

Netcool Agile Service Manager  
1.1.10

*Installation, Administration and User  
Guide  
04 June 2021*



**Note**

Before using this information and the product it supports, read the information in [“Notices” on page 417](#).

This edition applies to Version 1.1.9 of IBM Netcool Agile Service Manager (product number 5725-Q09) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 2016, 2021.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

# Contents

<b>Tables.....</b>	<b>vii</b>
<b>Preface.....</b>	<b>xi</b>
About the latest release.....	xi
<b>Chapter 1. Product overview.....</b>	<b>1</b>
Components.....	3
Glossary.....	6
<b>Chapter 2. Planning.....</b>	<b>9</b>
Hardware requirements.....	9
Software requirements.....	10
Sizing reference.....	11
<b>Chapter 3. Installing Agile Service Manager.....</b>	<b>13</b>
Installing on-prem.....	13
Installing the Agile Service Manager services.....	13
Adding new observers (on-prem).....	28
IBM Installation Manager.....	28
Installing the hybrid integration kit.....	30
Configuring DASH user roles.....	32
Uninstalling a legacy Agile Service Manager UI using the Installation Manager.....	33
Uninstalling the Netcool Agile Service Manager core services.....	34
Installing on OCP.....	35
Adding new observers (on OCP).....	35
Installing on OCP (Watson AIOps).....	36
Preparing the OCP cluster.....	36
Installing on OCP for Watson AIOps (without NOI).....	38
Configuring alternative storage.....	45
Uninstalling on OCP for Watson AIOps.....	45
Agile Service Manager Operator properties.....	47
Sizing reference (Watson AIOps only).....	50
<b>Chapter 4. Configuring components.....</b>	<b>53</b>
Configuring the Jenkins plugin.....	53
Refining Jenkins integration and visualization.....	57
Configuring application discovery service.....	61
Application discovery service configuration template.....	65
Supported application discovery accessListTypes and properties.....	68
Configuring Db2 for application discovery (OCP only).....	70
Configuring the network discovery services.....	85
Discoverable technologies and devices.....	90
Configuring EMS and SFTP for collectors.....	92
Network discovery service configuration template.....	94
Configuring a hybrid system.....	96
Configuring Humio integration (on-prem only).....	96
<b>Chapter 5. Running Observer jobs.....</b>	<b>99</b>
Defining observer security.....	99

Configuring encryption and authentication.....	99
Defining observer jobs using the UI.....	103
Configuring ALM Observer jobs.....	104
Configuring Ansible AWX Observer jobs.....	106
Configuring Application Discovery Observer jobs.....	108
Configuring AppDynamics Observer jobs.....	110
Configuring AWS Observer jobs.....	111
Configuring Azure Observer jobs.....	113
Configuring BigFix Inventory Observer jobs.....	115
Configuring Ciena Blue Planet Observer jobs.....	117
Configuring Cisco ACI Observer jobs.....	120
Configuring Contrail Observer jobs.....	123
Configuring DNS Observer jobs.....	127
Configuring Docker Observer jobs.....	129
Configuring Dynatrace Observer jobs.....	132
Configuring File Observer jobs.....	134
Configuring GoogleCloud Observer jobs.....	136
Configuring IBM Cloud Observer jobs.....	137
Configuring ITNM Observer jobs.....	139
Configuring Jenkins Observer jobs.....	141
Configuring Juniper CSO Observer jobs.....	142
Configuring Kubernetes Observer jobs.....	144
Configuring Network Discovery Observer jobs.....	148
Configuring New Relic Observer jobs.....	150
Configuring OpenStack Observer jobs.....	152
Configuring Rancher Observer jobs.....	157
Configuring REST Observer jobs.....	159
Configuring ServiceNow Observer jobs.....	161
Configuring TADDM Observer jobs.....	162
Configuring VMware NSX Observer jobs.....	164
Configuring VMware vCenter Observer jobs.....	166
Configuring Zabbix Observer jobs.....	169
Defining observer jobs manually.....	171
Defining ALM Observer jobs.....	172
Defining Ansible AWX Observer jobs.....	173
Defining AppDynamics Observer jobs.....	175
Defining Application Discovery Observer jobs.....	177
Defining AWS Observer jobs.....	179
Defining Azure Observer jobs.....	181
Defining BigFix Inventory Observer jobs.....	183
Defining Ciena Blue Planet Observer jobs.....	184
Defining Cisco ACI Observer jobs.....	187
Defining Contrail Observer jobs.....	190
Defining DNS Observer jobs.....	194
Defining Docker Observer jobs.....	196
Defining Dynatrace Observer jobs.....	200
Defining File Observer jobs.....	201
Defining GoogleCloud Observer jobs.....	203
Defining IBM Cloud Observer jobs.....	205
Defining ITNM Observer jobs.....	207
Defining Jenkins Observer jobs.....	208
Defining Juniper CSO Observer jobs.....	209
Defining Kubernetes Observer jobs.....	211
Defining Network Discovery Observer jobs.....	217
Defining New Relic Observer jobs.....	220
Defining OpenStack Observer jobs.....	222
Defining Rancher Observer jobs.....	225
Defining REST Observer jobs.....	227

Defining ServiceNow Observer jobs.....	232
Defining TADDM Observer jobs.....	234
Defining VMware NSX Observer jobs.....	237
Defining VMware vCenter Observer jobs.....	239
Defining Zabbix Observer jobs.....	242
<b>Chapter 6. Using Agile Service Manager.....</b>	<b>245</b>
Logging into the UI (OCP).....	245
Logging into the UI (Watson AIOps).....	246
Accessing the Topology Viewer in DASH (on-prem).....	246
Accessing topologies via direct-launch URL string.....	247
Rendering static topologies via UI-API.....	250
Rendering a topology.....	253
Viewing a topology.....	255
Viewing topology history.....	260
Rebuilding a topology.....	262
Performing topology administration.....	263
Using the topology dashboard.....	265
Using the Path tool.....	267
<b>Chapter 7. Administration.....</b>	<b>269</b>
Configuring authentication.....	269
Configuring access to core services.....	269
Changing default passwords.....	270
Customizing UI elements.....	271
Configuring custom tools.....	271
Defining custom icons.....	279
Editing resource type styles.....	281
Creating custom relationship type styles.....	283
Defining advanced topology settings.....	285
Configuring resource history TTL.....	287
Increasing default timeouts.....	288
Configuring alternate storage (OCP).....	289
Porting data for testing, backup and recovery.....	290
Backing up and restoring database data (on-prem).....	290
Backing up UI configuration data (on-prem).....	292
Restoring UI configuration data (on-prem).....	294
Backing up database data (OCP).....	295
Restoring database data (OCP).....	299
Backing up and restoring UI configuration data (OCP).....	305
Launching in context from Event Viewer.....	307
Updating a topology on the same DASH page.....	307
Updating a topology on a different DASH page.....	308
Launch-in-context parameters.....	308
Defining rules.....	309
Rules examples and samples.....	312
Using topology templates.....	315
Improving database performance.....	318
Changing gc_grace_seconds (OCP).....	318
Changing gc_grace_seconds (on-prem).....	320
Changing dlocal_read_repair_chance (OCP).....	321
Scaling (OCP).....	322
Scaling vertically.....	323
Scaling horizontally.....	323
System health and logging.....	336
Viewing the service logs (on-prem).....	337
Viewing the logs (OCP).....	338

<b>Chapter 8. Troubleshooting.....</b>	<b>341</b>
Installation troubleshooting (on-prem).....	341
Probe and gateway troubleshooting.....	342
Startup and UI troubleshooting.....	344
Search troubleshooting.....	345
Observer troubleshooting.....	346
Other troubleshooting.....	348
OCP troubleshooting.....	349
<b>Chapter 9. Reference.....</b>	<b>353</b>
Topology service reference.....	353
Properties.....	354
Edge labels.....	356
Edge types.....	357
Entity types.....	360
REST API.....	366
Status (and state).....	367
Timestamps.....	368
Cookbook.....	369
Virtual machine recipe.....	370
Physical device recipe.....	379
XML Gateway reference [deprecated from V. 1.1.6.1].....	382
Probe for Message Bus reference [deprecated from V 1.1.6.1].....	387
Example probe rules file.....	389
Status service reference.....	391
Topology viewer reference.....	393
Topology tools reference.....	404
Custom icons reference.....	408
Example sysctl.conf file.....	409
Swagger reference.....	410
Installation parameters.....	413
<b>Notices.....</b>	<b>417</b>
Trademarks.....	418

---

# Tables

1. Agile Service Manager core packages.....	3
2. Agile Service Manager observer packages .....	4
3. Netcool Agile Service Manager Core hardware requirements.....	9
4. Netcool Agile Service Manager software requirements.....	10
5. Agile Service Manager commands.....	26
6. Mandatory Agile Service Manager cloud properties.....	47
7. Full Agile Service Manager cloud properties.....	47
8. General sizing requirements.....	50
9. Total hardware requirements Agile Service Manager standalone Including OCP and infrastructure nodes.....	51
10. Total hardware requirements Agile Service Manager standalone services only.....	51
11. Hardware allocation and configuration.....	51
12. Agile Service Manager requirements for the Jenkins plugin.....	53
13. Application discovery model objects mapped to Agile Service Manager entity types.....	62
14. Supported accessListTypes and their properties.....	68
15. Db 2 configurable parameters and default values.....	80
16. Default network discovery users and roles.....	87
17. ALM Observer parameters for alm jobs.....	104
18. ALM Observer parameters for ALM rm (Resource Manager) jobs.....	105
19. Ansible AWX Observer Load parameters.....	106
20. Application Discovery Observer job parameters.....	109
21. AppDynamics Observer parameters.....	110
22. AWS Observer parameters.....	112

23. Azure Observer parameters.....	114
24. Bigfix Inventory Observer job parameters.....	115
25. Ciena Blue Planet Observer restapi Load parameters.....	118
26. Ciena Blue Planet Observer Websocket Listen parameters.....	118
27. Cisco ACI Observer restapi and websocket job parameters.....	121
28. Contrail Observer rabbitmq job parameters.....	124
29. Contrail Observer restapi job parameters.....	125
30. DNS Observer reverse job parameters.....	127
31. DNS Observer forward job parameters.....	128
32. Docker Observer job parameters.....	129
33. Dynatrace Observer job parameters.....	132
34. File Observer job parameters.....	135
35. GoogleCloud Observer parameters.....	136
36. IBM Cloud Observer job parameters.....	138
37. ITNM Observer load and listen job parameters.....	139
38. Jenkins Observer job parameters.....	141
39. Juniper CSO Observer job parameters.....	143
40. Kubernetes Observer load job parameters.....	145
41. Kubernetes Observer weave_scope job parameters.....	147
42. Network Discovery Observer job parameters.....	149
43. New Relic job parameters.....	151
44. OpenStack Observer restapi job parameters.....	153
45. OpenStack Observer rabbitmq job parameters.....	155
46. Rancher Observer parameters.....	157
47. REST Observer listen and bulk replace job parameters.....	160



48. ServiceNow Observer job parameters.....	162
49. TADDM Observer load job parameters.....	163
50. VMware NSX Observer job parameters.....	165
51. VMware vCenter Observer job parameters.....	167
52. Zabbix Observer parameters.....	169
53. Ciena Blue Planet Observer restapi Load parameters.....	185
54. Additional Ciena Blue Planet Observer Websocket Listen parameters.....	186
55. Encryption parameters required for ciscoaci_observer_common.sh.....	189
56. Mapping of Contrail object types to Agile Service Manager entity types:.....	191
57. Mapping IBM Cloud model objects to Agile Service Manager entity types.....	205
58. Mapping of ServiceNow object types to Agile Service Manager entity types:.....	233
59. Mapping TADDM model objects to Agile Service Manager entity types.....	235
60. Encryption parameters required for vmwarensx_observer_common.sh.....	238
61. Encryption parameters required for vmvcenter_observer_common.sh.....	241
62. Encryption parameters required for zabbix_observer_common.sh.....	243
63. Severity levels.....	258
64. TTL example for the 'sprocket' resource.....	288
65. Launch-in-context parameters.....	308
66. Log names and directories for Netcool Agile Service Manager services.....	337
67. Generic properties.....	354
68. Edge types for the Aggregation edge labels.....	357
69. Edge types for the Association edge labels.....	357
70. Edge types for the Data flow edge labels.....	358
71. Edge types for the Dependency edge labels.....	359
72. Predefined entity types and icons, where defined.....	360

73. General event state rules.....	385
74. Use of Netcool/OMNIbus alerts.status event fields by Agile Service Manager.....	385
75. Netcool/OMNIbus event data mapped onto Topology Service status.....	386
76. General event state rules.....	392
77. Use of Netcool/OMNIbus alerts.status event fields by Agile Service Manager.....	392
78. Netcool/OMNIbus event data mapped onto Topology Service status.....	393
79. Severity levels.....	400
80. Default Swagger URLs for Agile Service Manager services.....	411
81. Default Swagger URLs for Agile Service Manager observers.....	412
82. Helm installation parameters.....	413

# Preface

---

This PDF document contains topics from the Knowledge Center in a printable format.

## About the latest release

---

Agile Service Manager Version 1.1.10 is available.

### What's new in Version 1.1.10

**Documentation updated:** 04 June 2021

#### Updates to observer security configuration details

[“Configuring Ansible AWX Observer jobs” on page 106](#)

[“Configuring BigFix Inventory Observer jobs” on page 115](#)

[“Configuring Ciena Blue Planet Observer jobs” on page 117](#)

[“Configuring Cisco ACI Observer jobs” on page 120](#)

[“Configuring Docker Observer jobs” on page 129](#)

[“Configuring Dynatrace Observer jobs” on page 132](#)

[“Configuring ITNM Observer jobs” on page 139](#)

[“Defining ITNM Observer jobs” on page 207](#)

[“Configuring Juniper CSO Observer jobs” on page 142](#)

[“Configuring Kubernetes Observer jobs” on page 144](#)

[“Configuring OpenStack Observer jobs” on page 152](#)

[“Configuring Rancher Observer jobs” on page 157](#)

[“Configuring VMware NSX Observer jobs” on page 164](#)

[“Configuring VMware vCenter Observer jobs” on page 166](#)

[“Configuring Zabbix Observer jobs” on page 169](#)

**Documentation updated:** 17 December 2020

#### New Path tool topic

[“Using the Path tool” on page 267](#)

#### New Splunk custom tool example

[“Splunk integration example” on page 278](#)

#### New 'historical search' function

[Search for a resource](#)

#### New integration script to create Humio search tool

[“Configuring Humio integration \(on-prem only\)” on page 96](#)

#### New network discovery and devices overview

[“Discoverable technologies and devices” on page 90](#)

#### New application discovery sensor reference section

**Knowledge Center PDF:** [Application Discovery Sensor Reference](#)

**Knowledge Center HTML:** [Application discovery sensor reference](#)

#### New troubleshooting topics

[“Apply Elasticsearch patch during upgrade” on page 341](#)

[“System health check job fails \(after rollback\)” on page 351](#)



---

# Chapter 1. Product overview

IBM Netcool 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. **Agile Service Manager is available as both on-prem and RedHat OpenShift Container Platform (OCP) versions.**

## Benefits of Netcool Agile Service Manager

Services and applications are increasingly deployed in environments that take advantage of distributed and often virtualized infrastructure. For example, parts of a network might be cloud-based, with other connected elements contained within, or tethered to, legacy systems that exploit tried and tested on-prem capability. The result is often a highly distributed and increasingly complex hybrid network that requires an agile and dynamic operations management solution in order to leverage and exploit its rapidly evolving technologies.

Netcool Agile Service Manager allows the real-time view, support and management of highly dynamic infrastructures and services. By visualizing complex network topologies in real-time, updated dynamically or on-demand, and allowing further investigation of events, incidents and performance, operational efficiency is improved, problems are detected and solved faster, false alarms are reduced, and automation and collaboration between operational teams is improved. Also, data can be leveraged more efficiently both in real time and historically, thereby empowering teams and systems to create and nurture differentiated services for different customers.

IBM Netcool Agile Service Manager is cloud-born, and built on secure, robust and proven technologies. It is designed to be flexible and can be extended as needed using plug-in components and micro-services to cater for highly specific environments.

## Agile Service Manager observers


Observers are Agile Service Manager's means of obtaining topology information from a specific source, which may be inclusive of state. Typically, an observer aims to get a near-real-time stream of data from a source, but in cases where this is not possible or applicable, they can revert to other techniques, such as database or REST API queries.

Each observer minimally needs to run one 'job' and provides a set of REST APIs to enable job management. Once an observer has been given a job request on behalf of a tenant, it runs the job and sends data to the Agile Service Manager topology service; predominantly via Kafka, although some meta-data exchange is done RESTfully.

The data originating from an observer is considered to be name-spaced by the tenant ID running the job and the notion of a 'provider', which is up to the observer to determine. For example, the File Observer's provider is the string 'FILE.OBSERVER:file name to load'. This ensures that subsequent loads of data for the same file can be replaced or updated as needed.

## Basic deployment

Netcool Agile Service Manager is deployed with IBM Tivoli Netcool Operations Insight as part of an integrated solution. This figure depicts the basic Agile Service Manager **on-prem** architecture.



## Deployment scenarios

### Network Manager

You want to use Netcool Agile Service Manager to analyze the resource data discovered by Network Manager.

You configure the ITNM Observer to load topology data, and then monitor Network Manager for updates.

You define a seed resource in the Agile Service Manager UI, and then dynamically render a topology view centered around that resource, which can display linked resources up to four hops away.

You use this visualization to delve into the states, histories and relationships of the resources displayed.

New data is harvested continuously, which you can then analyze further.

### Netcool/OMNIBus

You want to extend your analysis of Netcool/OMNIBus events.

You configure the Agile Service Manager status service and the probe and gateway services to monitor the Netcool/OMNIBus ObjectServer for new events.

You configure the probe and gateway services to synchronize event views across the Netcool Agile Service Topology Viewer and the Netcool/OMNIBus Event Viewer.

You display a topology based on a specific resource (event), and then exploit Netcool Agile Service Manager's functionality to gain further insights into the displayed events.

### OpenStack

You use the OpenStack Observer to render detailed OpenStack topologies, and delve further into their states, histories and relationships.

### Bespoke topologies using the REST APIs

You want to load resource data from your own source in order to use the Netcool Agile Service Manager functionality to render topologies for analysis.

You use the REST APIs to configure a data source, load your data, and then use the Netcool Agile Service Manager UI to focus on a specific seed resource, before extending your topology outward.

## Components

Agile Service Manager consists of a number of services, and can be integrated into the IBM Netcool Operations Insight suite of products. You access Agile Service Manager through the IBM Dashboard Application Service Hub (DASH).

### Agile Service Manager core download packages

The Agile Service Manager core eAssembly consists of the following packages. Apart from the UI, which is installed using the IBM Installation Manager, all core packages are Docker containers.

Table 1. Agile Service Manager core packages	
Package	Details
nasm-cassandra	A distributed and robust database that is scalable while maintaining high performance
nasm-common	Component that contains the product licenses, common scripts and docker-compose
nasm-elasticsearch	A distributed search and analytics engine that is scalable and reliable
nasm-kafka	<p>A message bus that efficiently consolidates topology data from multiple sources</p> <p>In addition to the Kafka message bus, the nasm-kafka service also deploys the Kafka REST API, which verifies the existence of Kafka topics.</p>
nasm-layout	A service that lets you customize the way topologies are structured, providing a number of standard options, such as hierarchical, force-directed, and other views
nasm-merge	A service that lets you merge duplicate records of the same resource retrieved through different mechanisms into one composite resource
nasm-nginx	A service that manages access to all other Agile Service Manager micro-services
nasm-noi-gateway	A service that updates the Agile Service Manager status with Netcool/OMNIBus event data (via the status service)
nasm-noi-probe	A service that uses Agile Service Manager resource status information to generate events in the Netcool/OMNIBus Event Viewer.
nasm-observer	A service that you install on a Jenkins server, from where it updates the Agile Service Manager status (via the Jenkins Observer)
nasm-search	A service that inserts topology data into the Elasticsearch engine, and exposes REST APIs to search for resources
nasm-status	A service that extracts information from IBM Tivoli Netcool/OMNIBus events
nasm-topology	A service that lets you query networked resources, and retrieve both real-time and historical information about their state and relationships with other linked resources

<i>Table 1. Agile Service Manager core packages (continued)</i>	
<b>Package</b>	<b>Details</b>
nasm-ui-api	A service whose dedicated purpose is to provide topology-related data to the Agile Service Manager UI
nasm-ui	A user interface service accessed through DASH, which presents you with a topology view and lets you perform a number of further tasks in context
nasm-zookeeper	A robust, distributed and scalable synchronization service

## Agile Service Manager observer download packages

<i>Table 2. Agile Service Manager observer packages</i>	
<b>Package</b>	<b>Details</b>
nasm-alm-observer	A service that extracts information from the IBM Agile Lifecycle Manager
nasm-appdisco-observer	A service that retrieves discovered applications, their components, configurations and dependencies, from the Application Discovery database via REST API
nasm-appdynamics-observer	A service that obtains information from the AppDynamics Controller via the REST API
nasm-aws-observer	A service that reads data from the Amazon Web Services
nasm-azure-observer	A service that reads data from Azure cloud services through its REST APIs
nasm-bigfixinventory-observer	A service that reads data from a Bigfix Inventory instance through its REST API
nasm-cienablueplanet-observer	A service that retrieves topology data from the Blue Planet MCP instance via REST API
nasm-ciscoaci-observer	A service that makes REST calls to Cisco APIC in the Cisco ACI environment
nasm-contrail-observer	A service that makes REST calls to the Contrail API server to retrieve topology data from Juniper Network Contrail
nasm-dns-observer	A service that queries internal DNS servers, and returns response times and service addresses
nasm-docker-observer	A service that extracts information from Docker networks
nasm-dynatrace-observer	A service that queries a specified Dynatrace environment for information about its applications, services, process groups, and infrastructure entities
nasm-file-observer	A service that retrieves data written to a file in a specific format
nasm-google-observer	A service that reads data from the Google Cloud Platform's Compute Services through Google's Compute Services SDK.
nasm-ibmcloud-observer	A service that performs REST calls to the IBM Cloud REST API, which retrieve Cloud Foundry Apps information and services
nasm-itnm-observer	A service that extracts information from the IBM Tivoli Network Manager IP Edition database



<i>Table 2. Agile Service Manager observer packages (continued)</i>	
<b>Package</b>	<b>Details</b>
nasm-jenkins-observer	A service that receives build information generated by the Agile Service Manager plugin for Jenkins
nasm-juniperco-observer	A service that extracts information from Juniper Contrail Service Orchestration (CSO)
nasm-kubernetes-observer	A service that discovers Kubernetes services containers and maps relationships between them
asm-netdiso-observer	A service that retrieves discovered devices and relationships from the Network Discovery database via REST API
nasm-newrelic-observer	A service that loads New Relic Infrastructure resource data via a New Relic account with a New Relic Infrastructure subscription
nasm-openstack-observer	A service that extracts information from OpenStack
nasm-rancher-observer	A service that read data from Rancher through the REST API
nasm-rest-observer	A service that obtains topology data via REST endpoints
nasm-servicenow-observer	A service that performs REST calls to retrieve configuration management database (CMDB) data from ServiceNow
nasm-taddm-observer	A service that extracts information from the IBM Tivoli Application Dependency Discovery Manager database
nasm-vmvcenter-observer	A service that dynamically loads VMware vCenter data
nasm-vmwarensx-observer	A service that dynamically loads VMware NSX data
nasm-zabbix-observer	A service that extracts server information and its associated network resources from Zabbix via REST RPC

## **Agile Service Manager application and network discovery download packages**

In addition to the Network Discovery and Application Discovery observers, both listed in the observer packages table, Agile Service Manager contains the following application and network discovery download packages.

- Network Discovery Config
- Network Discovery Consul
- Network Discovery Control
- Network Discovery Schema Registry
- Network Discovery Security
- Network Discovery Sftp
- Network Discovery Sidecar
- Network Discovery Status
- Network Discovery Swagger Service
- Network Discovery Swagger UI
- Network Discovery Topogram
- Network Discovery Worker
- Network Discovery Collector
- Application Discovery Service
- Application Discovery Primary Storage

- DB2

### Related reference

[“Swagger reference” on page 410](#)

Specific links to Agile Service Manager Swagger documentation are included in many of the topics, as and when useful. This topic summarizes some of that information in a single location, for example by listing the default ports and Swagger URLs for each Agile Service Manager service.

### Related information

[IBM Netcool Agile Service Manager download document](#)

## Glossary

---

Refer to the following list of terms and definitions to learn about important Netcool Agile Service Manager concepts.

### Netcool Agile Service Manager terminology

#### CASE (Container Application Software for Enterprises)

The Container Application Software for Enterprises (CASE) specification defines metadata and structure for packaging a containerized application.

CASE provides a common packaging definition for the many ways to install, manage, operate, upgrade, and uninstall complex software using multiple roles and personas.

#### edge

An edge is a relationship between resources, also simply referred to as the 'link' between resources.

Edges have a *label*, which allocates them to a family of edges with specific behavior and governs how they are displayed in the UI, and an *edgeType*, which defines the relationship in real terms.

#### hop

A hop is a step along a single edge from one resource to another.

All resources that are connected directly to a seed resource are one hop removed, while those connected to the secondary resources are two hops removed from the seed resource, and so on.

Netcool Agile Service Manager displays topologies with resources up to four hops removed from the seed resource by default, which is configurable up to 30 hops.

**CAUTION:** Do not increase the hop count beyond your system's ability to cope. A large hop count can result in a very large topology, and rendering this can lead to timeout errors.

#### OCP

RedHat OpenShift container platform (OCP) is an automated Kubernetes container platform you can use to deploy and manage cloud applications.

#### observer

An observer is a service that extracts resource information and inserts it into the Agile Service Manager database.

Agile Service Manager includes a configuration UI to help you configure and run observer jobs.

#### observer job

The access details for a target system are defined in an observer job, which is triggered to retrieve data.

Observer jobs are configured and run from the Observer Configuration UI, and can be long-running or transient. For example, the Network Manager Observer topology 'load' job is a one-off, transient job (unless scheduled to run at set intervals), while the Network Manager 'listen' job is long-running and runs until explicitly stopped, or until the observer is stopped.

You can configure observer jobs manually by editing the configuration files for specific observer jobs, instead of using the Observer Configuration UI.

For the OCP version of Agile Service Manager, observer jobs are defined and run using Swagger.

## observer job script

In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## OLM (Operator Lifecycle Manager)

The Red Hat OpenShift Operator Lifecycle Manager (OLM) UI helps users install, update, and manage the lifecycle of all Operators and their associated services running across their clusters. It is part of the Operator Framework, an open source toolkit designed to manage Kubernetes native applications (Operators) in an effective, automated, and scalable way.

## provider

A provider is usually a single data source within the scope of a tenant.

**Note:** A provider's **uniqueId** property for a resource is unique only within the scope of a provider.

## proxy

The Agile Service Manager Nginx proxy server (nasm-nginx) manages access to all other Agile Service Manager micro-services by rewriting URLs.

## resource

A resource is a node in an interconnected topology, sometimes also referred to as a vertex, or simply a node. It can be anything in a user-specific topology that has been designated as such, for example a hardware or virtual device, a location, a user, or an application.

## scaling

You can scale up your system horizontally or vertically.

**Horizontal scaling** means increasing the replication factor of a particular service, and may also require adding additional hardware.

**Vertical scaling** means that you add more power (CPU or RAM) to an existing machine.

## seed

A seed is a single resource that has been chosen as the starting point of a topology. Once defined, the topology view is expanded one 'hop' at a time (the number of hops are configurable with a maximum of 30).

## status

Status is a property of one or more resources, and a single resource can have different types of status.

Each status can be in one of three states: 'open', 'clear' or 'closed'.

The severity can be one of the following: 'clear', 'information', 'indeterminate', 'warning', 'minor', 'major' or 'critical'.

The status of a resource can be derived from events, in the case of the resource having been retrieved via the status service, or it can be supplied when resources are posted to the topology service.

## Swagger

Agile Service Manager uses Swagger for automated documentation generation and utilizes a Swagger server for each micro-service.

You can access and explore the REST APIs of the topology service and observers using Swagger via the proxy service.

## tenant

A tenant is represented by a globally unique identifier, its tenant ID.

The default tenant ID is: cfd95b7e-3bc7-4006-a4a8-a73a79c71255

## topology

The arrangement of interconnected resources within a network, viewed in the Netcool Agile Service Manager UI.

## More information:

You can find additional information on the topology service in the [“Topology service reference”](#) on page 353 section.

More detailed information on the topology screen elements are in the [“Topology viewer reference”](#) on page 393 section.

More information on Swagger is included in the documentation where appropriate. You can find a list of the default Swagger URLs and ports here: [“Default Swagger URLs”](#) on page 411

## Chapter 2. Planning

This section helps you to plan your installation and use of Netcool Agile Service Manager by listing the minimum software and hardware requirements.

### Hardware requirements

This section lists the minimum hardware requirements.

Your minimum hardware requirements are determined by the needs of the components of your specific Netcool Operations Insight solution.

#### OCP hardware requirements:

See the following topic for information on hardware sizing for a Netcool Operations Insight deployment on Red Hat OpenShift: [https://www.ibm.com/support/knowledgecenter/SSTPTP\\_1.6.3/com.ibm.netcool\\_ops.doc/soc/integration/reference/soc\\_sizing\\_full.html](https://www.ibm.com/support/knowledgecenter/SSTPTP_1.6.3/com.ibm.netcool_ops.doc/soc/integration/reference/soc_sizing_full.html)

**Tip:** If you are installing IBM Common Services as part of an Agile Service Manager deployment on RedHat OpenShift Container Platform (OCP), you must add an additional 32 Gb of RAM to the worker node on which IBM Common Services is to be installed.

#### On-prem hardware requirements:

For its on-prem edition, the Agile Service Manager core components are deployed to a single server and the following physical or virtual hardware requirements must be met.

Specifically, a number of Kernel parameters must be configured to optimize Cassandra and Elasticsearch, which run better when Swap is either disabled or the Kernel **vm.swappiness** parameter is set to 1 (in the `sysctl.conf` file). This setting reduces the Kernel's tendency to swap and should not lead to swapping under normal circumstances, while still allowing the whole system to swap under emergency conditions. See the following topic in the reference section for an example of a `sysctl.conf` file: [“Example sysctl.conf file” on page 409](#)

Table 3. Netcool Agile Service Manager Core hardware requirements		
Requirement	Setting	Notes
CPU	32 cores	<b>Recommendation</b> You may need to increase the number of cores to 48 for heavy workloads, where an expected event input rate of 50 events per second (or more) is expected, and large numbers of resources (one million or more) are being managed.
Memory	64 GB	

Table 3. Netcool Agile Service Manager Core hardware requirements (continued)		
Requirement	Setting	Notes
Disk	500 GB	<b>Recommendations</b> Use disk arrays with redundancy, such as RAID10, with a minimum of 1000 IOPS. Have separate disks for the following services under the \$ASM_HOME/data directory (by creating mount-points in your Operating System): <ul style="list-style-type: none"> <li>• \$ASM_HOME/data/cassandra</li> <li>• \$ASM_HOME/data/elasticsearch</li> <li>• \$ASM_HOME/data/kafka</li> <li>• \$ASM_HOME/data/zookeeper</li> </ul>

#### Related information

Netcool Operations Insight Knowledge Center

## Software requirements

This section lists the minimum software requirements.

Netcool Agile Service Manager Core has the following requirements.

Table 4. Netcool Agile Service Manager software requirements	
Requirement	Details
Operating system	Red Hat Enterprise Linux 7 (x86-64) Apply the latest updates.
Docker	Docker for Red Hat Enterprise Linux Version 1.12 (or later) <a href="#">Installation</a> is described in the core installation topic. You can find more information about the Docker engine here: <a href="https://docs.docker.com/engine">https://docs.docker.com/engine</a>
WebSphere	WebSphere Application Server Version 8.5.5.17 (or later)
Java	IBM WebSphere SDK Java Technology Edition Version 7.0 (or later)
DASH	IBM Dashboard Application Service Hub (DASH) 3.1.3.7 (or later) Installed via Jazz for Service Management Version 1.1.3.9
Web GUI	Tivoli Netcool/OMNIBus Web GUI Version 8.1.0.21
IBM Installation Manager	Version 1.9.2

The Netcool Agile Service Manager User Interface is deployed into an existing DASH instance that has been deployed as part of a Netcool Operations Insight installation. See the [Netcool Operations Insight Version 1.6.3 Product and component version matrix](#) for more details on NOI components.

## Sizing reference

---

This topic includes references to the sizing (and other preparatory) requirements for the various Agile Service Manager deployment types.

### Sizing your system

#### **Agile Service Manager on OCP (with NOI)**

See the following NOI sizing topic: [Sizing for a Netcool Operations Insight on Red Hat OpenShift deployment](#)

#### **Agile Service Manager on-prem (with NOI)**

See the following NOI planning topics: [Planning for an on-premises installation](#)

#### **Agile Service Manager stand-alone on OCP (with Watson AIOps)**

See the following topic: [“Sizing reference for Agile Service Manager on OCP \(Watson AIOps\)” on page 50](#)





---

## Chapter 3. Installing Agile Service Manager

You can deploy Agile Service Manager either stand-alone or on RedHat OpenShift Container Platform (OCP). The installation and post-install configuration steps differ for these versions.

**Tip:** All versions of Agile Service Manager are installed together with Netcool Operations Insight.

**Important:** Agile Service Manager uses default passwords out of the box. For security reasons, change the default passwords after installation, as described in the [“Changing default passwords” on page 270](#) topic.

### Related concepts

[“Planning” on page 9](#)

This section helps you to plan your installation and use of Netcool Agile Service Manager by listing the minimum software and hardware requirements.

[“Configuring components” on page 53](#)

After installing Agile Service Manager, you configure the optional components, such as the Jenkins plugin. You can also configure a hybrid system consisting of a RedHat OpenShift Container Platform backend (core) and an Agile Service Manager UI

### Related tasks

[“Changing default passwords” on page 270](#)

Passwords are encrypted by default and stored in a number of environment (.env) files.

---

## Installing Agile Service Manager on-prem

To install Netcool Agile Service Manager, you complete a number of prerequisites tasks. You then install the Netcool Agile Service Manager core components, UI and observers, before deploying the Netcool Hybrid Deployment Option Integration Kit using the IBM Installation Manager.

**Updating your system:** If you have an existing installation of Agile Service Manager (Version 1.1.7 or earlier), follow this sequence to update your system:

1. Uninstall the old Agile Service Manager UI using IBM Installation Manager.
2. Install Agile Service Manager (core), including only the observers you require, and the optional application and network discovery components. Then run the `set-dash-config` script.
3. Install the Netcool Hybrid Deployment Option Integration Kit, then restart DASH.

## Installing the Agile Service Manager services

The Agile Service Manager application consists of several micro-services, which are provided as Docker containers. You can deploy these on a single server, though a multi-server deployment is required if the network and application discovery services are deployed in a production environment.

### Before you begin

Hardware and software requirements are listed in the [Planning](#) section.

**Updating your system:** If you are updating an existing installation with the latest version of Agile Service Manager, you may already have the prerequisites in place. Before updating an installation, **check** that you have the correct version of the prerequisites, and that you have applied any relevant upgrade steps documented in the release note upgrade topics.

**Tip:** You can use the `get_package_versions.sh` script to discover which Agile Service Manager packages are installed.

**Remember:** You must reinstall the Netcool Hybrid Deployment Option Integration Kit each time you reinstall or upgrade Agile Service Manager.

You must complete the following prerequisites before you can install Agile Service Manager.

1. Ensure that your Red Hat Enterprise Linux 7 (x86-64) operating system has the latest updates applied.
2. Ensure SELinux is disabled before performing the Agile Service Manager installation. To do so, edit the `/etc/selinux/config` file with a suitable editor and set `SELINUX=disabled`, before rebooting.
3. Enable the `rhel-x86_64-server-extras-7` and `rhel-7-server-optional-rpms` repositories so that the docker package can be installed. You can find more information on configuring the repository here: [https://access.redhat.com/documentation/en-US/Red\\_Hat\\_Enterprise\\_Linux/7/html/System\\_Administrators\\_Guide/sec-Configuring\\_Yum\\_and\\_Yum\\_Repositories.html#sec-Setting\\_main\\_Options](https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/System_Administrators_Guide/sec-Configuring_Yum_and_Yum_Repositories.html#sec-Setting_main_Options)
4. Using the following commands, create the Docker group and add your current user to that group to enable non-root access to the Docker daemon.

```
$ sudo groupadd docker
```

```
$ sudo usermod -aG docker $USER
```

More information on creating and administering Docker groups can be found here: <https://docs.docker.com/engine/installation/linux/linux-postinstall/>

**Important:** For the group changes to take effect, you must refresh your terminal session.

**Tip:** If you do not complete this step, you will either have to run the commands as the root user, or prefix your commands with the `sudo` command.

5. Obtain the Netcool Agile Service Manager installation image from the Passport Advantage site, and extract it to a temporary directory. More information can be found in the download document here: <http://www-01.ibm.com/support/docview.wss?uid=swg24043717>

**Note:** You need an IBM ID to access the download document.

6. Ensure you have IBM Dashboard Application Service Hub (DASH) 3.1.3.7 installed, and have configured access.

## About this task

When you install the Agile Service Manager components, they are loaded automatically.

You install **all** the core applications of Agile Service Manager including the UI, but **only** the observers that you require. The Docker Observer is a requirement, as it supplies Agile System Manager health view data in the Topology Viewer. In the unlikely event that you wish to install (and start up) all available observers, you must ensure that your system meets the minimum requirements listed here: [Table 3 on page 9](#)

If required, you also install the application and network discovery services. A single-server deployment of Agile Service Manager is only sufficient for a proof-of-concept deployment.

**Important:** For a production deployment, the application and network discovery services need to be deployed onto separate machines.

Finally, you configure communication between DASH and the Agile Service Manager UI by running the `set-dash-config.sh` script.

**Note:** The example data for software versions or directories used here may differ from your own scenario.

## Procedure

### Prepare the installation files

1. Move all Agile Service Manager packages to the installation target host.
  - Copy only the observers you intend to deploy to the installation directory, or delete any unwanted observer packages after you have downloaded them. Remember that the Docker Observer is required.

- Place the Agile Service Manager Base eAssembly and observer packages into the same directory.

**Important:** To prevent unwanted observers being installed and thereby placing unnecessary strain on your infrastructure, ensure that this directory contains **only** the observers you wish to install.

- For a multi-server production deployment of the application and (or) network discovery services, ensure that only the Network Discovery Observer and the Application Discovery Observer are present in the installation directory.

