileBasedRealm instead of LDAP. This means that the new user cannot be assigned to groups and therefore cannot be assigned roles in LDAP.

Certificate Verification popup

When logging in to DASH, an Information pop-up is displayed, requesting that certificates for some links be verified. These links do not work. This pop-up can be safely closed.

Operations Management on IBM Cloud Private

Incorrect path in `ibm-ea-asm-normaliser` section in *values.yaml*

Follow these instructions to correctly enable Topology Analytics during installation or upgrade.

1. Before your installation or upgrade, edit your configuration file, `values.yaml`.
2. Find the `ibm-ea-asm-normalizer` section:

```
ibm-ea-asm-normalizer:
  joinWindowSize: 15
  kafkaImage:
    name: ea/kafka
    tag: KAFKA_IMAGE_AUTO_UPDATE
```

3. Modify this section by removing `ea/` from the value of `ibm-ea-asm-normaliser.kafkaImage.name`, and by removing the tag `KAFKA_IMAGE_AUTO_UPDATE`. The modified section should be:

```
ibm-ea-asm-normalizer:
  joinWindowSize: 15
  kafkaImage:
    name: kafka
```

4. Ensure that you have accepted the license by selecting the check-box (UI install), or by setting `global.license=accept` in your `values.yaml` file (CLI install).
5. Continue with your installation or upgrade.

Note: When performing an upgrade, get your current values by running the command `helm get values` *`helm_release_name`* `--tls`, or use the `--reuse-values` option in the helm upgrade, where *helm_release_name* is the name of your Operations Management on IBM Cloud Private deployment.

Error message tells user to set incorrect license parameter

When performing a helm upgrade, the following error message is given if the license is not accepted: `You have not accepted the terms of the license agreement. You must review and accept the terms by setting license=accept.` The license parameter needs to be set as `global.license=accept`, not as `license=accept`.

Multiple icons displayed in Event viewer

The Event viewer displays multiple icons for the same correlation type against each event.

Countdown timer on Incident Viewer is not reset when refresh is clicked

When refresh is clicked the countdown timer is not reset to 0 as it should be. The view is still correctly refreshed.

Policy timeouts with error

When selecting a policy from **Insights->Manage Policies**, the policy will not load and displays the following error:

```
An error occurred while fetching data from the server. Please make sure you have an
active internet connection. More...
```

You may sometimes hit this issue if you open a policy which has a large number of event instances.

Analytics services generating large log files

Disks can fill up very fast with log files for each service instance. As a workaround, increase disk space in the docker containers space (the `/var/lib/docker` directory).

Blank window displayed when selecting Event Viewer from side bar of new tab in Internet Explorer V11

When you open the Incident Viewer and select a new tab, the new tab opens displaying the Incident view and a side bar with Event Viewer and Policies. If you select Event Viewer from the side bar, a blank window is displayed.

Default groups are missing from Tools Configuration > Access Criteria

When you log in to IBM Netcool/OMNIbus Web GUI and go to **Tools Configuration** > **Access Criteria**, the default groups, such as icpadmins, are missing under the "Access Criteria" section.

Scope-based grouping side panel displays incorrect times

The scope-based grouping side panel might display incorrect first occurrence and last occurrence times, compared to the actual occurrences in the Incident Viewer.

Authentication failure - cannot log in to Web GUI when external LDAP is configured

Unable to login to Web GUI using LDAP credentials. Configuring to an external LDAP server is not supported in this release.

Authentication failure - `ncobackup` cannot authenticate user "root" when `ncoprimary` is restarted

The following errors are listed in the `ncobackup` pod logs during startup:

```
Server 'AGG_B' initialised - entering RUN state.
Updating default root password
2019-06-13T22:11:43: Error: E-OBX-102-023: Failed to authenticate user root.
(-3600:Denied)
2019-06-13T22:11:43: Error: E-OBX-102-057: User root@<releasename>-ncobackup-0 failed to
login: Denied
Failed to connect
Error: Failed to get login token
Unable to update root password
```

The following error is seen in the `objectserver` logfile:

```
cat /opt/IBM/tivoli/netcool/omnibus/log/AGG_B_audit_file.log
2019-06-13T22:11:43: Error: E-SEC-010-002: authentication failure - cannot authenticate
user "root" : Denied
```

This error can be ignored. The ncobackup pod tries to change the default password and fails, because it has already been changed to the newer version.

UI throws "Pod Security Conflict" warning when selecting namespace during a deployment of Operations Management on IBM Cloud Private with Red Hat OpenShift

The following warnings are displayed:

Pod Security Conflict This chart requires a namespace with a ibm-priviliged-psp pod security policy Pod Security Warning Your ICP cluster is running all namespaces Unrestricted (ibm-anyuid-hostpath-psp) by default. This could pose a security risk

Red Hat OpenShift does not use pod security policies (PSP), and instead uses security constraint contexts (SCCs). This warning can be ignored, if you have correctly configured a security constraint context in “Configuring pod access control” on page 108.