### Install Agile Service Manager core, UI and observers

2. From the directory where you have placed the packages, install Agile Service Manager using the yum install command.

```
sudo yum install nasm-*.rpm
```

**Tip:** While it is possible to specify each individual installation image, it is recommended that you perform a wildcard (\*) installation to ensure that all components are installed. Remember that you **must** ensure that **only** the observers you wish to install are present. To install an individual observer later, see the information in the following topic: [“Adding new observers to an existing on-prem installation” on page 28](#)

Yum will install Docker and all other nasm-\* components as required, including all observers found in that directory. During the installation of the packages, the related Docker images are loaded. No data can be retrieved, however, until observer jobs are defined.

3. Required: During a first installation or during upgrades, you will be prompted to review and accept the license. You must do so **after installation has completed** using the following command:

```
/opt/ibm/netcool/asm/bin/license-review.sh
```

**Note:** Agile Service Manager can only start after the license has been accepted, and after all additional components have been installed, as described in the following steps.

### Install application discovery and network discovery

4. Install the application discovery and network discovery services, either as part of a single- or a multi-server deployment.

- **Single-server deployment**

**Only** install the application discovery and network discovery services on the same server as the other Agile Service Manager components for a non-production deployment, for example for proof-of-concept or testing purposes.

From the directory where you placed the application discovery or network discovery packages, install one or both of these services using the yum install command, as in the following **examples**:

- Install the application discovery components:

```
sudo yum install nasm-db2 nasm-app-disco-discovery nasm-app-disco-primarystorage  
nasm-appdisco-observer
```

- Install network discovery components:

```
sudo yum install nasm-net-disco-config nasm-net-disco-consul nasm-net-disco-control  
nasm-net-disco-sidecar nasm-net-disco-worker nasm-net-disco-schema-registry nasm-net-  
disco-security nasm-net-disco-status nasm-net-disco-swagger-service nasm-net-disco-  
swagger-ui nasm-net-disco-topogram nasm-netdisco-observer nasm-net-disco-collector  
nasm-net-disco-sftp
```

- **Multi-server deployment**

For **any** production environment, you **must** install the application discovery and network discovery services on separate servers.

In the following installation scenarios, Agile Service Manager servers are identified as follows:

**Agile Service Manager core server**

```
<host-asmcore>
```

### Application discovery server

<host-appdisco>

### Network discovery server

<host-netdisco>

- a. To allow the application and network discovery observers to connect to their remotely deployed services, edit the files that specify the connections between components.

#### For application discovery

Edit the environment variables in the \$ASM\_HOME/.env file. Add the following **host** and **port** details:

```
APPDISCO_SERVICE_HOST=<host-appdisco>
APPDISCO_SERVICE_PORT=<443>
```

#### For network discovery

Edit the environment variables in the \$ASM\_HOME/.env file. Add the following **host** details:

```
KAFKA_ROUTABLE_HOSTNAME=host-asmcore
NETDISCO_SWAGGER_SERVICE_HOST=<host-netdisco>
NETDISCO_SCHEMA_REGISTRY_HOST=<host-netdisco>
```

Edit the KAFKA\_SERVER **host** and **port** in the etc/nasm-netdisco-observer.yml file. Change the following line:

```
KAFKA_SERVER: <host-asmcore:19093>
```

- b. Install the components for the application and network discovery services:

#### For application discovery

On your application discovery host (<host-appdisco>):

```
yum install nasm-db2 nasm-app-disco-discovery nasm-app-disco-primarystorage nasm-nginx nasm-common
```

Remove the redundant references to the application discovery services (now installed on a separate server) from the \$ASM\_HOME/docker-compose.yml file. Run the following bash command from the command line:

```
for service in [ "cassandra" "layout" "merge" "search" "topology" "ui-api" "kafka"
"kafkarest" "zookeeper" "elasticsearch" ]; do
    /opt/ibm/netcool/asm/bin/update-docker-compose.py /opt/ibm/netcool/asm/docker-
compose.yml delete $service
done
```

#### For network discovery

On your network discovery host (<host-netdisco>):

```
yum install nasm-nginx nasm-net-disco-collector nasm-net-disco-config nasm-net-
disco-consul nasm-net-disco-control nasm-net-disco-sidecar nasm-net-disco-worker
nasm-net-disco-schema-registry nasm-net-disco-security nasm-net-disco-sftp nasm-
net-disco-status nasm-net-disco-swagger-service nasm-net-disco-swagger-ui nasm-net-
disco-topogram nasm-common nasm-elasticsearch nasm-kafka nasm-zookeeper
```

Remove the redundant references to the network discovery services (now installed on a separate server) from the \$ASM\_HOME/docker-compose.yml file. Run the following bash command from the command line:

```
for service in [ "cassandra" "layout" "merge" "search" "topology" "ui-api" "kafka"
"kafkarest" "zookeeper" ]; do
    /opt/ibm/netcool/asm/bin/update-docker-compose.py /opt/ibm/netcool/asm/docker-
compose.yml delete $service
done
```

Remove the following references to Kafka and Zookeeper from the \$ASM\_HOME/docker-compose.yml file:

```
# egrep '(kafka|zookeeper)' $ASM_HOME/docker-compose.yml
- kafka
- kafka
- zookeeper
- kafka
```

Add the following environment variables to the \$ASM\_HOME/.env file to set up a connection to the Agile Service Manager core version of Kafka:

```
SHERPA_ROUTABLE_HOSTNAME=<host-netdisco>
KAFKA_EXTERNAL_SERVER=<host-asmcore:19093>
```

During the installation of the packages, the related Docker images are loaded. However, the database required by the discovery services is only created and configured when Agile Service Manager is started. If in addition you are deploying the probe and gateway services, you must configure them before starting the services. No data can be retrieved or discovered until you have also defined observer jobs for the discovery services (and any other observers you are deploying).

5. Required: **Again** review and accept the license **after installation has completed** using the following command:

```
/opt/ibm/netcool/asm/bin/license-review.sh
```

**Note:** If you do not complete this step and accept the license, Agile Service Manager will not start. You only start Agile Service Manager once you have completed all installation tasks for your optional components (that is, one or more of the probe and gateway, application discovery, or network discovery services).

### Configure integration and verify deployment

6. To configure communication between DASH and the Agile Service Manager UI, run the following script:

```
bin/set-dash-config.sh
```

You are prompted to enter the following information:

#### DASH URL endpoint

DASH server root URL

Enter a complete path so that it can be reached from the Agile Service Manager server (including the protocol, host name, and port number).

Example: `https://dash-host.ibm.com:16311`

#### DASH administrator username

The DASH administrator user name.

Example: `smadmin`

#### DASH administrator password

The password for the DASH administrator user.

#### ASM hostname

Enter the fully-qualified host name of the server on which Agile Service Manager is installed.

Example: `asm-host.ibm.com`

This hostname must match the one entered when the Netcool Hybrid Deployment Option Integration Kit is installed, as documented [here](#).

7. Optional: You can verify that the images have been loaded using the `docker images` command.

### What to do next

**Note:** You configure the deployed probe and gateway services **after** installing the core Agile Service Manager containers (including the probe and gateway containers), but **before** starting the Agile Service Manager services.

You provide configuration details for the deployed application discovery service **after** installing the core Agile Service Manager and application discovery containers, and **after** ensuring that the installed Application Discovery Observer, Application Discovery engines and the provided Db2 database container are running.

You configure the deployed network service **after** installing the core Agile Service Manager containers, and **after** starting them.

**Remember:** You must reinstall the Netcool Hybrid Deployment Option Integration Kit each time you reinstall or upgrade Agile Service Manager.

**Important:** Agile Service Manager uses default passwords out of the box. For security reasons, change the default passwords after installation, as described in the [“Changing default passwords”](#) on page 270 topic.

### Related tasks

[“Changing default passwords”](#) on page 270

Passwords are encrypted by default and stored in a number of environment (.env) files.

[“Configuring the probe and gateway services”](#) on page 18

The Agile Service Manager probe and gateway containers are installed together with the other core components. Once you have configured these, Agile Service Manager can display resource information generated from Netcool/OMNIbus events.

[“Viewing the service logs \(on-prem\)”](#) on page 337

Logs for all Netcool Agile Service Manager services can be found in the \$ASM\_HOME/logs/<service> directories. You can set logging levels for the user-facing services, such as the observers and search, using scripts provided.

[“Installing the Netcool Hybrid Deployment Option Integration Kit”](#) on page 30

The Netcool Hybrid Deployment Option Integration Kit is installed into an existing DASH installation, and then configured to communicate with the Agile Service Manager services. The kit is shipped as part of the nginx container and you obtain it via HTTP by pointing to your proxy service during the installation process (rather than downloading it as a separate part).

## Configuring the probe and gateway services

The Agile Service Manager probe and gateway containers are installed together with the other core components. Once you have configured these, Agile Service Manager can display resource information generated from Netcool/OMNIbus events.

### Before you begin

#### Probe and gateway configuration overview:

You stop the probe and gateway, and restart Agile Service Manager, after each step:

1. You apply the \$ASM\_HOME/integrations/omnibus/\*.sql files to the Object Server(s).
2. You configure your ObjectServer(s) in the \$ASM\_HOME/integrations/omnibus/omni.dat file (do **not** add the gateway to this file).
3. Optionally, you add ObjectServer usernames and passwords.
4. Optionally, you add an ObjectServer certificate for TLS.

This topic describes these probe and gateway configuration steps in more detail.

The Agile Service Manager integration with an existing Netcool/OMNIbus system requires updates to the schema and automation (triggers) of that system. Agile Service Manager ships with a sample configuration, which a Netcool/OMNIbus administrator can reference to update their system.

The Agile Service Manager integration also requires connectivity information about the Netcool/OMNIbus system, which the Netcool/OMNIbus administrator should provide.

**Important:** To configure the probe and gateway services, the Agile Service Manager and Netcool/OMNIbus administrators should work together.

**Remember:** You configure the deployed probe and gateway services **after** installing the core Agile Service Manager containers (including the probe and gateway containers), but **before** starting the Agile Service Manager services.

**Upgrade note:** If you are running an existing Agile Service Manager deployment and have already configured earlier versions of the Event Observer, the XML Gateway and the Netcool/OMNIBus probe for Message Bus, and you intend to use the new services, you should perform the following prerequisite steps:

**Prerequisite steps before upgrading the probe:**

De-register the probe event sink (while the topology service is running):

```
$ASM_HOME/bin/topology_service_probe_deregister.sh --all
```

Stop and remove the currently configured HTTP probe.

**Prerequisite steps before upgrading the gateway:**

Stop and remove the currently configured HTTP gateway.

**Prerequisite step before upgrading the Event Observer**

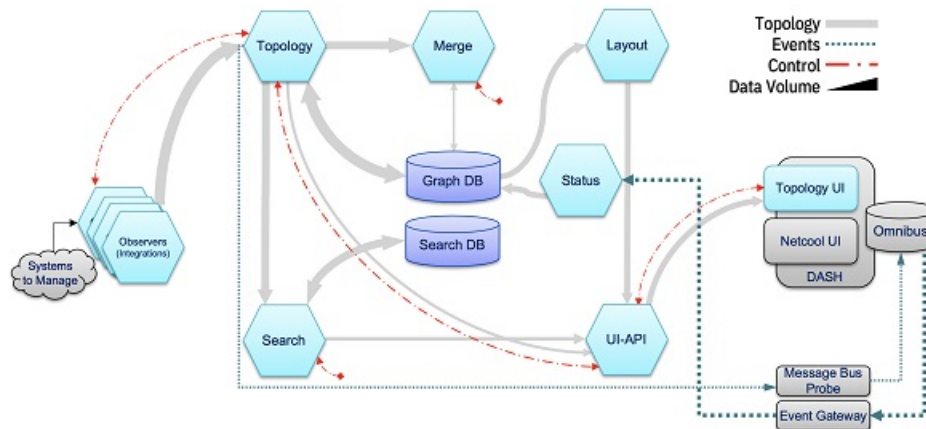
The Event Observer has been replaced with the new Agile Service Manager status service.

The upgrade process replaces the observer with the status service automatically, and no additional steps are required.

## About this task

The probe service receives status from Agile Service Manager, and generates corresponding events in the Netcool/OMNIBus Event Viewer. These events are then fed back to Agile Service Manager via the gateway service, which updates the Agile Service Manager status via the status service with the eventId.

The following diagram depicts how the probe and gateway services work together with the Agile Service Manager status service to keep the event status between Agile Service Manager and Netcool/OMNIBus synchronized.



Specifically, this diagram shows how the data flow from Agile Service Manager status generates Netcool/OMNIBus events. When Netcool/OMNIBus is the source of events, however, the data flow from the Event Observer to the topology service not only updates the status eventId with ServerName/ServerSerial, but generates the status itself.

**Restriction:** Agile Service Manager is **not** compatible with the model of multiple ObjectServer pairs in OCP.

The following nco tools are deployed with the gateway container, and are exposed as container scripts in the \$ASM\_HOME/bin directory. These tools accept the -h or -help options.

**nco\_aes\_crypt.sh**

You use this tool to encrypt the ObjectServer username and password.



**Important:** You have to configure the username and password for the probe and gateway **only if** the ObjectServer is running in secure mode. See the related link to the Netcool/OMNIBus Knowledge Center for more information.

Encryption requires a key file, as named by NOI\_KEY\_FILE in the \$ASM\_HOME/integrations/omnibus/secure.env file.

#### **nco\_keygen.sh**

You use this tool to generate a key file for encrypted usernames and passwords.

The key file is named by NOI\_KEY\_FILE in the \$ASM\_HOME/integrations/omnibus/secure.env file and is generated in the \$ASM\_HOME/security directory.

#### **nco\_ping.sh and nco\_sql.sh**

You use these tools to check access to a given ObjectServer; that is, to check connection status and verify TLS configuration.

The nco\_sql.sh tool can also be used to perform SQL queries or apply schemas and trigger changes.

Both tools require the \$ASM\_HOME/integrations/omnibus/omni.dat file, and, if TLS is configured, the \$ASM\_HOME/integrations/omnibus/secure.env file.

Examples:

```
$ nco_ping.sh NCOMS
Thu Sep 17 10:02:57 UTC 2020 No Object Server CA certificate configured for connection to
the Object Server
Pinging NCOMS
NCO_PING: Server available.
```

```
$ nco_sql.sh -server NCOMS -u asm
Thu Sep 17 10:06:44 UTC 2020 No Object Server CA certificate configured for connection to
the Object Server
Accessing NCOMS as user asm
Password:
1> select count(*) from alerts.status;
2> go
COUNT( * )
-----
          3449

(1 row affected)
1> exit
```

The following configuration files define the connection with Netcool/OMNIBus:

#### **omni.dat**

Netcool/OMNIBus connections data file (contains configuration [examples](#))

\$ASM\_HOME/integrations/omnibus/omni.dat

#### **G\_ASM.props**

Gateway properties file

\$ASM\_HOME/integrations/omnibus/kafka/gateway/G\_ASM.props

#### **probe.props**

Probe properties file

\$ASM\_HOME/integrations/omnibus/kafka/probe/probe.props

#### **secure.env**

Additional secure configuration file (optional)

\$ASM\_HOME/integrations/omnibus/secure.env

**Important:** After changing any of the configuration files in the \$ASM\_HOME/integrations/omnibus directory, you **must** restart the probe and gateway services before your changes take effect.

**Tip:** You can use the [included test scripts](#) to generate a test event.



## Procedure

### Configure target ObjectServer schema and triggers

1. Perform the following updates to the target Netcool/OMNIbus ObjectServers.

**Remember:** Work with the Netcool/OMNIbus administrator to apply these changes.

- a) Set the sub-second clearance mechanism.

#### Sub-second clearance mechanism

This mechanism allows the correct clearance of event updates that occur in the same second.

A new field is added to the `alerts.status` schema, which works in conjunction with the core Netcool/OMNIbus field `@LastOccurrence`; `@LastOccurrenceUSec`.

This mechanism is set via the Agile Service Manager probe rules file and referred to in an updated generic clear trigger.

- b) Define the Agile Service Manager status events clearance.

#### Agile Service Manager specific clearance

A new Netcool/OMNIbus SQL trigger handles the specific clearance of Agile Service Manager status events.

Examples of these updates are provided with Agile Service Manager and are located in the `$ASM_HOME/integrations/omnibus` directory.

#### asm-alert-fields.sql

Defines two new fields:

##### LastOccurrenceUSec

Allows sub-second clearing

##### AsmStatusId

Stores the topology service status ID

Without these fields, the probe and gateway services cannot connect.

#### asm-trigger.sql

Clears up events generated by Agile Service Manager when resources are deleted.

These events will not be cleared if this trigger has not been applied.

#### updated-generic-clear.sql

Updates `generic_clear` automation to allow sub-second clearing.



**Warning:** The sample updates supplied with Agile Service Manager should not be applied to an existing Netcool/OMNIbus deployment without a review by the Netcool/OMNIbus administrator, as they overwrite core Netcool/OMNIbus functions, which may have been customized in an existing Netcool/OMNIbus system. In this case the Netcool/OMNIbus administrator may need to develop custom updates for this integration.

**Hybrid scenario tip:** If Agile Service Manager is deployed on OCP and connecting to an existing Netcool/OMNIbus or NOI system, you can obtain the sample SQL files via the following steps:

- a. Log into the OCP system.
- b. Extract the `asm-trigger.sql` file:

```
oc get configmap asm-noi-gateway-config -o jsonpath="{.data.asm-trigger\.sql}" >
asm-trigger.sql
```

- c. Extract the `'updated-generic-clear.sql` file:

```
oc get configmap asm-noi-gateway-config -o jsonpath="{.data.updated-generic-clear\.sql }"
>
updated-generic-clear.sql
```

- d. Recreate `asm-alert-fields.sql`:

```
echo -e "alter table alerts.status add column AsmStatusId varchar(64);\nalter table
```

```
alerts.status add column LastOccurrenceUSec int;\ngo\n" > asm-alert-fields.sql
```

## Configure connectivity to the target ObjectServer(s)

**Note:** Completing this section of the procedure (step 2) is sufficient to create a connection to a single, unsecured ObjectServer. Its location should be provided by the Netcool/OMNIbus administrator.

2. Define the ObjectServer to which the probe and gateway services connect.

**Important:** The default file paths in the probe and gateway properties files **should not** be altered. These fixed paths identify the files within the containers mounted in the \$ASM\_HOME/etc configuration files, and **do not** correspond to the paths on the Agile Service Manager host.

- a) Identify the ObjectServer to which the probe and gateway services connect.

The ObjectServers must be identified in the \$ASM\_HOME/integrations/omnibus/omni.dat file. This location is fixed and will be mounted into the nasm-noi-probe and nasm-noi-gateway docker containers. In the following omni.dat sample a single ObjectServer is identified called 'AGG\_P'.

```
[AGG_P]
{
  Primary:    1.2.3.4    4100
}
```

Further ObjectServers may be configured in this file, for example to support failover or multi-tier systems. See the related link to the Netcool/OMNIbus documentation for more information on editing omni.dat files.

- b) Configure the probe service to connect to the identified ObjectServer.

The probe service is configured in the \$ASM\_HOME/integrations/omnibus/kafka/probe/probe.props file. This location is fixed and will be mounted into the nasm-noi-probe docker container. In the following probe.props sample the ObjectServer is specified by the 'Server' property, and must match the ObjectServer identified in the omni.dat file in the previous step, which in this example is 'AGG\_P'.

```
Server                : 'AGG_P'
```

- c) Configure the gateway service to connect to the identified ObjectServer.

The gateway service is configured in the \$ASM\_HOME/integrations/omnibus/kafka/gateway/G\_ASM.props file. This location is fixed and will be mounted into the nasm-noi-gateway docker container. In the following G\_ASM.props sample the ObjectServer is specified by the 'Gate.Reader.Server' property, and must match the ObjectServer identified in the omni.dat file.

```
Gate.Reader.Server    : 'AGG_P'
```

## Encrypt username and password

**Remember:** Completing this section of the procedure (steps three to six) is required **only if** you require ObjectServer authentication, that is, only if the ObjectServer is running in secure mode. For more information see the [Running the ObjectServer in secure mode](#) topic in the Netcool/OMNIbus Knowledge Center.

3. Use the nco\_keygen.sh tool to generate a key file.

For example:

```
$ASM_HOME/bin/nco_keygen.sh
Generating key file $ASM_HOME/security/noi_keyfile
```

4. Encrypt the username and password using the generated key.

In this example, the username and password are both asm:

```
$ASM_HOME/bin/nco_aes_crypt.sh asm
```

Example encrypted output:

```
@44:9nx51SfAdcPhyQ1mqcql00qHanR/wQUnZy943YI+TrQ=@
```

5. Add the encrypted username and password to the gateway property file:

\$ASM\_HOME/integrations/omnibus/kafka/gateway/G\_ASM.props

```
Gate.Reader.Server      : 'AGG_P'
ConfigCryptoAlg         : 'AES_FIPS'
ConfigKeyFile           : '$NCHOME/omnibus/asm/$NOI_KEY_FILE'
Gate.Reader.Username    : '@44:9nx51SfAdcPhyQ1mqcql00qHanR/wQUnZy943YI+TrQ=@'
Gate.Reader.Password    : '@44:9nx51SfAdcPhyQ1mqcql00qHanR/wQUnZy943YI+TrQ=@'
```

6. Add the encrypted username and password to the probe property file:

\$ASM\_HOME/integrations/omnibus/kafka/probe/probe.props

```
Server                  : 'AGG_P'
ConfigCryptoAlg         : 'AES_FIPS'
ConfigKeyFile           : '$NCHOME/omnibus/asm/$NOI_KEY_FILE'
AuthUserName            : '@44:9nx51SfAdcPhyQ1mqcql00qHanR/wQUnZy943YI+TrQ=@'
AuthPassword            : '@44:9nx51SfAdcPhyQ1mqcql00qHanR/wQUnZy943YI+TrQ=@'
```

## Define ObjectServer TLS

**Note:** Completing this section of the procedure (steps seven and eight) is required only if you require secure TLS communication with the ObjectServer.

7. Define the Object Server CA certificate by ensuring the following:

- a) The CA certificate file must exist (or copied to) in the \$ASM\_HOME/security directory.

Example:

```
cp <previous_location>/aCaCert.crt $ASM_HOME/security/aCaCert.crt
```

- b) The file is named by NOI\_CA\_CERTIFICATE in the \$ASM\_HOME/integrations/omnibus/secure.env file.

```
NOI_CA_CERTIFICATE=aCaCert.crt
```

8. Add TLS settings to the configuration files.

### omni.dat

Required

```
[AGG_P]
{
  Primary:    noi-omnibus  ssl  4100
}
```

### G\_ASM.props

Some configurations may require the gateway Gate.Reader.CommonNames property,

```
Gate.Reader.Server      : 'AGG_P'
#Gate.Reader.CommonNames : May be required
```

### probe.props

Some configurations may require the probe SSLServerCommonName property.

```
Server                  : 'AGG_P'
#SSLServerCommonName    : May be required
```

9. Required: After changing any of the configuration files in the \$ASM\_HOME/integrations/omnibus directory, you have to restart the probe and gateway services before your changes take effect.

## Results

The gateway service is ready to supply Netcool/OMNIBus event data to the Agile Service Manager topology service, and the probe service is ready to receive status from Agile Service Manager, and then pass it on to the Netcool/OMNIBus Event Viewer.

## Troubleshooting:

If the **probe runs, but the gateway continually restarts**, this may be due to the Kafka topic that is required by the gateway not having been created yet.

The topic **will** be created at startup by the status service (for Agile Service Manager Version 1.1.8 or later), or by the Event Observer (for Agile Service Manager Version 1.1.7 or earlier).

The gateway service will continue to restart until the topic becomes available.

## Example

The omni.dat file contains configuration examples:

```
# Example omni.dat file. Copy your own omni.dat file in here, and configure the
# probe and gateway to connect to the appropriate object servers in their props files:
#
# - $ASM_HOME/integrations/omnibus/kafka/probe/probe.props
# - $ASM_HOME/integrations/omnibus/kafka/gateway/G_ASM.props
#
[AGG_P]
{
    Primary:      noi-omnibus      4100
}

#
# Example failover pair config. The gateway and probe properties would require:
#
# - Probe
# Server: 'AGG_V'
#
# - Gateway
# Gate.Reader.Server: 'AGG_V'
#
[AGG_V]
{
    Primary:      primary-host.example.com    4100
    Backup:       backup-host.example.com     4100
}

#
# Example of connecting the probe and gateway to different layers of a multi-tiered architecture.
# The gateway and probe properties would require:
#
# - Probe
# Server: 'COL_P'
#
# - Gateway
# Gate.Reader.Server: 'AGG_P'
#
[AGG_P]
{
    Primary:      aggregation-layer-host.example.com    4100
}

[COL_P]
{
    Primary:      collection-layer-host.example.com     4100
}

#
# Example TLS config. The gateway and probe properties could use:
#
# - Probe
# Server: 'TLS_AGG_P'
#
# - Gateway
# Gate.Reader.Server: 'TLS_AGG_P'
#
# or can make use of the gateway Gate.Reader.CommonNames and probe SSLServerCommonName
# properties
#
# The Object Server CA certificate or key database file + password must be configured
# in $ASM_HOME/integrations/omnibus/secure.env
#
[TLS_AGG_P]
{
    Primary:      tls-host.example.com    ssl 4100
}
```

**Remember:** After changing any Netcool/OMNibus configuration files, such as the omni.dat file as described in this example, you have to restart the probe and gateway services before your changes take effect.

## What to do next

Next, you start Agile Service Manager.

**Advanced Configuration:** A Netcool/OMNIBus administrator can further configure the probe and gateway services by applying custom rules and mapping. Probe and gateway configuration is accessible here:

- To improve performance and prevent unnecessary events from being displayed in the topology viewer, you can filter out events. You can also pass additional alerts.status fields to Agile Service Manager.
- Use the following gateway configuration files to apply advanced configuration:

```
${ASM_HOME}/integrations/omnibus/kafka/gateway/row_filter.def
```

```
${ASM_HOME}/integrations/omnibus/kafka/gateway/field_filter.map
```

- Probe rules transform the input into events suitable for the Netcool/OMNIBus alerts.status table. The name of the file must be given as a probe property or command line option (which in this example is 'probe.rules').
- Use (or copy and edit) the supplied [sample probe rules file](#).
- Use the following probe configuration file to apply advanced configuration:

```
${ASM_HOME}/integrations/omnibus/kafka/probe/probe.rules
```

You can review the probe and gateway log files (asm\_probe.log and gateway.log). See the related links for more information.

**Generating a test event:** To validate the probe and gateway configuration, you can use the following test scripts (included in the \$ASM\_HOME/bin directory). This script creates a resource with status, which then generates an event in the Netcool/OMNIBus ObjectServer via the probe service, which in turn updates the status in Agile Service Manager via the gateway service.

```
nco_test_create.sh  
nco_test_verify.sh  
nco_test_delete.sh
```

You can print the help text for each of these scripts using the -h command line option.

### Related concepts

[“Probe and gateway troubleshooting” on page 342](#)

See the following information to troubleshoot probe and gateway issues.

### Related tasks

[“Starting the Netcool Agile Service Manager core services” on page 26](#)

You start the Agile Service Manager services using the asm\_start.sh command.

[“Viewing the service logs \(on-prem\)” on page 337](#)

Logs for all Netcool Agile Service Manager services can be found in the \$ASM\_HOME/logs/<service> directories. You can set logging levels for the user-facing services, such as the observers and search, using scripts provided.

### Related reference

[“Status service reference” on page 391](#)

This topic contains reference information for the Agile Service Manager status service.

### Related information

[Running the IBM Tivoli Netcool/OMNIBus ObjectServer in secure mode](#)

[Manually editing the IBM Tivoli Netcool/OMNIBus connections data file \(omni.dat\)](#)

## Starting the Netcool Agile Service Manager core services

You start the Agile Service Manager services using the `asm_start.sh` command.

### Before you begin

**Note:** The Agile Service Manager core installation includes the probe and gateway containers, as well as the discovery services containers. You only start Agile Service Manager once you have completed all installation tasks for your optional components (that is, one or more of the probe and gateway, application discovery, or network discovery services).

### About this task

**Restriction:** Do not start or stop the on-premise version of Agile Service Manager using the Docker commands, as that may result in errors when retrieving required environment variables from the `$ASM_HOME/.env` file.

Table 5. Agile Service Manager commands	
Action	Agile Service Manager command
Verify services	<code>asm_status.sh</code>
Start services	<code>asm_start.sh</code>
Stop services	<code>asm_stop.sh</code>

### Procedure

Start the Agile Service Manager services on the Agile Service Manager server by running the `asm_start.sh` command from the home directory.

```
$ASM_HOME/bin/asm_start.sh
```

**Note:** If you are deploying the discovery services, the start command will first create and configure the required database, before starting the services. This may take some time.

Example system output:

```
Creating network "asm_default" with driver "bridge"
Creating asm_elasticsearch_1
Creating asm_zookeeper_1
Creating asm_cassandra_1
Creating asm_kafka_1
Creating asm_kafkarest_1
Creating asm_topology_1
Creating asm_noi-probe_1
Creating asm_dns-observer_1
Creating asm_layout_1
Creating asm_ui_1
Creating asm_rest-observer_1
Creating asm_noi-gateway_1
Creating asm_kubernetes-observer_1
Creating asm_ui-api_1
Creating asm_merge_1
Creating asm_search_1
Creating asm_file-observer_1
Creating asm_docker-observer_1
Creating asm_proxy_1
Creating asm_observer_1
```

**Tip:** You can check the status of the discovery services.

## Check application discovery status

To view the status of the application discovery services, see the logs.

Check the database status:

```
$ tail -f /opt/ibm/netcool/asm/logs/db2/create_database.log
```

Check the database schema status:

```
$ grep DbInit /opt/ibm/netcool/asm/logs/app-disco-pss/pss/wlp.log
```

Check the storage status:

```
$ docker exec -ti asm_app-disco-primarystorage_1 /opt/IBM/taddm/dist/bin/control status
```

Check the discovery services status:

```
$ docker exec -ti asm_app-disco-discovery_1 /opt/IBM/taddm/dist/bin/control status
```

## Check network discovery status

To view the status of the network discovery services, launch the Consul UI on port 8501, for example:

```
https://<your-deployment>:8501/ui/sherpa/services
```

## Results

You can also check the state of the Agile Service Manager deployment using the `asm_status.sh` command.

```
/opt/ibm/netcool/asm/bin/asm_status.sh
```

Example system output:

CONTAINER ID	IMAGE	COMMAND	CREATED	STATUS	PORTS
22090abaa914	nasm-nginx:1.14.1.41	"/opt/ibm/netcool/..."	57 seconds ago	Up 54 seconds	0.0.0.0:443->8443/tcp
asm_proxy_1	nasm-file-observer:1.0.13.11613	"sh -c /opt/ibm/\$A..."	About a minute ago	Up 57 seconds	9098-9099/tcp
0c18fd0d70b7	nasm-file-observer_1	"sh -c /opt/ibm/\$A..."	About a minute ago	Up 57 seconds	7082-7083/tcp
6c6f86877a3f	nasm-merge:1.0.6.11649	"sh -c /opt/ibm/\$A..."	About a minute ago	Up 57 seconds	9086-9087/tcp
asm_merge_1	nasm-docker-observer:1.0.12.11614	"sh -c /opt/ibm/\$A..."	About a minute ago	Up 57 seconds	7080-7081/tcp
fe2c78739419	nasm-docker-observer_1	"sh -c /opt/ibm/\$A..."	About a minute ago	Up 58 seconds	3080/tcp
5d821c375e6d	nasm-search:1.0.10.11620	"sh -c /opt/ibm/\$A..."	About a minute ago	Up 57 seconds	9108-9109/tcp
asm_search_1	nasm-ui-api:1.0.2-11609	"sh -c ./start.sh"	About a minute ago	Up 57 seconds	8080/tcp
de5243327d30	nasm-kubernetes-observer:1.1.22.11617	"sh -c /opt/ibm/\$A..."	About a minute ago	Up 57 seconds	9104-9105/tcp
asm_ui-api_1	nasm-noi-gateway:10.0.2.1.15	"/docker-entrypoint..."	About a minute ago	Up 57 seconds	7084-7085/tcp
be9b4eb7107f	nasm-rest-observer:1.0.8.11614	"sh -c /opt/ibm/\$A..."	About a minute ago	Up 57 seconds	9106-9107/tcp
asm_kubernetes-observer_1	nasm-layout:1.0.5.11635	"sh -c /opt/ibm/\$A..."	About a minute ago	Up 59 seconds	80/tcp, 443/tcp, [...]
f1d3c88cdfa7	nasm-dns-observer:1.0.10.11618	"sh -c /opt/ibm/\$A..."	About a minute ago	Up About a minute	8080-8081/tcp
asm_noi-gateway_1	nasm-topology:1.1.23.11636	"sh -c /opt/ibm/\$A..."	About a minute ago	Up About a minute	9092-9093/tcp, [...]
bda9dc00a8d7	nasm-kafkarest:2.1.1.82	"/bin/sh -c /opt/k..."	About a minute ago	Up About a minute	9200/tcp, 9300/tcp
asm_rest-observer_1	nasm-cassandra:3.11.4.76	"/opt/ibm/start-ca..."	About a minute ago	Up About a minute	2181/tcp, 2888/tcp, 3888/tcp
2761da1495f1	nasm-elasticsearch:6.5.4.89	"/opt/elasticsearc..."	About a minute ago	Up 57 seconds	7086-7087/tcp
asm_layout_1	nasm-zookeeper:3.4.14.56	"/bin/sh -c /opt/s..."	About a minute ago	Up 57 seconds	
4df8bdc8b849	nasm-observer:1.1.24.11717	"sh -c /opt/ibm/\$A..."	About a minute ago	Up 57 seconds	
asm_dns-observer_1	asm_observer_1				

## What to do next

Next, you install the Netcool Agile Service Manager UI.

**Tip:** You can stop the Agile Service Manager service using the `asm_stop.sh` command.

```
/opt/ibm/netcool/asm/bin/asm_stop.sh
```

Example output:

```
Stopping asm_proxy_1 ... done
Stopping asm_kubernetes-observer_1 ... done
Stopping asm_noi-gateway_1 ... done
Stopping asm_noi-probe_1 ... done
Stopping asm_rest-observer_1 ... done
...
```

### Related tasks

[“Configuring the probe and gateway services” on page 18](#)

The Agile Service Manager probe and gateway containers are installed together with the other core components. Once you have configured these, Agile Service Manager can display resource information generated from Netcool/OMNIbus events.

## Adding new observers to an existing on-prem installation

You can add new observers to an existing on-prem installation of Agile Service Manager using the command line interface (CLI).

### About this task

To add a new observer to an existing installation, you install the new observer using the yum command, then restart Agile Service Manager.

### Procedure

1. Run the following command to install the observer:

```
yum install <observer docker rpm file>
```

#### Example:

```
yum install nasm-alm-observer
```

2. Restart Agile Service Manager:

```
<ASM_HOME>/bin/asm_start.sh
```

## IBM Installation Manager

You use IBM Installation Manager to install the Agile Service Manager Hybrid Integration Kit. Installation Manager is available for download from the IBM Fix Central website.

### Before you begin

The required version of Installation Manager is Version 1.9.2 .

**Tip:** If you are deploying Netcool Agile Service Manager as part of another IBM solution such as Netcool Operations Insight (NOI) you will already have IBM Installation Manager on your system.

You must have an IBM ID to download software from IBM Fix Central. You can register for an IBM ID at <http://www.ibm.com>.

### About this task

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.

### Procedure

1. Open the IBM Fix Central website at the following URL:

<http://www.ibm.com/support/fixcentral/>



2. On the **Find product** tab:
  - a) Enter IBM Installation Manager in the **Product selector** field.
  - b) Select **Version 1.9.2** from the **Installed Version** list.
  - c) Select your intended host operating system from the **Platform** list and click **Continue**.
3. On the **Identify Fixes** page, choose **Browse for fixes** and **Show fixes that apply to this version**. Click **Continue**.
4. On the **Select Fixes** page, select the installation file appropriate to your intended host operating system and click **Continue**.
5. When prompted, enter your IBM ID user name and password.
6. If your browser has Java enabled, choose the Download Director option. Otherwise, select the HTTP download option.
7. Start the installation file download. Make a note of the download location.

## What to do next

Install the IBM Installation Manager (GUI, console, or silent installation).

## Installing IBM Installation Manager (via GUI or console)

You can install Installation Manager with a wizard-style GUI or an interactive console.