`loadSampleData.sh` script fails on IBM Cloud Private with OpenShift

The `loadSampleData.sh` fails to run on IBM Cloud Private with OpenShift as a result of the aggregation dedup service pod timing out it's redis connection. After a while, the service stops connecting if redis is down. To work around this issue, restart the stateless aggregation dedup pods. Run the **`kubectl get pod | grep dedup`** command to get a list of dedup pods, then run the **`kubectl delete pod`** command for each dedup pod.

Auto-deploy mode is enabled after running the `loadSampleData.sh` script.

If you have installed Operations Management on IBM Cloud Private in manual deploy mode, running the `loadSampleData.sh` script switches your configuration to auto-deploy mode. For more information about manually deploying policies, see the *Temporal Group Policies Deploy First* parameter in "Configuring Installation Parameters for Operations Management on IBM Cloud Private" on page 125.

To disable auto-deploy mode, re-run a portion of the training by running the `runTraining.sh` script.

Get the start and end times of the sample data as in the following example:

```
kubectl delete pod ingesnoi3;
kubectl run ingesnoi3 -i --restart=Never --env=LICENSE=accept
--image=<repository>/ea/ea-events-tooling: getTimeRange.sh samples/
demoTrainingData.json.gz
pod "ingesnoi3" deleted
{"minfirstoccurence":{"epoc":1552023064603,"formatted":
"2019-03-08T05:31:04.603Z"},"maxlastoccurrence":{"epoc":
1559729860924,"formatted":"2019-06-05T10:17:40.924Z"}}
```

Re-run the training with seasonality disabled, as in the following example:

```
kubectl delete pod ingesnoi3;
kubectl run ingesnoi3 -i --restart=Never --env=LICENSE=accept
--image-pull-policy=Always --image=<repository>
/ea/ea-events-tooling: runTraining.sh -- -r test-install -a SEASONALITY
-s 1552023064603 -e 1559729860924 -d false
```

Where *<repository>* is the repository name and *-s* and *-e* are the start and end times returned by the **`getTimeRange()`** command and **`-d false`** disables live policy updates.

Authentication errors reported in logs when ncoprimary and ncobackup pods are restarted.

The following errors can be seen in the ncoprimary and ncobackup pod logs when the ncoprimary and ncobackup pods startup:

```
Information: I-SEC-104-002: Cannot authenticate user "root": Denied
Error: E-OBX-102-023: Failed to authenticate user root. (-3600:Denied)
Error: E-OBX-102-057: User root@dev401-ncoprimary-0 failed to login: Denied
Failed to connect
Error: Failed to get login token
Unable to update root password
```

These errors are caused by the objectserver starting up and attempting to change the default root password, which has already been changed.

Unable to make TLS connections using administrator GUI versions earlier than V8.1.0.18

Versions of administrator GUI **`nco-config`** earlier than V8.1.0.18 fail to make a TLS connection. This is a known issue.

If you want to use the secrets that are automatically created by Operations Management on IBM Cloud Private, you must disable FIPS mode

If the certificate is automatically generated and FIPS mode is specified with the `$NCHOME/etc/security/fips.conf` file, the Netcool/OMNIbus client fails to make a connection. The `CT-LIBRARY error` error message is displayed. If you want to use the secrets that are automatically created by Operations Management on IBM Cloud Private then you must disable FIPS mode. Do this by removing the following file: `$NCHOME/etc/security/fips.conf file`.

Additional Insight Pack installation does not persist

After you install an additional Insight Pack with Operations Analytics - Log Analysis and restart your pod, the Insight Pack disappears. This is a known issue. As a workaround, reinstall the Insight Pack.

Operations Analytics - Log Analysis dashboard does not persist

After you create a Operations Analytics - Log Analysis dashboard with a search result widget and restart your pod, the dashboard disappears. This is a known issue. As a workaround, recreate the dashboard.

Outage of Event Analytics pages or features

Following an Impact Server failover, Event Analytics pages or features might not be available for a short time. This is a known issue.

Cannot change openldap suffix or base distinguished name

When Netcool Operations Insight is running on IBM Cloud Private, and you are using the built-in LDAP server container `openldap` that is provided with Operations Management on IBM Cloud Private then you must not change the LDAP suffix, which is part if the LDAP base DN, or the LDAP bind distinguished name (bind DN). These are hardcoded based on the values that you specified during installation. If you do try to change these values then the Dashboard Application Services Hub pod, `webgui`, will not start.

Storage requirements

You must have sufficient storage, as described in "Requirements for an installation on IBM Cloud Private" on page 98. When you install Operations Management on IBM Cloud Private, you must specify the **Enable Data Persistence** option, as described in "Configuring Installation Parameters for Operations Management on IBM Cloud Private" on page 125.

Remote access to Db2 container

It is not possible to remotely access the Db2 container (`db2ese`) that is running within an IBM Cloud Private cluster. You might want to do this if you are using a database client, rather than using the command line within the Db2 container.

Remote access to Netcool/Impact Derby database

It is not possible to remotely access the Netcool/Impact Derby database that is running within an IBM Cloud Private cluster.

New method to access the status information of Netcool/Impact Name Server instances

The method to access the status information of Netcool/Impact Name Server instances has changed from the method documented in the Netcool/Impact documentation at https://www-03preprod.ibm.com/support/knowledgecenter/SSSHYH_7.1.0.16/com.ibm.netcoolimpact.doc/admin/imag_gui_server_viewing_nameserver_status_c.html. If you are running Netcool/Impact under Operations Management on IBM Cloud Private v1.5.0.1, then use the following method to access this status information:

1. Run the following command:

   ```
   helm status release-name --tls
   ```

   Where *release-name* is the release name for your Operations Management on IBM Cloud Private deployment.

   This command generates output to your console.

2. In the output navigate to the **Impact Servers** section, which looks similar to the following snippet:

   ```
   Impact Servers:
     Update your hosts file (On the machine you are running your Browser) or your DNS
   settings with this mapping
     $NODE_IP nci-0.release-name.master-node
     $NODE_IP nci-1.release-name.master-node

     firefox https://nci-0.release-name.master-node/nameserver/services
     firefox https://nci-1.release-name.master-node/nameserver/services
   ```

3. Copy and paste the URLs into your browser to retrieve the Netcool/Impact name server instance status.

Device Dashboard

Device Dashboard is unable to differentiate between anomaly thresholds configured for poll definitions that have the same metric name

The **Device Dashboard** is unable to differentiate between anomaly thresholds configured for poll definitions that have the same metric name. To ensure that metrics are correctly distinguished in the **Device Dashboard Performance Insights** portlet ensure that any poll definitions created have a unique metric name.

However, this workaround does not apply to the Default Chassis Ping and Default Interface Ping poll definitions, which by default both use a metric with the name PingTime. The consequences of this are best illustrated using an example:

1. Assume you set and enable an anomaly threshold on the Default Chassis Ping poll definition, with the intention of viewing anomalies against the chassis ping time metric.
2. Assume also that you do not set an anomaly threshold on the Default Interface Ping poll definition, as you do not want to view anomalies against the interface ping time metric.
3. However, the threshold you set and enabled on the Default Chassis Ping poll definition will apply to both the chassis and the interface ping time metrics. Consequently in the Interfaces tab of the **Performance Insights** portlet, selecting **PingTime** from the **Metric** drop-down list will return unexpected anomalies in the portlet.

The workaround is to set and enable the same thresholds on both the Default Chassis Ping and Default Interface Ping poll definitions.

Tooltip on sparkline in Performance Insights portlet displays unnecessary scroll bar

The sparklines in the **Performance Insights** portlet each have an associated tooltip. Some of these tooltips display an unnecessary vertical scrollbar. This issue can safely be ignored.

Performance Insights portlet: sorting the data by Value does not work

In the **Interfaces** tab within the **Performance Insights** portlet sorting the data by clicking the **Value** column does not work.

When launching a new Device Dashboard by right-clicking an event in the Event Viewerwithin an existing Device Dashboard the Device Dashboard does not refresh with the new entity identifier

When launching a new **Device Dashboard** by right-clicking an event in the **Event Viewer** within an existing **Device Dashboard** the newly launched **Device Dashboard** does not refresh with the entity identifier corresponding to the selected event. To work around this issue, select **Event Viewer** from the Dashboard Application Services Hub navigation, select the relevant event and launch the **Device Dashboard** from there.

Integration problem: Widgets do not open when selecting "Show Traffic"

When right-clicking and selecting **Show Traffic** from any network topology (for example, **Network Hop View**), widgets do not open. To work around this issue, find the device in **Event Viewer** , right-click and select **Show Traffic**.

Device Dashboard does not launch from a Network Hop View opened within a new browser tab

If you launch a **Network Hop View** within a new browser tab by running the **Find in** > **Network Hop View** from any other GUI, then you will not be able to launch the **Device Dashboard** from that **Network Hop View**.

Performance Insights portlet filter values are cleared on refresh of Device Dashboard

You can apply filters in the **Device Dashboard Performance Insights** portlet, in both the **Device** and **Interfaces** tabs. However, when the **Device Dashboard** is refreshed, any values entered into the **Interfaces** tab filter field are automatically cleared. You will need to enter filter values again after the refresh.

Note: This issue does not affect values entered into the **Device** tab filter field.

Performance Insights portlet filter values are cleared when switching tabs

When you switch from the **Interfaces** to the **Device** tab in the **Device Dashboard Performance Insights** portlet, any values entered into the **Interfaces** tab filter field are automatically cleared. You will need to enter filter values again.

Note: This issue does not apply when you switch from the **Device** to the **Interfaces** tab.

When portlet preferences are saved Device Dashboard portlets are not displayed properly

After changing portlet preferences in the following **Device Dashboard** portlets, unexpected results occur in the following portlets:

- **Performance Insights** portlet: metric and count data disappear from the portlet, and the preferences do not take effect.
- **Performance Timeline** portlet: timeline data disappears; however, the portlet preferences are saved and are automatically applied when graph is next rendered, either automatically on the next refresh or by clicking a sparkline.