## 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.

## 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. To 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`  
We recommend using the `icosgrp` for Netcool Agile Service Manager.
  - 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
```

Where *data\_location* specifies the data directory. You must specify a data directory that all members of the group can access. 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 panel, 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 the Agile Service Manager Hybrid Integration Kit.

### **What to do next**

If required, add the Installation Manager installation directory path to your PATH environment variable.

## **Installing the Netcool Hybrid Deployment Option Integration Kit**

The Netcool Hybrid Deployment Option Integration Kit is installed into an existing DASH installation, and then configured to communicate with the Agile Service Manager services. The kit is shipped as part of the nginx container and you obtain it via HTTP by pointing to your proxy service during the installation process (rather than downloading it as a separate part).

### **Before you begin**

**Remember:** You must reinstall the Netcool Hybrid Deployment Option Integration Kit each time you reinstall or upgrade Agile Service Manager.

**Note:** Any tasks relating to Web GUI that are mentioned here only apply if you have Web GUI installed (which must be Tivoli Netcool/OMNIbus Web GUI Version 8.1.0.21).

**Note:** Ensure that you have the correct versions of all required NOI product components installed. See the [Netcool Operations Insight Version 1.6.3 Product and component version matrix](#) for more details on NOI components.

You must complete the following prerequisites before you can install the Netcool Hybrid Deployment Option Integration Kit.

1. Ensure that Agile Service Manager core has been installed and is running.
2. Ensure that WebSphere Application Server Version 8.5.5.17 has been installed. Follow the IBM Knowledge Center instructions here: [https://www.ibm.com/support/knowledgecenter/SSAW57\\_8.5.5/com.ibm.websphere.nd.multiplatform.doc/ae/welcome\\_ndmp.html](https://www.ibm.com/support/knowledgecenter/SSAW57_8.5.5/com.ibm.websphere.nd.multiplatform.doc/ae/welcome_ndmp.html)
3. Ensure that WebSphere Application Server is running during the installation of the Netcool Hybrid Deployment Option Integration Kit.

4. Ensure that a compatible version of the IBM WebSphere Java SDK has been installed (for example together with WebSphere Application Server Version 8.5.5.17), for example at least IBM WebSphere SDK Java Technology Edition Version 7.0. Follow the IBM Knowledge Center instructions here: [https://www.ibm.com/support/knowledgecenter/SSAW57\\_8.5.5/com.ibm.websphere.installation.nd.doc/ae/tins\\_installation\\_jdk7\\_gui.html](https://www.ibm.com/support/knowledgecenter/SSAW57_8.5.5/com.ibm.websphere.installation.nd.doc/ae/tins_installation_jdk7_gui.html)
5. Ensure that IBM Dashboard Application Service Hub (DASH) 3.1.3.7 or later has been installed. Follow the IBM Knowledge Center instructions for installing Jazz for Service Management (Version 1.1.3.9) here: [https://www.ibm.com/support/knowledgecenter/en/SSEKCU\\_1.1.3.0/com.ibm.psc.doc/install/psc\\_c\\_install.html](https://www.ibm.com/support/knowledgecenter/en/SSEKCU_1.1.3.0/com.ibm.psc.doc/install/psc_c_install.html)
6. Configure the Installation Manager repository to point at the following URL `https://<ASM_HOST>/im/repository/repository.xml`.
7. The IBM Installation Manager analyzes existing installations to determine defaults, which it then presents to you during installation. To verify these defaults and change them if necessary, ensure you have the required environment information about your installation and the DASH environment to hand.

## About this task

You add the Netcool Hybrid Deployment Option Integration Kit to an existing DASH installation and configure the application to communicate with the previously installed application. You use the default client ID and secret.

**Important:** Ensure that you are logged in as the same user who installed DASH.

The steps for starting Installation Manager differ depending on the user mode in which it was installed. The steps for installing with the Installation Manager console are common to all user modes and operating systems. Take note of the following information about permissions on the supported operating systems:

- The 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 Nonadministrator 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.

**Note:** It is recommended to use Install Manager Group Mode for installation.

Respond to each Installation configuration option to ensure it matches your pre-defined Installation Information Checklist.

## Procedure

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

```
./IBMIM
```

2. From the Installation Manager Main Menu, select **Preferences**.
3. In the Preferences menu, select **Repositories**.
4. In the Repositories menu, select **Add Repository**.
5. Enter the path to the repository file on the Agile Service Manager proxy, and click **OK**.
6. In the main **Installation Manager** window, click **Install** for a fresh installation.
7. Select the **Hybrid Deployment Option Integration Kit** package, and click **Next**.
8. Complete the following installation steps:

**Tips:**

- The Hybrid Deployment Option Integration Kit is also used in other deployments, such as OCP/CNEA installations. If you encounter optional steps or terminology not specific to your Agile Service Manager deployment of the kit as documented here, you can ignore these.
- Remember to use the default client ID and secret.
  - a) On the Installation Manager **Licenses** tab, read and accept the license agreement, then click **Next**.
  - b) Select the default location, then click **Next**.  
The installation process moves onto the **Features** tab, which displays your selected version of the Netcool Hybrid Deployment Option Integration Kit, as well as the features you can install.
  - c) Select **Jazz for Service Management**, enter the required information for the WebSphere Application Server, then click **Next**.
  - d) Select **OAuth 2.0 Configuration**, and set the **Redirect URL** to point to <ASM\_HOST>/users/api/authprovider/v1/was/return, for example https://asm-host.ibm.com/users/api/authprovider/v1/was/return
  - e) Click **Install** to complete the installation.  
The IBM Installation Manager installs the Netcool Hybrid Deployment Option Integration Kit into your existing DASH installation.
  - f) Click **Finish** to exit the installer.

9. Restart DASH.

## What to do next

Next, you configure DASH user roles to allow users to access Agile Service Manager.

## Configuring DASH user roles

You configure DASH user roles so that users can use the Netcool Agile Service Manager UI. This task is the same for both on-prem and OCP deployments of Agile Service Manager.

### About this task

You can assign the following DASH user roles to users:

#### **inasm\_operator**

A user with the inasm\_operator role can access the Netcool Agile Service Manager UI, and use it to search for and visualize the resources in the Netcool Agile Service Manager core application.

#### **inasm\_editor**

The same as for inasm\_operator.

In addition, a user with the inasm\_editor role can add comments to resources from the Topology Viewer Context (right-click) menu. (A user with the inasm\_operator role can view comments, but not add new ones.)

#### **inasm\_admin**

The same as for inasm\_editor.

In addition, a user with the inasm\_admin role has access to a number of administrator tools, where they can define custom UI elements for the Topology Viewer. See the [“Customizing UI elements” on page 271](#) topic for more information.

To configure DASH user roles, you must log into DASH with admin user credentials.

**Tip:** You can also assign roles to a user indirectly by assigning them to a group of which the user is a member.

## Procedure

1. As the admin user, log into your DASH web application.

### For on-prem

If you have used the default root location of `/ibm/console`, use the following logon URL:

```
https://<DASH-HOST>:<DASH-PORT>/ibm/console/logon.jsp
```

### For OCP

You login to the Agile Service Manager OCP installation using a URL of the following format (example):

```
http://netcool.noi.apps.<your-ocp-cluster>/ibm/console
```

Where *noi* is the Netcool Operations Insight Helm release name. Use the following command to retrieve the DASH URL:

```
helm status NOI helm release name --tls
```

2. Select **Console Settings** from the DASH menu.
3. Select **User Roles** from the Console Settings menu (under the Roles heading).
4. Select the **User Roles** tab, and then click **Search**.  
Known users are displayed in the Search results table.
5. **For each user** requiring access to the Netcool Agile Service Manager UI, perform the following actions:
  - a) Click the required user ID in the Search results table.  
All roles that are available to the selected user are displayed in the Available Roles table.
  - b) Select the required roles, as appropriate.
  - c) Click **Save**.

## Results

Once you have saved your changes, the user role changes take effect. All users with their newly assigned roles are now able to log into DASH and access the Netcool Agile Service Manager UI. From there, users can search and visualize resources from the Netcool Agile Service Manager topology service.

**Remember:** You can also assign roles to a user by assigning them to a group to which the user belongs.

## Uninstalling a legacy Agile Service Manager UI using the Installation Manager

To uninstall legacy (that is, pre-1.1.8 versions) of the Agile Service Manager UI, you use IBM Installation Manager.

### Before you begin

**Remember:** You only use this process to uninstall 1.1.7 (or earlier) versions of the Agile Service Manager UI, which were also installed with IBM Installation Manager.

**Tip:** When uninstalling a complete Agile Service Manager installation (Version 1.1.7 or earlier), you remove both the core services and the user interface. The recommended sequence of removal is to uninstall the UI first using the IBM Installation Manager, and then remove the core, which is described in a related topic.

### About this task

Use IBM Installation Manager to remove Netcool Agile Service Manager UI.

**Important:** Ensure that you are logged in as the same user who installed DASH.

## Procedure

1. Change to the `/eclipse` subdirectory of the Installation Manager installation directory.
2. Use the following command to start the Installation Manager wizard:  
`./IBMIM`
3. On the main **Installation Manager window**, click **Uninstall**.
4. Select **IBM Netcool Agile Service Manager**, then click **Next**.  
The installed package groups are displayed.
5. Under IBM Netcool GUI Components, select **IBM Netcool Agile Service Manager**, then click **Uninstall**.  
The user interface is uninstalled.
6. Click **Finish** to exit the Installation Manager.
7. After using the Installation Manager to uninstall Netcool Agile Service Manager, you must restart DASH to ensure that it recognizes the removal of Netcool Agile Service Manager.

## Results

The Installation Manager removes the files and directories that it installed, leaving behind the application configuration file and log files.

## What to do next

After uninstalling the UI, you remove the core services.

### Related tasks

[“Uninstalling the Netcool Agile Service Manager core services” on page 34](#)

To uninstall the Netcool Agile Service Manager core, you remove `nasm-common`. Due to package dependencies, this will stop and remove **all** of the Netcool Agile Service Manager core Docker containers and images, and then remove the installation files from your server.

## Uninstalling the Netcool Agile Service Manager core services

To uninstall the Netcool Agile Service Manager core, you remove `nasm-common`. Due to package dependencies, this will stop and remove **all** of the Netcool Agile Service Manager core Docker containers and images, and then remove the installation files from your server.

## Before you begin

To uninstall Netcool Agile Service Manager, you remove the core services and then perform cleanup of any remaining components.

## About this task

The following procedure sequentially stops and removes all Netcool Agile Service Manager Docker packages, before removing the images and then the installation packages from the server. Any other Docker components are not affected, and the Docker service is not stopped.

**Note:** Although you can remove individual packages, this is the recommended uninstall procedure for Netcool Agile Service Manager core components.

**Tip:** Use the following command to see a list of the installed docker containers:

```
docker ps -a
```

## Procedure

Use the following command to stop the server, remove the installation images from Docker, and then remove the rpm packages from the server.

```
sudo yum remove nasm-common
```

## Results

The Netcool Agile Service Manager core services have been removed.

# Installing Agile Service Manager on RedHat OpenShift Container Platform

---

To install Netcool Agile Service Manager on RedHat OpenShift Container Platform, you follow the NOI documentation installation steps. You can add new observers to an existing Agile Service Manager OCP installation.

## NOI documentation

[https://www.ibm.com/support/knowledgecenter/SSTPTP\\_1.6.3/com.ibm.netcool\\_ops.doc/soc/integration/task/int\\_installing-on-rhocp.html](https://www.ibm.com/support/knowledgecenter/SSTPTP_1.6.3/com.ibm.netcool_ops.doc/soc/integration/task/int_installing-on-rhocp.html)

## Adding new observers to an existing OCP installation

You can add new observers to an existing Red Hat OpenShift Container Platform (OCP) installation of Agile Service Manager using either the command line interface (CLI), or the Red Hat OpenShift Operator Lifecycle Management console (UI).

### About this task

OCP installations of Agile Service Manager are deployed as part of IBM Netcool Operations Insight (NOI). To add new observers to an existing OCP installation you edit the yaml file and then run a Helm upgrade.

### Procedure

#### For a CLI NOI installation

See the following topic for more information on NOI installations via CLI: [Installing operators with the CLI using PPA](#)

1. Enable the new observers in the `deploy/crds/noi.ibm.com_nois_cr.yaml` file.  
Change the value for an observers to 'true' to enable it, as in the following sample:

```
...
  observers:
    alm: false
    appdynamics: false
    aws: true
    azure: false
...
  docker: true
  dynatrace: false
  file: true
...
```

2. Run the following command to apply the changes made:

```
kubectl apply -f noi.ibm.com_nois_cr.yaml
```

#### For a UI NOI installation

See the following topic for more information on NOI installations via UI: [Installing operators with the Operator Lifecycle Management console](#)

3. In the Operator Lifecycle Management console, go to **Operators > Installed Operators**.
4. Choose the correct project from the **Project** drop-down list.
5. Click the **Netcool Operations Insight operator**.
6. Click the **YAML** tab and edit the file. To enable an observer, set it to 'true', for example:

```
...
  observers:
    alm: false
```

```
... appdynamics: true
```

7. Click **Save** to apply the changes.

#### Related information

## Installing Agile Service Manager on OCP for Watson AIOps

To install a stand-alone version of Agile Service Manager on RedHat OpenShift Container Platform for use with IBM Watson AIOps (and without NOI), you follow these installation steps.

#### From the glossary:

##### **CASE (Container Application Software for Enterprises)**

The Container Application Software for Enterprises (CASE) specification defines metadata and structure for packaging a containerized application.

CASE provides a common packaging definition for the many ways to install, manage, operate, upgrade, and uninstall complex software using multiple roles and personas.

##### **OLM (Operator Lifecycle Manager)**

The Red Hat OpenShift Operator Lifecycle Manager (OLM) UI helps users install, update, and manage the lifecycle of all Operators and their associated services running across their clusters. It is part of the Operator Framework, an open source toolkit designed to manage Kubernetes native applications (Operators) in an effective, automated, and scalable way.

#### Related tasks

[“Adding new observers to an existing OCP installation” on page 35](#)

You can add new observers to an existing Red Hat OpenShift Container Platform (OCP) installation of Agile Service Manager using either the command line interface (CLI), or the Red Hat OpenShift Operator Lifecycle Management console (UI).

[“Configuring alternative storage \(on OCP\)” on page 45](#)

You can deploy Agile Service Manager on OCP (for Watson AIOps) with alternative storage.

#### Related information

[Product Documentation for OpenShift Container Platform 4.6](#)

[Installing Netcool Operations Insight on OpenShift](#)

## Preparing the OCP cluster for Agile Service Manager installation (with Watson AIOps)

Prepare your cluster for the installation of a stand-alone OCP version of Agile Service Manager for deployment with Watson AIOps.

### Before you begin

Before preparing your cluster, ensure you have all of the prerequisites in place.

#### **RedHat OpenShift Container Platform (OCP)**

Download and install RedHat OpenShift Container Platform (OCP) Version 4.6:

[Product Documentation for OpenShift Container Platform 4.6](#)

#### **OpenShift CLI (oc)**

Download and install the required version of the OpenShift CLI (oc), which must be the same (or later) as the version of OCP you are using.

[Getting started with the OpenShift CLI \(oc\)](#)



## Procedure

1. Provision the required machines. The hardware architecture on which Agile Service Manager is installed must be AMD64.
2. Ensure that you have access to an administrator account on the target Red Hat OpenShift cluster. Agile Service Manager must be installed by a user with administrative access on the cluster.
3. Create a custom namespace into which to deploy:

```
oc create namespace $NAMESPACE
```

Where:

- `$NAMESPACE` is the name of the custom namespace that you want to create

4. Create a service account named `${RELEASE_NAME}-service-account` over which the Agile Service Manager instance named `${RELEASE_NAME}` would run.

```
oc create serviceaccount ${RELEASE_NAME}-service-account -n $NAMESPACE
```

5. Obtain the entitlement key that is assigned to your IBM ID, which will grant you access to the IBM Entitled Registry.

- a) Log into the following site with the account (username and password) that has entitlement to IBM software.

<https://myibm.ibm.com/products-services/containerlibrary>

The key that is displayed (180 characters) is the password that will be used when accessing the Entitled Registry.

- b) In the Entitlement keys section, select **Copy key** to copy the entitlement key to the clipboard.
- c) Create a docker secret to enable your deployment to pull Agile Service Manager images from the IBM Entitled Registry.

```
oc create secret docker-registry asm-registry-secret \
--namespace=$NAMESPACE \
--docker-server=cp.icr.io \
--docker-username=cp \
--docker-password=entitlement_key
```

Where:

- `asm-registry-secret` is the name of the secret that you are creating
- `$NAMESPACE` is the namespace into which you want to deploy Agile Service Manager
- `entitlement_key` is the entitlement key that you copied previously

- d) Add the registry secret to your service account,

```
oc patch serviceaccount ${RELEASE_NAME}-service-account -p '{"imagePullSecrets": [{"name": "asm-registry-secret"}]}' -n $NAMESPACE
```

Where:

- `asm-registry-secret` is the name of the Docker registry secret that you created previously
- `$NAMESPACE` is the namespace into which you want to deploy Agile Service Manager

## Related information

[Red Hat OpenShift videos](#)

## Installing Agile Service Manager on OCP (Watson AIOps)

A stand-alone OCP version of Agile Service Manager is supported for deployment with Watson AIOps, without NOI.

### Before you begin

1. Review your system sizing needs: [“Sizing reference for Agile Service Manager on OCP \(Watson AIOps\)” on page 50](#)
2. Ensure you have prepared your cluster for installation: [“Preparing the OCP cluster for Agile Service Manager installation \(with Watson AIOps\)” on page 36](#)
3. Obtain Version 3.4.4 (or later) of the cloudctl utility, give it executable permissions, and ensure that it is in your PATH. You can obtain cloudctl here: <https://github.com/IBM/cloud-pak-cli/releases>
4. Download the IBM Netcool Agile Service Manager CASE bundle (ibm-asm-operator-case) to your Red Hat OpenShift cluster, and extract the bundle to a temporary folder (<destination\_dir>).  
Download (example):

```
cloudctl case save --case ibm-asm-operator-case --outputdir <destination_dir> -repo https://raw.githubusercontent.com/IBM/cloud-pak/master/repo/case
```

Extract (example):

```
tar -xvf <destination_dir>/ibm-asm-operator-case.tgz
```

## Installing Agile Service Manager on OCP for Watson AIOps (via CASE with OLM UI)

A stand-alone OCP version of Agile Service Manager is supported for deployment with Watson AIOps. You can install this version using both the Container Application Software for Enterprises (CASE) and the Red Hat OpenShift Operator Lifecycle Manager (OLM) UI.

### Before you begin

Ensure you have completed the [prerequisite tasks](#).

### About this task

You use CASE to install the Agile Service Manager catalog source and operator. You then use the OLM UI to verify that the catalog and operator are present, and to create an Agile Service Manager instance.

### Procedure

1. Install the Agile Service Manager Catalog using CASE.

```
cloudctl case launch \
  --case ibm-asm-operator-case \
  --namespace ${NAMESPACE} \
  --inventory asmOperatorSetup \
  --action install-catalog
```

Where:

- `${NAMESPACE}` is the custom namespace to be used for your deployment, created when you prepared your cluster

To verify the Agile Service Manager Catalog source:

- From the OLM UI, navigate to **Administration > Cluster Settings**, and then select the **OperatorHub** configuration resource under the **Global Configurations** tab.  
Verify that the `ibm-asm-catalog` catalog source is present.

2. Install the IBM Netcool Agile Service Manager operator using CASE.

```
cloudctl case launch \
  --case ibm-asm-operator-case \
  --namespace ${NAMESPACE} \
  --inventory asmOperatorSetup \
  --action install-operator
```

Verify the IBM Netcool Agile Service Manager operator:

- From the OLM UI, navigate to **Operators > Installed Operators**.  
Verify that the status of the IBM Netcool Agile Service Manager operator is **Succeeded**.

### 3. Create an Agile Service Manager instance.

- From the OLM UI, navigate to **Operators > Installed Operators**, and select the **IBM Netcool Agile Service Manager Operator**.
- Under **Provided APIs**, select **Cloud Deployment**, and then select **Create Instance**.
- From the OLM UI, use the YAML view or the form view to configure the properties for the Agile Service Manager deployment. For more information about configurable properties for a this deployment, see the [Operator properties](#) topic.
- Click **Create**.

Under the **Cloud deployment** tab, the Agile Service Manager instance is displayed. You can view the status of each of the updates on the installation. The deployment is ready when the status is **OK**.

## What to do next

If you want to enable or disable a feature or observer after your installation, you can edit the Agile Service Manager instance by running the following command:

```
oc edit ASM <asm-instance-name>
```

You can then select to enable or disable the feature or observer.

**Tip:** When you disable features post install, the resource is not automatically deleted. To find out if the feature is deleted or not, check the operator log.

### Limitations:

- Changing an existing deployment from its original deployment type to any other deployment type is not supported.
- Changing an instance's deployment parameters in the form view is not supported after deployment.
- If you update custom secrets in the OLM console, the crypto key becomes corrupted and the command to encrypt passwords does not work. **Only update custom secrets with the CLI.**

### Related tasks

[“Installing Agile Service Manager on OCP for Watson AIOps \(via CASE without OLM UI\)” on page 40](#)

A stand-alone OCP version of Agile Service Manager is supported for deployment with Watson AIOps. You can install this version using Container Application Software for Enterprises (CASE) **without** also using the Red Hat OpenShift Operator Lifecycle Manager (OLM) UI.

[“Configuring the Helm chart to use alternate storage \(OCP\)” on page 289](#)

Agile Service Manager by default supports local volumes for storage on Kubernetes. To configure an alternate storage backend, you must set the storage class. This requires a manual command line installation, and **not** an installation via Helm.

### Related reference

[“Agile Service Manager Operator cloud properties” on page 47](#)

This topic lists the configurable properties for a deployment of Agile Service Manager on OCP for Watson AIOps (via CASE and CLI, or the OLM UI), as well as on an **offline** Red Hat OpenShift cluster that has no internet connectivity (via airgapped environment).

### Related information

[Installing the IBM Watson AIOps service](#)

## Installing Agile Service Manager on OCP for Watson AIOps (via CASE without OLM UI)

A stand-alone OCP version of Agile Service Manager is supported for deployment with Watson AIOps. You can install this version using Container Application Software for Enterprises (CASE) **without** also using the Red Hat OpenShift Operator Lifecycle Manager (OLM) UI.

### Before you begin

Ensure you have completed the [prerequisite tasks](#).

### About this task

You use CASE to create the catalog source, install the Agile Service Manager operator, and create an Agile Service Manager instance.

### Procedure

1. Install the Agile Service Manager Catalog and Operator.

- a) Install the Catalog using CASE

```
cloudctl case launch
--case ibm-asm-operator-case
--namespace ${NAMESPACE}
--inventory asmOperatorSetup
--action install-operator-native
```

Where:

- `${NAMESPACE}` is the custom namespace to be used for your deployment, created when you prepared your cluster

- b) Verify the Agile Service Manager Operator.

```
oc get deployment asm-operator -n ${NAMESPACE}
```

System output example:

NAME	READY	UP-TO-DATE	AVAILABLE	AGE
asm-operator	1/1	1	1	55s

2. Create an Agile Service Manager instance.

- a. Edit the sample Agile Service Manager Custom Resource and update the required values, such as 'license acceptance', 'entitlement secret', 'Image registry', and 'release name'. For more information about configurable properties for a this deployment, see the [Operator properties](#) topic.

```
vi ibm-asm-operator-case/inventory/asmOperator/files/asm.ibm.com_asm_cr.yaml
```

- b. Apply the save configuration:

```
oc apply -f ibm-asm-operator-case/inventory/asmOperator/files/asm.ibm.com_asm_cr.yaml -n ${NAMESPACE}
```

- c. Verify the instance being loaded:

```
oc get ASM ${RELEASE_NAME} -n ${NAMESPACE}
```

System output should return an instance with the release name specified earlier. If the release name is `example-asm`, then the output should be:

NAME	AGE
example-asm	28s

d. To view the running configuration of Agile Service Manager Instance:

```
oc describe asm ${RELEASE_NAME} -n ${NAMESPACE}
```

## What to do next

If you want to enable or disable a feature or observer after your installation, you can edit the Agile Service Manager instance by running the following command:

```
oc edit ASM <asm-instance-name>
```

You can then select to enable or disable the feature or observer.

**Tip:** When you disable features post install, the resource is not automatically deleted. To find out if the feature is deleted or not, check the operator log.

### Limitations:

- Changing an existing deployment from its original deployment type to any other deployment type is not supported.
- Changing an instance's deployment parameters in the form view is not supported after deployment.
- If you update custom secrets in the OLM console, the crypto key becomes corrupted and the command to encrypt passwords does not work. **Only update custom secrets with the CLI.**

### Related tasks

[“Installing Agile Service Manager on OCP for Watson AIOps \(via CASE with OLM UI\)” on page 38](#)

A stand-alone OCP version of Agile Service Manager is supported for deployment with Watson AIOps. You can install this version using both the Container Application Software for Enterprises (CASE) and the Red Hat OpenShift Operator Lifecycle Manager (OLM) UI.

[“Configuring the Helm chart to use alternate storage \(OCP\)” on page 289](#)

Agile Service Manager by default supports local volumes for storage on Kubernetes. To configure an alternate storage backend, you must set the storage class. This requires a manual command line installation, and **not** an installation via Helm.

### Related reference

[“Agile Service Manager Operator cloud properties” on page 47](#)

This topic lists the configurable properties for a deployment of Agile Service Manager on OCP for Watson AIOps (via CASE and CLI, or the OLM UI), as well as on an **offline** Red Hat OpenShift cluster that has no internet connectivity (via airgapped environment).

### Related information

[Installing the IBM Watson AIOps service](#)

[Download Watson AIOps](#)

## Installing Agile Service Manager on OCP for Watson AIOps (offline on airgap)

A stand-alone OCP version of Agile Service Manager is supported for deployment with Watson AIOps. You can install this version on an **offline** Red Hat OpenShift cluster that has no internet connectivity by using an airgapped environment using Container Application Software for Enterprises (CASE) and the Red Hat OpenShift Operator Lifecycle Manager (OLM) UI.

## Before you begin

Ensure you have completed the [prerequisite tasks](#).

## About this task

You install Agile Service Manager on an offline Red Hat OpenShift cluster with no internet connectivity by creating an online bastion host that can download the Agile Service Manager CASE bundle, access the required images in the IBM Entitled Registry, and mirror them to a registry on the Red Hat OpenShift cluster. The cluster can then be used to install the Agile Service Manager operator, and create an Agile Service Manager instance.

## Procedure

### Create a target registry to store all the images locally

1. Install and start a production grade Docker V2 compatible registry, such as Quay Enterprise, JFrog Artifactory, or Docker Registry.

The target registry must be accessible by the Red Hat OpenShift cluster and the bastion host. The Red Hat OpenShift internal registry is not supported. Create a secret for access to the target registry.

```
oc create secret docker-registry target-registry-secret \
  --docker-server=target_registry \
  --docker-username=user \
  --docker-password=password \
  --namespace=target_namespace
```

Where:

- **target\_registry** is the target registry that you created
- **target-registry-secret** is the name of the secret that you are creating
- **user** and **password** are the credentials to access your target registry
- **target\_namespace** is the namespace into which you want to deploy Agile Service Manager

### Prepare the bastion host

2. Logon to the bastion machine and verify that the bastion server has access to:

- The public internet to download CASE and images
- The target registry where the images will be mirrored
- The target cluster to install the Agile Service Manager operator
- The cloudctl and oc utilities

**Remember:** You can obtain your entitlement key here: <https://myibm.ibm.com/products-services/containerlibrary>

### Download the CASE bundle onto the bastion server

3. Download the Agile Service Manager CASE bundle, (**ibm-asm-operator-case**), into a local directory on your bastion server.

```
cloudctl case save --case ibm-asm-operator-case --outputdir destination_dir --repo https://raw.githubusercontent.com/IBM/cloud-pak/master/repo/case
```

Where:

- **<destination\_dir>** is a directory of your choosing, for example **/tmp/cases**
4. Verify that the CASE bundle, **images.csv**, and **charts.csv** have been successfully downloaded on your bastion server, with the following command:

```
find <destination_dir> -type f
```

Where:

- **<destination\_dir>** is a directory of your choosing, for example **/tmp/cases**

### Configure bastion server authentication

5. Set up access to the IBM Entitled Registry, **cp.icr.io**, from which you will be pulling images. Run the following command on your bastion server:

```
cloudctl case launch
--case ibm-asm-operator-case
--namespace ${NAMESPACE}
--inventory asmOperatorSetup
--action configure-creds-airgap
--args "--registry cp.icr.io --user cp --pass password"
```

Where:

- `${NAMESPACE}` is the custom namespace into which you want to deploy Agile Service Manager
- `cp` is the default user
- `password` is your IBM Entitled Registry entitlement key, as found when you prepared your cluster

#### 6. Set the target registry environment variable `$TARGET_REGISTRY`.

Run the following command on your bastion server:

```
export TARGET_REGISTRY=target_registry
```

Where:

- `target_registry` is the URL of the target registry that you created

**Important:** Do not use a forward slash (/) in the target registry. Airgap credentials are not saved if the target registry contains a /.

#### 7. Set up access to the target image registry, into which you will be copying images.

Run the following command on your bastion server:

```
cloudctl case launch
--case ibm-asm-operator-case
--namespace ${NAMESPACE}
--inventory asmOperatorSetup
--action configure-creds-airgap
--args "--registry $TARGET_REGISTRY --user USERNAME --pass PASSWORD"
```

Where:

- `${NAMESPACE}` is the custom namespace into which you want to deploy Agile Service Manager
- `USERNAME` and `PASSWORD` are the credentials for accessing the target registry that you created

The credentials are saved to `~/.airgap/secrets/$TARGET_REGISTRY.json`

### Mirror images from CASE to the target registry in the airgap environment

#### 8. Mirror images from CASE to the target registry on Red Hat OpenShift. This can take up to 15 minutes.

Run the following command on your bastion server:

```
cloudctl case launch
--case ibm-asm-operator-case
--namespace ${NAMESPACE}
--inventory asmOperatorSetup
--action mirror-images
--args "--registry $TARGET_REGISTRY/$TARGET_REGISTRY_NAMESPACE --inputDir /tmp/cases"
```

Where:

- `$TARGET_REGISTRY_NAMESPACE` is the custom namespace of the target registry under which the images will be copied

The images listed in the downloaded CASE (`images.csv`) are copied to the target registry in the airgap environment.

### Configure your Red Hat OpenShift Cluster for airgap

#### 9. Configure your Red Hat OpenShift Cluster for airgap. This step can take 30+ minutes.

Run the following command on your Red Hat OpenShift cluster to create a global image pull secret for the target registry, and create a `imagesourcecontentpolicy`.

```
$ cloudctl case launch
--case ibm-asm-operator-case
--namespace ${NAMESPACE}
--inventory asmOperatorSetup
--action configure-cluster-airgap
--args "--registry $TARGET_REGISTRY --inputDir /tmp/cases"
```

Where:

- `${NAMESPACE}` is the custom namespace to be used for your deployment
- `/tmp/cases` is the input directory containing the CASE bundle



**Warning:**

- Cluster resources must adjust to the new pull secret, which can temporarily limit the usability of the cluster. Authorization credentials are stored in `$HOME/.airgap/secrets` and `/tmp/airgap*` to support this action.
- Applying `imagesourcecontentpolicy` causes cluster nodes to recycle.

## Install the catalog and operator on your Red Hat OpenShift cluster

10. Install the Catalog on your cluster using CASE:

```
cloudctl case launch
--case ibm-asm-operator-case
--namespace ${NAMESPACE}
--inventory asmOperatorSetup
--action install-operator-catalog
```

Where:

- `${NAMESPACE}` is the custom namespace to be used for your deployment.

11. Verify the Agile Service Manager Catalog Source.

- From the Red Hat OpenShift OLM UI, navigate to **Administration > Cluster Settings**, and then select the **OperatorHub** configuration resource under the **Global Configurations** tab.
- Verify that the **ibm-asm-catalog** catalog source is present.

12. Verify the Agile Service Manager Operator.

- From the Red Hat OpenShift OLM UI, navigate to **Operators > Installed Operators**.
- Verify that the status of the Agile Service Manager operator is Succeeded.

## Create Agile Service Manager instance

13. Create the instance.

- From the Red Hat OpenShift OLM UI, navigate to **Operators > Installed Operators**, and select the Agile Service Manager operator.
- Under Provided APIs, select **Cloud Deployment**, and then select **Create Instance**.

14. Define the operator properties.

- From the Red Hat OpenShift OLM UI, use the YAML view to configure the properties for the Agile Service Manager deployment.
- Edit the properties to provide access to the target registry.

Add the following:

```
spec:
  helmValuesASM:
    "global.image.useTag": false
```

For more information about configurable properties for a cloud-only deployment, see [“Agile Service Manager Operator cloud properties” on page 47](#).

15. Click **Create**.

Under the **All Instances** tab, the Agile Service Manager instance appears. View the status of each of the updates on the installation. When the instance's status shows **OK** then the deployment is ready.



## What to do next

If you want to enable or disable a feature or observer after your installation, you can edit the Agile Service Manager instance by running the following command:

```
oc edit ASM <asm-instance-name>
```

You can then select to enable or disable the feature or observer.

**Tip:** When you disable features post install, the resource is not automatically deleted. To find out if the feature is deleted or not, check the operator log.

### Limitations:

- Changing an existing deployment from its original deployment type to any other deployment type is not supported.
- Changing an instance's deployment parameters in the form view is not supported after deployment.
- If you update custom secrets in the OLM console, the crypto key becomes corrupted and the command to encrypt passwords does not work. **Only update custom secrets with the CLI.**

## Configuring alternative storage (on OCP)

You can deploy Agile Service Manager on OCP (for Watson AIOps) with alternative storage.

### About this task

See the related link for more information on deciding which storage type meets your requirements.

#### Related information

[Persistent storage using local volumes](#)

[Persistent storage using VMware vSphere volumes](#)

## Uninstalling Agile Service Manager on OCP for Watson AIOps

You can uninstall a stand-alone installation of Agile Service Manager on OCP deployed with Watson AIOps using Container Application Software for Enterprises (CASE), either with or without the Red Hat OpenShift Operator Lifecycle Manager (OLM) UI.

### About this task

Choose one of the following methods to uninstall the stand-alone OCP version of Agile Service Manager.

#### Uninstalling *with* OLM UI

Complete [this procedure](#).

#### Uninstalling *without* OLM UI

Complete [this procedure](#).

## Uninstalling Agile Service Manager with OLM UI

### Procedure

1. Delete the Agile Service Manager instance.
  - a) In the OLM UI, go to **Operator > Installed Operators** and select the project where you installed Agile Service Manager.
  - b) Click **IBM Netcool Agile Service Manager > All Instances**.
  - c) Select **Delete ASM** from the menu.
2. Delete the service account created for above operator instance.
  - a) Go to **User Management > Service Accounts** and select the project where you installed Agile Service Manager.

- b) Delete the service-account with name `${RELEASE_NAME}-service-account`, where `${RELEASE_NAME}` is the Agile Service Manager instance name.
3. Delete the Agile Service Manager operator.
  - a) Go to **Operator > Installed Operators** and select the project where you installed Agile Service Manager.
  - b) Select the options menu for **IBM Netcool Agile Service Manager** and click **Uninstall Operator**.
4. Delete the Custom Resource Definitions (CRDs).



**CAUTION:** Do not delete CRDs in scenarios where multiple instances of Agile Service Manager exist on the same cluster.

- a) Go to **Administration > Custom Resource Definitions**.
  - b) Select the CRDs that were created by the Agile Service Manager installation.
  - c) Delete all the CRDs that start with 'asm' and 'asmformation'.
5. Delete the secrets that were created for your deployment.
  - a) Go to **Workloads > Secrets** and select the project where you installed Agile Service Manager.
  - b) Delete all secrets that start with `${RELEASE_NAME}`, where `${RELEASE_NAME}` is the Agile Service Manager instance name.
6. Delete the ConfigMaps that were created by Agile Service Manager.
  - a) Go to **Workloads > Config Maps** and select the project where you installed Agile Service Manager.
  - b) Delete all config maps that start with `${RELEASE_NAME}`, where `${RELEASE_NAME}` is the Agile Service Manager instance name.
7. Delete the persistent volume claims and persistent volumes:



**CAUTION:** Confirm that you are not about to delete important data before deleting persistent volume claims and persistent volumes.

- a) Go to **Storage > Persistent Volume Claims** and delete all Persistent Volume Claims for Agile Service Manager.
  - b) Go to **Storage > Persistent Volume** and delete all Persistent Volumes for Agile Service Manager.

## Uninstalling Agile Service Manager without OLM UI

### Procedure

1. Delete the Agile Service Manager operator instance.

```
oc delete asm ${RELEASE_NAME} -n ${NAMESPACE}
```

2. Delete the service account created for this instance.

```
oc create serviceaccount ${RELEASE_NAME}-service-account -n ${NAMESPACE}
```

3. Delete the Agile Service Manager operator.

```
cloudctl case launch \
  --case ibm-asm-operator-case \
  --namespace ${NAMESPACE} \
  --inventory asmOperatorSetup \
  --action uninstall-operator-native
```

Append `--args "--deleteCRDs"` to above command if CRDs need to be deleted as part of uninstall.



**CAUTION:** Do not delete CRDs in scenarios where multiple instances of Agile Service Manager exist on the same cluster.

4. Delete the secrets that were created for your deployment.

```
oc delete secret $REGISTRY_SECRET -n ${NAMESPACE}
```

Where \$REGISTRY\_SECRET is the secret created for image registry access during installation.

5. Delete all secrets that start with '\${RELEASE\_NAME}'.

```
oc get secret -o name | grep ${RELEASE_NAME}| xargs oc delete -n ${NAMESPACE}
```

6. Delete the ConfigMaps that were created by Agile Service Manager.