To resolve this issue, do the following:

- **Performance Insights** portlet: click a device in the Topology portlet. This forces a manual refresh of the **Performance Insights** portlet.
- **Performance Timeline** portlet: click any sparkline in the **Performance Insights** portlet. This reinstates the **Performance Timeline** portlet and refreshes the content of the portlet.

Occasionally the content of the Device Dashboard Performance Insights portlet does not update when you select a different interface metric

When selecting a new interface metric from the **Metric** drop-down list in the **Performance Insights** portlet **Interfaces** tab, the content of portlet occasionally does not update. To work around this issue, select the metric twice from the drop-down list.

Event Analytics

If your problem is not listed in this section, then refer to Troubleshooting Event Analytics for additional issues.

If one group in a pattern is moved to a new state, the other groups in the pattern disappear from the UI

Symptom:

If one group in a pattern is moved to a new state, the other groups in the pattern disappear from the UI. The pattern and selected group are moved to the new state, but the other groups in the pattern are not visible.

Resolution:

If this situation occurs, to display all the patterns groups, move the group or pattern back to the original state.

If it is necessary to move some groups allocated to a pattern to a new state, you should delete the pattern before the groups are moved. Then, the pattern can be re-created for the same Event Types as before. This method re-allocates the remaining groups that were previously in the pattern.

In the View Related Events portlet the Related Event Groups panel right-click menu commands momentarily appear in capital letters

If you perform the following sequence of steps in the View Related Events portlet then the commands in the Related Event Groups panel right-click menu momentarily appear in full upper-case.

1. Run a configuration.
2. Navigate to the View Related Events portlet.
3. On the Related Event Groups panel, right-click a group and select to either deploy, watch, or archive the group.

4. As soon as **Deploy**, **Watch**, or **Archive** is selected and before the group has moved, right click the related events group again. The right-click menu appears with capital letters.

Unable to add aggregate fields using the Event Analytics Setup Wizard

When running the Event Analytics Setup Wizard, if you add aggregate fields in the **Report fields** section of the wizard, and save the wizard data, subsequent configuration runs do not display the new fields in seasonal or related event reports. In order to add these fields you must also restart the Netcool/Impact GUI server.

It is not possible to determine which event triggers a pattern from the View Related Events portlet

If you run a configuration and then deploy a pattern, then trigger the pattern so that the trigger statistics are updated, when you subsequently display the View Related Events portlet, it is not possible to determine which event caused the pattern to trigger.

During installation of Event Analytics if the Derby database becomes corrupted reset it to default state

During the installation of Event Analytics, if the Derby database becomes corrupted you must reset it to default state by performing the following steps:

1. Navigate to `$IMPACT_HOME/install/dbcore` and find the zip archive named `ImpactDB_NOI_FP15.zip`.
2. Stop of the Netcool/Impact servers in the cluster, including the primary and secondary servers running the Derby database.
3. Back up the existing directory structure in `$IMPACT_HOME/db/`*`Server_Name`*`/derby`, where *`Server_Name`* is the name of the Netcool/Impact server. Once the backup is complete, remove the directory structure.
4. Copy the `ImpactDB_NOI_FP15.zip` zip file found in step 1 to `$IMPACT_HOME/db/`*`Server_Name`*`/derby` and unzip tin there.
5. Start the primary Netcool/Impact server and allow it time to fully initialize.
6. Start the secondary Netcool/Impact server and allow it to resynchronize from the primary server.

Ensure that the property `impact.featureToStringAsSource.enabled` is set to `true`

The property `impact.featureToStringAsSource.enabled` must be set to `true` in order for Event Analytics to work correctly. Set the following value in the Netcool/Impact properties file and run the `nci_trigger` command to activate the property, as described in the following topic: Netcool/Impact Knowledge Center: `nci_trigger`.

```
impact.featureToStringAsSource.enabled=true
```

Event Analytics Configuration wizard reports error following failover of `nciserver-0` pod

When Netcool Operations Insight is running on IBM Cloud Private, and the `nciserver-0` container fails over and back again, the Event Analytics Configuration wizard will subsequently report a `java.io.FileNotFoundException` error. To resolve this issue, regenerate the Netcool/Impact policy list by running the `createPolicyList` script, as described in the following topic: Netcool/Impact Knowledge Center: createPolicyList.

Seasonality rule generates errors if only Minute of hour is set

If you create a seasonality rule where the only clause set is **Minute of hour**, then the rule will not fire and will generate errors. An example of this is the following:

```
minute of hour = 2, 3, 4
```

To work around this issue, add to the rule a clause that will always be true; for example:

```
day of week = Mon, Tue, Wed, Thurs, Fri, Sat, Sun
minute of hour = 2, 3, 4
```

Attempt to edit seasonality rule returns blank screen

When deploying a seasonality rule and then attempting to edit the rule, sometimes the Modify Existing Rule screen opens but is blank, and you are unable to edit the rule.