```
oc get configmap -o name | grep ${RELEASE_NAME}| xargs oc delete -n ${NAMESPACE}
```

7. Delete the persistent volume claims and persistent volumes:



**CAUTION:** Confirm that you are not about to delete important data before deleting persistent volume claims and persistent volumes.

- a) Delete the persistent volume claims:

```
oc delete pvc <pvc-name>
```

- b) Delete the persistent volumes:

```
oc delete pv <pv-name>
```

## Agile Service Manager Operator cloud properties

This topic lists the configurable properties for a deployment of Agile Service Manager on OCP for Watson AIOps (via CASE and CLI, or the OLM UI), as well as on an **offline** Red Hat OpenShift cluster that has no internet connectivity (via airgapped environment).

**ASM Operator Version:** Ensure you choose the latest version, which by default is 1.1.0. **Do not** install version 1.0.0.

### Mandatory Agile Service Manager cloud properties

To create Agile Service Manager cloud instances, the following properties are mandatory and must be provided with valid values.

Table 6. Mandatory Agile Service Manager cloud properties	
Property	Description
clusterDomain	Specify the fully qualified domain name (FQDN)
deploymentType	Deployment type ('trial', 'smallProduction', or 'production')
entitlementSecret	Entitlement secret to pull images
license.accept	"I have read and agreed to the license agreement"

### Full Agile Service Manager cloud properties

The following table lists all the properties that are available to create Agile Service Manager cloud instances.

Table 7. Full Agile Service Manager cloud properties	
Property	Description
clusterDomain	Specify the fully qualified domain name (FQDN)

Table 7. Full Agile Service Manager cloud properties (continued)

Property	Description
deploymentType	Deployment type ('trial', 'smallProduction', or 'production')
entitlementSecret	Entitlement secret to pull images
helmValuesASM	To be used as an attribute of your 'Go Spec' object example
license.accept	Set the following to <b>true</b> to accept the license: "I have read and agreed to the license agreement"
advanced.imagePullPolicy	The default pull policy is IfNotPresent, which causes the Kubelet to skip pulling an image if it already exists.
advanced.imagePullRepository	Docker registry from which all component images are pulled
persistence.enabled	Enable persistence storage
persistence.storageClassCassandraBackup	Storage class for Cassandra backup
persistence.storageClassCassandraData	Storage class for Cassandra data
persistence.storageClassElastic	Storage class for Elasticsearch
persistence.storageClassKafka	Storage class for Kafka
persistence.storageClassZookeeper	Storage class for Zookeeper
persistence.storageSizeCassandraBackup	Storage size for Cassandra backup
persistence.storageSizeCassandraData	Storage size for Cassandra data
persistence.storageSizeElastic	Storage size for Elasticsearch
persistence.storageSizeKafka	Storage size for Kafka
persistence.storageSizeZookeeper	Storage size for Zookeeper
storageClassFileObserver	Storage class for File Observer
storageSizeFileObserver	Storage size for File Observer
asm.cassandra	Deploy Cassandra
asm.kafka	Deploy Kafka
asm.zookeeper	Deploy Zookeeper

*Table 7. Full Agile Service Manager cloud properties (continued)*

<b>Property</b>	<b>Description</b>
asm.elasticsearch	Deploy Elasticsearch
asm.almObserver	Deploy ALM (IBM Agile Lifecycle Manager) Observer
asm.appdynamicsObserver	Deploy AppDynamics Observer
asm.awsObserver	Deploy AWS Observer
asm.ansibleawxObserver	Deploy Ansible AWX Observer
asm.azureObserver	Deploy Microsoft Azure Observer
asm.bigfixinventoryObserver	Deploy BigFix Inventory Observer
asm.cienablueplanetObserver	Deploy Ciena Blue Planet Observer
asm.ciscoaciObserver	Deploy Cisco ACI Observer
asm.contrailObserver	Deploy Juniper Contrail Observer
asm.commonUi	Deploy Common UI
asm.dnsObserver	Deploy DNS Observer
asm.dockerObserver	Deploy Docker Observer
asm.dynatraceObserver	Deploy Dynatrace Observer
asm.fileObserver	Deploy File Observer
asm.googlecloudObserver	Deploy Google Cloud Platform Observer
asm.ibmcloudObserver	Deploy IBM Cloud Observer
asm.itnmObserver	Deploy ITNM (IBM Tivoli Network Manager) Observer
asm.jenkinsObserver	Deploy Jenkins Observer
asm.junipercsoObserver	Deploy Juniper Networks CSO Observer
asm.kubernetesObserver	Deploy Kubernetes Observer
asm.newrelicObserver	Deploy New Relic Observer
asm.openstackObserver	Deploy OpenStack Observer
asm.restObserver	Deploy REST Observer
asm.servicenowObserver	Deploy ServiceNow Observer

Table 7. Full Agile Service Manager cloud properties (continued)	
Property	Description
asm.taddmObserver	Deploy TADDM (IBM Tivoli Application Dependency Discovery Manager) Observer
asm.vmvcenterObserver	Deploy VMware VCenter Observer
asm.vmwarensxObserver	Deploy VMware NSX Observer
asm.zabbixObserver	Deploy Zabbix Observer
asm.rancherObserver	Deploy Rancher Observer

### Related tasks

“Installing Agile Service Manager on OCP for Watson AIOps (via CASE without OLM UI)” on page 40

A stand-alone OCP version of Agile Service Manager is supported for deployment with Watson AIOps. You can install this version using Container Application Software for Enterprises (CASE) **without** also using the Red Hat OpenShift Operator Lifecycle Manager (OLM) UI.

“Installing Agile Service Manager on OCP for Watson AIOps (via CASE with OLM UI)” on page 38

A stand-alone OCP version of Agile Service Manager is supported for deployment with Watson AIOps. You can install this version using both the Container Application Software for Enterprises (CASE) and the Red Hat OpenShift Operator Lifecycle Manager (OLM) UI.

## Sizing reference for Agile Service Manager on OCP (Watson AIOps)

This topic **only** describes the sizing requirements for a stand-alone Private Cloud deployment of Agile Service Manager on RedHat OpenShift Container Platform for IBM Watson AIOps.

### Agile Service Manager stand-alone on RedHat OpenShift Container Platform (for Watson AIOps)

**Important:** The following sizing only applies to a stand-alone deployment of Agile Service Manager on RedHat OpenShift Container Platform for use with IBM Watson AIOps. For Agile Service Manager deployment with Netcool Operations Insight, see the NOI Knowledge Center topics.

Table 8. General sizing requirements				
Category	Resource	Trial	Production	smallProduction
<b>Event Management sizing information</b>				
<b>Topology Management sizing information</b>				
System size	Approx. resources	200,000	1,000,000	70,000
Environment options	Container size	<b>Trial</b>	<b>Production</b>	<b>smallProduction</b>
	High Availability	NO	YES	YES

The following table shows the total hardware requirements for an Agile Service Manager standalone system deployed on a OCP cluster, including both the Agile Service Manager standalone and OCP related hardware needs. This information is intended for customers that need to create a new OCP cluster to deploy the full **OCP+ASMstandalone** stack.

*Table 9. Total hardware requirements Agile Service Manager standalone Including OCP and infrastructure nodes*

Category	Resource	Demo/PoC/Trial	Production	smallProduction
Minimum CPU/MEM/DISK	x86 CPUs	8	45	20
	Memory (GB)	32	120	108
	Disk (GB)	240	720	720
Persistent storage Requirements (Gi)		820	2740	480
Requirements		1300	7450	7450

The following table shows the hardware requirements attributed to the Agile Service Manager standalone footprint deployed on OCP. This information is useful for customers who already have an OCP cluster installed, but need to add worker/compute nodes to it to accommodate Agile Service Manager standalone.

*Table 10. Total hardware requirements Agile Service Manager standalone services only*

Category	Resource	Demo/PoC/Trial	Production	smallProduction
Minimum CPU/MEM/DISK	x86 CPUs (Min)	4	33	8
	Memory (GB) (Min)	16	72	60
	Disk (GB) (Min)	120	360	360
Persistent storage Requirements (Gi)		820	2740	480
Requirements		1300	7450	7450

The following table illustrates the recommended resource allocation for the OCP Master and Worker nodes, along with the recommended configuration for the disk volumes associated with each persisted storage resource. Worker nodes may be larger when Agile Service Manager standalone is deployed along with other applications.

*Table 11. Hardware allocation and configuration*

Category	Resource	Demo/PoC/Trial	Production	smallProduction
OCP Control Plane (Master) Node(s)	Minimum Nodes Count	1	3	3
CPU cores and memory requirements	x86 CPUs	4	4	4
	Memory (GB) (Min)	16	16	16
	Disk (GB) (Min)	120	120	120
<b>Agile Service Manager standalone components suggested configuration</b>				
OCP Compute (Worker) Nodes	Minimum Nodes Count	1	3	3

<i>Table 11. Hardware allocation and configuration (continued)</i>				
<b>Category</b>	<b>Resource</b>	<b>Demo/PoC/Trial</b>	<b>Production</b>	<b>smallProduction</b>
CPU cores and memory requirements	x86 CPUs	4	11	3
	Memory (GB)	16	24	20
	Disk (GB)	120	120	120
Persistent storage min Requirements (Gi)	Cassandra	200	1500	200
	Kafka	150	300	60
	Zookeeper	15	30	15
	Elasticsearch	450	900	200
	File Observer	5	10	5
Persistent storage recommended IOPS Requirements	Cassandra	600	2400	2400
	Kafka	300	2400	2400
	Zookeeper	50	150	150
	Elasticsearch	300	2400	2400
	File Observer	50	100	100



## Chapter 4. Configuring components

After installing Agile Service Manager, you configure the optional components, such as the Jenkins plugin. You can also configure a hybrid system consisting of a RedHat OpenShift Container Platform backend (core) and an Agile Service Manager UI

### Related concepts

[“Installing Agile Service Manager” on page 13](#)

You can deploy Agile Service Manager either stand-alone or on RedHat OpenShift Container Platform (OCP). The installation and post-install configuration steps differ for these versions.

## Configuring the Jenkins plugin

The Agile Service Manager software includes the Jenkins plugin, which you install on your Jenkins server using the Jenkins Plugin Manager Advanced installation wizard. From the Jenkins server, the plugin gathers and sends information to the Jenkins Observer.

### Before you begin

Ensure you have all the required administrator or management permissions to let you deploy the Agile Service Manager Jenkins component onto the Jenkins server.

Table 12. Agile Service Manager requirements for the Jenkins plugin

Plugin	Version	Description
Jenkins core	2.150.1	Jenkins server minimum version
apache-httpcomponents-client-4-api	4.5.10-2.0	Provided by Jenkins core
workflow-step-api	2.22	
workflow-cps	1.9	
credentials	2.3.1	Compatible version usually provided by Jenkins core
Git	2.6.5	
artifactory	2.15.0	Optional (for using the Artifactory integration only)

### Obtain the Jenkins plugin:

#### On-prem

The Jenkins plugin is available on the Agile Service Manager server after the Jenkins Observer is installed. Find it as in the following example:

```
$ find /opt/ibm/netcool/asm/integrations/jenkins/*plugin*
```

Sample result:

```
/opt/ibm/netcool/asm/integrations/jenkins/asm-observer-plugin-<version>.hpi
```

#### On OCP

You can find the Jenkins plugin included with the Jenkins Observer image. For example:

```
$ JENKINS_OBSERVER_POD=$(oc get pod -l app=jenkins-observer -o jsonpath='{.items[0].metadata.name}')
```

```
$ echo ${JENKINS_OBSERVER_POD}
asm-jenkins-observer-84c456f58c-jxlrj
```

To copy the plugin from the image:

```
$ oc cp $JENKINS_OBSERVER_POD:/opt/ibm/asm-observer-plugin-<version>.hpi asm-observer-
plugin-<version>.hpi
```

```
$ ls asm-observer-plugin-<version>.hpi
asm-observer-plugin-<version>.hpi
```

## Obtain credentials:

### On-prem

Obtain the credentials from ASM\_HOME / .env

The default is asm and asm

### On OCP

The Agile Service Manager UI API credentials are dynamically created. If Agile Service Manager is installed with a release name of asm, then:

```
ASM_RELEASE_NAME=asm
ASM_USER=$(oc get secret ${ASM_RELEASE_NAME}-asm-credentials -o jsonpath={.data.username} |
base64 -d)
ASM_PASS=$(oc get secret ${ASM_RELEASE_NAME}-asm-credentials -o jsonpath={.data.password} |
base64 -d)
```

Use the following command to obtain the UI API credentials:

```
echo $ASM_USER $ASM_PASS
```

Example user and password credentials returned:

```
asm-netcool-user N8gJJGEfmmDnF6Q/1zg8NyAGKgQ9PmZQLhUSKd9/j54=
```

## Obtain certificate:

### On-prem

If your installation uses the proxy service, generate the required certificate as in the following example:

```
openssl pkcs12 -export -out <your_cert_file>.p12 -inkey asm-nginx.key
-in asm-nginx.crt
```

To generate the certificate from your installation truststore, use the following example:

```
keytool -importkeystore -srckeystore {path to keystore} -srcstorepass
{encrypted keystore password} -destkeystore <your_cert_file>.p12
-deststoretype PKCS12 -deststorepass {password you want to set}
-destkeypass {password you want to set}
```

### On OCP

The following sample steps apply to the default router deployed on OCP. If additional routers are used, adjust these steps accordingly. You perform these steps as **cluster administrator**.

Obtain the certificate:

```
oc get secret -n openshift-ingress router-certs-default -o
jsonpath='{.data.tls\.crt}' | base64 -d > ocp-default-router.crt
```

Obtain the certificate key:

```
oc get secret -n openshift-ingress router-certs-default -o
jsonpath='{.data.tls\.key}' | base64 -d > ocp-default-router.key
```

Convert to PKCS12:

```
openssl pkcs12 -export -out ocp-default-router.p12 -inkey
ocp-default-router.key -in ocp-default-router.crt
Enter Export Password: enter-suitably-secure-password
Verifying - Enter Export Password: enter-suitably-secure-password
```

Ensure that the files you have extracted from OCP are securely deleted.

### Obtain observer URL:

#### On-prem

The expected observer URL is `https://<your-asm-onprem-host>/`

#### On OCP

Obtain the hostname as in the following example:

```
# oc get routes -l release=asm -l app=jenkins-observer -o jsonpath='{.items[*].spec.host}'
asm.netcool.apps.<your-ocp-cluster>
```

## About this task

The Jenkins plugin defines Jenkins pipeline DSL steps that you use to instruct the Jenkins builds to generate and then send topology data to the Jenkins Observer. The DSL steps to send the notification lets you include Artifactory BuildInfo modules information in order for the build products to be modeled. You create an observer instance, which then triggers the sending of the notification to the observer.

For information on configuring Jenkins Observer jobs, see the related tasks.

**Important:** You perform the following steps as **Jenkins administrator**.

## Procedure

### Install the Jenkins plugin

1. Log in to your Jenkins server with a user ID that has administration (management) permissions.
2. Navigate to **Manage Jenkins > Manage Plugins** and select the **Advanced** tab.
3. In the Upload Plugin section, upload and install the provided plugin (asm-observer-plugin-<version>.hpi)

### Configure the Jenkins plugin

4. Navigate to **Manage Jenkins > Configure System** and select the **Agile Service Manager plugin**.
5. Enter the Jenkins Observer URL obtained earlier in the **Observer URL** field.
6. Create credentials for Agile Service Manager using the **Manage Jenkins > Credentials** configuration page.

Enter the credentials and the certificate details obtained earlier from the (or as the) cluster administrator.

#### Auth Credentials

Enter the username and password to authenticate with the Jenkins Observer at API level.

Click **ADD**.

#### Certificate

Enter the PKCS12 format certificate to secure the communication between the Jenkins plugin and the observer using SSL.

Click **ADD**.

See the Jenkins site for more detailed information on using the Jenkins Credentials feature: <https://jenkins.io/doc/book/using/using-credentials/>

7. Optional: Enter a **Default Tenant ID**.

The tenant to be used as default when sending information to Agile Service Manager (if it is not overridden in the build).

8. Optional: Enter a **Default Job ID**.

The default job ID to be used when sending information to Agile Service Manager (if it is not overridden in the build). This represents the ID of the job the observer is running to process your information.

9. Click **TEST CONNECTION** to attempt to establish a connection with the observer using the configured credentials.

If successful, a Connection successful! message is displayed. If unsuccessful, an error is returned.

**Note:** The connection test does not verify the default job and tenant IDs.

### Configuring your Jenkins build to notify Agile Service Manager

After configuring the installed plugin in Jenkins, you integrate the provided functionality into your existing Jenkins build. The following two steps describe how to use scripted pipeline syntax to send your build information to the Agile Service Manager Jenkins Observer.

**Remember:** You perform the following two steps on your Jenkins build (and **not** the Jenkins Observer).

10. Create an instance of the Agile Service Manager observer communication object using the default credentials set during the [configuration](#) steps.

```
def <my observer> = ASM.newObserver()
```

#### Tip:

##### Override credentials

To override any configured credentials (that is, the Jenkins credentials ID rather than the actual credentials themselves), you can use the following syntax:

```
def <my observer> = ASM.newObserver  
authCredsId: 'asm-auth-credentials-id', certCredsId: 'asm-cert-credentials-id'
```

11. Send a build notification using the following syntax:

```
ASM.notifyASM asmObserver: <my observer>
```

#### Tip:

You can override the default values for the tenant and job IDs:

```
ASM.notifyASM tenantId: '<tenantId>', jobId: '<jobId>', asmObserver: <my  
observer>
```

You can integrate the Jenkins plugin with the Artifactory component using the modules information to model any artifacts the build has published to the repository:

```
ASM.notifyASM asmObserver: <my observer>, artModules:  
<buildInfoObject>.getModules()
```

### Example

The following Jenkinsfile code sample shows the integration with Agile Service Manager. It starts with the declaration of the observer instance, and ends with the notification to Agile Service Manager. This notification should be added as a `post / always` block, so that the information is sent as the last operation of the build, and is sent regardless of the result of the build.

The optional `try / catch` around the notification to Agile Service Manager is included in this sample to avoid the build failing if the notification to Agile Service Manager fails.

The following sample also depicts the use of Artifactory to publish the build products, and how that information can be included in the notification to Agile Service Manager using the **artModules** parameter.

```
def asm = ASM.newObserver()  
def buildInfo = Artifactory.newBuildInfo()
```

```

buildInfo.env.capture = true

... your build logic here ...

post {
    always {
        script {
            try {
                ASM.notifyASM asmObserver: asm, artModules: buildInfo.getModules()
            } catch (err) {
                echo "An error occurred sending a notification to ASM: ${err}"
            }
        }
    }
}

```

**Tip:** You can further refine your integration and visualization, as described [here](#).

## What to do next

After you have installed and configured the Jenkins plugin, you configure the Jenkins Observer listen job.

You can further refine your Agile Service Manager system to take advantage of its topology and other customization functionality.

### Related tasks

[“Configuring Jenkins Observer jobs” on page 141](#)

Using the Jenkins Observer, you can define listen jobs that receive build information generated by the Agile Service Manager plugin for Jenkins.

[“Defining Jenkins Observer jobs” on page 208](#)

Using the Jenkins Observer, you can define listen jobs that receive build information generated by the Agile Service Manager plugin for Jenkins.

[“Refining Jenkins integration and visualization” on page 57](#)

You can extend your Jenkins integration with rules to merge data from different sources, custom topology display conventions, and the use of templates for the automated generation of topologies.

### Related information

[“Jenkins Observer troubleshooting” on page 347](#)

## Refining Jenkins integration and visualization

You can extend your Jenkins integration with rules to merge data from different sources, custom topology display conventions, and the use of templates for the automated generation of topologies.

### Before you begin

You first install and configure the Jenkins plugin, and then configure your Jenkins build to notify Agile Service Manager, before extending its functionality as described here.

### About this task



**Trouble:** See the Jenkins Observer [troubleshooting](#) section for useful Jenkins troubleshooting information.

### Related concepts

[“Customizing UI elements” on page 271](#)

You can customize a number of Agile Service Manager UI elements for your deployment, such as tooltips, link styles and icons. You can also create custom tools which users can access through a topology's context menu.

[“Porting data for testing, backup and recovery” on page 290](#)

You can create backups of your Agile Service Manager UI configuration data in order to run a test configuration, or simply to safeguard your custom settings. You can also back up and restore your topology data.

### **Related tasks**

[“Configuring Jenkins Observer jobs” on page 141](#)

Using the Jenkins Observer, you can define listen jobs that receive build information generated by the Agile Service Manager plugin for Jenkins.

[“Defining Jenkins Observer jobs” on page 208](#)

Using the Jenkins Observer, you can define listen jobs that receive build information generated by the Agile Service Manager plugin for Jenkins.

[“Configuring the Jenkins plugin” on page 53](#)

The Agile Service Manager software includes the Jenkins plugin, which you install on your Jenkins server using the Jenkins Plugin Manager Advanced installation wizard. From the Jenkins server, the plugin gathers and sends information to the Jenkins Observer.

[“Using templates to generate defined topologies” on page 315](#)

You can create topology templates to generate defined topologies, which search your topology database for instances that match its conditions. These defined topologies can then be searched for in the Topology Viewer and displayed.

[“Defining rules” on page 309](#)

Rules help streamline topologies and conserve system resources, for example by merging different observer records of the same resource into a single composite resource, or by excluding specific changes from being recorded against a resource history.

### **Related information**

[“Jenkins Observer troubleshooting” on page 347](#)

## **Defining Jenkins data merge rules**

You can extend your Jenkins integration with rules to merge data from different sources.

### **About this task**

You can create merge rules that merge data from different sources as described in the [“Defining rules” on page 309](#) topic.

### **Procedure**

Define a merge rule that merges Docker images information provided by the Docker Observer.

If you plan to include Artifactory modules information as part of your build notifications using the Agile Service Manager plugin, you can define a merge rule that will merge Artifactory's published Docker modules information with the information provided by the Docker Observer. For example:

```
{
  "name": "dockerFromJenkins",
  "ruleType": "mergeRule",
  "tokens": [
    "${docker.image.id}", "sha256:${docker.image.id}"
  ],
  "ruleStatus": "enabled",
  "entityTypes": [
    "image"
  ],
  "observers": [
    "jenkins-observer"
  ],
  "providers": [
    "*"
  ]
}
```

## Customizing the Jenkins topologies

You can customize how your Jenkins topologies are displayed in a number of ways, either by using the Jenkins customizations supplied with Agile Service Manager, or by creating your own custom display elements.

### About this task

You can use the Agile Service Manager backup and restore functionality to import preset Jenkins UI right-click URL tools definitions included with Agile Service Manager.

The Jenkins-specific UI customizations are contained in the `{ASM_HOME}/data/tools/jenkins_observer_config.json` configuration file.

The configuration file also includes some styling changes to the builds relationship type.

See the [example](#) at the end of this topic for more details.

**Tip:** The information obtained from Jenkins via the Agile Service Manager plugin for Jenkins and the Jenkins Observer includes URLs for several Jenkins elements. You can create topology tools that use these URLs to provide in-context navigation to Jenkins or your Git repositories from within the topologies displayed in Agile Service Manager.

### Procedure

1. Run the following command to generate types and tools appropriate for your Jenkins topologies:

```
{ASM_HOME}/bin/import_ui_config.sh -file jenkins_observer_config.json
```

2. Should the provided customizations conflict with existing definitions, you can add the `-overwrite` flag:

```
{ASM_HOME}/bin/import_ui_config.sh -file jenkins_observer_config.json -overwrite
```

### Example

**Sample topology tools for Jenkins topologies (provided in the `{ASM_HOME}/data/tools/jenkins_observer_config.json` file.)**

**Opens a GitHub repository commits log**

```
function url2Http(url) {
  if(url.indexOf("http") != -1) {
    return url;
  }
  var res = url.split(":");
  var ghUrl = res[0].replace("git@", "https://");
  var repoUrl = res[1].replace(".git", "");
  return ghUrl + "/" + repoUrl;
}

if (asmProperties.repositoryUrl && asmProperties.revision) {
  var normUrl = url2Http(asmProperties.repositoryUrl);
  window.open(normUrl + '/ commits/'
    + asmProperties.revision);
}
```

Relationship type is 'builds'

**Opens the specific commit**

```
function url2Http(url) {
  if(url.indexOf("http") != -1) {
    return url;
  }
  var res = url.split(":");
  var ghUrl = res[0].replace("git@", "https://");
  var repoUrl = res[1].replace(".git", "");
  return ghUrl + "/" + repoUrl;
```

```

}

if (asmProperties.repositoryUrl && asmProperties.revision) {
    var normUrl = url2Http(asmProperties.repositoryUrl);
    window.open(normUrl + '/commit/'
        + asmProperties.revision);
}

```

Relationship type is 'builds'

### Opens the Jenkins resources records

```

if (asmProperties.url) {
    window.open(asmProperties.url);
}

```

Resource types are 'person', 'build', 'job'

### Opens a GitHub repository tree view from the commit

```

function url2Http(url) {
    if(url.indexOf("http") != -1) {
        return url;
    }
    var res = url.split(":");
    var ghUrl = res[0].replace("git@", "https://");
    var repoUrl = res[1].replace(".git", "");
    return ghUrl + "/" + repoUrl;
}

if (asmProperties.repositoryUrl && asmProperties.revision) {
    var normUrl = url2Http(asmProperties.repositoryUrl);
    window.open(normUrl + '/tree/'
        + asmProperties.revision);
}

```

Relationship type is 'builds'

### Opens the build log console view in Jenkins

```

if (asmProperties.buildConsoleLog) {
    window.open(asmProperties.buildConsoleLog);
}

```

Resource type is 'build'

## What to do next

If required, you can customize how your Jenkins topologies are displayed further using the [“Customizing UI elements”](#) on page 271 topics.

## Creating Jenkins topology templates

Over time topologies generated via a specific Jenkins Observer job may grow, including the information from different Jenkins builds. You can use the Agile Service Manager topology template functionality to create meaningful, scoped views per build.

### About this task

This topic describes two suggested build templates:

- Topology per build
- Current pipeline status

The use of the Agile Service Manager templates is described in the [“Using templates to generate defined topologies”](#) on page 315 topic.

## Procedure

### Create 'per build' topology templates



The following template lets you visualize individual topologies for each build.

1. As the admin user, log into your DASH web application, select **Administration** from the DASH menu, then select **Topology Templates** under the Agile Service Management heading.
2. Search for build by its entity type, or by a specific build name (for example, sample-pipeline/master), and then select **View Topology** next to a result.
3. Select the **Dynamic** template type to ensure that all matching builds will generate topologies like this.
4. On the **Preview** pane, right-click the build resource and select **Follow Relationship**, then repeat this for all the connected resources in your build topology.

**Important:** Do not expand the 'lastRunOf' relationship.

5. Optionally, define tags for the template, then click **Save template and generate topologies**.

#### Create a 'current pipeline status' topology template

The following template generates topologies that show the current status of your pipeline, that is, the latest build for which a notification was received.

6. As the admin user, log into your DASH web application, select **Administration** from the DASH menu, then select **Topology Templates** under the Agile Service Management heading.
7. Search for job by its entity type, or by a specific build name (for example, sample-pipeline), and then select **View Topology** next to a result.
8. Select the **Dynamic** template type to ensure that all matching builds will generate topologies like this.
9. On the **Preview** pane, right-click the job resource, select **Follow Relationship**, then click **lastRunOf**. The latest build resource for which a topology was generated is displayed.
10. Expand any relationships you want to include in this template.

**Important:** Do not expand the 'RunOf' relationship.

11. Optionally, define tags for the template, then click **Save template and generate topologies**.

## Configuring application discovery service

The Agile Service Manager application discovery service is installed together with the other core components. Once you have provided the discovery details, Agile Service Manager can discover application data.

### Before you begin

Ensure you have all the required application discovery information available, such as the application discovery username and password, scope and accessList details.

**Remember:** You provide configuration details for the deployed application discovery service **after** installing the core Agile Service Manager and application discovery containers, and **after** ensuring that the installed Application Discovery Observer, Application Discovery engines and the provided Db2 database container are running.

Verify the Application Discovery deployment before configuring the observer jobs. Use the following command to verify that the docker containers are running:

```
ASM_HOME/bin/docker-compose ps
```

The system should return output indicating that the following containers have a state of Up and therefore are running:

```
...
asm_appdisco-observer_1
asm_app-disco-discovery_1
asm_app-disco-primarystorage_1
asm_db2_1
```

## About this task

You first upload the configuration details in a commonconfig.json file via **POST /file** command. You then run the **PUT /config** command for the Application Discovery Observer, which first validates the files, and then saves them in the topology management artifact and application discovery.

The application discovery service discovers configuration information about the entire application infrastructure, identifying deployed software components, physical servers, network devices, virtual systems, and host data. It then saves the information in the provided database. The Application Discovery Observer then retrieves network topology data from the database via REST API, which it uses to create topologies within the Agile Service Manager topology service.

Table 13. Application discovery model objects mapped to Agile Service Manager entity types	
Application Discovery model object	Agile Service Manager entity types
AppServer	application
ComputerSystem	host, server, router, switch
CPU	cpu
L2Interface	networkinterface
IpInterface	networkinterface
IpAddress	ipaddress
OperatingSystem	os
Service	service
StorageExtent	storage
Function	service
SAPSystem	application
Collection	Group
Router	Router

### Terminology:

- A computer system can be a host, server, router or switch
- A computer system contains CPU, L2Interface and storage
- Operating system, application server and service run on a computer system
- A computer system can connect to another computer system
- A SAPSystem contains a collection
- An application server can be a member of a collection

## Procedure

1. Edit the example configuration file included with the Application Discovery installation, or create a new configuration file.

### For on-prem

Edit the included configuration file:

```
ASM_HOME/data/appdisco-observer/ExampleCommonconfig.json
```

Save the file as commonconfig.json

Alternatively, in the ASM\_HOME/bin directory, generate the configuration file using the generate\_common\_config.sh script:

```
./generate_common_config.sh -d
```

Expected output includes:

```
...Successfully generated commonconfig.json in /opt/ibm/netcool/asm/bin/../data/appdisco-observer folder...
```

#### For OCP

Copy the text from the following example file: [“Application discovery service configuration template” on page 65](#)

Save the file as commonconfig.json

2. Edit or populate the following required properties in the configurations file for Application Discovery observer.

See the related links for the full text of the included example configuration file.

#### configId

Discovery profile ID

#### username

Application Discovery username

Value is administrator

#### password

Application Discovery encrypted password

**For on-prem:** Value is collation, which **must** be encrypted.

**For OCP:** No update is required. The Application Discovery Observer retrieves the password from appdisco secret:

```
{releasename}-topology-nasm-app-disco-admin-secret
```

#### discovery\_profile

Application Discovery profile

#### scope

Application Discovery scope

#### accessListType

Application Discovery access list objects and their properties

See the following topic for all accessListTypes and properties: [“Supported application discovery accessListTypes and properties” on page 68](#)

**Note:** The Application Discovery user gets locked out after three failed login attempts. This lockout is cleared after 30 minutes.

3. Encrypt the following parameters:

- password
- privatePassword
- enablePassword
- passPhrase
- communitystring

#### For on-prem

To encrypt these parameters, run the encrypt\_password.sh script in the ASM\_HOME/bin directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the parameters to be encrypted, in turn. The encryption utility will return encrypted versions.

#### For OCP

For OCP, encrypt these parameters with the following command:

```
kubectl exec -ti asm-topology-pods -- java -jar /opt/ibm/topology-service/topology-  
service.jar  
encrypt_password --password '<password to encrypt>'
```

4. Upload the configurations file (that is, 'commonconfig.json') via the Application Discovery Observer's Swagger UI **POST /file** command, or cURL command.

Example cURL command:

```
curl --insecure --header 'Content-Type: multipart/form-data' --header  
'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --  
header  
'Authorization: Basic YXNtOmFzbQ==' --request POST --url 'https://localhost/1.0/appdisco-  
observer/files' --form  
'job_file=@/tmp/commonconfig.json'
```

The configurations file will be saved in the ASM\_HOME/data/appdisco-observer/tenant-id directory.

5. Run **PUT /config** via Swagger UI or cURL command to create/update/replace the discovery details from the configurations file to the topology artifacts and Application Discovery.

Example cURL command:

```
curl --insecure --header 'Content-Type: application/json' --header  
'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --  
header  
'Authorization: Basic YXNtOmFzbQ==' --request PUT --url 'https://localhost/1.0/appdisco-  
observer/config?filename=commonconfig.json'
```

**Success notice:** The file undergoes validation, and has been successfully saved when you receive response code 200 in the Swagger UI.

**Note:** To view the discovery configuration, you can run **GET /config/{id}** via Swagger UI or cURL command. Example cURL command:

```
curl --insecure --header 'Accept: application/json' --header  
'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --header 'Authorization: Basic  
YXNtOmFzbQ=='  
'https://localhost/1.0/appdisco-observer/config/default'
```

6. Use the Swagger UI or a cURL command to add an anchor or a gateway.

#### Example anchor cURL command

```
curl -X POST --header 'Content-Type: application/json' --header 'Accept: application/  
json' --header 'Authorization: Basic YXNtOmFzbQ==' 'https://<ASM host>/1.0/appdisco-  
observer/anchor?configId=user6&IP%20Address=9.46.74.13&Port=3400'
```

#### Example gateway cURL command

```
curl -X POST --header 'Content-Type: application/json' --header 'Accept: application/  
json' --header 'Authorization: Basic YXNtOmFzbQ==' 'https://<ASM host>/1.0/appdisco-  
observer/gateway?configId=user6&IP%20Address=9.46.75.179'
```

**Note:** Anchor and gateway are global configurations. The value for the configId must be the same as defined for the configId property in the commonconfig.json file.

## Results

The discovery configurations are saved in topology management artifacts and Application Discovery. You can now run the full load job.

A **Response code 400** error can occur when the incorrect **componentType** has been provided:

**Example error message:**

```
{
  "_error": {
    "message": "Failed to configure Application Discovery access list object",
    "level": "error",
    "description": "Failed to configure Application Discovery access list componentType:
Computer Systems, vendor: none and accessListName: ComputerSystem_cvt_host1 via PUT or POST,
aborting operation... Refer to Application Discovery observer logs for more details"
  }
}
```

**Workaround**

Correct the configuration file, then re-upload it, and then rerun the **PUT /config** command.

## What to do next

Next, you define an Application Discovery Observer job.

**Related tasks**

[“Defining Application Discovery Observer jobs” on page 177](#)

You use the Application Discovery Observer to retrieve network topology data from the Application Discovery database via REST API, and then use this data to create topologies within the topology service.

[“Viewing the service logs \(on-prem\)” on page 337](#)

Logs for all Netcool Agile Service Manager services can be found in the \$ASM\_HOME/logs/<service> directories. You can set logging levels for the user-facing services, such as the observers and search, using scripts provided.

**Related reference**

[“Supported application discovery accessListTypes and properties” on page 68](#)

The following accessListTypes and their properties are supported by the Agile Service Manager application discovery service, and must be defined in the application discovery configuration file.

[“Application discovery service configuration template” on page 65](#)

The following application discovery configuration template is provided (in json format).

## Application discovery service configuration template

The following application discovery configuration template is provided (in json format).

### Application discovery configuration template

An example configuration file is included here: ASM\_HOME/data/appdisco-observer/ExampleCommonconfig.json

```
{
  "configId": "default-appdisco",
  "username": "username",
  "password": "4ihmyEBH6pETt14yIxcQ9w==",
  "disco_status_retry_delay_s": 20,
  "connect_read_timeout_ms": 50000,
  "scopes": [
    {
      "type": "Host | Subnet",
      "address": "10.160.162.87 | 2.2.2.2/255.255.255.224",
      "protocol": "IPV4"
    }
  ],
  "seeds": [
    {
      "address": "2.2.2.2/255.255.255.224",
      "timeout": 250,
      "retries": 1
    }
  ],
  "snmppasswords": [
    {

```

```

        "communitystring": "itnmlondon",
        "address": "2.2.2.2/255.255.255.224",
        "snmpversion": "V2",
        "snmpv3": null,
        "snmpport": 161,
        "timeout": 3000,
        "retries": 3
    },
    ],
    "accessListType": [
        {
            "accesslistName": "ComputerSystem",
            "componentType": "Computer System | Computer System (Windows)",
            "vendor": "none",
            "communitystring": "",
            "username": "username",
            "password": "4ihmyE",
            "privatePassword": "",
            "enablePassword": "",
            "authenticationType": "Default | Password | Public Key Infrastructure(PKI)",
            "authenticationProtocol": "",
            "encryptionAlgo": "",
            "clientId": "",
            "secretKey": "",
            "sslType": "",
            "walletAuth": {
                "cert": ""
            },
            "keystoreAuth": {
                "cert": "",
                "passPhrase": ""
            },
            "truststoreAuth": {
                "cert": "",
                "passPhrase": ""
            }
        },
        {
            "accesslistName": "NetworkElementSNMP",
            "componentType": "Network Element (SNMP)",
            "vendor": "none",
            "communitystring": "4ihmyE",
            "username": "",
            "password": "",
            "privatePassword": "",
            "enablePassword": "",
            "authenticationType": "",
            "authenticationProtocol": "",
            "encryptionAlgo": "",
            "clientId": "",
            "secretKey": "",
            "sslType": "",
            "walletAuth": {
                "cert": ""
            },
            "keystoreAuth": {
                "cert": "",
                "passPhrase": ""
            },
            "truststoreAuth": {
                "cert": "",
                "passPhrase": ""
            }
        },
        {
            "accesslistName": "NetworkElementSNMPV3",
            "componentType": "Network Element (SNMPV3)",
            "vendor": "none",
            "communitystring": "",
            "username": "username",
            "password": "rVpP==",
            "privatePassword": "",
            "enablePassword": "",
            "authenticationType": "",
            "authenticationProtocol": "MD5",
            "encryptionAlgo": "None",
            "clientId": "",
            "secretKey": "",
            "sslType": "",
            "walletAuth": {
                "cert": ""
            },
            "keystoreAuth": {
                "cert": ""
            }
        }
    ]
}

```

```

        "passPhrase":""
    },
    "truststoreAuth":{
        "cert":"",
        "passPhrase":""
    }
},
{
    "accesslistName":"Database",
    "componentType": "Database",
    "vendor": "DB2",
    "communitystring": "",
    "username":"db2inst1",
    "password":"Z/Xz==",
    "privatePassword":"",
    "enablePassword":"",
    "authenticationType":"",
    "authenticationProtocol":"",
    "encryptionAlgo":"",
    "clientId":"",
    "secretKey":"",
    "sslType":"",
    "walletAuth":{
        "cert":""
    },
    "keystoreAuth":{
        "cert":"",
        "passPhrase":""
    },
    {
        "accesslistName":"CiscoDevice",
        "componentType": "Cisco Device",
        "vendor": "none",
        "communitystring": "",
        "username":"username",
        "password":"9lqfxgje==",
        "privatePassword":"",
        "enablePassword":"zxrdft4d",
        "authenticationType":"",
        "authenticationProtocol":"",
        "encryptionAlgo":"",
        "clientId":"",
        "secretKey":"",
        "sslType":"",
        "walletAuth":{
            "cert":""
        },
        "keystoreAuth":{
            "cert":"",
            "passPhrase":""
        },
        "truststoreAuth":{
            "cert":"",
            "passPhrase":""
        }
    },
    {
        "accesslistName":"genericTaddmComponent",
        "componentType": "Cisco Works",
        "vendor": "none",
        "communitystring": "",
        "username":"username",
        "password":"9lqfxgje==",
        "privatePassword":"",
        "enablePassword":"",
        "authenticationType":"",
        "authenticationProtocol":"",
        "encryptionAlgo":"",
        "clientId":"",
        "secretKey":"",
        "sslType":"",
        "walletAuth":{
            "cert":""
        },
        "keystoreAuth":{
            "cert":"",
            "passPhrase":""
        },
        "truststoreAuth":{
            "cert":"",
            "passPhrase":""
        }
    }
}

```