Order of columns in the Configure Analytics screen is inconsistent

The order of columns in the **Configure Analytics** screen is inconsistent, and is as follows. The columns are numbered here for ease of reference; they are not numbered in the GUI.

1. Name
2. Event Identity
3. Seasonality Status
4. Related Event Status
5. Start Time
6. End Time
7. Seasonality Phase
8. Seasonality Phase Progress
9. Related Event Phase
10. Related Event Phase Progress
11. Scheduled
12. Related Event Last Run Duration
13. Seasonal Event Last Run Duration

Most of the columns follow the order Seasonality/Seasonal Event, then Related Event. For example, see column pairs 3-4, 7-9, 8-10. However columns 12-13 follow the opposite order (Related Event then Seasonal Event). If you have lots of rows in the table and you scroll down so that the column titles are no longer visible, you might get the Related Event Last Run Duration and Seasonal Event Last Run Duration mixed up. Be aware that the order for these two column is inconsistent with the other columns.

Seasonal event not matching original selection when opting to trigger rule off events related to the seasonal event

When you create a rule based on a seasonal event, if that seasonal event has associated related events then you have the option to trigger the rule based on one or more of those related events, by selecting the **Edit Selection...** button in the **Create Rule** screen. When you select this option, the originally selected seasonal event should be in the list and should be selected. Occasionally you might find that the event selected is not the originally selected seasonal event. In this case, you need to search for the desired seasonal event and reselect it.

Related events configuration run errors following adding of aggregate fields

After you have configured the system to add an aggregate field to the analytics reports using the **Report fields** panel in the **Configure Event Analytics** wizard screen, you might subsequently encounter errors when you run a related events configuration.

If this is the case then use `nci_trigger` to export the configuration and check what the `Actionable` attribute for your new aggregate field is set to.

- If the `Actionable` attribute is set to "false", then contact IBM Support.
- If the `Actionable` attribute is set to "true", then set it to "false" and then run `nci_trigger` to configure the system using the existing settings.

For more information on using `nci_trigger` to configure the system, see “Configuring Event Analytics using the command line” on page 259.

Error on first use of the Event Analytics Configuration wizard

On first use of the Event Analytics Configuration wizard, you might encounter an error similar to the following:

```
ATKRST132E An error occurred while transferring a request to the following remote
provider:
 - '404:Cannot access data provider - Impact_NCICLUSTER
```

To resolve this issue, restart the Netcool/Impact server and Netcool/Impact UI server, then restart the wizard.

Blank screen in Event Analytics Configuration wizard Report fields screen

If, while working through the Event Analytics configuration wizard, you encounter a blank **Report fields** screen, this might be because you have custom fields in your Historical Event database, and you did one of the following:

- You did not create a database view to map the custom fields to the corresponding standard field names in Netcool/Impact, as described in "Mapping customized field names" on page 279. In this case, perform the following steps:
  1. Exit the wizard.
  2. Create a database view, as described in "Mapping customized field names" on page 279.
  3. Restart the wizard. When you get to the **Configure historical event database** screen, make sure to select the database view from the **History table** drop-down list, as described in "Configuring the historical event database" on page 251.
- You created the database view, but when configuring the Historical Event database in the wizard, you specified the unmapped history table instead of the database view.
  1. Click the **< Back** button in the wizard until you get to the **Configure historical event database** screen.
  2. In the **Configure historical event database** screen, make sure to select the database view from the **History table** drop-down list, as described in "Configuring the historical event database" on page 251.

Unable to export analytics data

When trying to export analytics data, you encounter a "500" error, as follows:

```
Error 500: java.io.IOException: Unable to validate authorization
```

If this happens then perform the following steps:

1. Edit the `$IMPACT_HOME/etc/<NCI>_server.props` file.
2. Set the `impact.noi.export.hostname` variable to a recognized hostname.

Unable to configure an event pattern in the Event Patterns GUI due to empty Trigger Action section

This occurs if, during the configuration of event pattern processing that is performed as part of Event Analytics system setup, the historical event database column that was chosen for the default event type does not contain any data.

View Related Events > Groups: right-click menu items displayed in capital letters

If you **Deploy**, **Watch**, or **Archive** a related event group from the right-click menu in the **View Related Events** > **Groups** panel and then immediately right-click again to access the menu before the group has moved, the menu appears with capital letters. There is no impact on moving groups in this manner.

Time window for Suggested Patterns displays large values

The time window displayed for suggested patterns may occasionally display a very large value in minutes (more than one day in minutes). This is due to anomalies in the data being processed by the event analytics algorithms. If this situation is seen for a suggested pattern, the advised course of action is to either edit the pattern so that the time window is reduced and save this pattern, or to rerun the Event Analytics configuration in the **Configure Analytics** portlet. This should resolve the issue.

Related events groups and patterns from the Watched tab not moving to the Expired tab

You can modify the expiry time for **Watched** or **Active** related events groups. When the expiry time is reached, expired groups and related events from the **Watched** tab are not moved to the **Expired** tab within the View Related Events portlet.

Parts of the Events Pattern screen not displaying because an invalid group is selected

If you attempt to create an event pattern with groups that have *EventID* as null, the check boxes for **Trigger Action**, **Event Type** and **Resource Columns** do not appear on the **Pattern Criteria** tab. An error message is displayed if you try to save the pattern. To avoid this issue, ensure that only valid groups are used to create patterns.

View Related Events error: "Failed to load data"

The error "Failed to load data" can appear if you select **All** from the **Configuration** pane, when the state is **Active** or **Watched**, and there is more than one set of related event configurations with data. To avoid this issue, select the individual configurations instead.

Chart data retrieved from the server is incorrectly formatted

Problems can occur when displaying Seasonal Event Graphs if you are experiencing network latency or a slow network connection between your browser and the DASH/Impact machine. In such a scenario, an incomplete graph page might be displayed with the following error message:

```
Chart data retrieved from the server is incorrectly formatted
TypeError: Cannot read property 'grammar' of null
```

Related Events with a medium or weak relationship profile displayed as strong

In the **Related Events Details** view under **More Information**, related events with a medium or weak relationship profile might be incorrectly displayed as a strong relationship profile.

Event Analytics configurations are deleted on Netcool/Impact cluster node failover

Following a Netcool/Impact cluster node failover, the running configuration will be correctly stopped but any queued configurations will be lost on the secondary node. These configurations cannot be retrieved and must be recreated on the secondary node.

For more information, see the following technote: http://www.ibm.com/support/docview.wss?uid=swg22012656.

Suggested pattern with a blank Event Type parameter field

When editing a suggested pattern, the **Event Type** parameter field will be empty if the property `type.0.eventtype` is set to a value that is empty in the database. To avoid this issue, ensure that the `type.0.eventtype` property is not set to an empty value in the Event History Database. Selecting an event type field that contains all null values in the history database will result in the pattern criteria section of the create or edit pattern screen appearing blank.

Blank fields during creation of synthetic event on non-occurrence.

Users are given the option to supply one or more values for additional columns by selecting **Set additional fields**. With the exception of the values specified in the **Create Event** window, only the values of Node and `Summary` are copied into the synthetic event.

Removing Netcool/Impact data models that are no longer needed.

This issue refers to the following Netcool/Impact data models that are no longer needed after you upgrade to the IBM Netcool Operations Insight fix pack. Use the Netcool/Impact GUI to remove these data models:

ObjectServerHistoryDb2ForSeasonality
ObjectServerHistoryOrclForSeasonality
ObjectServerHistoryMSSQLForSeasonality

Note: If you choose not to remove these data models, there is no impact to the Event Analytics functionality. You would remove these data models only for cosmetic purposes.

Event summary is truncated

The event summary is occasionally truncated in the **Related Events Details** portlet Timeline view. To view the event summary, modify the screen resolution temporarily.

Seasonal event rule with time condition does not run

A seasonal event rule, with an action to run after a specific time that is more than 25 seconds, is not run. To ensure that a seasonal event rule that is scheduled to run after a specific time runs correctly, select a time condition of less than or equal to 25 seconds.

Edit selection window hangs when 'selecting all' for a large number of related events during seasonal event rule creation

When creating a seasonal event rule for an event with a large number of related events, you can check the **Select all related events** check box to associate all the related events with the seasonal event rule. The problem occurs when a large number of related events, on the order of 1000 or more, are selected and **Edit selection** is clicked. The **Edit selection** window is displayed but it remains in a loading state.

To avoid this issue, split the report into smaller reports and create rules around reports with fewer related events.

Incorrectly displaying 'No event selected' error message

When creating a pattern for related events and clicking **Use Selected Event as Template**, if you have not selected an event, the system correctly displays the 'No event selected' error message. However, if you then do select an event and click **Use Selected Event as Template** again, the error message persists

In this case, you can disregard and close the error message. It will not affect the creation of the pattern.

When testing an event pattern collapsed sections do not display correctly when reopened

Following testing of an event pattern, if you collapse the **Groups**, **Group Instances**, and **Events** sections in the **Test** tab of the **Events Patten GUI** using the splitters provided and then you expand the sections again, the data columns in the **Groups** and **Events** panels become very narrow and the data cannot be read.

To work around this issue, you can do one of the following:

- Resize the window frame. This causes the columns to resize so that the data becomes visible.
- Manually resize the column widths or refresh the page. Either of these actions causes the columns to be displayed correctly again

Event Search

Unable to access Event Search from the Event Viewer Information right-click command

When Netcool Operations Insight is running on IBM Cloud Private, if you are in the Event Viewer, and you open the Event Properties window by double-clicking an event or by right-clicking an event, then clicking **Information** > **Event Search**, then you get a network error indicating that the request to access Operations Analytics - Log Analysis failed. As a workaround, you should access Event Search using an alternative right-click method.

Severity colors in Operations Analytics - Log Analysis charts use random colors for event severity values

Note: This issue is fixed in version 1.4.1.2.

Operations Analytics - Log Analysis charts, such as Event Trend by Severity, Event Storm by AlertGroup, and Severity Distribution, use random colors for chart elements that represent event severity values, and do not use the standard event severity colors used in Web GUI Event Viewer. For example, a pie chart in the Severity Distribution chart portlet might contain a blue pie chart segment representing events of critical severity, even though the standard color for critical severity in the Web GUI Event Viewer is red.

Furthermore, if two or more charts are presented on a dashboard, then it is likely that the color used for a given event severity in one chart portlet on that dashboard page will differ from the same event severity in another chart portlet on that same dashboard page.

Note: Each chart has its own severity color legend. Refer to the severity color legend on each chart to determine which colors are used on that chart to denote different severity values.

This use of random colors that differ from the standard event severity colors used in Web GUI Event Viewer is intentional, and is meant to enable you to distinguish between event severity values on a chart.

HeatMap chart types do not render correctly

Note: This issue is fixed in version 1.4.1.2.

The HeatMap type charts fail to render when using IBM Operations Analytics - Log Analysis V1.3.3, displaying the following error: `Parameter category is invalid for Heat Map`. The error is caused by a difference in the way charts are defined in Operations Analytics - Log Analysis 1.3.2 and 1.3.3.

The following charts are affected:

- In **OMNIbusInsightPack** > **Last Day** > **Dynamic Event Dashboard**:
  - Hotspots by Node and Alert Group
  - Hotsport by AlertGroup and Severity
- In **OMNIbus_Static_Dashboard**, opened from the Web GUI **EventSearch** tab by right-clicking and selecting **Show Event Dashboard by Node**:
  - Hotspots by Node and AlertGroup
  - Hotspots by AlertGroup and Severity

To correct this problem, edit the following chart specification files, and search for all instances of `category` and change them to `categories`:

- `$UNITY_HOME/AppFramework/Apps/OMNIbusInsightPack_v1.3.0.2/Last_Day/OMNIbus_Dynamic_Dashboard.app`
- `$UNITY_HOME/AppFramework/Apps/OMNIbusInsightPack_V1.3.0.2/OMNIbus_Static_Dashboard.app`

Right-click keyword search results hidden by default

This issue affects Event Search prior to Operations Analytics - Log Analysis version 1.3.5. When using the Event Search right-click menu within Web GUI and selecting **Show keywords and event count**, the search results in Event Search are empty. The problem is caused by the **Search Patterns** pane being hidden by default in the Operations Analytics - Log Analysis user interface.

To view the **Show keywords and event count** results, click the **Restore Section** triangle icon on the left side of the Event Search user interface.

Hotspots by Node, AlertKey not displaying.

If the **Hotspots by Node, Alert Group and AlertKey** chart fails to display in the Last_Month->Operational Status Dashboard the SOLR heap size might need increasing.

Note: The **Hotspots by Node, Alert Group and AlertKey** chart is CPU intensive and can be slow to render for systems with a large amount of data.

More information: Search runtime exceptions and IBM SmartCloud® Analytics - Log Analysis Performance and Tuning Guide.

Hover values for the Event Trend by Severity charts do not appear to match the axes.

When hovering over a point on a particular severity the values returned might not appear to match the axes on the chart. It is because the hover values represent that severity only, whereas the values on the axes are cumulative. For example, if there are 20 Intermediate severity events and 26 Major severity events displayed on the line above, the Major events will appear against 46 on the Y-axis.

Drill down does not return results.
The drill down function is not available for the type Omnibus Tree Map. This affects some of the charts in the Operational Efficiency dashboard in the Last Month folder, and in the OMNIbus Operations Manager and OMNIbus Spot Important Alerts dashboards in the Last Hour folder.

If drill down is required for these charts you can use the default **Tree Map** specification instead. To change specification, click the chart settings icon on the right of the chart and change **Chart Type** from **OMNIbus Tree Map** to **Tree Map**.

Help text not rendered correctly for Event Analysis and Reduction in bi-directional locales.
This affects the help text in the **Event Analysis and Reduction** dashboards for **Analyze and Reduce Event Volumes**, **Introduction to the Apps**, and **Helpful Links**. The help text is not rendered correctly in the Arabic and Hebrew locales.

You can view text directly as an HTML file in a browser that supports bi-directional locales. The relevant files are `Analyze_and_reduce_event_volumes.html`, `Introduction_to_the_Apps.html`,and `Helpful_links.html`. The files are located in `$UNITY_HOME/AppFramework/Apps/OMNIbusInsightPack_v1.3.0.2/locale/<LOCALE>/LC_MESSAGES`.

Globalization
The following issues affect the non-English language versions of Netcool Operations Insight.

Event Search: count and time stamp hover help are not translated in some charts
This issue affects the OMNIbusInsightPack_v1.3.1. The following texts are not translated and appear in English only:

- Hover help for time stamp and count in stacked bar charts, heat maps and pie charts
- Legend for the **Count** field in the "**Hotspots by Node and AlertGroup**" chart and "**Hotspots by AlertGroup and Severity**" chart of the xxxOMNIbus Static Dashboard custom app
- Legend for the **Count** field in the "**Last Day - Hotspots by AlertGroup and Severity**" chart, "**'Last Day - Event Storm by Node**" chart, and "**Last Day - Hotspots by Node and AlertGroup**" of the OMNIbus Dynamic Dashboard custom app

Network Health Dashboard: In the Top Performers widget, the Japanese translation of the time to refresh is not displayed correctly
In the **Top Performers** widget, the Japanese translation of the time to refresh is not displayed correctly. The correct word order in Japanese has the number preceding the text; however, the word order displayed has the number following the text.

Topology Search: error message only partially translated
This issue affects the Network Manager Insight Pack V1.3.0.0. If you attempt to run a topology search between two nodes on the same device, the error message that is displayed is only partially translated. The error message in full is as follows:

```
An error occurred calculating the path between 'X' and 'X',
The source and destination Node's cannot be the same
```

(Where *X* is the value of the **NmosObjInst** for the device. The first half of the message, `An error occurred calculating the path between 'X' and 'X'`, is translated. The second half of the message `The source and destination Node's cannot be the same` is not translated and always appears in English.

IBM Installation Manager

Console mode: cannot install Netcool/OMNIbus core and Web GUI or Netcool/Impact at the same time.
When you install Netcool/OMNIbus core and Web GUI or Netcool/Impact at the same time with Installation Manager - console mode, the installation paths for Web GUI and Netcool/Impact are not prompted for and the installation fails. If you are performing the installation with Installation Manager - console mode, you must install the components separately.

Note: Installing Jazz for Service Management and IBM WebSphere Application Server is not supported for Installation Manager - console mode.

Network Health Dashboard

The known problems vary depending on which version of Network Manager is installed.

If you have Network Manager 4.2.0.4 installed

There are no known problems for the **Network Health Dashboard**.

If you have Network Manager 4.2.0.3 installed

The **Network Health Dashboard** has the following known problems.

Must install Network Health Dashboard with the Network Manager GUI Components

For Fix Packs 1, 2, and 3, if you want to install the **Network Health Dashboard**, you must install it in the same IBM Installation Manager session as the Network Manager GUI Components. Otherwise, some interaction between widgets and pages might not work.

If you upgrade the Network Manager GUI Components from Fix Pack 1 to 2, or 2 to 3, you must upgrade the **Network Health Dashboard** in the same IBM Installation Manager session.

If you roll back the Network Manager GUI Components from Fix Pack 2 to Fix Pack 1, you must roll back the **Network Health Dashboard** in the same IBM Installation Manager session.

If you roll back the Network Manager GUI Components from Fix Pack 3 to Fix Pack 2, you do not need to roll back the **Network Health Dashboard** in the same IBM Installation Manager session.

If you encounter problems with upgrading or rolling back the Network Manager GUI Components, apply the following workaround: Uninstall the **Network Health Dashboard**, retry the GUI upgrade or rollback, then reinstall the **Network Health Dashboard**.

Network Health Dashboard does not work after rolling back from Fix Pack 3

If you update from Network Manager 4.2 to 4.2 Fix Pack 3 and then roll back to 4.2, the **Network Health Dashboard** does not function. Take a backup of your system before upgrading, and restore the backup instead of rolling back the product.

You must install the Network Health Dashboard with the Network Manager GUI Components

If you want to install the **Network Health Dashboard**. you must install it in the same IBM Installation Manager session as the Network Manager GUI Components. Otherwise, some interaction between widgets and pages might not work.

You must roll back the Network Manager GUI Components and the Network Health Dashboard packages together

If you have installed the Network Manager 4.2 GA release and the **Network Health Dashboard**, and you then updated to 4.2 Fix Pack 1, and at this point you want to roll back one or more packages, then you must roll back the Network Manager GUI Components and the **Network Health Dashboard** packages together. If you roll back one of these packages without the other, there is a risk that your system will become corrupted.

Operations Analytics - Log Analysis

Operations Analytics - Log Analysis installation fails

Operations Analytics - Log Analysis V1.3.x requires a minimum 5 GB disk space on all supported operating systems. If less than 5 GB is found, the installer hangs. No error message is displayed. If this problem occurs, terminate the installer.

Support

IBM Electronic Support offers a portfolio of online support tools and resources that provides comprehensive technical information to diagnose and resolve problems and maintain your IBM products. IBM has many smart online tools and proactive features that can help you prevent problems from occurring in the first place, or quickly and easily troubleshoot problems when they occur. For more information, see:

http://www.ibm.com/support/electronicsupport

Related concepts

Solution overview

Read about key concepts and capabilities of IBM Netcool Operations Insight.

Appendix B. Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement might not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758 U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

If you are viewing this information in softcopy form, the photographs and color illustrations might not be displayed.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml.

Adobe, Acrobat, PostScript and all Adobe-based trademarks are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other product and service names might be trademarks of IBM or other companies.

Part Number:

Printed in the Republic of Ireland

(1P) P/N:

SC27-8601-11