```

    },
    ],
    "telnetaccess":null,
    "dns":null
}

```

## Supported application discovery accessListTypes and properties

The following accessListTypes and their properties are supported by the Agile Service Manager application discovery service, and must be defined in the application discovery configuration file.

### accessListTypes and properties

**Note:** Encrypt the following parameters using the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

- password
- privatePassword
- enablePassword
- passPhrase
- communitystring

Table 14. Supported accessListTypes and their properties

accessListType	Properties
<b>computerSystem</b> <b>Description</b> Contains properties for Computer System	<b>componentType</b> Computer System   Computer System (Windows) <i>enum</i> <b>Vendor</b> none <i>enum</i> <b>username</b> User name to access the server <i>string</i> <b>Example:</b> windows <b>password</b> Password to access the server <b>Example:</b> 1234 <b>authenticationType</b> Type of authentication for the computer system <i>string</i> <b>Example:</b> Default   Password   Public Key Infrastructure (PKI)
<b>snmp</b> <b>Description</b> Contains properties for SNMP	<b>componentType</b> Network Element (SNMP) <i>enum</i> <b>vendor</b> none <i>enum</i> <b>communitystring</b> Community string for the network device <b>Example:</b> 12345



Table 14. Supported accessListTypes and their properties (continued)

accessListType	Properties
<p>networkElement</p> <p><b>Description</b> Contains properties for SNMP NetworkElement (SNMPV3)</p>	<p><b>componentType</b> Network Element (SNMPV3) <i>enum</i></p> <p><b>vendor</b> none <i>enum</i></p> <p><b>username</b> User name to access the server <i>string</i> <b>Example:</b> SNMPV3</p> <p><b>password</b> Password to access the server <b>Example:</b> 1234</p> <p><b>privatePassword</b> Password used if data encryption is set for SNMP <b>Example:</b> 1234</p> <p><b>authenticationType</b> Type of authentication protocol used by SNMP <i>string</i> <b>Example:</b> None   MD5   SHA</p> <p><b>encryptionAlgo</b> Type of encryption algorithm used by SNMP <i>string</i> <b>Example:</b> None   DES   AES_128   AES_192   AES_256</p>
<p>dataBase</p> <p><b>Description</b> Contains properties for database</p>	<p><b>componentType</b> Database <i>enum</i></p> <p><b>vendor</b> DB2 <i>enum</i></p> <p><b>username</b> User name to access the server <i>string</i> <b>Example:</b> Db2</p> <p><b>password*</b> Password to access the server <b>Example:</b> 1234</p>

Table 14. Supported accessListTypes and their properties (continued)

accessListType	Properties
<b>ciscoDevice</b> <b>Description</b> Contains properties for Cisco Device	<b>componentType</b> Cisco Device <i>enum</i> <b>vendor</b> none <i>enum</i> <b>username</b> User name to access the server <i>string</i> <b>Example:</b> WebSphere <b>password</b> Password to access the server <b>Example:</b> 1234 <b>enablePassword</b> Enable password for the Cisco device Telnet protocol SSH1 or SSH2 is used <b>Example:</b> zxrdf4d
<b>genericTaddmComponent</b> <b>Description</b> Contains properties for general Tivoli Application Dependency Discovery Manager (TADDMM) component type	<b>componentType</b> Cisco Works <i>enum</i> <b>vendor</b> None <i>enum</i> <b>username</b> User name to access the server <i>string</i> <b>Example:</b> Sybase DB <b>password</b> Password to access the server <b>Example:</b> 1234

## Configuring Db2 for application discovery on OCP

For on-prem Agile Service Manager deployments, a Db2 database for use by the application discovery services is deployed by default. For OCP, you can either deploy an OCP instance of Db2, or connect your OCP application discovery service to an on-prem Db2 instance.

## Installing Db2 Community Edition on OCP

To deploy the Agile Service Manager application discovery service on OCP, you install Db2 Community Edition into an OpenShift cluster.

### Before you begin

You need the following prerequisites:

## Cluster

OpenShift Version 3.11, or Version 4.3

An IBM Cloud account with Cluster Administrator permissions

## Kubernetes

Version 1.11.0 (or later)

## Helm

**Power: At least** Version 2.12.x (or later), **but not** Version 3.x

**X86: At least** Version 2.14.x (or later), **but not** Version 3.x

## Storage

**One** of the following:

- NFS
- IBM Cloud File Storage (gold storage class)
- Portworx
- Red Hat OpenShift Container Storage 4.3 (or later)
- A hostPath PV that is a mounted clustered filesystem

## Resources

One worker node:

- **Cores:** 5.7 (2 for the Db2 engine and 3.7 for Db2 auxiliary services)
- **Memory:** 10.4 GiB (4 GiB for the Db2 engine and 6.4 GiB for Db2 auxiliary services)

## About this task

See for more Db2 installation and configuration information.

## SELinux considerations

For information on using Red Hat OpenShift with SELinux in enforcing mode, see the following Db2 Knowledge Center page:

[NFS storage requirements for Db2 on SELinux](#)

## Enabling Db2 HADR

To enable Db2 HADR (High Availability Disaster Recovery) in a db2u standalone deployment the following option has to be added to the helm options file:

```
hadr.enabled="true"
```

## Enabling Db2 REST

To enable Db2 REST in a db2u standalone deployment the following option to be added to the helm options file:

```
rest.enabled="true"
```

To enable the REST server to communicate outside the cluster, you create an external route and apply it.

Example route yaml file:

```
apiVersion: route.openshift.io/v1
kind: Route
metadata:
  labels:
    app: db2oltp-rest
    app.kubernetes.io/managed-by: Helm
    chart: ibm-db2
    component: db2oltp
    heritage: Helm
    release: db2oltp-rest
  name: db2oltp-rest-db2u-rest-svc
spec:
  port:
```

```

    targetPort: rest-server
  tls:
    termination: passthrough
  to:
    kind: Service
    name: db2oltp-rest-db2u-rest-svc
    weight: 100
  wildcardPolicy: None

```

You can create a yaml file using the following command:

```
oc create -f <YAML_FILE>
```

## Results

To verify the Db2 installation, see the instructions found in the `Notes.txt` file that is included in the Db2 installation package. Alternatively, view the instructions using the following command:

```
helm status my-release
```

## Example

### Installation with single storage class:

```

$ cd ./ibm_cloud_pak/pak_extensions/common
$ ./db2u-install \
  --db-type db2oltp \
  --db-name MYDB \
  --namespace db2u-project \
  --release-name db2u-release-1 \
  --storage-class managed-nfs-storage

```

### Installation with advanced storage options specified in a configuration file:

A configuration file can be passed using the `--helm-opt-file=FILENAME` argument to set the advanced helm options.

For example:

```

$ cd ./ibm_cloud_pak/pak_extensions/common
$ ./db2u-install \
  --db-type db2oltp \
  --namespace db2u-project \
  --release-name db2u-release-2 \
  --helm-opt-file ./helm_options

```

A sample configuration file is included with this chart here: `./ibm_cloud_pak/pak_extensions/common/helm_options`

```

storage.useDynamicProvisioning="true"
storage.enableVolumeClaimTemplates="true"
storage.storageLocation.dataStorage.enablePodLevelClaim="true"
storage.storageLocation.dataStorage.enabled="true"
storage.storageLocation.dataStorage.volumeType="pvc"
storage.storageLocation.dataStorage.pvc.claim.storageClassName="ocs-storagecluster-ceph-rbd"
storage.storageLocation.dataStorage.pvc.claim.size="40Gi"
storage.storageLocation.metaStorage.enabled="true"
storage.storageLocation.metaStorage.volumeType="pvc"
storage.storageLocation.metaStorage.pvc.claim.storageClassName="ocs-storagecluster-cephfs"
storage.storageLocation.metaStorage.pvc.claim.size="40Gi"

```

**Note:** The values used for the `storageClassName` parameter must be valid names from the `oc get sc` output on your system. For more information, see [Table 15 on page 80](#)

## Related reference

[“Db2 reference” on page 80](#)

This topic lists the configurable parameters and their default values for the Db2 Community Edition Version 11.5.4.0 deployment on OCP. It also contains sample code for the `create_database.sh` script.

## Related information

## Configuring Db2 Community Edition on OCP

After installing Db2 Community Edition into an OpenShift cluster, you configure it to connect to the Agile Service Manager application discovery service.

### Before you begin

Ensure you have [installed Db2 Community Edition on OCP](#).

### About this task

You first [prepare the Db2 server](#), then [configure the connection to the application discovery service on OCP](#).

After installation and configuration, you can verify that the application discovery service has been successfully configured with DB2 on OCP by [checking the service status](#).

## Procedure

### Db2 server prerequisites

1. Capture the primary storage and the database pods.

Example of a primary storage pod:

```
# kubectl get pod | grep pss
fvt-topology-nasm-app-disco-pss-deployment-6c5b55594f-gbphr 1/2 Running
5 41m
```

Example of the database pod:

```
kubectl get pod | grep db2u-0
Example system output:
asm204-db2u-0 1/1 Running 0 27h
```

2. Copy the `create_db2_database.sh` script from the primary storage pod to the OCP `/tmp` directory.

Example:

```
kubectl cp fvt-topology-nasm-app-disco-pss-deployment-6c5b55594f-gbphr:/opt/IBM/taddm/
utils/create_db2_database.sh /tmp/create_db2_database.sh
```

3. Copy the `create_db2_database.sh` script from the previous step to the `tmp` directory in the database pod.

Example:

```
kubectl cp create_database.sh asm204-db2u-0:/tmp
```

4. Enter the database pod and provide execute permissions to the script.

Example:

```
kubectl exec -it asm204-db2u-0 /bin/bash
sudo chmod a+x /tmp/create_db2_database.sh
```

5. Create a new user 'archuser'.

Example:

```
sudo useradd -g db2iadm1 -u 1010 -d /home/archuser -m archuser
sudo chage -I -1 -m 0 -M 99999 -E -1 archuser
sudo echo archuser:archuser | chpasswd
```

6. As the **db2inst1** user, create a database instance.

Run the following command from the tmp directory. In the following example, the database is called 'taddm':

```
./create_db2_database.sh taddm
```

**Tip:** If the system output consists of the following message, a database called 'taddm' already exists:

```
DB2 version 1.5.4
INFO: Database with same name taddm already exists.
```

If this occurs, run the 'drop database' command, then recreate the 'taddm' database, as in the following example:

```
db2 drop database taddm
```

System output if successful:

```
DB20000I The DROP DATABASE command completed successfully.
```

Rerun the following command:

```
./create_db2_database.sh taddm
```

#### 7. Confirm that the database has been created:

```
db2 list database directory
```

Example system output if the database exists:

```
System Database Directory
Number of entries in the directory = 1
Database 1 entry:
Database alias           = TADDM
Database name            = TADDM
Local database directory = /mnt/blumeta0/db2/databases
Database release level   = 15.00
Comment                  =
Directory entry type      = Indirect
Catalog database partition number = 0
Alternate server hostname =
Alternate server port number =
```

Test the database connection:

```
db2 connect to TADDM user db2inst1 using db2inst1
```

Example system output:

```
Database Connection Information
Database server           = DB2/LINUX8664 11.5.4.0
SQL authorization ID      = DB2INST1
Local database alias      = TADDM
```

#### 8. Exit pod.

### OCF application discovery server configuration

#### 9. On the OCP application discovery server, create a configuration file called db2-secret.yaml, as in the following example:

```
apiVersion: v1
kind: Secret
metadata:
  name: db2-secret
  namespace: production
data:
  DBARCHIVEUSER: YXJjaHVzZXI=
  DBUSER: ZGIyaW5zdDE=
```

```
DB_ARCH_PASSWORD: YXJjaHVzZXI=  
DB_PASSWORD: ZGIyaW5zdDE=  
type: Opaque
```

**Note:** The values of the following parameters have to be base64 encoded:

**DBARCHIVEUSER**

Default value: archuser

**DBUSER**

Default value: db2inst1

**DB\_ARCH\_PASSWORD**

Default value: archuser

**DB\_PASSWORD**

Default value: db2inst1

10. Create a secret that stores DB details by running the following command:

```
kubectl apply -f db2-secret.yaml
```

11. Populate the database details in the Netcool Operations Insight on Red Hat OpenShift Operator properties.

The following sample shows the database parameters to connect the application discovery service with the pre-configured database:

```
appDisco:  
  enabled: true  
  dburl : 'db2hostname.com'  
  dbsecret: 'db2-secret'  
  secure: false  
  certSecret: ''
```

**dburl**

The URL of the server on which DB2 is installed.

For DB2 on OCP, the format of dburl is:

```
<Service Name>.<Namespace Name>.svc.cluster.local
```

Example:

```
asm204-db2u.production.svc.cluster.local
```

Run '**oc get svc**' to obtain the service name, for example:

```
oc get svc | grep asm204-db2u  
asm204-db2u  
ClusterIP 172.30.200.131 <none>  
50000/TCP,50001/TCP,25000/TCP,25001/TCP,25002/TCP,25003/TCP,25004/TCP,25005/TCP 27h
```

**dbsecret**

The name of the secret in the same namespace as the instance to be installed.

It must contain the following keys:

**DBUSER**

The name of the DB2 user.

**DB\_PASSWORD**

The password of the DB2 user.

**DBARCHIVEUSER**

The name of the DB2 archive user.

**DB\_ARCH\_PASSWORD**

The password of the DB2 archive user.

**secure**

'False' if the connection to DB2 is insecure.

## certSecret

Applicable only if 'secure' property is 'true'.

## What to do next

After the installation and configuration of the application discovery service and Db2, check the service status. Run the following command on the application discovery OCP server:

```
# kubectl exec -it <releasename>-topology-nasm-app-disco-pss-deployment-<pod-id> /opt/IBM/taddm/dist/bin/control status
```

The following example shows that the application discovery service has been successfully configured with Db2 on OCP:

```
kubectl exec -it appd204-topology-nasm-app-disco-pss-deployment-84dbdb89c9-sjk9m /opt/IBM/taddm/dist/bin/control status
```

```
Defaulting container name to pss.  
Use 'kubectl describe pod/appd204-topology-nasm-app-disco-pss-deployment-84dbdb89c9-sjk9m -n production' to see all of the containers in this pod.  
DbInit: Started  
Liberty: Started  
StorageService: Started  
  
TADDM: Running
```

## Configuring on-prem Db2 with OCP application discovery (TCPIP)

You can connect the application discovery service of an OCP deployment of Agile Service Manager with an on-prem Db2 installation. This topic describes how to do so using TCPIP. You must complete the TCPIP setup before completing the SSL setup.

## Before you begin

Before setting up the application discovery service for Agile Service Manager on OCP to run with an on-prem version of Db2, you must download and install DB2 Version 11.5 Enterprise Server Edition. See the Db2 Knowledge Center for more installation and configuration information.

## About this task

**Note:** TCPIP (or insecure) mode is enabled by default when DB2 is installed.

## Procedure

### DB2 server prerequisites

1. On the Db2 server, and as the **root** user, configure 'archuser' by running the following command:

```
useradd -g db2iadm1 -u 1010 -d /home/archuser -m archuser  
chage -I -1 -m 0 -M 99999 -E -1 archuser  
echo archuser:archuser | chpasswd
```

2. Obtain the sample code from [here](#) and paste it into the `create_database.sh` script.
3. As the **db2inst1** user, create a database instance.

Run the following command:

```
./create_database.sh <name_of_instance>
```

In the following example, the database is called 'taddm':

```
./create_database.sh taddm
```

## OCP application discovery server configuration



4. On the OCP application discovery server, create a configuration file called `db2-secret.yaml`, as in the following example:

```
apiVersion: v1
kind: Secret
metadata:
  name: db2-secret
  namespace: production
data:
  DBARCHIVEUSER: YXJjaHVzZXI=
  DBUSER: ZGIyaW5zdDE=
  DB_ARCH_PASSWORD: YXJjaHVzZXI=
  DB_PASSWORD: ZGIyaW5zdDE=
type: Opaque
```

**Note:** The values of the following parameters have to be base64 encoded:

**DBARCHIVEUSER**

Default value: `archuser`

**DBUSER**

Default value: `db2inst1`

**DB\_ARCH\_PASSWORD**

Default value: `archuser`

**DB\_PASSWORD**

Default value: `db2inst1`

5. Create a secret that stores DB details by running the following command:

```
kubectl apply -f db2-secret.yaml
```

6. Populate the database details in the Netcool Operations Insight on Red Hat OpenShift Operator properties.

The following sample shows the database parameters to connect the application discovery service with the pre-configured database:

```
appDisco:
  enabled: true
  dburl : 'db2hostname.com'
  dbsecret: 'db2-secret'
  secure: false
  certSecret: ''
```

**dburl**

The URL of the server on which DB2 is installed.

**dbsecret**

The name of the secret in the same namespace as the instance to be installed.

It must contain the following keys:

**DBUSER**

The name of the DB2 user.

**DB\_PASSWORD**

The password of the DB2 user.

**DBARCHIVEUSER**

The name of the DB2 archive user.

**DB\_ARCH\_PASSWORD**

The password of the DB2 archive user.

**secure**

'False' if the connection to DB2 is insecure.

**certSecret**

Applicable only if 'secure' property is 'true'.

## Related information

### Configuring on-prem Db2 with OCP application discovery (SSL)

You can connect the application discovery service of an OCP deployment of Agile Service Manager with an on-prem Db2 installation. This topic describes how to do so using SSL.

### Before you begin

Before setting up the application discovery service for Agile Service Manager on OCP to run with an on-prem version of Db2, you must download and install DB2 Version 11.5 Enterprise Server Edition. See the Db2 Knowledge Center for more installation and configuration information.

**Important:** Before configuring a secure (SSL) connection between the OCP application discovery service and the on-prem Db2 server, you must complete steps 1 - 5 of the [TCPIP setup](#).

### About this task

**Note:** TCPIP (or insecure) mode is enabled by default when DB2 is installed.

### Procedure

#### DB2 server prerequisites

1. On the Db2 server, and as the **root** user, create a key database called `mydbserver.kdb` and a stash file called `mydbserver.sth` using the GSKCapiCmd tool.

**Tip:** The GSKCapiCmd tool is packaged with DB2 and can be found in the `/opt/ibm/db2/V11.5/gskit/bin` directory.

Run the following commands:

```
export LD_LIBRARY_PATH=/opt/ibm/db2/V11.5/lib64/gskit_db2
cd /opt/ibm/db2/V11.5/gskit/bin
./gsk8capicmd_64 -keydb -create -db "mydbserver.kdb" -pw "myServerPassw0rdpw0" -stash
```

2. As the **root** user, add (create) a certificate for your server to the key database.

The server sends this certificate to clients during the SSL handshake to provide authentication for the server.

Example command to create a self-signed certificate:

```
./gsk8capicmd_64 -cert -create -db "mydbserver.kdb" -pw "myServerPassw0rdpw0" -label "myselfsigned" -dn "CN=myhost.mycompany.com,O=myOrganization,OU=myOrganizationUnit,L=myLocation,ST=ON,C=CA"
```

3. As the **root** user, extract the certificate created to a file.

This certificate can be distributed to computers running clients. In this case, the application discovery service can use this certificate to establish SSL connections to the DB2 server. The following command extracts the certificate to the `mydbserver.arm` file.

```
./gsk8capicmd_64 -cert -extract -db "mydbserver.kdb" -pw "myServerPassw0rdpw0" -label "myselfsigned" -target "mydbserver.arm" -format ascii -fips
```

4. Set up the Db2 server for SSL support:

```
chmod 744 mydbserver.kdb mydbserver.sth
su db2inst1
db2 update dbm cfg using SSL_SVR_KEYDB /opt/ibm/db2/V11.5/gskit/bin/mydbserver.kdb
db2 update dbm cfg using SSL_SVR_STASH /opt/ibm/db2/V11.5/gskit/bin/mydbserver.sth
db2 update dbm cfg using SSL_SVR_LABEL myselfsigned
```

```
db2 update dbm cfg using SSL_SVCENAME 60000
db2set -i db2inst1 DB2COMM=SSL,TCPIP
db2stop
db2start
```

5. As the **db2inst1** user, verify that both SSL and TCPIP modes are enabled.

a) Run the following command:

```
[db2inst1@taddmvm162095 root]$ db2 get dbm cfg | grep SVCE
```

System output:

```
TCP/IP Service name      (SVCENAME) = 50000
SSL service name         (SSL_SVCENAME) = 60000
```

The system output shows that the SSL and TCPIP ports are enabled for connection with the DB2 server.

b) Run the following command:

```
[db2inst1@taddmvm162095 root]$ db2set -i
```

System output:

```
DB2_EVALUNCOMMITTED=YES
DB2_SELECTIVITY=DSCC
DB2COMM=SSL,TCPIP
DB2AUTOSTART=YES
```

The system output for DB2COMM shows that SSL and TCPIP modes are enabled.

### OCF application discovery server configuration

6. On the OCF application discovery server, perform the following steps:

a) Copy the previously created `mydbserver.arm` certificate to the application discovery client server.

```
ssh root@MASTER-1-IP
```

b) Rename the `mydbserver.arm` certificate 'tls.crt'.

```
cp mydbserver.arm tls.crt
```

c) Create a secret in the same namespace in which the application discovery service is installed using the `tls.crt` certificate.

```
oc create secret generic "db-cert" --from-file=/root/tls.crt -n app-disco-namespace
```

**Remember:** The following step requires that the first five steps of the TCPIP setup have been completed.

7. Populate the database details in the Netcool Operations Insight on Red Hat OpenShift Operator properties.

The following sample shows the database parameters to connect the application discovery service with the pre-configured database in secure (SSL) mode:

```
appDisco:
  enabled: true
  dburl : 'db2hostname.com'
  dbsecret: 'db2-secret'
  secure: true
  certSecret: 'db-cert'
```

#### dburl

The URL of the server on which DB2 is installed.

#### dbsecret

The name of the secret in the same namespace as the instance to be installed.

It must contain the following keys:

**DBUSER**

The name of the DB2 user.

**DB\_PASSWORD**

The password of the DB2 user.

**DBARCHIVEUSER**

The name of the DB2 archive user.

**DB\_ARCH\_PASSWORD**

The password of the DB2 archive user.

**secure**

'True' if the connection to DB2 is secure.

**certSecret**

Applicable only if 'secure' property is 'true'.

This is the secret name created earlier.

This secret **must** contain the DB2 certificate 'tls.crt'.

**Related information****Db2 reference**

This topic lists the configurable parameters and their default values for the Db2 Community Edition Version 11.5.4.0 deployment on OCP. It also contains sample code for the `create_database.sh` script.

**Configurable parameters and default values for Db2 Community Edition (on OCP)**

The following table lists the configurable parameters and default values for Db2 Community Edition Version 11.5.4.0.

Table 15. Db 2 configurable parameters and default values		
Parameter	Description	
dedicated	Enforce that Db2 is scheduled on worker nodes that have been labeled and tainted with <code>icp4data=\${value}</code> . Use with option <code>global.nodeLabel.value</code> below.	true
subType	Database subtype can be either [smp, mpp] for db2wh. For db2oltp, the subtype is smp.	smp
arch	The helm chart will try to detect the architecture based on the master node. Choose an explicit architecture here to overwrite it.	x86_64
mln.total	The total number of MLNs to be used and evenly spread out to the number of worker nodes.	1
global.dbType	Database type can be either [db2wh, db2oltp]	db2oltp

Table 15. Db 2 configurable parameters and default values (continued)

Parameter	Description	
global.nodeLabel.value	Value for the node label in order to deploy Db2 on a dedicated node. The node must be labeled and tainted with icp4data=\${value}	
images.universalTag	The tag specified here will be used by all images except those that explicitly specify a tag.	latest
images.pullPolicy	Always, Never, or IfNotPresent. Defaults to Always	IfNotPresent
images.auth.image.tag	The tag for the version of the LDAP authentication container	11.5.4.0-56
images.auth.image.repository	The container is deployed to serve as an LDAP server for database authentication if ldap.enabled is true and no ldap server is not specified	icr.io/obs/hdm/db2u/db2u.auxiliary.auth
images.rest.image.tag	The container tag that is used to pull the version of the rest container.	11.5.4.0-56
images.rest.image.repository	The REST container image, which is used to perform host the REST sever	icr.io/obs/hdm/db2u/db2u.rest
images.etcd.image.tag	The etcd container tag that used to pull the version of the needed container.	3.3.10-56
images.etcd.image.repository	The etcd container is used in MPP configurations to automate the failover of MPP nodes	icr.io/obs/hdm/db2u/etcd
images.instdb.image.tag	The container tag that is used to pull the version of the Database and instance payload container	11.5.4.0-56
images.instdb.image.repository	This container carries the payload required to restore a database and instance into a newly deployed release.	icr.io/obs/hdm/db2u/db2u.instdb
images.db2u.replicas	The number of Db2 Warehouse pods that will serve the database. Replica count of 1 signifies SMP and 2 and more is an MPP configuration.	1
images.db2u.image.tag	The container tag that is used to pull the version of Db2 Warehouse main engine container.	11.5.4.0-56

Table 15. Db 2 configurable parameters and default values (continued)

Parameter	Description	
images.db2u.image.repository	The main database engine for the Db2 Warehouse release	icr.io/obs/hdm/db2u/db2u
images.tools.image.tag	The container tag that is used to pull the version of the tools container.	11.5.4.0-56
images.tools.image.repository	The tools container image, which is used to perform outside of the engine operations	icr.io/obs/hdm/db2u/db2u.tools
database.name	The name of the database. Defaults to BLUDB	BLUDB
database.pageSize	The default database page size. Defaults to 32768.	32768
database.bludadminPwd	Password for the LDAP database administrator which is the main LDAP user	
database.codeset	The default database codeset. Defaults to UTF-8.	UTF-8
database.collation	The default database collation sequence. Defaults to IDENTITY.	IDENTITY
database.territory	The default database territory. Defaults to US.	US
storage.storageClassName	Choose a specific storage class name to use during deployment. A storage class offers the foundation for dynamic provisioning.	
storage.existingClaimName	Name of an existing Persistent Volume Claim that references a Persistent Volume	
storage.useDynamicProvisioning	If dynamic provisioning is available in the cluster this option will automatically provision the requested volume if set to true.	False
instance.db2Support4K	Db2 supports storage devices that use a 4KB sector size in production environments. Default to false.	False
limit.cpu	CPU cores limit to apply to Db2. Db2 won't be able to exceed the provided value.	2
limit.memory	Memory limit to apply to Db2. Db2 won't be able to exceed the provided value.	4.3Gi

## Sample code for the create\_database.sh script

When setting up a Db2 database for the application discovery service on OCP, you have to copy and paste the following code into the create\_database.sh script.

```
#!/bin/sh

# /*****
#*
#* Copyright © Aricent Holdings Luxembourg S.a.r.l. 2020. All rights reserved.
#*
#*****/
#//
#//////// End Standard Header //////////*/

##### Begin Standard Header - Do not add comments here #####
#
# File:      %W%
# Version:   %I%
# Modified:  %G% %U%
# Build:     %R% %L%
#
# Licensed Materials - Property of IBM
#
# Restricted Materials of IBM
#
# 5724-N55
#
# (C) COPYRIGHT IBM CORP. 2007, 2011. All Rights Reserved.
#
# US Government Users Restricted Rights - Use, duplication or
# disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
#
##### End Standard Header #####

# set -x

if test -z "$1"
then
    echo "Usage: $0 <database_name>"
    exit 1
fi

tune_db_noauto() {
echo "SUCCESS: Create database $1 completed."
db2 <<EOF
connect to $1
alter bufferpool IBMDEFAULTBP SIZE 240000
create bufferpool BUF8K immediate size 384 pagesize 8k
create bufferpool BUF32K immediate size 384 pagesize 32k
connect reset

connect to $1
create tablespace userspace2 pagesize 32k managed by system using ('usercont2') bufferpool
BUF32K
create large tablespace largespace2 pagesize 32k managed by database using (file 'largedata'
2560) bufferpool BUF32K
create system temporary tablespace largetemp2 pagesize 32k managed by system using ('systemp2')
bufferpool BUF32K

create tablespace USERSPACE3 pagesize 8k managed by system using ('usercont3') bufferpool BUF8K

alter bufferpool BUF8K SIZE 40000
alter bufferpool BUF32K SIZE 8000

update db cfg using LOGFILSIZ 8192
update db cfg using LOGBUFSZ 1024
update db cfg using LOGPRIMARY 12
update db cfg using LOGSECOND 40
update db cfg using CATALOGCACHE_SZ 4000
update db cfg using PCKCACHESZ 4096
update db cfg using AVG_APPLS 3
update db cfg using NUM_IOCLEANERS automatic
update db cfg using NUM_IOSERVERS 10
update db cfg using DFT_DEGREE ANY
update db cfg using APP_CTL_HEAP_SZ 4000
update db cfg using MAXLOCKS 35 LOCKLIST 20000
update db cfg using SHEAPTHRES_SHR MANUAL SORTHEAP 8000
update db cfg using AUTO_MAINT OFF
```

```

update db cfg using UTIL_HEAP_SZ 8000
update db cfg using DBHEAP 8000
update db cfg using APPLHEAPSZ 2000
disconnect $1
quit
EOF
db2set DB2_EVALUNCOMMITTED=YES
db2set DB2_SELECTIVITY=DSCC
}

tune_db_auto() {
echo "SUCCESS: Create database $1 completed."
if test "$version" -ne "9"
then
    USE_AUTOMATIC=YES
    PARAM_VAL=AUTOMATIC
else
    USE_AUTOMATIC=
fi

db2 connect to $1

if [ -z "$USE_AUTOMATIC" ]; then PARAM_VAL=240000; fi
db2 alter bufferpool IBMDEFAULTBP SIZE $PARAM_VAL

db2 <<EOF
create bufferpool BUF8K immediate size 384 pagesize 8k
create bufferpool BUF32K immediate size 384 pagesize 32k
connect reset

connect to $1
create tablespace userspace2 pagesize 32k managed by automatic storage bufferpool BUF32K
create large tablespace largespace2 pagesize 32k managed by automatic storage bufferpool BUF32K
create system temporary tablespace targetemp2 pagesize 32k managed by automatic storage
bufferpool BUF32K
create tablespace USERSPACE3 pagesize 8k managed by automatic storage bufferpool BUF8K

update db cfg using LOGFILSIZ 8192
update db cfg using LOGBUFSZ 1024
update db cfg using LOGPRIMARY 12
update db cfg using LOGSECOND 40
update db cfg using CATALOGCACHE_SZ 4000
update db cfg using PCKCACHESZ 4096
update db cfg using AVG_APPLS 3
update db cfg using NUM_IOCLEANERS automatic
update db cfg using NUM_IOSERVERS 10
update db cfg using DFT_DEGREE ANY
update db cfg using APP_CTL_HEAP_SZ 4000
update db cfg using MAXLOCKS 35 LOCKLIST 20000
update db cfg using AUTO_MAINT OFF

EOF

if [ -z "$USE_AUTOMATIC" ]; then PARAM_VAL=40000; fi
db2 alter bufferpool BUF8K SIZE $PARAM_VAL

if [ -z "$USE_AUTOMATIC" ]; then PARAM_VAL=8000; fi
db2 alter bufferpool BUF32K SIZE $PARAM_VAL

if [ -z "$USE_AUTOMATIC" ]
then
    db2 update db cfg using SHEAPTHRES_SHR MANUAL SORTHEAP 8000
else
    db2 update db cfg using SHEAPTHRES_SHR AUTOMATIC SORTHEAP AUTOMATIC
fi

if [ -z "$USE_AUTOMATIC" ]; then PARAM_VAL=8000; fi
db2 update db cfg using UTIL_HEAP_SZ $PARAM_VAL

if [ -z "$USE_AUTOMATIC" ]; then PARAM_VAL=8000; fi
db2 update db cfg using DBHEAP $PARAM_VAL

if [ -z "$USE_AUTOMATIC" ]; then PARAM_VAL=2000; fi
db2 update db cfg using APPLHEAPSZ $PARAM_VAL

db2 disconnect $1
db2 quit

echo ""
echo "-----$1 Instance Created Successfully-----"
echo ""

```



```

db2set DB2_EVALUNCOMMITTED=YES
db2set -immediate DB2_SELECTIVITY=DSCC 2>&1 >/dev/null || db2set DB2_SELECTIVITY=DSCC
}

db2cfexp /tmp/db2cfg.$$ maintain
output=`cat /tmp/db2cfg.$$ | grep -i "^[DB>.*:$1\]"`
rm -f /tmp/db2cfg.$$

# Get db2 version from db2cfexp command. Donot use db2level because
# output could be translated
db2cfexp /tmp/db2cfg.$$ maintain
version_str=`cat /tmp/db2cfg.$$ | grep "APPLICATION=" | awk '{print $2}'`
echo "DB2 version $version_str"

# Check if db already exists
dbName=`db2 list db directory | grep "Database name" | awk '{print $4}' | grep -w -i $1`
dbName=`echo $dbName | tr '[:upper:]' '[:lower:]'`
userDbName=`echo $1 | tr '[:upper:]' '[:lower:]'`
if [ "$userDbName" = "$dbName" ]
then
    echo ""
    echo "INFO: Database with same name $1 already exists."
    exit 3
fi

release=`echo $version_str | cut -c 3`
version=`echo $version_str | cut -c 1`
rm -f /tmp/db2cfg.$$

if test "$version" -eq 9 -a "$release" -eq 8 -o "$version" -eq 0 -a "$release" -eq 1 -o
"$version" -eq 0 -a "$release" -eq 5 -o "$version" -eq 1 -a "$release" -eq 1 -o "$version" -eq
1 -a "$release" -eq 5
then
    db2 create db $1 using codeset UTF8 territory US
    if test "$?" -eq 0
    then
        tune_db_auto $1
    else
        echo "ERROR: Create database $1 failed."
        exit 2
    fi
else
    if test "$version" -eq 9
    then
        db2 create db $1 automatic storage no using codeset UTF8 territory US
        if test "$?" -eq 0
        then
            tune_db_noauto $1
        else
            echo "ERROR: Create database $1 failed."
            exit 2
        fi
    else
        echo "Unsupported DB2 version $version_str"
        exit 1
    fi
fi
fi

```

## Configuring the network discovery services

The Agile Service Manager network discovery service is installed together with the other core components. Once you have configured it, Agile Service Manager can discover network data.

### Before you begin

Ensure you have the required network data available before you configure the Agile Service Manager network discovery service. When creating a discovery profile, you may need information about the following:

#### Scope

You can specify subnet ranges or specific devices for the discovery as a 'zone'. A zone can be as granular as an individual router.

You can specify zones within zones to include or exclude them from the discovery.

## Seeds

Seeds specify the locations from which to begin discovering devices. Discovery seeds can be IP addresses or subnet addresses.

## Collector discovery

You can further configure a discovery with the included Alcatel5620Sam Java collector, which collects topology data from Element Management Systems (EMS) and integrates this data into the discovered topology.

Collector discovery is activated using the **collector** tag (by default, set to 'null').

**Note:** When network discovery is enabled, so is the collector service, which will fail to start unless the SFTP secret is created, as described in the “Configuring EMS and SSH File Transfer Protocol (SFTP) Network Connectivity for Collector Discovery on OCP” on page 92 topic.

A Java Virtual Machine must be available on the target machine, and your EMS devices must be within the configured discovery scope.

For configuration information pertaining to this specific collector, see the IBM Network Manager IP Edition documentation here: [https://www.ibm.com/support/knowledgecenter/en/SSSHRK\\_4.2.0/disco/task/dsc\\_confalcatel5620samjava.html](https://www.ibm.com/support/knowledgecenter/en/SSSHRK_4.2.0/disco/task/dsc_confalcatel5620samjava.html)

**Remember:** You configure the deployed network service **after** installing the core Agile Service Manager containers, and **after** starting them.

**OCP requirement:** On the OCP hosts, network discovery requires that `pids_limit` be set at least to 44406 inside the `crio.conf` file.

For information about changing the values in the `crio.conf` file using Machine Configs, see <https://access.redhat.com/solutions/4717701>

For reference information about Machine Configs, see <https://docs.openshift.com/container-platform/4.2/architecture/architecture-rhcos.html>

## About this task

You first upload the configuration details in a `commonconfig.json` file via **POST /file** command. You then run the **POST /config** command for the Network Discovery observer, which first validates the files, and then saves them.

Network discovery service discovers configuration information about the entire network infrastructure, identifying deployed software components, physical servers, network devices, virtual systems, and host data. It then saves the information in the provided database. The Network Discovery Observer then retrieves network topology data from the database via REST API, which it uses to create topologies within the Agile Service Manager topology service.

## Procedure

1. Edit the example configuration file included with the network discovery installation, or create a new configuration file.

### For on-prem

Edit the included configuration file:

```
ASM_HOME/data/netdisco-observer/ExampleCommonconfig.json
```

Save the file as `commonconfig.json`

Alternatively, in the `ASM_HOME/bin` directory, generate the configuration file using the `generate_common_config.sh` script:

```
./generate_common_config.sh -d
```

Expected output includes:

```
...Successfully generated commonconfig.json in /opt/ibm/netcool/asm/bin/../data/netdisco-observer folder...
```

## For OCP

Copy the text from the following example file: [“Network discovery service configuration template”](#) on page 94

Save the file as `commonconfig.json`

2. Populate the following required properties in the configurations file for Network Discovery observer. See the related links for the full text of the included example configuration file.

### **configId**

Discovery profile ID

### **username**

Network Discovery username

The default is `admin`

To use a new default user, you must first create the user using **POST /user**. The following default users, passwords and roles are available:

Table 16. Default network discovery users and roles		
User	Password	Role
admin	admin@123	ADMIN
superuser	superuser@123	ADMIN
operator	operator@123	OPERATOR

### **password**

Network Discovery encrypted password

**For on-prem:** Run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory to encrypt the parameter:

```
./bin/encrypt_password.sh
```

Enter and then confirm the parameters to be encrypted, in turn. The encryption utility will return encrypted versions.

**For OCP:** Run the following command to encrypt the parameters:

```
kubectl exec -ti asm-topology-pods -- java -jar /opt/ibm/topology-service/topology-service.jar encrypt_password --password '<password to encrypt>'
```

### **snmppasswords**

SNMP community strings, which must be in encrypted form

**For on-prem:** Run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory to encrypt the parameter:

```
./bin/encrypt_password.sh
```

Enter and then confirm the parameters to be encrypted, in turn. The encryption utility will return encrypted versions.

**For OCP:** Run the following command to encrypt the parameters:

```
kubectl exec -ti asm-topology-pods -- java -jar /opt/ibm/topology-service/topology-service.jar encrypt_password --password '<password to encrypt>'
```

### **scope**

Network Discovery scope for IP-based devices and subnets

Uses host and subnet

You must supply the netmask for both host and subnet

**seeds**

Provides the starting points from which to look for devices

**telnetaccess**

By default set to null. You must supply the password in encrypted form.

**dns**

By default set to null.

**collector**

By default set to null.

Collector configuration is used only during a collector discovery.

You must supply the password in encrypted form.

**Note:** When network discovery is enabled, so is the collector service, which will fail to start unless the SFTP secret is created, as described in the [“Configuring EMS and SSH File Transfer Protocol \(SFTP\) Network Connectivity for Collector Discovery on OCP”](#) on page 92 topic.

### 3. Optional: Configure automatic restarting of discovery.

The network discovery service polls active discovery processes at intervals to check that they are still running. If a running discovery fails (for example, because the worker on which it is running stops), the discovery is run again from the beginning when the worker is restarted and any other active discoveries running on that worker have finished.

Use the ISO 8601 standard notation when defining durations.

- To configure the polling and discovery restart intervals:

**On-prem**

Edit the `$ASM_HOME/etc/nasm-net-disco-status.yml` file.

Change the values for the **SHERPA\_STATUS\_SERVICE\_WORKER\_CHECK\_DURATION** property.

**OCP**

Edit the `${releaseName}-nasm-net-disco-status` pod for Kubernetes and OCP.

Change the values for the **SHERPA\_STATUS\_SERVICE\_WORKER\_CHECK\_DURATION** property.

- To configure the intervals at which discovery attempts to restart:

**On-prem**

Edit the `$ASM_HOME/etc/nasm-net-disco-status.yml` file.

Change the value for the **SHERPA\_RETRY\_INTERVALS** property.

**OCP**

Edit the `${releaseName}-nasm-net-disco-status` pod for Kubernetes and OCP.

Change the value for the **SHERPA\_RETRY\_INTERVALS** property.

### 4. Upload the configurations file (that is, 'commonconfig.json') via the Network Discovery Observer's Swagger UI **POST /file** command, or cURL command.

Example cURL command:

```
curl --insecure --header 'Content-Type: multipart/form-data' --header
'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --
header
'Authorization: Basic YXNtOmFzbQ==' --request POST --url 'https://localhost/1.0/netdisco-
observer/files' --form
'job_file=@/tmp/commonconfig.json'
```

The configurations file will be saved in the `ASM_HOME/data/netdisco-observer/{tenant-id}` directory.

### 5. Run **POST /config** via Swagger UI or cURL command to create the discovery details from the configurations file to the topology management artifact and Network Discovery.

Example cURL command:

```
curl -X POST --header 'Content-Type: application/json' --header
```

```
'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --header 'Authorization: Basic YXNtOmFzbQ==' 'https://localhost/1.0/netdisco-observer/config?filename=commonconfig.json'
```

**Tip:** When you need to update or replace the discovery details, use **PUT /config** via the Swagger UI, or use the cURL command, as in the following example:

```
curl --insecure --header 'Content-Type: application/json' --header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --header 'Authorization: Basic YXNtOmFzbQ==' --request PUT --url 'https://localhost/1.0/netdisco-observer/config?filename=commonconfig.json'
```

**Success notice:** The file undergoes validation, and has been successfully saved when you receive response code 200 in the Swagger UI.

**Note:** To view the discovery configuration, you can run **GET /config/{configId}** via Swagger UI or cURL command. Example cURL command:

```
curl --insecure --header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --header 'Authorization: Basic YXNtOmFzbQ==' 'https://localhost/1.0/netdisco-observer/config/default'
```

## Results

The discovery configurations are saved and you can now run the full load job.

A **Response code 400** error can occur when the incorrect scope has been provided:

### Example error message:

```
{
  "_error": {
    "message": "Invalid netmask",
    "level": "error",
    "description": "Netmask cannot be null or empty"
  }
}
```

## Workaround

Correct the configuration file, then re-upload it, and then rerun the **PUT /config** command.

## What to do next

Next, you define a Network Discovery Observer job.

### Related tasks

[“Defining Network Discovery Observer jobs” on page 217](#)

You use the Network Discovery Observer to retrieve network topology data from the Network Discovery database via REST API and use this data to create topologies within the topology service.

[“Viewing the service logs \(on-prem\)” on page 337](#)

Logs for all Netcool Agile Service Manager services can be found in the \$ASM\_HOME/logs/<service> directories. You can set logging levels for the user-facing services, such as the observers and search, using scripts provided.

### Related reference

[“Network discovery service configuration template” on page 94](#)

The following network discovery configuration template is provided (in json format).

### Related information

## Discoverable technologies and devices

Network discovery is capable of discovering different network technologies and devices.

Network discovery can discover the following:

### OSI layers

- Layer 2 connectivity
- Layer 3 connectivity

### Element Management Systems (EMS)

- Alcatel 5620 SAM

### Multi-Protocol Label Switching (MPLS)

- MPLS Virtual Private Networks (VPNs)
- VPN Route Forwarding (VRF)
- Virtual Private LAN Service (VPLS)
- MPLS Virtual Route Forwarding (VRF)
- MPLS Label-Switched Paths (LSP)
- MPLS pseudowires

### Standards-based support

Certain devices are not supported explicitly but can be discovered through support for Internet Engineering Task Force (IETF) standards:

- Interface and IfStack table
- ENTITY mib
- VRRP
- LLDP
- IP Inet routing
- dot1d dot1q switching
- SONMP

### Devices

- 3Com stacked switch
- Alteon routers and switches
  - Virtual Router Redundancy Protocol (VRRP)
  - Spanning Tree Protocol
- Avaya (formerly Nortel) devices
  - Nortel Passport
  - SynOptics Network Management Protocol (SONMP)
  - Containment
- Brocade devices
  - Containment
  - Foundry Discovery Protocol (FDP)
- Centillion switches

- VLAN
- Cisco devices
  - Cisco Discovery Protocol (CDP)
  - Cisco Carrier Routing System Link Aggregation Group (CRS LAG)
  - IEEE 802.1 QinQ configuration
  - IP SLA
  - Spanning Tree Protocol
  - Virtual Device Context (VDC) from Cisco Nexus devices
  - Virtual Port Channel (vPC)
  - VLAN Trunking Protocol (VTP)
- Check Point VSX devices
  - Virtual firewalls
  - Context data
- CoreBuilder 9000 switches
- Dasan switches
  - Layer 2 connectivity
  - VLANs
  - Trunk ports
- Enterasys switches
  - Cabletron Discovery Protocol
  - FDB
  - VLAN
- Extreme switches
  - Layer 2 connectivity
  - EDP
  - VLAN
- F5 switches
- Foundry switches
- HP ProCurve switches
- Huawei switches
  - Network Quality Analyser (NQA)
  - Quidway
- IBM System Networking switches
  - Layer 2 connectivity
  - VLAN containment
  - VLAN tags
  - VLAN Trunk
  - Trunk Group
- Juniper
  - Juniper Box Anatomy
  - Virtual router and VRF context-sensitive stacking information
  - LAG

- VLAN tagging
- IEEE 802.1 QinQ
- IP SLA
- ARP
- Marconi ES-3810 switches
- Nokia routers
- VRRP

## Other protocols

- Hot Standby Routing Protocol (HSRP)
- Link Layer Discovery Protocol (LLDP)
- Spanning Tree Protocol (STP)

## Configuring EMS and SSH File Transfer Protocol (SFTP) Network Connectivity for Collector Discovery on OCP

You set up SSH File Transfer Protocol for use by the network discovery service.

### Before you begin

SFTP is available on most Linux system out-of-the-box and can be readily configured.

#### Prerequisite for collectors:

#### Configuring SFTP in an OCP installation

As part of the collectors feature of the network discovery service, a Secure File Transfer Protocol (SFTP) server is required for using the Alcatel5620Sam Collector. If you do not want to use this collector, you do not need an SFTP server.

On OCP, no SFTP server is supplied. You must install your own SFTP server, and then configure network discovery to use it. The SFTP server can be inside or outside the OCP cluster.

To configure network discovery collectors to use your SFTP server, issue the following command:

```
oc create secret generic ${releaseName}-sftp-secret --from-literal=host=${IP} --from-literal=port=${PORT} --from-literal=username=${USERNAME} --from-literal=password=${PASSWORD} --from-literal=directory=${DIRECTORY}
```

Where:

#### Secret Name

The name of the secret is the release name of the Agile Service Manager installation, followed by `sftp-secret`.

#### IP

IP is the SFTP server IP address or hostname, which should be reachable from the OCP cluster.

#### PORT

PORT is the SFTP server port number.

#### USERNAME

USERNAME is the SFTP server user name to use to connect to the SFTP server from the collector.

#### PASSWORD

PASSWORD is the SFTP server password to use to connect to the SFTP server from the collector.

#### DIRECTORY

DIRECTORY is the location where the collector looks for the files.



**Note:** Restart the Network Discovery Collector pods to enable the changes. If you change the SFTP details in the future, you must re-create the SFTP secret and restart the Network Discovery Collector pods each time.

### Configure EMS certificates

To enable https connection between Network Discovery Collectors and the EMS, issue the following command:

```
oc create configmap ${releaseName}-net-disco-ems-ca-certs --from-file /temp/collector-ems-certs/
```

Where:

#### Configmap Name

The name of configmap is the release name of the Agile Service Manager installation, followed by net-disco-ems-ca-certs

#### /temp/collector-ems-certs/

The folder that contains all EMS certificates (the filename extension of the certificate should be .crt).

**Note:** Restart the Network Discovery Collector pods to enable the changes. If you change the certificates in the future, you must reconfigure the EMS certificates and restart the Network Discovery Collector pods each time.

## About this task

SFTP is required to connect to an EMS during a collector discovery.

## Procedure

### SFTP server setup

1. Create a user on the desired server.

The user and password in the following examples here are 'benedict'.

```
[root@server ~]# useradd benedict  
[root@server ~]# passwd benedict
```

### Client setup

2. From the client environment, copy the ssh key of the newly created user

```
[home@client ~]$ ssh-copy-id benedict@server
```

3. Verify the ssh connection by logging into the machine.

```
[home@client ~]$ ssh benedict@server  
[benedict@server ~]$ exit  
logout  
Connection to server closed
```

4. Verify the sftp connection.

```
[home@client ~]$ sftp benedict@server  
Connected to server  
sftp> quit
```

### SFTP server setup

5. After verifying both the sftp and ssh connections, return to the same server and create a new group to manage the chroot users.

```
[root@server ~]# groupadd sftpusers
```

6. Create a common directory for the chroot users.

```
[root@server ~]# mkdir /sftp
```

7. Create a subdirectory for each individual user.

```
[root@server ~]# mkdir /sftp/benedict
```

8. Create the 'home' directory for the user.

```
[root@server ~]# mkdir /sftp/benedict/home
```

9. Modify the user to add them to the user group.

```
[root@server ~]# usermod -aG sftpusers benedict
```

10. Change permission for the user's chrooted 'home' directory only.

```
[root@server ~]# chown benedict:sftpusers /sftp/benedict/home/
```

11. Modify the /etc/ssh/sshd\_config to include the following lines towards the end of the file:

```
Match Group sftpusers
  ForceCommand internal-sftp -d /home
  ChrootDirectory /sftp/%u
```

12. Restart the sshd service.

```
[root@server ~]# systemctl restart sshd
```

### Client verification

13. Verify that the ssh connection no longer works.

```
[home@client ~]$ ssh benedict@server
This service allows sftp connections only.
Connection to server closed.
```

14. Verify that sftp access is still valid.

```
[home@client ~]$ sftp benedict@server
Connected to server.
sftp> pwd
Remote working directory: /home
sftp> cd /root
Couldn't stat remote file: No such file or directory
```

15. Make sure that a regular user can still log in via ssh without issues.

## Network discovery service configuration template

The following network discovery configuration template is provided (in json format).

### Network discovery configuration template

An example configuration file is included here: ASM\_HOME/data/netdisco-observer/ExampleCommonconfig.json

```
{
  "configId": "default-netdisco",
  "username": "username",
  "password": "4ihmyEBH6pETt14yIxcQ9w==",
  "disco_status_retry_delay_s": 20,
  "connect_read_timeout_ms": 500000,
  "discovery_profile": "N/A",
  "scopes": [
    {
      "type": "Subnet",
      "address": "172.20.120.0/255.255.248.0",
      "protocol": "IPv4"
    }
  ],
  "seeds": [
    {
      "address": "172.20.120.0/255.255.248.0",
      "timeout": 250,

```

```

        "retries":1
    }
],
"snmppasswords":[
    {
        "communitystring":"4ihmyEBH6pETt14yIxcQ9w==",
        "address":"172.20.120.0/255.255.248.0",
        "snmpversion":"V2",
        "snmpv3":{
            "securityname":"shades",
            "authpassword":"4ihmyEBH6pETt14yIxcQ9w==",
            "privtype":"des",
            "privpassword":"4ihmyEBH6pETt14yIxcQ9w==",
            "level":"authPriv",
            "authType":"sha1"
        },
        "snmpport":161,
        "timeout":3000,
        "retries":3
    }
],
"accessListType":null,
"anchors":null,
"gateways":null,
"telnetaccess":[
    {
        "address":"172.20.120.0/255.255.255.255",
        "usernameprompt":"^[[[:space:]]*[[A-Za-z0-9()&<@ /:~. _-]+[$%>#] [[[:space:]]*$",
        "username":"privlusr",
        "passwordprompt":"^[[[:space:]]*[[A-Za-z0-9()&<@ /:~. _-]+[$%>#] [[[:space:]]*$",
        "password":"4ihmyEBH6pETt14yIxcQ9w==",
        "consoleprompt":"^[[[:space:]]*[[A-Za-z0-9()&<@ /:~. _-]+[$%>#] [[[:space:]]*$",
        "accessport":23,
        "timeout":1000,
        "usessh":false,
        "advanced":
            {
                "command":"pwd",
                "passwordprompt":"^[[[:space:]]*[[A-Za-z0-9()&<@ /:~. _-]+[$%>#] [[[:space:]]*$",
                "password":"4ihmyEBH6pETt14yIxcQ9w==",
                "consoleprompt":"^[[[:space:]]*[[A-Za-z0-9()&<@ /:~. _-]+[$%>#] [[[:space:]]*$",
                "selectedcommands":[
                    "ls",
                    "pwd"
                ]
            }
    }
],
"dns":[
    {
        "servicename":"google",
        "system":false,
        "domainsuffix": ".com",
        "iptoname":[
            {
                "ip":"172.20.120.0",
                "name":"ibm"
            }
        ]
    }
],
"collector":[
    {
        "collectorid":"AlcatelSAM-IE",
        "collectortype":"Alcatel5620SamCollector",
        "httpaccess":{
            "url":"http://alcatelsam-host:27080/",
            "username":"username",
            "password":"4ihmyEBH6pETt14yIxcQ9w==",
            "md5password":"4ihmyEBH6pETt14yIxcQ9w==",
            "token":"4ihmyEBH6pETt14yIxcQ9w=="
        }
    }
]
}

```

## Collector sample

The following example code depicts a possible collector deployment with **sample** values.

The target EMS is running on host `nasm-net-disco-collectorsim.host.example.com` via HTTP port 27080, and will be accessed using the supplied username and password.

The scope is set to `null`, meaning all devices are accepted. If you configure a scope, you **must** ensure that your EMS devices are included.

You **must** supply the password in encrypted form.

```
{
  "configId": "default-netdisco",
  "username": "username",
  "password": "4ihmyEBH6pETt14yIxcQ9w==",
  "disco_status_retry_delay_s": 20,
  "connect_read_timeout_ms": 50000,
  "discovery_profile": "N/A",
  "scopes": null,
  "seeds": null,
  "snmppasswords": null,
  "accessListType": null,
  "telnetaccess": null,
  "dns": null,
  "collector": [
    {
      "collectorid": "AlcatelSAM-IE",
      "collectortype": "Alcatel5620SamCollector",
      "httpaccess": {
        "url": "http://alcatelsam-host:27080/",
        "username": "username",
        "password": "4ihmyEBH6pETt14yIxcQ9w==",
        "md5password": "4ihmyEBH6pETt14yIxcQ9w==",
        "token": "4ihmyEBH6pETt14yIxcQ9w=="
      }
    }
  ]
}
```

## Configuring a hybrid system

For information on configuring a hybrid system with NOI deployed on-prem and Agile Service Manager on OCP, See the relevant NOI documentation [here](#).

**Assumption:** The hybrid system deployment described here is intended specifically for managing event data extracted from Netcool/OMNIbus. For installation and other configuration information, see the appropriate installation, configuration and administration topics.

### Related information

## Configuring Humio integration (on-prem only)

If you have deployed a Humio system, you can perform searches on it from within the Topology Viewer. This topic describes the integration steps required (for **on-prem deployments** only).

### Before you begin

#### Required Humio details

Base URL of your Humio instance, for example `https://my--humio-host.com:8080`

Name of the Humio repository you want to search

### About this task

To integrate with your existing Humio system, you run the provided integration script, providing information when prompted, and then restart Agile Service Manager. This creates the **Search in Humio** right-click tool for use from within the Topology Viewer.

## Procedure

1. Log into Agile Service Manager.
2. Run the Humio integration script:

```
$ASM_HOME/bin/set-humio-config.sh
```

3. Enter the following information when prompted:

- Humio instance
- Humio repository name

4. When the script has completed, run the Agile Service Manager start script:

```
$ASM_HOME/bin/asm_start.sh
```

## Results

The UI API service creates the **Search in Humio** tool, which you can launch from the Topology Viewer context menu.

### Tip:

- If required, you can delete the tool from the Topology Tools page.
- The UI API service checks whether the Humio search tool exists each time Agile Service Manager starts, and creates it if it does not. To prevent the tool from being recreated, an administrator user can edit the `.humio-env` file and set the two Humio parameters to the value of **notSet**.

## Related tasks

[“Configuring custom tools” on page 271](#)

As an administrator or advanced user, you can create custom topology tools, which users can then access from within a topology's context menu. This functionality allows you to access properties of a selected item (such as a resource or relationship) and execute some custom functionality within the context of that item.



---

## Chapter 5. Running Observer jobs

Agile Service Manager is available with a large number of observers, and can be deployed as on-prem or RedHat OpenShift Container Platform (OCP) versions. Not all observers are available on OCP.

### observer

An observer is a service that extracts resource information and inserts it into the Agile Service Manager database.

Agile Service Manager includes a configuration UI to help you configure and run observer jobs.

### Observer job names

The characters that you can use when defining the Unique IDs (**unique\_id**) of Agile Service Manager observer jobs have been restricted to the following:

ABCDEFGHIJKLMNOPQRSTUVWXYZ abcdefghijklmnopqrstuvwxyz0123456789-.\_~:#[\]@!  
\$&'()\*+,-;=

If you have used any characters other than these (such as '/'), you must recreate the job definition.

### Observer job types

Observer jobs are either Listen jobs or Load jobs.

#### Observer Listen jobs

Listen jobs run constantly waiting for resource update data.

#### Observer Load jobs

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

### Observer service

Requests to start or stop jobs made through the UI are processed by the observer service, which functions as a proxy between the UI and the observers (which in turn run the observer jobs).

The Observer Service REST API provides a single endpoint for start, stop and list jobs for all observers. When it receives a request to start a new job, the observer service validates the job parameters directly with the observer, and if the job parameters pass validation, it passes the job request to the observer to be executed.

**Remember:** Swagger documentation for the observer service is available at the following default location: `https://<your host>1.0/observer/swagger`

---

## Defining observer security

All observer jobs require password encryption, and in addition some observer jobs require authentication credentials. This topic describes such configuration tasks for both OCP and on-prem versions of Agile Service Manager.

For the OCP version of Agile Service Manager, observer jobs are defined and run using Swagger. For information on how to customize and deploy observers on OCP, also see the included Swagger documentation.

**Remember:** Swagger links to specific observers are in “Swagger reference” on page 410, and more information on specific observers is located in the subtopics under [“Defining observer jobs” on page 171](#).

## Configuring password encryption, authentication certificates and keystores

All observer jobs require password encryption, and in addition some observers require authentication credentials such as certificates, keystores, or both. This topic describes such configuration tasks for OCP

and on-prem versions of Agile Service Manager, and also describes how to post an observer job using Swagger (or cURL).

## About this task

The following steps are described:

- Encrypt the passwords for all observer load or listen jobs
- Obtain an authentication certificate
- Store that certificate as a secret
- Post an observer job

## Procedure

### Encrypt the passwords for all observer load and listen jobs

1. The jobs for all observers require the password in the configuration file to be encrypted.

To encrypt the password:

#### For OCP

Use the commands in the following example:

```
kubectl exec -ti asm-topology-pods -- java -jar /opt/ibm/topology-service/topology-  
service.jar  
encrypt_password --password 'password'
```

Where the value of *asm-topology-pods* can be obtained using the following command:

```
kubectl get pods | grep <release name>-topology  
<release name>-topology-xxxxx-xxxxx
```

The encryption utility will return an encrypted password.

#### For on-prem

Run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

### To acquire an SSL certificate and copy the certificate to the security folder (on-prem)

2. Use the following **Cisco ACI Observer example** to acquire an SSL certificate.

In the following example, you use OpenSSL to connect to Cisco APIC over port 443, and extract a SSL Certificate from Cisco APIC to a *<certificate\_file\_name>.crt* file.

```
echo -n | openssl s_client -connect {Cisco APIC IPAddress}:443 | sed -ne '/-BEGIN  
CERTIFICATE-/,/-END CERTIFICATE-/p' > ./ {certificate_file_name}.crt
```

3. Copy the certificate file *<certificate\_file\_name>.crt* to the `$ASM_HOME/security` directory to complete the SSL setup.

### Managing authentication certificates and storing them as secrets (OCP)

4. Obtain the authentication certificate using OpenSSL.

```
echo -n | openssl s_client -connect {ipAddress}:{port} | sed -ne '/-BEGIN CERTIFICATE-/,/  
-END CERTIFICATE-/p' | base64 -w 0 > target_system.crt
```

Where `target_system.crt` contains the encoded certificate, and `{ipAddress}` could be the IP address of any of the following target systems:

- Bigfix Inventory
- Ciena Blue Planet



- Cisco ACI
- Juniper CSO
- Kubernetes master node
- OpenStack
- Rancher
- VMware NSX
- VMware vCenter
- Zabbix

**Example** target\_system.crt:

```
[root@localhost ~]# cat target_system.crt
LS0tLS1CRUdJTiBDRVJUSUZJQ0FURSB0tLS0tCk1JSUN3RENDQWFnQ0NRGRuMENqU3BXZXhUQU5CZ2txaGtpRz13
MEJBVVVGURBaU1RMHdDd11EVlFRFRERBUKIKVUVsRE1SRXdEd11KS29aSWhtY05BUWtCRmdKV1V6QWVGdzb4
TmptBeE1qRXd0ekV5TWpWYU1DSXhEVEFMQmd0VkJBTU1CRUZRU1VNeEVUQVBCZ2txa
GtpRz13MEJDUUVXQWxwVE1JSUJJakFOckJna3Foa2lHOXcwQkFRRUZBQU9DQVE4QU1JSUJDZ0tDQVFFQW10b0dxd
FI0RlFpWkd0UWftand1YmxRYjRobU0KTzJwOGtjbGwUwL2NuUno3cSttWGYzWlRQYTZsWEk2MG9BbmVPSGowZEVA
MkhwRWFfb1BUBWJmWUF6Y0ZQd11EVlFRFRERBUKIKVUVsRE1SRXdEd11KS29aSWhtY05BUWtCRmdKV1V6QWVGdzb4
WV5NVhhaFBtdkZDbDBtM3Y3Cisxa11FMFRNVlBnTk56R0ZSUxU1RVlGc3FZWZGZGZlU1Ua1F6cks3YnE0Rk
JiMw1kVjFsYnVOMWhISzd2SFEKS3ZUNHBGbgGx1NTRHU0JhZ2RSbUdad0dta0tNZHRGUkEvc3pBWEMreJQ0cHN3T05yd
TJnbDR3bG5MZWZwFQUx1THM4UDgrOUx0eFN3YVJvb0VMcHRjb3pKdEpUb2E4QS9zZXRaSi81RUJQNmhj
Nk1yUWxHQktRSURBUUFCck1BMEduDU3FHU01iM0RRRUJCUVVBQTRJQkFRQkJuZz1JK2pBdjhNUjBYemM1SUUxd
TBkK0JweW90ZGVGRk14T2sKZWFsNzNubmktZmMh4QUQzd1QzenZSE1SUe0d31xMWJqQ05LY3BZOGVcbVJvVzh0Sn1
NdG9vcU9hN1JMWNPTAoyeVZub1Vna092THRvVjM5eFNQ1B0MzV4YXJjdG9yYdE9NZWJRw1c1ZC9Hc1lPZUFLTL1
Nr1TlQwRmtreDE0UzJFC1pBV19IUUVHaVpUR0tQNkx1czYzLzJiTEJVNHDGUjg3bjNkdFJFVUp5eGQ4
ZDJDTFA4MkE2UTNOT210ZEdkam0KSnfQZXNEaWwXW5Gd09xUk1XOWFGWTVUSUt0L25PQzhqc10cVFM
ZTJZc1lnZ242N0crLyBQy9kV21JSVQ2dgpBWTVMejhwOWQwSzZUaGxLeVpNZkdYVkJvNmM1vTms1ajQ4ckJlZ2J5c
FhTm1J2SnIKLS0tLS1FTkQgQ0VSVE1GSUNBVEU0tLS0tLQo=

[root@localhost ~]#
```

**Tip:** To get the ipaddress and port for each respective observer, see “Swagger reference” on page 410 or the observer subtopics under “Defining observer jobs” on page 171

5. **For the GoogleCloud Observer**, encrypt the contents of the service account key file using base64. For example:

```
cat {project_id_file_name}.json | base64 -w 0 > googlecloud.json
```

Where googlecloud.json contains the encoded service account key file, and {project\_id\_file\_name}.json is the service account key file downloaded from 'Credentials' under 'API & Services' in the Google Cloud Platform dashboard.

Next, store the encoded service account key file as a secret.

6. Store a certificate as a secret.

Each installed Agile Service Manager release has a single special secrets file. Data added to that is made available to the appropriate observer containers. Run the following command, assuming **<release-name>-topology** is the Helm release name for Agile Service Manager when installed with NOI.

```
$ kubectl edit secret <release-name>-topology-custom-secrets
```

Paste in the encoded certificate generated in the previous step.

- a) Find the correct secrets file using the following command:

```
$ kubectl get secrets -l app=<release-name>-topology-custom-secrets
NAME TYPE DATA AGE
asm-custom-secrets Opaque 2 29d
```

- b) Edit the appropriate file for your release.

```
$ kubectl edit secret <release-name>-topology-custom-secrets
```

- c) Add a name and value pair to the data section.

The value is the certificate generated earlier. The name is what you enter as the certificate file name to run the observer job.

```
data:
{name}:{value}
```

**Example** of expected content in the secret file after adding vcenter.crt is as follows (where the data section is between the 'apiVersion' and 'kind' sections).

**Note:** This VMware vCenter Observer example registers the vcenter.crt SSL certificate in OCP Secret, and vcenter.crt is the job parameter value for the VMware vCenter Observer. Define a new **{name}** parameter in the same file for other observers that require SSL certificates. You provide the certificate in OCP Secret and the settings you provide for truststore/truststore password will be used to generate a new truststore automatically.

```
apiVersion: v1
data:
vcenter.crt:LS0tLS1CRUdJTiBDRVJUSUZJQ0FURS0tLS0tCk1JSUN3RENDQWFnQ0NRRGRuMENqU3BXZXhUQU5CZ2
txaGtpRz13
MEJBVVVGQURBaU1RMHdDd1lEVlFRFRERBUkIKVUVsRE1SRXdEd1lKS29aSWh2Y05BUWtCRmdKVlV6QWVGdzb4
TmpBeE1qRXd0ekV5TWpwYU1DSXhEVEFMQmd0VkJBTU1CRUZRU1VNeEVUQVBCZ2
txa
GtpRz13MEJDUUVXQWxwVE1JSUJJakFOckJna3Foa2lHOXcwQkFRRUZBQU9DQVE4QU1JSUJDZ0tDQVFFQW10b0dxd
FI0R1FPWkd0UWFtand1YmxRYjRobU0KTzJwOGtjbGUwL2NuUno3cSttWGYzWlRQYTZsWEk2MG9BbmVPSGowZEVA
MkhwRFFb1BUBWJmWUF6Y0ZQdJNVWApMWjM3VWVoMDZXTjMxS29tSSs2czJtSk1IiWM0Mw44M1RiUU5uWUNjYjZjd1
ZLc
WV5NVhhaFBtdkZDbDBtM3Y3Cisxa1lFMFRNVlBnTk56R0ZSUXU1RVlGc3FZWZGbfZhz0lUa1F6cks3YnE0Rk
JiMw1kVjFsYnVOMWhISzd2SFEKS3ZUNHBGbgx1NTRHU0JhZ2RSbUdad0dta0tNZHRGUKEv3pBWEMrejQ0cHN3T05y
d
TJnbDR3bG5MZTVvM2NWZwpFQUx1THM4UDgrOUx0eFN3YWJvb0VMcHRjb3pKdEpUb2E4QS9zZXRaSi81RUJQNmhj
Nk1yUWxHQkRSURBUUFCCK1BMEdDU3FHU0l1m0RRRUJCUVVBQTRJQkFRQkJuZz1JK2pBdjhhNUjBYemM1SUUxd
TBkK0JweW90ZGVrbk14T2sKZWFsNzNUbmKzWmh4QUQzd1QzenZNSE1SUEc0d3lxMWJqQ005LY3BZOGVcbVJuVzh0Sn1
NdG9vcU9hN1JMWGNPTAoyeVZub1Vna092THRvPjM5eFNfQ1B0MzV4YXJJdGYyde9NZWJRWlcl1ZC9Hc11PZUFLT1
NzTlQWRmtreDE0UzJFC1pBVi9IUUVHaVpUR0tQNkx1czYzLzJiTEJVNhdGUjg3bjNkdFJFVUp5eGQ4
ZDJDTFA4Mke2UTNOT210ZEdkam0KSnfQZXNEawXWE5Gd09xUk1XOWFGWTVUSUt0L25PQzhqczI0cVFm
ZTJZcllnZ242N0crLytBQy9kV21JSVQ2dgpBWTVMjeHw0WQwSzZUaGxLeVpNZkdYVknNMF1vTms1ajQ4ckJlZ2J5c
FhTM1J2SnIKLS0tLS1FTkQgQ0VSVElGSUNBVEU0tLS0tLQo=
kind:Secret
```

If the edit is successful, the following message will be displayed:

```
secret "<release-name>-topology-custom-secrets" edited
```

7. In the OCP GUI, you can view the configured secret under the **Menu > Workload > Secrets** option, where the **Name** is 'asm-custom-secrets'. Within asm-custom-secrets, all data configured earlier is displayed.

### Posting a job

8. Post the job via the Swagger UI or cURL.

**Note:** The default value for the password\_ssl\_truststore\_file property is **password** and has to be encrypted.

Example cURL command:

```
curl --location --insecure --header 'Content-Type: application/json' --header
'Accept: application/json' --header 'X-TenantID:
cfd95b7e-3bc7-4006-a4a8-a73a79c71255' -d '{
  "unique_id": "my job",
  "type": "query",
  "parameters": {
    "data_center": "LondonDC1",
    "vcenter_username": "admin",
    "vcenter_password": "RW+w==",
    "vcenter_api_url": "https://localhost/rest",
    "vcenter_certificate": "vcenter.crt",
    "ssl_truststore_file": "localhost.jks",
    "password_ssl_truststore_file": "IxcQ9w==",
    "connect_read_timeout_ms": 5000
  }
}' 'https://<master-ip address>/1.0/vmcenter-observer/jobs/restapi'
```

**Note:** When using cURL, you may need to add `--location` so that it will follow redirects, and `--insecure` as the proxy server is using HTTPS.

## What to do next

For a repeating job, you can wrap the cURL in a script and use a normal cron job.

## Defining observer jobs using the Observer Configuration UI

You configure observer jobs using the Observer Configuration UI, which you access through DASH.

**Note:** To schedule jobs and configure truststore certificates, use the information in the following sections: “Defining observer security” on page 99 and “Defining observer jobs” on page 171

### To access the observer dashboard:

1. Using a compatible browser, open DASH using the DASH URL. For example:

```
https://<DASH HOST>:<DASH PORT>/ibm/console/
```

2. Login using your user credentials.
3. In DASH, open the **Administration** menu.
4. Under the Agile Service Management heading, click **Observer jobs** to display the **Observer Configuration UI**. From here, you can search for an existing job, or open either the **Existing jobs** or **Add a new job** expandable sections.

### Existing jobs

The **Existing jobs** panel displays all jobs as tiles.

Each job state is indicated (Pending, Running, Stopping, Stopped, Finished, Scheduled). You can hover over the Schedule icon to view additional details. The Schedule icon is either a clock icon if scheduled to run at set intervals, or an hourglass icon if it is scheduled to run only once at a future time).

From here, you can run a job manually, or switch it On or Off, depending on the job type.

You can also use the **List of options** drop-down to **View & Edit** or **Delete** a job, or **View History** for it (View History is available for Load jobs only). If a schedule has been set for a Load job, you can cancel the schedule using **Remove Schedule**, which will revert the job to a default, on-demand Load job.

When you select **View History** for an observer, the **Observer job history** window opens listing the last ten jobs, when they were run, their observation count, duration and status. You can refine the list using the **All**, **Complete**, or **Incomplete** tabs.

### Add a new job

The **Add a new job** panel displays all jobs that can be configured in tile format.

Click the **Configure** button under a specific observer to open its job configuration window.

The Observer job configuration UI lists each job parameter that you have to configure.

### Tip:

If an observer has been stopped or removed, you will be unable to run existing jobs, or add new jobs. Stopped or removed observers and jobs that are listed in the Observer Configuration UI will be disabled (grayed out) or removed in progressive (housekeeping) steps. If you are reinstalling or reactivating an observer, the jobs and the observer will again become available.

1. **Up to 5 minutes** after removal, observers and jobs still appear as normal until the housekeeping process runs, but cannot be used.
2. **Up to 60 minutes** after removal, the observer is still listed, but jobs are grayed out and marked offline until the next housekeeping process runs. You can delete existing jobs, but cannot view, add or edit jobs.
3. **After 60 minutes** the removed observer is no longer listed, **but jobs remain**, though they are grayed out and marked offline. You can delete existing jobs, but cannot view, add or edit jobs.

4. If **at any time** you reinstall or reactivate the observer, it reappears in the UI, and existing (previously active) jobs are no longer grayed out. You can delete, view or edit existing jobs, or add new jobs.

## Configuring ALM Observer jobs

Using the Agile Lifecycle Manager Observer, you can define jobs that dynamically load data associated with intent from the Agile Lifecycle Manager for analysis by Netcool Agile Service Manager.

### Before you begin

**Important:** The ALM Observer supports IBM Agile Lifecycle Manager Version 2.2.

Ensure you have the Agile Lifecycle Manager Kafka server host and topics to hand, such as the Agile Lifecycle Manager server, the Kafka port, and the topics used for lifecycle events.

**Important:** To access Agile Lifecycle Manager remotely, you must ensure that the Agile Lifecycle Manager installation has been configured with the **KAFKA\_ADVERTISED\_HOST\_NAME** so as to allow remote connections. For more information, see the Configuration reference topic in the Agile Lifecycle Manager Knowledge center at the following location: [https://www.ibm.com/support/knowledgecenter/SS8HQ3\\_2.2.0/GettingStarted/r\\_alm\\_quickreference.html](https://www.ibm.com/support/knowledgecenter/SS8HQ3_2.2.0/GettingStarted/r_alm_quickreference.html)

The Agile Lifecycle Manager Observer is installed as part of the core installation procedure.

### About this task

The Agile Lifecycle Manager Observer jobs listen to the Kafka 'state change' topics of Agile Lifecycle Manager, as well as the Agile Lifecycle Manager Resource Manager. Information is extracted from Agile Lifecycle Manager about Assemblies and Resources and a topology is created.

After installation, you define and start the following two jobs.

#### Listen for lifecycle events ('alm' job)

The **alm** job is a long-running job that monitors its source for updates and runs until it is explicitly stopped, or until the observer is stopped.

#### Listen for Resource Manager events ('rm' job)

The **rm** job is a long-running job that monitors its source for updates and runs until it is explicitly stopped, or until the observer is stopped.

Table 17. ALM Observer parameters for <i>alm</i> jobs		
Parameter	Action	Details
Unique ID	Enter a unique name for the job	Required
IBM Agile Lifecycle Manager instance name	Use this to identify the Agile Lifecycle Manager installation and any associated Resource Managers.	Required
Topic	Use this to identify the Agile Lifecycle Manager Kafka topic.	Required
Group ID	Use this to identify the Kafka group ID to be used.	Required
Connection	Use this to specify the Kafka Host and Port to be used.	Required
Observer job description	Enter additional information to describe the job.	Optional

Table 18. ALM Observer parameters for ALM **rm** (Resource Manager) jobs

Parameter	Action	Details
Unique ID	Enter a unique name for the job	Required
IBM Agile Lifecycle Manager instance name	Use this to identify the Agile Lifecycle Manager installation and any associated Resource Managers.	Required
Topic	Use this to identify the Agile Lifecycle Manager <b>Resource Manager</b> Kafka topic.	Required
Group ID	Use this to identify the Kafka group ID to be used.	Required
Connection	Use this to specify the Kafka Host and Port to be used.	Required
Observer job description	Enter additional information to describe the job.	Optional

## Procedure

1. From the **Observer Configuration UI**, click **Configure** under the IBM Agile Lifecycle Manager icon, or select an existing ALM job to be edited.
2. Choose either **alm** or **rm** from the job type drop-down.

### To configure an alm job

3. Enter or edit the following parameters:
  - Unique ID
  - IBM Agile Lifecycle Manager instance name
  - Topic (the Kafka topic for the Agile Lifecycle Manager lifecycle events)
  - Group ID
  - Connection

### To configure an rm job

4. Enter or edit the following parameters:
  - Unique ID
  - IBM Agile Lifecycle Manager instance name
  - Topic (the Kafka topic for the Agile Lifecycle Manager Resource Manager lifecycle events)
  - Group ID
  - Connection

**Important:** The value of the **IBM Agile Lifecycle Manager instance name** parameter needs to be the same for both jobs to allow for the topology to be combined.

5. Enter an **Observer job description** to explain the purpose of the job in more detail.
6. Click **Run job** to save your job and begin retrieving information.

## Configuring Ansible AWX Observer jobs

Using the Ansible AWX Observer, you can define load jobs that will gather and read data from a defined Ansible AWX instance. When a load job is run, it will load a baseline topology data view from Ansible AWX through its REST APIs.

### Before you begin

**Important:** The Ansible AWX Observer supports the on-premise Ansible AWX version 11.1.0 and Tower API version 3.6.1.

The Ansible AWX Observer is installed as part of the core installation procedure.

Ensure you have the Ansible AWX details to hand, such as the host, port, user and token before defining and running the observer job.

**Required:** Before defining the observer load job, you must obtain a personal access token (PAT) from AWX. The token authentication is recommended for the usage of Tower's API.

1. Access the AWX UI.
2. From the menu bar, navigate to the **Users** page that is under the **Access** section.
3. Select a username, then click the **Token** option.
4. Click on the **+** (plus) icon.
5. Leave the application field blank, and choose **read** for the token's access.

### About this task

The Ansible AWX Observer imports ITSM Resource Topology Service data to Agile Service Manager.

You define and start the following job. You must edit the parameters in the configuration file before running this job.

#### Full Topology Upload job (via restapi)

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

Table 19. Ansible AWX Observer Load parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job.	Required
Ansible AWX hostname	Specify the Ansible AWX hostname.	Required
Ansible AWX host port	Specify the Ansible AWX host port.	Required
Username (owner) of the personal access token (PAT)	Specify the username (owner) of the personal access token (PAT).	Required
Encrypted personal access token (PAT)	Specify the personal access token (PAT).	Required. Must be encrypted.
HTTPS trustStore file name	Specify the trustStore file name	For <b>both on-prem and OCP</b> , provide the JKS name, and the observer will then create the JKS file accordingly.

Table 19. Ansible AWX Observer Load parameters (continued)

Parameter	Action	Details
trustStore file password	<p>Specify the trustStore file password. You must supply the trustStore file password in encrypted form.</p> <p><b>To encrypt the password on premise</b> Run the <code>encrypt_password.sh</code> script in the <code>\$ASM_HOME/bin</code> directory</p> <p><b>To encrypt the password on OCP</b> Run the relevant <code>kubect</code> command, as described <a href="#">here</a>.</p>	Required. Must be encrypted.
Target system certificate	Specify a certificate by name to load into the trustStore.	<p><b>On-prem</b> To create and store the certificate, see <a href="#">“Configuring password encryption, authentication certificates and keystores”</a> on page 99</p> <p><b>OCP</b> To obtain the authentication certificate using OpenSSL and store them as secrets, see <a href="#">“Defining observer security”</a> on page 99</p>
SSL verification	Choose whether SSL validation is true or false. Turning SSL validation to <code>false</code> will use HTTPS without host verification.	Optional
Number of jobs to display for each template	Specify the number of recent jobs that will be modeled on the topology for each of the job template. Max value is 50. Larger value might result in the topology being less perceptible	Optional
Connection timeout (milliseconds)	Enter a connection timeout in milliseconds.	Optional. Must be a value greater than 0 (zero), and the default is 5000 (5 seconds).
Read timeout (milliseconds)	Enter a read timeout in milliseconds	Optional
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

**Encryption requirement:** The Load job requires the password in the configuration file in encrypted form. To encrypt, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

**SSL requirements:** To acquire SSL certificates and build SSL truststores, use the relevant instructions for both on-prem and OCP in the following section: [“Defining observer security” on page 99](#)

## Procedure

1. From the **Observer Configuration UI**, click **Configure** under the AWX icon, or select an existing job to be edited.
2. Enter or edit the following parameters:
  - Unique ID
  - Ansible AWX hostname
  - Ansible AWX host port
  - Username (owner) of the personal access token (PAT)
  - Encrypted personal access token (PAT)
  - HTTPS trustStore file name
  - trustStore file password
  - Target system certificate (optional for on-prem, required for OCP)
  - SSL verification (optional)
  - Number of jobs to display for each template (optional)
  - Connection timeout (milliseconds) (optional)
  - Read timeout (milliseconds) (optional)
3. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
4. Enter an **Observer job description** to explain the purpose of the job in more detail.
5. Click **Run job** to save your job and begin retrieving information.

## Configuring Application Discovery Observer jobs

The Application Discovery Observer job retrieves network topology data, including discovered applications, their components, configurations and dependencies, from the Application Discovery database via REST API, and uses this data to create topologies within the Agile Service Manager topology service.

### Before you begin

The Application Discovery services are installed as part of the core installation procedure. This includes the Application Discovery Observer, Application Discovery engines and Db2 docker containers, which have been installed and should be running.

Before you configure an application discovery job, you must configure the discovery details. This procedure is described here: [“Configuring application discovery service” on page 61](#)

When successful, the configuration details will be saved in topology management artifacts.

### About this task

The Application Discovery Observer job retrieves topology data from the Application Discovery database via REST API. The observer loads and updates the resources and their relationships within the Netcool Agile Service Manager core topology service.



You define and start the following job.

### Load job

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

Table 20. Application Discovery Observer job parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job	Required
Profile ID for common configuration schema	Specify the Application Discovery configuration ID	Required
Flag to trigger AppDisco discovery	Set to 'true' to trigger both AppDisco and AppDisco observer discovery, and 'false' to trigger AppDisco observer discovery only.	Required
Date	Extend the result set returned to Agile Service Manager based on the date given. Set to empty for snapshot discovery.	Optional. Date must be in dd-mm-yyyy HH:mm:ss format
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

**Note:** You can only run **one** job per profile.

### Procedure

1. From the Observer Configuration UI, click **Configure** under the AppDisco icon, or select an existing AppDisco job to be edited.
2. Enter or edit the following parameters:
  - Unique ID
  - Profile ID for common configuration schema
  - Flag to trigger AppDisco discovery
  - Date in dd-mm-yyyy HH:mm:ss format
3. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
4. Enter an **Observer job description** to explain the purpose of the job in more detail.
5. Click **Run job** to save your job and begin retrieving information.

### Results

The Application Discovery Observer job discovers and then uploads the application services data to the Agile Service Manager topology service.

### Related tasks

[“Configuring application discovery service” on page 61](#)

The Agile Service Manager application discovery service is installed together with the other core components. Once you have provided the discovery details, Agile Service Manager can discover application data.

[“Installing Db2 Community Edition on OCP” on page 70](#)

To deploy the Agile Service Manager application discovery service on OCP, you install Db2 Community Edition into an OpenShift cluster.

#### Related reference

[“Application discovery service configuration template” on page 65](#)

The following application discovery configuration template is provided (in json format).

[“Supported application discovery accessListTypes and properties” on page 68](#)

The following accessListTypes and their properties are supported by the Agile Service Manager application discovery service, and must be defined in the application discovery configuration file.

## Configuring AppDynamics Observer jobs

Using the AppDynamics Observer, you can define a full load job that will read data from the AppDynamics Controller via the REST API. This job can provide, for example, business applications, nodes, and tiers to Agile Service Manager.

### Before you begin

**Important:** The AppDynamics Observer supports the cloud/SaaS AppDynamics version 4.5.12 and API version 4.5.x.

Ensure you have the AppDynamics details to hand, such as the instance, account, username and password before running the observer job.

The AppDynamics Observer is installed as part of the core installation procedure.

**Tip:** Before defining the observer load job, you must create an AppDynamics user with the correct permissions. This is required for REST API authentication.

1. On the AppDynamics **Administration** page, click the gear icon (top right).
2. Under the Admin group, select **Administration**, then select the **Users** tab.
3. Select the **AppDynamics** option from the **Display users** drop-down list.
4. Assign all available roles to the user by selecting **Add from Roles > Select All > Done**, then click **Save**.

### About this task

The AppDynamics Observer imports ITSM Resource Topology Service data to Agile Service Manager.

You define and start the following job. You must edit the parameters in the configuration file before running this job.

#### Full Topology Upload job

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

Table 21. AppDynamics Observer parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job.	Required
Instance	Specify the AppDynamics Controller instance.	Required

Table 21. AppDynamics Observer parameters (continued)

Parameter	Action	Details
Account	Specify the Tenant account name.	Required
Username	Specify the name of the user of the specified tenant account.	Required
Password	Specify the password for the specified user.	Required. Must be encrypted.
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

### Encryption requirement:

The Load job requires the password in the configuration file in encrypted form. To encrypt the password, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

### Procedure

1. From the **Observer Configuration UI**, click **Configure** under the AppDynamics icon, or select an existing job to be edited.
2. Enter or edit the following parameters:
  - Unique ID
  - Instance
  - Account
  - Username
  - Password
3. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
4. Enter an **Observer job description** to explain the purpose of the job in more detail.
5. Click **Run job** to save your job and begin retrieving information.

## Configuring AWS Observer jobs

Using the AWS Observer, you can define jobs that read services data from the Amazon Web Services (AWS) through AWS SDK and generate a topology. It is installed as part of the core installation procedure.

### Before you begin

**Important:** The AWS Observer supports the cloud/SaaS AWS version 1.11.

Ensure you have the AWS details to hand, such as AWS Region, Access Key ID and Access Secret Key.

## About this task

The AWS Observer supports multiple Amazon web services such as EC2 for its 'elastic compute' services.

You define and start the following job. You must edit the parameters in the configuration file before running this job.

### Full Topology Upload job

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

Table 22. AWS Observer parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job	Required
Access Key	Specify the AWS access key.	Required
Secret Key	Specify the AWS secret key.	Required. Must be encrypted.
Region	Specify the AWS region or multiple regions to discover.	Required
Property to exclude	Single or multiple properties to exclude	Optional
Proxy	Specify the proxy ('true' or 'false')	Required. If set to 'true', you must supply the proxy host, port and protocol.
Proxy Host	Specify the proxy host	Optional (unless proxy is set to 'true')
Proxy Port	Specify the proxy port	Optional (unless proxy is set to 'true')
Proxy Protocol	Specify the proxy protocol	Optional (unless proxy is set to 'true')
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

### Encryption requirement:

The Load job requires the **secretKey** in the configuration file in encrypted form. To encrypt them, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the secret key. The encryption utility will return an encrypted **secretKey**.

## Procedure

### To find your Access Key and Secret Access Key:

1. Log into your AWS Management Console.
2. Click on your user name at the top right of the page.
3. Click on the **Security Credentials** link from the drop-down menu.
4. Find the **Access Credentials** section, and copy the latest Access Key ID.

5. Click on the **Show link** in the same row, and copy the Secret Access Key.

#### To find the region

6. Check the region at the following location:

<https://docs.aws.amazon.com/general/latest/gr/rande.html>

**Note:** To discover more than one region, select the check boxes from the Amazon Region drop-down.

#### To configure the AWS job

7. From the **Observer Configuration UI**, click **Configure** under the AWS icon, or select an existing job to be edited.
8. Enter or edit the following parameters:
  - Unique ID
  - Access key
  - Secret access key
  - Region
  - Proxy
9. Enter one (single) or more (multiple) properties to be excluded.
10. Enter a host to define the Proxy Host.
11. Enter a port to define the Proxy Port.
12. Enter a protocol to define the Proxy Protocol.
13. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
14. Enter an **Observer job description** to explain the purpose of the job in more detail.
15. Click **Run job** to save your job and begin retrieving information.

## Configuring Azure Observer jobs

Using the Azure Observer, you can define a full load job that will read data from Azure cloud services through its REST APIs and generate a topology.

### Before you begin

**Important:** The Azure Observer supports the cloud/SaaS Azure version.

Ensure you have the Azure details to hand, such as the Tenant ID, Client ID, and client password before running the observer job.

The observer is installed as part of the core installation procedure.

### About this task

The Azure Observer retrieves topology data from Azure cloud services via REST APIs exposed by the Azure API server.

You define and start the following job. You must edit the parameters in the configuration file before running this job.

#### Full Topology Upload job

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

Table 23. Azure Observer parameters

Parameter	Action	Details
Unique ID	Enter a unique name for the job	Required
Data center	Enter the data center of the Azure service	Required
Tenant ID	Specify the Azure account tenant ID	Required
Client ID	Specify the client ID to access the registered app	Required
Client password	Enter the client password to access the registered app	Required. Must be encrypted.
Connection timeout (milliseconds)	Define the Azure connection timeout	Optional
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

#### Encryption requirement:

The Load job requires the client password in the configuration file in encrypted form. To encrypt the client password, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

#### Procedure

1. From the **Observer Configuration UI**, click **Configure** under the Azure icon, or select an existing job to be edited.
2. Enter or edit the following parameters:
  - Unique ID
  - Data center
  - Tenant ID
  - Client ID
  - Client password
3. Enter a **Connection timeout** (in milliseconds).
4. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
 Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
5. Enter an **Observer job description** to explain the purpose of the job in more detail.
6. Click **Run job** to save your job and begin retrieving information.

## Configuring BigFix Inventory Observer jobs

You configure BigFix Inventory Observer jobs to read data from a Bigfix Inventory instance through its REST APIs, and generate a topology.

### Before you begin

**Important:** The BigFix Inventory Observer supports the on-premise BigFix Inventory version 9.5.

The Bigfix Inventory Observer is installed as part of the core installation procedure.

Before configuring a Bigfix Inventory job, ensure you have the Bigfix Inventory details to hand such as the Bigfix Inventory URL, API token and SSL trustStore.

**Important:** The Bigfix Inventory Observer supports Bigfix Inventory Version 9.5.0.

### About this task

You define and start the following job.

#### Bigfix Inventory Observer job (full topology load)

This job loads a baseline of topology data from an Bigfix Inventory environment.

Run this job whenever you need Bigfix Inventory topology data refreshed.

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

Table 24. Bigfix Inventory Observer job parameters

Parameter	Action	Details
Unique ID	Enter a unique name for the job.	Required
BigFix Inventory API token	Enter the BigFix token for authentication.	Required. Must be encrypted.
BigFix Inventory instance URL	Specify the API URL of the BigFix Inventory endpoint (including port).	Required. Usually in the following format: <code>https://&lt;hostname or IP address&gt;:&lt;port&gt;</code>
Bigfix Inventory resources	Specify the resources to be discovered.	Optional. Lists supported values such as software, hardware or *. If left blank, all available resources are discovered.
Bigfix Inventory certificate	Specify the name of the certificate to be loaded into the trust store.	<b>On-prem</b> To create and store the certificate, see <a href="#">“Configuring password encryption, authentication certificates and keystores”</a> on page 99 <b>OCP</b> To obtain the authentication certificate using OpenSSL and store them as secrets, see <a href="#">“Defining observer security”</a> on page 99

Table 24. Bigfix Inventory Observer job parameters (continued)

Parameter	Action	Details
HTTPS trustStore file name	Specify the trustStore file name.	For <b>both on-prem and OCP</b> , provide the JKS name, and the observer will then create the JKS file accordingly.
trustStore file password	Specify the trustStore password to decrypt the HTTPS trustStore file.	Required. Must be encrypted.
SSL Validation	Choose whether SSL validation is on or off. Turning SSL validation off will use HTTPS without host verification.	Optional
Bigfix Inventory connection timeout (ms)	Enter the time at which the connection times out.	Optional. Must be a value greater than 0 (zero), and the default is 5000 (5 seconds).
Data Center	Specify the data center(s) in which the Bigfix Inventory instance runs.	Required. If more than one, list them (comma-separated).
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

**Encryption requirement:** The job requires passwords in encrypted form. To encrypt the Bigfix Inventory token and SSL trustStore file password, run the `encrypt_password.sh` script in the `$ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

## Procedure

### To configure and run Bigfix Inventory Observer jobs

1. From the **Observer Configuration UI**, click **Configure** under the Bigfix Inventory icon, or select an existing Bigfix Inventory job to be edited.
2. Enter or edit the following parameters:
  - Unique ID
  - BigFix Inventory API token (must be encrypted)
  - BigFix Inventory instance URL
  - Bigfix Inventory resources (optional)
  - Bigfix Inventory certificate (optional)
  - HTTPS trustStore file name
  - trustStore file password (must be encrypted)
  - SSL Validation (optional)
  - Bigfix Inventory connection timeout (ms) (optional)
  - data\_center
3. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.



Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.

4. Enter an **Observer job description** to explain the purpose of the job in more detail.
5. Click **Run job** to save your job and begin retrieving information.

#### To acquire Bigfix Inventory SSL certificate and build SSL truststore

6. Use the following command to use OpenSSL to connect to Bigfix Inventory, and extract a SSL Certificate from Bigfix Inventory to a `<certificate_file_name>.cert` file.

```
echo -n | openssl s_client -connect {Bigfix Inventory IpAddress}:{port} | sed
-ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > ./{certificate_file_name}.cert
```

7. Use the following Java keytool command to import the Bigfix Inventory certificate file into a keystore and encrypt the keystore with a given password.

```
keytool -import -v -trustcacerts -alias {Bigfix Inventory Hostname}
-file {certificate_file_name}.cert -keystore {keystore file name}
-storepass {your password to encrypt keystore}
```

8. Copy the keystore file (`{keystore file name}`) to the `$ASM_HOME/security` directory to complete the SSL setup.

### What to do next

Run this job whenever you need Bigfix Inventory topology data refreshed.

## Configuring Ciena Blue Planet Observer jobs

Using the Ciena Blue Planet Observer, you can define jobs that will gather and read all topology data from the Blue Planet MCP instance by REST API and generate a topology.

### Before you begin

**Important:** The Ciena Blue Planet Observer supports MCP 4.0.

The Ciena Blue Planet Observer is installed as part of the core installation procedure.

Ensure you have the Ciena Blue Planet details to hand, such as API username, API password, MCP URL, MCP certificate, truststore file and truststore password.

### About this task

The Ciena Blue Planet Observer has two jobs, the `restapi load` and `websocket listen` jobs.

- When a load job is run, it loads baseline topology data through Blue Planet MCP APIs: Network Elements (constructs), EquipmentHolder, Equipment, TPE (Terminating Point Encapsulation), and FRE (Forwarding Relationship Encapsulation).
- When a listening job is run, the observer listens to changes in resource's status from the BluePlanet MCP instance through a websocket connection.

**Tip:** Defining observer jobs using the UI is the same for both on-premise and OCP.

#### Full Topology Upload job (via `restapi`)

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

Table 25. Ciena Blue Planet Observer restapi Load parameters

Parameter	Action	Details
Unique ID	Enter a unique name for the job.	Required
Server URL	The URL of the MCP server instance	Required
MCP Certificate	Certificate name to load into the trust store	<b>On-prem</b> To create and store the certificate, see <a href="#">“Configuring password encryption, authentication certificates and keystores”</a> on page 99  <b>OCP</b> To obtain the authentication certificate using OpenSSL and store them as secrets, see <a href="#">“Defining observer security”</a> on page 99
SSL truststore file	Exact HTTPS trust store file name	For <b>both on-prem and OCP</b> , provide the JKS name, and the observer will then create the JKS file accordingly.
SSL truststore password	The password to decrypt HTTPS trust store file	Required. Must be encrypted.
SSL Validation	Choose whether SSL validation is on or off. Turning SSL validation off will use HTTPS without host verification.	Optional
Connection timeout	Sets the connection and read timeout in milliseconds	Optional. Must be a value greater than 0 (zero), and the default is 5000 (5 seconds).
Username	MCP API username	Required
Password	MCP API password	Required
Data Center	The data center the MCP instance is running in	Required
Tenant Name	The tenant to use	Required
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

#### Websocket Listen job

A long-running job that monitors its source for updates and runs until it is explicitly stopped, or until the observer is stopped.

Table 26. Ciena Blue Planet Observer Websocket Listen parameters

Parameter	Action	Details
Unique ID	Enter a unique name for the job.	Required

Table 26. Ciena Blue Planet Observer Websocket Listen parameters (continued)

Parameter	Action	Details
Websocket URL	The MCP websocket URL to connect to	Required
Websocket Topics Subscriptions	The MCP topics to subscribe to	Required
Server URL	The URL of the MCP server instance	Required
MCP Certificate	Certificate name to load into the trust store	Optional for on-prem. If used, must be in the /opt/ibm/netcool/asm/security directory.  Required for OCP. Use the instructions in the following topic to obtain the authentication certificate using OpenSSL and store them as secrets: <a href="#">“Defining observer security” on page 99</a>
SSL truststore file	Exact HTTPS trust store file name	Required. The supported format is JKS and the file is relative to \$ASM_HOME/security
SSL truststore password	The password to decrypt HTTPS trust store file	Required. Must be encrypted.
SSL Validation	Choose whether SSL validation is on or off. Turning SSL validation off will use HTTPS without host verification	Optional
Connection timeout	Sets the connection and read timeout in milliseconds	Optional. Must be a value greater than 0 (zero), and the default is 5000 (5 seconds).
Username	MCP API username	Required
Password	MCP API password	Required
Data Center	The data center the MCP instance is running in	Required
Tenant Name	The tenant to use	Required
Observer job description	Enter additional information to describe the job.	Optional

**SSL requirements:** To acquire SSL certificates and build SSL truststores, use the relevant instructions for both on-prem and OCP in the following section: [“Defining observer security” on page 99](#)

## Procedure

1. From the **Observer Configuration UI**, click **Configure** under the Ciena Blue Planet icon, or select an existing Ciena Blue Planet job to be edited.
2. Enter or edit the following parameters for both restapi full load and websocket listen jobs:
  - Unique ID
  - Server URL

- MCP certificate
  - SSL truststore file
  - SSL truststore password (must be encrypted)
  - SSL Validation (optional)
  - Connection timeout
  - Username
  - Password
  - Data Center
  - Tenant Name
3. Enter or edit the following additional parameters for the websocket listen job:
    - Websocket URL
    - Websocket topic subscriptions
  4. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
 Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
  5. Enter an **Observer job description** to explain the purpose of the job in more detail.
  6. Click **Run job** to save your job and begin retrieving information.

## Configuring Cisco ACI Observer jobs

You use the Cisco ACI Observer when you have a Cisco ACI environment with Cisco Application Policy Infrastructure Controller (APIC) in your environment. The Observer interfaces with Cisco APIC and makes active REST calls to Cisco APIC in the Cisco ACI environment. You configure observer jobs that dynamically load Cisco ACI data for analysis by Netcool Agile Service Manager from the **Observer Configuration UI**.

### Before you begin

**Important:** The Cisco ACI Observer supports the on-premise Cisco ACI version 4.1.

Ensure you have the Cisco ACI service details to hand, such as the Cisco APIC username, Cisco APIC password, Cisco APIC SSL TrustStore and Cisco APIC URL.

The Cisco Application Centric Infrastructure (ACI) Observer is installed as part of the core installation procedure.

**Optional: Install and configure a proxy:** The Cisco ACI Observer can (optionally) connect to Cisco ACI via a proxy. See the [Proxy Host](#) and [Proxy Port](#) parameters listed in the table. A proxy can be used to create a tunnel to the target system.

#### Example: Squid proxy

1. Install the squid proxy:

```
$ sudo yum install squid
```

2. Edit the `/etc/squid/squid.conf` file to add `acl` for the source and target system. Example of a line added:

```
acl localnet src 1.2.3.0/19 # ASM Host
```

3. Initialize the squid directories:

```
$ sudo squid -z
```

#### 4. Configure the squid proxy for auto-start:

```
$ sudo systemctl start squid
$ sudo systemctl enable squid
```

See the related links for more information about the squid proxy.

### About this task

A Cisco ACI Observer job extracts Cisco ACI resources from Cisco APIC via REST. The Observer loads and updates the resources and their relationships within the Netcool Agile Service Manager core topology service.

You define and start the following jobs.

#### Restapi Load job

A transient (one-off) job that loads all requested topology data using Cisco APIC REST APIs to build a tenant logical construct topology or a fabric topology, and then exits.

A 'restapi' job loads initial topology data, and can resynchronize topology data from Cisco ACI into the Agile Service Manager topology.

You assign 'restapi' as the job type for /jobs/restapi observer endpoint.

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

#### Websocket Listen job

A long-running job that listens for notifications from Cisco APIC to build the topology and runs until it is explicitly stopped, or until the observer is stopped.

A 'websocket' job monitors changes from Cisco APIC object notification and updates the Agile Service Manager topology.

You always run a 'websocket' job **after** running a 'restapi' job type.

You assign 'websocket' as the job type for /jobs/websocket observer endpoint.

Table 27. Cisco ACI Observer <b>restapi</b> and <b>websocket</b> job parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job.	Required
Cisco APIC password	Enter the password for Cisco APIC authentication.	Required. Must be in encrypted text.
Cisco APIC endpoint	Specify the API URL of the Cisco APIC endpoint.	Required. Usually in the following format: <code>https://[hostname or IP address]/api</code>
Cisco APIC certificate	Specify a certificate by name to load into the trustStore.	<b>On-prem</b> To create and store the certificate, see <a href="#">“Configuring password encryption, authentication certificates and keystores”</a> on page 99  <b>OCP</b> To obtain the authentication certificate using OpenSSL and store them as secrets, see <a href="#">“Defining observer security”</a> on page 99

Table 27. Cisco ACI Observer **restapi** and **websocket** job parameters (continued)

Parameter	Action	Details
HTTPS trustStore file name	Specify the trustStore file name.	For <b>both on-prem and OCP</b> , provide the JKS name, and the observer will then create the JKS file accordingly.
HTTPS trustStore file password	Specify the trustStore password to decrypt the HTTPS trustStore file.	Required
SSL Validation	Choose whether SSL validation is on or off. Turning SSL validation off will use HTTPS without host verification.	Optional
Cisco APIC username	Specify the username to connect as, or listen to.	Required
Tenant name	Use this to identify the tenant.	Required. Set to 'admin' if there is no specific tenant. Set to '' to load Fabric Topology resources.
Proxy Host	Specify the proxy host through which to connect	Optional
Proxy Port	Specify the proxy port, defaults to 8080	Optional
Connection timeout	Specify the connection timeout in ms (default is 5000)	Optional
Read timeout	Specify the read timeout in ms (default is 2000)	Optional
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

#### Encryption requirement:

Both jobs require passwords in encrypted form. To encrypt the password and file name, run the `encrypt_password.sh` script in the `$ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

To acquire a Cisco APIC SSL certificate and build the SSL truststore, see the instructions from [this step onwards](#).

## Procedure

### To configure Cisco ACI Observer jobs

1. From the **Observer Configuration UI**, click **Configure** under the Cisco ACI icon, or select an existing Cisco ACI job to be edited.
2. Choose either **restapi** or **websocket** from the job type drop-down.
3. Enter or edit the following parameters for both job types:
  - Unique ID

- Cisco APIC password (must be encrypted)
  - Cisco APIC endpoint
  - Cisco APIC certificate (optional)
  - HTTPS trustStore file name
  - HTTPS trustStore file password (must be encrypted)
  - SSL Validation (optional)
  - Cisco APIC username
  - Tenant name
4. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
  5. Enter an **Observer job description** to explain the purpose of the job in more detail.
  6. Click **Run job** to save your job and begin retrieving information.

#### To acquire a Cisco APIC SSL certificate and build the SSL truststore

7. For **OCP** Agile Service Manager deployments, use the relevant instructions in the following topic:  
[“Defining observer security” on page 99](#)
8. For **on-prem** Agile Service Manager deployments, use the relevant instructions in the following topic:  
[Defining Cisco ACI Observer jobs \(on-prem\)](#)

#### Related information

## Configuring Contrail Observer jobs

Using the Contrail Observer, you can retrieve topology data from Juniper Network Contrail Release 4.1 via REST APIs exposed by the Contrail API server. This observer is developed against Juniper Network Contrail that integrates with OpenStack orchestration platform (Ubuntu 18.04 + Contrail Cloud - Ocata).

### Before you begin

**Important:** The Contrail Observer supports the on-premise Contrail version 4.1.0.

Ensure you have the Contrail API Server and OpenStack credentials details to hand. For rabbitmq jobs, you must also specify the location of the RabbitMQ queue and its authentication details.

The Contrail observer is installed as part of the core installation procedure.

### About this task

Contrail Observer jobs retrieve topology data from Juniper Network Contrail Release 4.1 via REST APIs exposed by the Contrail API server. The observer loads and updates the resources and their relationships within the Netcool Agile Service Manager core topology service.

You define and start the following jobs.

#### Rabbitmq Listen job

A long-running job that monitors its source for updates and runs until it is explicitly stopped, or until the Observer is stopped.

This job loads all supported resources during startup, and listens to RabbitMQ messages from 'vnc\_config.object-update' fanout exchange.

There is no need to run the restapi job before running the rabbitmq job, because the rabbitmq job performs a restapi job during initialization.

## Restapi Load job

This job loads all supported resources.

Run this job whenever you need the Contrail topology data refreshed.

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

Table 28. Contrail Observer <i>rabbitmq</i> job parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job.	Required
RabbitMQ hostname or IP address	Enter a hostname or IP address for the RabbitMQ server.	Required
RabbitMQ port	Specify the port for connection.	Required
RabbitMQ username	Specify the username for authentication with RabbitMQ.	Required
RabbitMQ password	Enter the password for authentication with RabbitMQ.	Required. Must be encrypted.
RabbitMQ virtual host	Enter the RabbitMQ virtual hostname	Optional. Default is <code>/</code> .
Contrail API URL	Specify the URL for the Contrail API server.	Required
OpenStack Authentication URL	Enter the authentication URL for the identity service.	Required
OpenStack username	Specify the OpenStack user name to connect as (or to).	Required
OpenStack password	Specify the OpenStack password with which to authenticate.	Required
Authentication type	Specify the authentication type used.	Optional. The default is <b>Keystone</b> , the other option is <b>None</b> .
OpenStack project domain name	Specify the OpenStack project domain name.	Optional
OpenStack user domain name	Specify the OpenStack project user domain name.	Optional
OpenStack project name	Enter the OpenStack project name for version 3 authentication	Optional
OpenStack tenant name	Enter the OpenStack tenant name for version 2 authentication	Optional
OpenStack identity API version	Select an option from the dropdown list.	Optional
Connection and read timeout (milliseconds)	Choose the time at which the connection and read action times out.	Optional. Must be a value greater than 0 (zero). The default is 5000 (5 seconds).



Table 28. Contrail Observer **rabbitmq** job parameters (continued)

Parameter	Action	Details
Observer job description	Enter additional information to describe the job.	Optional

Table 29. Contrail Observer **restapi** job parameters

Parameter	Action	Details
Unique ID	Enter a unique name for the job.	Required
Contrail API URL	Specify the URL for the Contrail API server.	Required
OpenStack Authentication URL	Enter the authentication URL for the identity service.	Required
OpenStack username	Specify the OpenStack user name to connect as (or to).	Required
OpenStack password	Specify the OpenStack password with which to authenticate.	Required
Authentication type	Specify the authentication type used.	Optional. The default is <b>Keystone</b> , the other option is <b>None</b> .
OpenStack project domain name	Specify the OpenStack project domain name.	Optional
OpenStack user domain name	Specify the OpenStack project user domain name.	Optional
OpenStack project name	Enter the OpenStack project name for version 3 authentication	Optional
OpenStack tenant name	Enter the OpenStack tenant name for version 2 authentication	Optional
OpenStack identity API version	Select an option from the dropdown list.	Optional
Connection and read timeout (milliseconds)	Choose the time at which the connection and read action times out.	Optional. Must be a value greater than 0 (zero). The default is 5000 (5 seconds).
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

### Encryption requirement:

Both jobs require the Contrail token in encrypted form. To encrypt the token, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

## Procedure

1. From the **Observer Configuration UI**, click **Configure** under the Contrail icon, or select an existing Contrail job to be edited.
2. Choose either **restapi** or **rabbitmq** from the job type drop-down.

### To configure a restapi job

3. Enter or edit the following **required** parameters:
  - Unique ID
  - Contrail API URL
  - Openstack Authentication URL
  - Openstack username
  - Openstack password
4. Enter or edit the following **optional** parameters:
  - Authentication type
  - OpenStack project domain name
  - OpenStack user domain name
  - OpenStack project name
  - OpenStack tenant name
  - OpenStack identity API version
  - Connection and read timeout (milliseconds)
5. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
6. Enter an **Observer job description** to explain the purpose of the job in more detail.

### To configure a rabbitmq job

7. Enter or edit the following parameters:
  - Unique ID
  - RabbitMQ hostname or IP address
  - RabbitMQ port
  - RabbitMQ username
  - RabbitMQ password (must be encrypted)
  - RabbitMQ virtual host (**optional**)
  - Contrail API URL
  - Openstack Authentication URL
  - Openstack username
  - Openstack password
8. Enter or edit the following **optional** parameters:
  - Authentication type
  - OpenStack project domain name
  - OpenStack user domain name
  - OpenStack project name
  - OpenStack tenant name

- OpenStack identity API version
  - Connection and read timeout (milliseconds)
9. Enter an **Observer job description** to explain the purpose of the job in more detail.
  10. Click **Run job** to save your job and begin retrieving information.

## Configuring DNS Observer jobs

Using the DNS Observer, you can query internal DNS server performance, and use the returned information on response times and service addresses to create topologies within the topology service. The DNS Observer supports forward and reverse job types, with 'recurse' or 'no recurse' options.

### Before you begin

Ensure you have the DNS access details to hand, such as DNS server, address types and port numbers.

The DNS Observer is installed as part of the core installation procedure.

### About this task

The DNS Observer provides DNS query services and topological insight into how a specified DNS server is performing forward (name-to-IP address) or reverse (IP address-to-name) lookups. Query results include a list of addresses, information on how long it takes the DNS server to resolve a lookup, and, optionally (with the maximum number of recursive calls set at 200) how the DNS server is recursively resolving a given name or IP address.

**Tip:** The relationship types can be customized with line color, width and pattern functions. See the [“Creating custom relationship type styles” on page 283](#) topic for more information.

You define and start the following jobs.

#### Reverse lookup job

A transient (one-off) Load job that loads all requested DNS reverse (IP address-to-name) lookup topology data.

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

#### Forward lookup job

A transient (one-off) Load job that loads all requested DNS forward (name-to-IP address) lookup topology data.

Table 30. DNS Observer <b>reverse</b> job parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job.	Required
ip_address	Enter the host or internet name to lookup.	Required
Address Types	Specify the address types to be observed.	Required. Select either IPv4 or IPv6.
Server	Specify the DNS server.	Required
Port	Specify the UDP DNS port.	Optional
Recursive query	Toggle <b>True</b> or <b>False</b> .	Optional. If set to True, the maximum number of calls is set at 200.

Table 30. DNS Observer **reverse** job parameters (continued)

Parameter	Action	Details
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

Table 31. DNS Observer **forward** job parameters

Parameter	Action	Details
Unique ID	Enter a unique name for the job.	Required
Domain name	Specify a domain.	Required
Address Types	Specify the address types to be observed.	Required. Select either IPv4 or IPv6.
Server	Specify the DNS server.	Required
Port	Specify the UDP DNS port.	Optional
Recursive query	Toggle <b>True</b> or <b>False</b> .	Optional. If set to True, the maximum number of calls is set at 200.
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

## Procedure

1. From the **Observer Configuration UI**, click **Configure** under the DNS icon, or select an existing DNS job to be edited.
2. Choose either **reverse** or **forward** from the job type drop-down.

### Configure a 'reverse' job

3. Enter or edit the following required parameters:

- Unique ID
- ip\_address
- Address Types
- ServerPort (optional)
- Recursive query (optional)

4. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.

Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.

5. Enter an **Observer job description** to explain the purpose of the job in more detail.

### Configure a 'forward' job

6. Enter or edit the following parameters:

- Unique ID
- Domain name
- Address Types

- ServerPort (optional)
  - Recursive query (optional)
7. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
  8. Enter an **Observer job description** to explain the purpose of the job in more detail.
  9. Click **Run job** to save your job and begin retrieving information.

## Configuring Docker Observer jobs

Using the Docker Observer, you can discover Docker network resources, including Docker Swarm clusters, and then visualize (or model) this data as a topology view in the Agile Service Manager UI. You configure observer jobs from the **Observer Configuration UI**.

### Before you begin

**Important:** The Docker Observer supports Docker version 3.1.0.

**Note:** Docker UCP v3.1.0 supports only TLS 1.2 for SSL negotiation and has removed support for TLS 1 and TLS 1.1.

Ensure you have the details for your Docker job to hand, specifically your Docker system's Unix socket, and / or host and port number.

The Docker Observer is installed as part of the core installation procedure.

**Update Notes:** If you have updated a previous version of Agile Service Manager with existing Docker Observer job data, you must run a data migration script (as documented in the release notes on-prem update topic) before running new observer jobs.

### About this task

Using the Observer Configuration UI you configure observer jobs that query the Docker REST API to retrieve data and display it as a topology in the Topology Viewer. The Docker Observer can model external Docker systems, and it can also provide a System health view of the Docker system on which Agile Service Manager runs.

The job parameters determine whether to connect to a local Docker on the same (UNIX) host as the observer using the **unix\_socket** parameter, or to a remote Docker using the **host** and **port** parameters.

Table 32. Docker Observer job parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job	Required
Host	Use this to identify the TCP host socket (HTTP or HTTPS) on which to access the remote Docker system.	Required for <b>remote</b> Docker access only
Username	Specify the username of the remote Docker environment with HTTPS.	Required for remote Docker with HTTPS access only.
Password	Specify the password of the remote Docker environment with HTTPS.	Required for remote Docker with HTTPS access only. Must be encrypted.

Table 32. Docker Observer job parameters (continued)

Parameter	Action	Details
Docker SSL Certificate	Specify the certificate file name.	<b>On-prem</b> To create and store the certificate, see <a href="#">“Configuring password encryption, authentication certificates and keystores” on page 99</a>  <b>OCP</b> To obtain the authentication certificate using OpenSSL and store them as secrets, see <a href="#">“Defining observer security” on page 99</a>
Docker SSL TrustStore File	Specify the trustStore file name.	For both <b>on-prem</b> and <b>OCP</b> , required for remote Docker with HTTPs access only. Provide the JKS name, and the observer will then create the JKS file accordingly.
SSL TrustStore File Password	Specify the trustStore password.	Required for remote Docker with HTTPs access only. Must be encrypted.
Port	Use this to identify the TCP port (HTTP or HTTPS) on which to access the remote Docker system.	Required for <b>remote</b> Docker access only
Unix Socket	Use this to access local docker environments using the complete path.	Required for <b>local</b> Docker access only. Host and port parameters <b>must</b> be empty.
View	Use this to select which resources are modeled in the topology view.	Optional. The Default displays running resources only. Options are:  <b>Container</b> All running containers  <b>Image</b> Images used by running containers  <b>Task</b> Running tasks only
Containers to exclude	List container you want to exclude.	Optional
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

**SSL requirements:** To acquire SSL certificates and build SSL truststores, use the relevant instructions for both on-prem and OCP in the following section: [“Defining observer security” on page 99](#)

## Procedure

1. From the **Observer Configuration UI**, click **Configure** under the Docker icon, or select an existing Docker job to be edited.
2. Configure one of the following job types:
  - To discover **remote** Docker network resources **through TCP port exposure**, enter or edit the following parameters:
    - Unique ID
    - Host
    - Port
    - View (optional)
    - Containers to exclude (optional)
    - Job description (optional)
  - To discover **remote** Docker network resources **through HTTPS**, enter or edit the following parameters:
    - Unique ID
    - Host
    - Port
    - Username
    - Password
    - Docker SSL Certificate (optional)
    - Docker SSL TrustStore File
    - SSL TrustStore File Password
    - View (optional)
    - Containers to exclude (optional)
    - Job description (optional)
  - To discover **local** Docker networks (if the Unix socket is accessible via the Docker container), enter or edit the following parameters:
    - Unique ID
    - Unix socket
    - View (optional)
    - Containers to exclude (optional)
    - Job description (optional)
3. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.

Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
4. Click **Run job** to save your job and begin retrieving information.

## Configuring Dynatrace Observer jobs

Using the Dynatrace Observer, you can query a specified Dynatrace environment for information about its applications, services, process groups, and infrastructure entities.

### Before you begin

**Important:** The Dynatrace Observer supports Dynatrace cloud/SaaS and Dynatrace onPrem v1.182.102.

Ensure you have generated a Dynatrace token to access your Dynatrace environment. You also need topology access scope to access the Dynatrace resources.

Ensure you have the Dynatrace access details to hand, such as Dynatrace API URL and API token.

The Dynatrace Observer is installed as part of the core installation procedure.

### About this task

You define and start the following job.

#### Dynatrace Load job

A transient (one-off) job that loads all requested Dynatrace resource data.

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

Table 33. Dynatrace Observer job parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job.	Required
Base URL	<p>Specify the API URL of the Dynatrace endpoint (including version).</p> <p>Example of a Dynatrace cloud-based URL:</p> <pre>https://&lt;your-environment-id&gt;.live.dynatrace.com/api/v1/</pre> <p>Example of a Dynatrace managed or onPrem URL:</p> <pre>https://&lt;your-domain&gt;/e/&lt;your-environment-id&gt;/api/v1/</pre>	Required
API Token	Enter the Dynatrace token for authentication.	Required. Must be encrypted.
Enable/Disable Secure Connection to CSO's host	Select to either enable or disable a secure connection to the CSO's host.	Set to <b>true</b> to secure the connection, or <b>false</b> to bypass



Table 33. Dynatrace Observer job parameters (continued)

Parameter	Action	Details
Dynatrace certificate	If you have set the <b>enable_secure_host_connection</b> parameter to <b>true</b> , you must specify a certificate by name to load into the trustStore.	<b>On-prem</b> To create and store the certificate, see <a href="#">“Configuring password encryption, authentication certificates and keystores” on page 99</a> <b>OCP</b> To obtain the authentication certificate using OpenSSL and store them as secrets, see <a href="#">“Defining observer security” on page 99</a>
SSL TrustStore File	If you have set the <b>enable_secure_host_connection</b> parameter to <b>true</b> , you must specify the SSL TrustStore filename.	For <b>both on-prem and OCP</b> , provide the JKS name, and the observer will then create the JKS file accordingly.
SSL TrustStore Password	If you have set the <b>enable_secure_host_connection</b> parameter to <b>true</b> , you must specify the SSL TrustStore password. You must supply the trustStore file password in encrypted form. <b>To encrypt the password on premise</b> Run the encrypt_password.sh script in the \$ASM_HOME/bin directory <b>To encrypt the password on OCP</b> Run the relevant kubectl command, as described here.	Required. Must be encrypted.
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

**Encryption requirement:** The job requires the API token in encrypted form. To encrypt the Dynatrace token, run the encrypt\_password.sh script in the \$ASM\_HOME/bin directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

## Procedure

1. From the **Observer Configuration UI**, click **Configure** under the Dynatrace icon, or select an existing Dynatrace job to be edited.
2. Enter or edit the following job parameters:
  - Unique ID

- Base URL
  - API Token
3. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
  4. Enter an **Observer job description** to explain the purpose of the job in more detail.
  5. Click **Run job** to save your job and begin retrieving information.

## Configuring File Observer jobs

Using the File Observer functionality, you can write bespoke data to a file in a specific format, upload this data to the topology service, and then visualize this data as a topology view in the Agile Service Manager UI.

### Before you begin

The File Observer is installed as part of the core installation procedure.

### About this task

The File Observer reads topology data from files located in the `$ASM_HOME/data/file-observer/` directory, and uploads it. You must create these files manually.

Topology data in a file is comprised of vertices (nodes) and edges. A vertex represents an object (resource), while an edge represents the relationship between two objects.

Each line of the file you create should be in one of the formats below, loading a single resource vertex (including optional relationships in the `_references` field) or a single edge, deleting a single vertex, or pausing execution.

Lines starting with V: (vertex), E: (edge), D: (delete) or W: (wait) are treated as instruction lines to be processed. Other lines, for example lines that are empty or commented out, are ignored.

#### Line format

##### V:

Load a resource vertex, with a JSON representation as documented for the body of the topology service API method: `POST /resources`

If specifying the **\_status** element, acceptable state values are open, closed, or clear, and acceptable severity values are clear, indeterminate, warning, minor, major, or critical.

##### E:

Load an edge, with a JSON representation as documented for the **\_references** section of the body of the topology service API method `POST /resources`

##### D:

Delete a resource vertex, identified by its `uniqueId`

##### W:

Pause for the given duration (for testing purposes only).

Takes an integer period followed by a string specifying the units.

#### Tip:

- An example file is available in the `$ASM_HOME/data/file-observer` directory.
- See the related links for more information on available timestamps formats.

**Important:** Ensure that the file is structured correctly. For each line of the file, information included after the closing `}` that matches an opening `{` is ignored, and no error is recorded.

## File Load job

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

Table 34. File Observer job parameters

Parameter	Action	Details
Unique ID	Enter a unique name for the job	Required
File Name	Specify the name of the file to be loaded.	Required. Must be relative to the \$ASM_HOME/data/file-observer/ directory (rather than absolute).
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

## Procedure

1. From the **Observer Configuration UI**, click **Configure** under the File icon, or select an existing File job to be edited.
2. Enter or edit the following parameters:
  - Unique ID
  - File Name
3. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
4. Enter an **Observer job description** to explain the purpose of the job in more detail.
5. Click **Run job** to save your job and begin retrieving information.

## Results

The File Observer job loads all requested topology data from the file specified. This job runs only once.

## What to do next

Run this job whenever the content in your file has been updated.

**Optional configuration:** The size of file posted to the File Observer is set to 8Mb by default for on-prem deployments. A limit is useful in guarding against, for example, some types of denial-of-service (DOS) attacks. You can change the default.

### On-prem

Access the `/opt/ibm/netcool/asm/etc/nginx/conf.d/nasm-file-observer.rules` file and change the following property:

```
client_max_body_size 8m
```

### On OCP

No limits exist.

## Related information

[Timestamps formats](#)

## Configuring GoogleCloud Observer jobs

Using the GoogleCloud Observer, you can define a full load job that will read services data from the Google Cloud Platform's Compute Services through Google's Compute Services SDK, and then generate a topology.

### Before you begin

**Important:** The Google Cloud Observer supports the cloud/SaaS Google Cloud version.

The GoogleCloud Observer is installed as part of the core installation procedure.

The GoogleCloud Observer supports GoogleCloud's compute services. Ensure you have the GoogleCloud details in hand, such as the Project ID, Service Account Key File and Zone, before running the observer job.

### About this task

The GoogleCloud Observer supports a transient (one-off) Load job that loads all requested topology data via Google's Compute Services SDK to build the topology, and then exit.

You define and start the following job. You must edit the parameters in the configuration file before running this job.

#### Full Topology Upload job

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

Table 35. GoogleCloud Observer parameters		
Parameter	Action	Details
Project ID	Enter the Google Cloud Platform Project ID	Required
Service Account Key File	Supply the Google Cloud Platform Service Account Key File.  Copy the json file to the \$ASM_HOME/security directory for on-prem.	Required
Zone	Specify the Google Cloud Platform Zones.	Required.
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

**Note:** You must create a service account key file or use an existing one to allow the GoogleCloud Observer to discover resources from GoogleCloud.

### Procedure

#### To configure the GoogleCloud job

1. From the **Observer Configuration UI**, click **Configure** under the GoogleCloud icon, or select an existing job to be edited.

#### To create a service account key file

2. From the Google Cloud Platform dashboard, under your 'Project ID', go to **APIs and Services** and then choose **Credentials**.

The **Credentials** page is displayed listing a number of authentication methods.

3. Select the **Service account** authentication service
4. From **Create Credentials**, choose **Service account**.
5. Complete the service account details, then click **Create** to create a new service account.
6. Select the **Compute Engine** > **Compute Admin** role, then click **Continue** and **Done**.
7. To download the .json file from the **Credentials** page, click **Service account** > **Manage Service Accounts**.
8. Select the new service account, then choose **Create key** from the overflow menu (the three dots under the Actions column).
9. Choose a key type of **JSON**, then click **Create**. A .json file will be downloaded.
  - **For on-prem**, store the .json file under /opt/ibm/netcool/asm/security
  - **For ICP/OCF**, follow [these steps](#) to store the service account key file as a secret.

The filename will be used in the observer parameter (**service\_account\_key\_file**) for the full load job.

10. Enter or edit the following parameters:
  - Project ID
  - Service Account Key File
  - Zone
11. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.

Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
12. Enter an **Observer job description** to explain the purpose of the job in more detail.
13. Click **Run job** to save your job and begin retrieving information.

## Results

The job gathers information and updates the topology.



**Trouble:** While the job is running, the status of discovered resources may appear as 'indeterminate' in the topology until the full upload is complete.

## Configuring IBM Cloud Observer jobs

Use the IBM Cloud Observer when you have IBM Cloud installed in your environment to run jobs that read data from an IBM cloud instance. These jobs retrieve Cloud Foundry Apps information and services, and then dynamically load the retrieved data for analysis by Netcool Agile Service Manager.

### Before you begin

**Important:** The IBM Cloud Observer supports the cloud/SaaS IBM Cloud version.

**Important:** The IBM Cloud Observer supports Cloud Foundry API version 2.92.

Ensure you have the IBM Cloud access details to hand to specify and access the cloud instance, such as the instance ID, credentials, and region.

The IBM Cloud Observer is installed as part of the core installation procedure.

## About this task

You define and start the following job.

### Full Topology Upload job

The IBM Cloud Observer imports ITSM Resource Topology Service data to Agile Service Manager.

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

Table 36. IBM Cloud Observer job parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job	Required
Instance	Enter the name of the IBM cloud instance.	Required
User email	The email address used to access the instance.	Required
Encrypted password	Enter the password used to access the instance.	Required. Must be in encrypted text.
IBM Cloud Region	Choose the cloud instance region from the drop-down list: <ul style="list-style-type: none"><li>• US_S</li><li>• UK</li><li>• EU</li><li>• AP</li></ul>	Required. Each region has its own URI, and only a single region is discovered in a full load job. To discover different regions, a full load job needs to be triggered for each region.
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

### Encryption requirement:

Both jobs require passwords in encrypted form. To encrypt the password and file name, run the `encrypt_password.sh` script in the `$ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

## Procedure

1. From the **Observer Configuration UI**, click **Configure** under the IBM Cloud icon, or select an existing IBM Cloud job to be edited.
2. Enter or edit the following parameters:
  - Unique ID
  - Instance
  - User email
  - Encrypted password
  - IBM Cloud Region

3. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
4. Enter an **Observer job description** to explain the purpose of the job in more detail.
5. Click **Run job** to save your job and begin retrieving information.

## Results

The IBM Cloud Observer job loads all requested topology data. Run this job whenever you need the IBM Cloud topology data refreshed.

## Configuring ITNM Observer jobs

Using the IBM Tivoli Network Manager (ITNM) Observer, you can define jobs that dynamically load data discovered by Network Manager for analysis by Netcool Agile Service Manager.

### Before you begin

**Important:** The ITNM Observer supports the on-premise ITNM version 4.2.

Ensure you have the IBM Tivoli Network Manager service details to hand, such as the domain, host and port number.

The ITNM Observer is installed as part of the core installation procedure.

### About this task

The ITNM Observer jobs extract Network Manager resources using an Object Query Language JDBC driver. The Observer loads and updates the resources and their relationships within the Agile Service Manager core topology service.

You configure the following two jobs.

#### Load

A transient (one-off) job that queries Network Manager for topology data, and performs a complete upload for a single ITNM domain.

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

#### Listen

A long-running job that monitors the Network Manager message bus for changes and update the topology service accordingly. When the job is started, the observer creates an OQL connection that listen for changes in the ITNM network. Any resources added, changed or deleted are passed on by the OQL connection and the Agile Service Manager topology service is updated.

The listen job runs until it is explicitly stopped, or until the observer is stopped.

Table 37. ITNM Observer <b>load</b> and <b>listen</b> job parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job	Required
ITNM instance name	Specify the ITNM instance name.	Required. Specify the same instance name for all jobs to enable connectivity across domains.

Table 37. ITNM Observer <b>load</b> and <b>listen</b> job parameters (continued)		
Parameter	Action	Details
ITNM domain	Specify the ITNM domain.	Required
Hostname or Server IP	Specify the ITNM host name or server IP on which the domain is running.	Required
ITNM domain port	Specify the ITNM port for the specified domain.	Required. For more information, see <a href="#">Tip (ITNM port)</a> .
OQL connection timeout (ms)	Specify the OQL Connection timeout value.	Optional. The default is 3000 (30 seconds).
Exclude resources without connections	Select whether to display disconnected resources.	Optional. The choices are <b>true</b> and <b>false</b> . The default is <b>true</b> .
Edge Type Map	Map ITNM layers to topology relationship types.	Optional  This parameter accepts the default or custom Agile Service Manager edge types. For custom edge types, register the edge in the topology service Swagger under the POST /types/edge API.  If left blank, Agile Service Manager will auto-map. For more information, see the <a href="#">edge type mapping tip</a> .
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

**Tip (ITNM port):** The value of **port** will vary if multiple domains exist. To identify which port is associated with a specific domain in your Network Manager host, open the \$NCHOME/etc/precision/ServiceData.cfg file and locate the line that specifies which ncp\_config service binds to the domain, for example:

```
SERVICE: ncp_config DOMAIN: NCOMS ADDRESS: 172.17.0.4 PORT: 7968 SERVERNAME:
core.ibm.com DYNAMIC: NO
```

The **port** identified in this example is 7968 (while the **domain** is NCOMS, and the **host** (ITNM Server IP) is 172.17.0.4).

**Tip (edge type mapping):** To identify topology relationship types, see the following file: \$NCHOME/precision/disco/stitchers/DNCIM/PopulatedDNCIMTopologies.stch  
Alternatively, run the following OQL statement against the model service to list the available topology types:

```
select ENTITYNAME from ncimCache.entityData where METAClass='Topology'
```

## Procedure

1. From the **Observer Configuration UI**, click **Configure** under the ITNM icon, or select an existing ITNM job to be edited.
2. Choose either **load** or **listen** from the job type drop-down.



3. Configure the following **required** parameters for both **load** and **listen** jobs:
  - Unique ID
  - ITNM instance name
  - ITNM domain
  - Hostname or Server IP
  - ITNM domain port
4. Configure the following **optional** parameters for both **load** and **listen** jobs:
  - OQL connection timeout (ms)
  - Exclude resources without connections
  - Edge Type Map
5. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
 Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
6. Enter an **Observer job description** to explain the purpose of the job in more detail.
7. Click **Run job** to save your job and begin retrieving information.

## Configuring Jenkins Observer jobs

Using the Jenkins Observer, you can define listen jobs that receive build information generated by the Agile Service Manager plugin for Jenkins.

### Before you begin

Ensure you have all required Jenkins details to hand.

The Agile Service Manager Version 1.1.8 release targets a typical CI/CD pipeline implementation in Jenkins that makes use of a Git-based repository as the SCM, and JFrog Artifactory as the build products repository.

Before configuring the Jenkins Observer, you must install and configure the Jenkins plugin provided with the Agile Service Manager installation images (\*.hpi). The plugin sends your instrumented build data to the Jenkins Observer in order to generate the appropriate topology information:  
[“Configuring the Jenkins plugin” on page 53](#)

### About this task

You define and start the following job.

#### Listen job

The standalone listen job receives Jenkins build notification data for a specified namespace and processes it as a topology.

The listen job is long-running, and runs until it is explicitly stopped or until the observer is stopped.

Table 38. Jenkins Observer job parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job.	Required
Jenkins observation namespace	Specify a provider namespace with which to associate all the topology data generated by this job.	Required

Table 38. Jenkins Observer job parameters (continued)		
Parameter	Action	Details
Observer job description	Enter additional information to describe the job.	Optional

## Procedure

1. From the **Observer Configuration UI**, click **Configure** under the JNK icon, or select an existing job to be edited.
2. Enter or edit the following Listen job parameters:
  - Unique ID
  - Namespace
  - Observer job description (optional)
3. Click **Run job** to save your job and begin retrieving information.

## Related tasks

[“Configuring the Jenkins plugin” on page 53](#)

The Agile Service Manager software includes the Jenkins plugin, which you install on your Jenkins server using the Jenkins Plugin Manager Advanced installation wizard. From the Jenkins server, the plugin gathers and sends information to the Jenkins Observer.

[“Refining Jenkins integration and visualization” on page 57](#)

You can extend your Jenkins integration with rules to merge data from different sources, custom topology display conventions, and the use of templates for the automated generation of topologies.

[“Defining Jenkins Observer jobs” on page 208](#)

Using the Jenkins Observer, you can define listen jobs that receive build information generated by the Agile Service Manager plugin for Jenkins.

## Related information

[“Jenkins Observer troubleshooting” on page 347](#)

# Configuring Juniper CSO Observer jobs

Using the Juniper CSO Observer, you can define a full load job that will gather and read data about topology data from Juniper CSO. The observer is installed as part of the core installation procedure.

## Before you begin

**Important:** The Juniper CSO observer supports the on-premise Juniper CSO version 4.1.0.

Ensure you have the Juniper CSO details to hand, such as details of the Juniper CSO API Server and its credentials.

## About this task

The Juniper CSO Observer retrieves topology data from Juniper CSO Release 4.1 via REST APIs exposed by CSO API server.

You define and start the following job.

### Full Topology Upload job

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

Run this job whenever you need the Juniper CSO topology data refreshed.

Table 39. Juniper CSO Observer job parameters

Parameter	Action	Details
Unique ID	Enter a unique name for the job.	Required
Juniper CSO Central MS URL	Specify the Juniper CSO Central micro-service URL.	Required
CSO Keystone Authentication URL	Enter the authentication URL for the identity service.	Required
CSO user domain name	Enter the CSO user domain name.	Required
CSO domain or project or tenant name	Enter the CSO domain or project or tenant name.	Required
CSO Authentication username	Specify the CSO authentication user name.	Required
CSO Authentication password	Specify the CSO authentication password.	Required. Must be encrypted.
Enable/Disable Secure Connection to CSO Host	Set to 'true' to secure the connection, otherwise set to 'false' to bypass.	Required
SSL Truststore File	If enable_secure_host_connection is set to 'true', then supply the HTTPS trust store filename.	For <b>both on-prem and OCP</b> , provide the JKS name, and the observer will then create the JKS file accordingly.
SSL Truststore Password	If enable_secure_host_connection is set to 'true', then supply a password to decrypt the HTTPS trust store file.	Optional. Must be encrypted.
SSL Validation	Choose whether SSL validation is on or off. Turning SSL validation off will use HTTPS without host verification.	Optional
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

**Encryption requirement:** The Load job requires the password in the configuration file in encrypted form. To encrypt, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

**SSL requirements:** To acquire SSL certificates and build SSL truststores, use the relevant instructions for both on-prem and OCP in the following section: [“Defining observer security” on page 99](#)

## Procedure

1. From the **Observer Configuration UI**, click **Configure** under the Juniper CSO icon, or select an existing job to be edited.
2. Enter or edit the following parameters:
  - Unique ID
  - Juniper CSO Central MS URL
  - CSO Keystone Authentication URL

- CSO Authentication username
  - CSO Authentication password
  - Enable/Disable Secure Connection to CSO's host
  - CSO user domain name
  - CSO domain or project or tenant name
  - SSL Truststore file (optional)
  - SSL Truststore password (optional)
  - SSL Validation (optional)
3. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
  4. Enter an **Observer job description** to explain the purpose of the job in more detail.
  5. Click **Run job** to save your job and begin retrieving information.

## Configuring Kubernetes Observer jobs

Using this observer, you can configure jobs that discover the structure of your Kubernetes clusters, including pods, worker nodes and containers.

### Before you begin

**Important:** The Kubernetes Observer supports Kubernetes version 1.18.

For Kubernetes load jobs, ensure you have the Kubernetes service details to hand, such as the Kubernetes host IP and SSL Certificate details. For Weave Scope listen jobs, first install Weave Scope, and then configure a job using the Weave Scope IP and port parameters.

Existing Load job functionality has been divided into two separate jobs, **Load** and **Local**. Local observations now run on the `jobs/local` endpoint. Existing scripts that use `jobs/local` to trigger a local observation of the Kubernetes environment will need to change the endpoint to `jobs/local`. With the introduction of new local job, the system-health job has been automatically migrated to use the same endpoint by the Kubernetes cron job. Any existing local jobs using the load API that are not automatically migrated, can be seen in the UI, deleted, and created using the correct job type.

The Kubernetes Observer is installed as part of the core installation procedure.

### About this task

The observer reads topology data from Kubernetes through its REST APIs, or Weave Scope.

You can run the following jobs:

#### Load

A transient (one-off) job that loads all requested topology data from a Kubernetes environment.

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

#### Local

Performs a local observation of the Kubernetes REST API for available resources, and loads them in the topology service.

**Not supported for on-premise installation.**

## Weave\_scope Listen

A standalone job that listens to the Weave Scope agent and continues to stream topology and state data to Agile Service Manager.

The Weave Scope listen job provides visibility of your Kubernetes services, pods, containers, deployments, stateful sets, Cron Jobs and processes for a specified namespace.

A long-running job that monitors its source for updates and runs until it is explicitly stopped, or until the Observer is stopped

You must install Weave Scope and then use the Weave Scope Master IP and Node port parameters. For more information on Weave Scope, see the following location: <https://www.weave.works/docs/scope/latest/introducing/>

### For OCP

1. Create Namespace 'weave' with 'ibm-privileged-psp'.

```
kubectl create namespace weave
kubectl -n weave create rolebinding weave-clusterrole-rolebinding --
clusterrole=ibm-privileged-clusterrole --group=system:serviceaccounts:
weave
```

2. Install Weave Scope using the following command:

```
kubectl apply -f "https://cloud.weave.works/k8s/scope.yaml?k8s-service-
type=NodePort&k8s-version=$(kubectl version | base64 | tr -d '\n')"
```

This will result in a port being opened that the Observer can use.

3. You can discover the NodePort using the following command:

```
kubectl -n weave describe service weave-scope-app
```

4. Launch the Weave Scope UI using the following URL:

```
https://<master ip>:<NodePort>
```

Table 40. Kubernetes Observer <b>load</b> job parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job	Required
Encrypted Kubernetes token	The service account token for kubernetes.	Required. Must be encrypted.
Kubernetes Master IP address	Enter the Kubernetes Master IP address.	Required
Kubernetes API port	Enter the Kubernetes API port number.	Required
Trust all certificate by bypassing certificate verification	Enter <b>true</b> if you want to connect to Kubernetes without certificate	Required

Table 40. Kubernetes Observer **load** job parameters (continued)

Parameter	Action	Details
Exact HTTPS certificate file name	If 'Trust all certificate' is set to <b>false</b> , then this parameter is Required.	<b>On-prem</b> To create and store the certificate, see “Configuring password encryption, authentication certificates and keystores” on page 99  <b>OCP</b> To obtain the authentication certificate using OpenSSL and store them as secrets, see “Defining observer security” on page 99
data_center	Specify the name of the data center in which the Kubernetes instance is running.	Required
API query timeout (ms)	Specify the Kubernetes REST API query timeout.	Optional. The default is 5000 ms (that is, 5 seconds)
Correlate	If 'true', enables the Event Analytics correlation on the namespace groups.	Optional
Namespace	Specify the Kubernetes namespace.	Optional. If left empty, all namespaces are observed.
Terminated pods	Choose whether terminated pods should be hidden ( <b>true</b> or <b>false</b> ).	Optional. The default is false.
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

**Encryption requirement:** The Load job requires the token to be encrypted. You encrypt the Kubernetes token using the `encrypt_password.sh` script in the `$ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password, which you enter in the **Encrypted Kubernetes token** field when configuring the Load job.

**SSL certificate requirement:** The Load job requires an SSL Certificate, and for it to be in a specific location:

1. Get the kubernetes master IP and its API port using:

```
kubectl cluster-info
```

2. Run the following OpenSSL command:

```
echo -n | openssl s_client -connect {master ip}:{api} | sed -ne
'/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > ./certificate_file_name.crt
```

The certificate is saved as `certificate_file_name.crt`

3. Copy the certificate file to the `$ASM_HOME/security` directory.

- When configuring the Load job, enter the certificate file name in the **Exact HTTPS certificate file name** field.

Table 41. Kubernetes Observer <b>weave_scope</b> job parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job	Required
Host	Enter the Weave Scope host name (or IP address) of the web socket to be observed.	Required
Port	Enter the Weave Scope port number of the web socket to be observed.	Required
Cluster Name	Enter the name of the cluster or data center to be observed.	Required
Namespaces	Enter a list of namespaces to be observed.	Optional. If left empty, all namespaces will be observed.
Resource types	Select the Weave Scope resource types to observe.	Optional.
Resources to exclude	List resources to be excluded by ID, label, rank or namespace.	Optional. Containers named 'pod' are excluded by default.
Observer job description	Enter additional information to describe the job.	Optional

## Procedure

- From the **Observer Configuration UI**, click **Configure** under the Kubernetes icon, or select an existing Kubernetes job to be edited.
- Choose either **load**, **local** or **weave\_scope** from the job type drop-down.

### Configure the load or local jobs

- Enter or edit the following **required** parameters:
  - Unique ID
  - Encrypted Kubernetes token
  - Kubernetes Master IP address
  - Kubernetes API port
  - Exact HTTPS certificate file name
  - data\_center
- Enter or edit the following **optional** parameters:
  - Namespaces
  - API query timeout (ms)
  - Terminated pods
- Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
- Enter an **Observer job description** to explain the purpose of the job in more detail.

### Configure the weave\_scope job

7. Enter the following **required** parameters:

- Unique ID
- Host
- Port

**Tip:** The NodePort can be obtained using the following command:

```
kubectl -n weave describe service weave-scope-app
```

- Cluster Name

**Note:** The **host** and **port** parameter fields must be empty.

8. Enter the following **optional** parameters:

- Namespaces

**Tip:** Run the following command in the Kubernetes environment to get a list of namespaces:

```
kubectl get namespaces
```

- Resource types
- Resources to exclude
- Observer job description

9. Click **Run job** to save your job and begin retrieving information.

## Configuring Network Discovery Observer jobs

The Network Discovery Observer job retrieves network topology data, including discovered devices and relationships, from the Network Discovery database via REST API, and uses this data to create topologies within the Agile Service Manager topology service.

### Before you begin

The Network Discovery Services are installed as part of the core installation procedure. This includes the Network Discovery Observer, Network Discovery Collector and Network Discovery Engine, which have been installed and should be running, as well as the required scripts to manage jobs.

Before you configure a network discovery job, you must configure the discovery details. This procedure is described here: [“Configuring the network discovery services” on page 85](#)

When successful, the configuration details will be saved in topology management artifacts.

**Known issue and workaround:** The Kafka service may not yet be fully initialized when the nasm-net-disco-schema-registry is started during an Agile Service Manager start up or restart, and it may subsequently fail a consul health check.

### Workaround

1. Verify via consul that the nasm-net-disco-schema-registry has failed its healthcheck:

```
https://localhost:8501/ui/nasm_net_disco/services
```

2. Restart the nasm-net-disco-schema-registry:

```
docker restart nasm-net-disco-schema-registry
```

3. Wait until all Agile Service Manager services have registered with consul, and only then run the network discovery.



## About this task

You define and start the following job. You must edit the parameters in the configuration file before running this job.

### Load job

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

The Network Discovery Observer receives network data from the network services. Once the discovery configuration has successfully been sent to the topology management artifact and Network Discovery engine, you can run a full load job to discover Network Discovery data.

Table 42. Network Discovery Observer job parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job	Required
Config ID	Specify the Network Discovery configuration ID	Required
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

## Procedure

1. From the **Observer Configuration UI**, click **Configure** under the NetDisco icon, or select an existing job to be edited.
2. Enter or edit the following parameters:
  - Unique ID
  - Config ID
3. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
4. Enter an **Observer job description** to explain the purpose of the job in more detail.
5. Click **Run job** to save your job and begin retrieving information.

## Results

The Network Discovery Observer job discovers and then uploads the network services data to the Agile Service Manager topology service.

## What to do next

**Tip:** Stopping a job using `DELETE /jobs/{id}` will stop the discovery first and then the job. It will take a while for the discovery engine and observer job to stop. If you want to resubmit another job, ensure that both the discovery and job have stopped by checking `GET /jobs/load/status` and `GET /jobs/{id}`.

### Related tasks

[“Configuring the network discovery services” on page 85](#)

The Agile Service Manager network discovery service is installed together with the other core components. Once you have configured it, Agile Service Manager can discover network data.

[“Configuring EMS and SSH File Transfer Protocol \(SFTP\) Network Connectivity for Collector Discovery on OCP” on page 92](#)

You set up SSH File Transfer Protocol for use by the network discovery service.

[“Defining Network Discovery Observer jobs” on page 217](#)

You use the Network Discovery Observer to retrieve network topology data from the Network Discovery database via REST API and use this data to create topologies within the topology service.

#### **Related reference**

[“Network discovery service configuration template” on page 94](#)

The following network discovery configuration template is provided (in json format).

## **Configuring New Relic Observer jobs**

Use New Relic Observer when you have a New Relic account with a New Relic Infrastructure subscription. Using New Relic Observer, you can configure jobs that dynamically load New Relic Infrastructure resource data via New Relic for analysis by Netcool Agile Service Manager.

### **Before you begin**

**Important:** The New Relic Observer supports the cloud/SaaS New Relic version.

Ensure you have the New Relic account and New Relic Infrastructure subscription details to hand, such as the account name, account ID, and New Relic Insights API query key.

The New Relic Observer is installed as part of the core installation procedure.

**Restriction:** New Relic applies a 1000 results limit on all New Relic Query Language (NRQL) queries. To accommodate this limit when retrieving data from the SystemSample, StorageSample, ProcessSample and NetworkSample event tables, the New Relic Observer uses the following NRQL query time clause:

```
"SINCE 4 hours ago LIMIT 1000"
```

### **About this task**

The Observer uses the New Relic Infrastructure subscription and makes active New Relic Query Language (NRQL) calls over REST to New Relic Insights to download New Relic Infrastructure resource data.

The New Relic Observer loads the following New Relic Infrastructure resources and their relationships to the Agile Service Manager core topology service:

- Host
- Storage
- OS
- Network Interfaces
- Processes

The New Relic Observer job extracts New Relic Infrastructure resources from New Relic using New Relic Query Language (NRQL) over REST. The observer loads and updates the resources and their relationships within the Agile Service Manager core topology service.

You configure the following job.

#### **Load**

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

Table 43. New Relic job parameters

Parameter	Action	Details
Unique ID	Enter a unique name for the job.	Required
New Relic Name	Specify the New Relic account name or tenant name.	Required
New Relic account ID	Specify the New Relic account ID.	Required. For more information, see <a href="#">account ID tip</a>
New Relic Insights Query API key	Specify the New Relic Insights Query API key.	Required. Must be encrypted. For more information, see <a href="#">query API key tip</a>
filterCriteria	Extend the result set returned to Agile Service Manager.	Optional. The default value is 'SINCE 4 hours ago LIMIT 1000'. For more information, see the documentation for New Relic Query Language.
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

**Tip (New Relic account ID):** To obtain the account ID, first log into the New Relic login page: <https://login.newrelic.com/login> and then obtain the account ID from this URL: <https://rpm.newrelic.com/accounts/<accountId>>

**Tip (New Relic Insights Query API key):** A new Relic user with a new Relic Infrastructure subscription is required to generate a new Relic Insights query API Key as outlined here: <https://docs.newrelic.com/docs/insights/insights-api/get-data/query-insights-event-data-api>

#### Encryption requirement

The load job requires the `insightsQueryAPIKey` in encrypted form. To encrypt the `insightsQueryAPIKey`, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

#### Procedure

1. From the **Observer Configuration UI**, click **Configure** under the New Relic icon, or select an existing New Relic job to be edited.
2. Configure (at least) the following parameters:
  - Unique ID
  - New Relic account name or tenant name
  - New Relic account ID
  - New Relic Insights Query API Key (must be encrypted)
3. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
4. Enter an **Observer job description** to explain the purpose of the job in more detail.

5. Click **Run job** to save your job and begin retrieving information.

## Results

This job loads all requested topology data. Run this job whenever you need New Relic topology data refreshed.

## Configuring OpenStack Observer jobs

Using the OpenStack Observer, you can configure jobs that dynamically load OpenStack data for analysis by Agile Service Manager.

### Before you begin

**Important:** The OpenStack Observer supports the on-premise OpenStack version Stein.

Ensure you have the OpenStack service details to hand, such as the parameters for its APIs or RabbitMQ message bus. If you are configuring a query job, have OpenStack location and authorisation details to hand. If you are configuring a rabbitmq job, you must also identify and provide access to the RabbitMQ message bus.

### OpenStack installation requirements:

If you have installed OpenStack using DevStack, you must add the code specified here to the end of the `local.conf` file, and reinstall it. If you have installed OpenStack using another installation method, you must add the code specified here to the `nova.conf` file, and then restart the Nova (compute) service.

#### If you have already installed OpenStack using DevStack

Add the following code to the end of the `local.conf` file, and then reinstall OpenStack.

#### If you are planning to install OpenStack using DevStack

Add the following code to the end of the `local.conf` file before installation.

```
[[post-config|$NOVA_CONF]]
[DEFAULT]
notification_topics = notifications,com.ibm.asm.obs.nova.notify
notification_driver=messagingv2
notify_on_state_change=vm_and_task_state
notify_on_any_change=True
```

### For standard (or any other) OpenStack installations

Add the following code under the `[DEFAULT]` section of the `nova.conf` file, and then restart the Nova (compute) service.

```
notification_topics = notifications,com.ibm.asm.obs.nova.notify
notification_driver=messagingv2
notify_on_state_change=vm_and_task_state
notify_on_any_change=True
```

The OpenStack Observer is installed as part of the core installation procedure.

**Note:** OpenStack uses RBAC-based protection of its API by defining policy rules based on an RBAC approach. Availability of resources retrieved by the observer is also governed by the same policy. For example, a VM created in project A by users with the admin role may only be available to other users with the same admin role. This can be configured or modified according to user requirements in the OpenStack's policy configuration.



**Trouble:** A **Certificate Chaining Error** can occur when launching an OpenStack Observer job. See the following troubleshooting topic for more information: [“OpenStack Observer certificate chaining error” on page 346](#)

## About this task

The OpenStack Observer jobs extract OpenStack resources via REST or RabbitMQ. The Observer loads and updates the resources and their relationships within the Netcool Agile Service Manager core topology service.

You configure and run the following two jobs.

### Restapi Load job

A transient (one-off) job that loads all requested topology data from the OpenStack instance by REST API.

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

The job loads baseline topology data through the following OpenStack's APIs:

- Keystone (identity)
- Cinder (block storage)
- Glance (image)
- Heat (orchestration)
- Neutron (network)
- Nova (compute)

**Restriction:** An OpenStack environment that has a list of endpoints whereby the 'heat-cfn' service comes first (before the 'heat' service) will encounter a JSON parsing error recorded in the logs due to a known issue in the openstack4j library. When this happens, the full load for the heat service will be skipped entirely. Other services will run as normal.

### Rabbitmq Listen job

A long-running job that reads messages on OpenStack's RabbitMQ message bus for activity from the Cinder (block storage), Heat (orchestration), Neutron (network) and Nova (compute) components continually, until it is explicitly stopped, or until the Observer is stopped.

The rabbitmq job should only be run after an initial restapi job has been completed.

**Restriction:** Only one rabbitmq job should be listening to one queue (or sets of queues) at any one time. If you need to listen to multiple projects, then separate queues must be set up in OpenStack, with appropriate names, before separate listening jobs are submitted for each. For example, for Nova via the **rmq\_nova\_notify** attribute, for Neutron via the **rmq\_neutron\_notify** attribute.

Table 44. OpenStack Observer <b>restapi</b> job parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job	Required
OpenStack authentication type	Specify the OpenStack connection authentication technique to use.	Required. Choose either <b>V2_Tenant</b> , <b>V3_Unscoped</b> , <b>V3_Project</b> , <b>V3_Domain</b> , or <b>V3_ProjectDomain</b> .
OpenStack password	Specify the OpenStack password with which to authenticate.	Required. Must be encrypted.
OpenStack identity endpoint	Specify the authentication URL.	Required. Must include the port and version.

Table 44. OpenStack Observer *restapi* job parameters (continued)

Parameter	Action	Details
Data center name	Specify the name of the data center in which the OpenStack instance is running.	Required. If more than one OpenStack instance is run, and duplicate project or tenant names exist, you must disambiguate them here.
OpenStack username	Specify the OpenStack user name to connect as (or to).	Required
OpenStack project name	Specify the OpenStack project name.	Required
OpenStack user domain name	Specify the user domain name to use, for example the value of OS_USER_DOMAIN_NAME	Optional
OpenStack project domain name	Specify the project domain name to use, for example the value of OS_PROJECT_DOMAIN_NAME	Optional
OpenStack region name	Specify the OpenStack region.	Optional
OpenStack perspective	Select the URL perspective the API accesses data from.	Optional. Choose from <b>Admin</b> , <b>Public</b> , and <b>Internal</b> .
Connection and read timeout (ms)	Choose the timeout setting for the connection and read actions.	Optional. The default is 5000 (5 seconds).
SSL Verification	Choose whether to use SSL verification (true or false). If false, HTTPS is used, but <b>without</b> hostname validation.	Optional
OpenStack host certificate	Specify a certificate name to load into the trust store.	<b>On-prem</b> To create and store the certificate, see <a href="#">“Configuring password encryption, authentication certificates and keystores”</a> on page 99  <b>OCP</b> To obtain the authentication certificate using OpenSSL and store them as secrets, see <a href="#">“Defining observer security”</a> on page 99
SSL truststore file name	Specify a truststore file name.	For <b>both on-prem and OCP</b> , provide the JKS name, and the observer will then create the JKS file accordingly.
SSL truststore file password	Specify a truststore file password.	Required. Must be encrypted.
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

Table 45. OpenStack Observer **rabbitmq** job parameters

Parameter	Action	Details
Unique ID	Enter a unique name for the job	Required
RabbitMQ username	Specify the AMQP user name to connect to the broker.	Required
RabbitMQ password	Specify the password to use to connect to the broker.	Required. Must be encrypted.
RabbitMQ hosts	Enter a (comma-separated) list of hosts in the RabbitMQ cluster.	Required. Add a comma-separated list of host(s) (for example, host1:5672, host2:5672) to connect to in the RabbitMQ cluster. The first host connected to successfully is used.
Data center name	Specify the name of the data center in which the OpenStack instance is running.	Required. If more than one OpenStack instance is run, and duplicate project or tenant names exist, you must disambiguate them here.
OpenStack username	Specify the OpenStack user name to connect as (or to).	Required
OpenStack project name	Specify the OpenStack project.	Optional
RabbitMQ virtual host name	Specify the virtual host to connect to the broker.	Optional
Use SSL?	Choose whether to use an SSL connection.	Optional. Choose <b>true</b> or <b>false</b> . For RabbitMQ, you must choose <b>true</b> .
Nova v2 Oslo message queue	Specify the Nova v2 Oslo message queue.	Optional
Neutron v2 Oslo message queue	Specify the Neutron v2 Oslo message queue.	Optional
Cinder v2 Oslo message queue	Specify the Cinder v2 Oslo message queue.	Optional
Heat v2 Oslo message queue	Specify the Heat v2 Oslo message queue.	Optional
Number of consumer instances	Specify the number of consumer instances to create for each API queue type.	Optional
Observer job description	Enter additional information to describe the job.	Optional

**Important:** You **must** specify the following properties consistently for both the restapi and rabbitmq jobs:

- Data center name
- OpenStack project name
- OpenStack username

**Encryption requirement:**

The restapi and rabbitmq jobs require passwords in the configuration file to be encrypted. To encrypt the OpenStack or RabbitMQ passwords, run the `encrypt_password.sh` script in the `$ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password, for example:

```
2IuExvqz5SGnGgROYGLAQg==
```

**SSL requirements:** To acquire SSL certificates and build SSL truststores, use the relevant instructions for both on-prem and OCP in the following section: [“Defining observer security” on page 99](#)

## Procedure

1. From the **Observer Configuration UI**, click **Configure** under the OpenStack icon, or select an existing OpenStack job to be edited.
2. Choose either **restapi** or **rabbitmq** from the job type drop-down.

### Configure the OpenStack Observer restapi job

3. Enter or edit the following **required** parameters:
  - Unique ID
  - OpenStack authentication type
  - OpenStack password (must be encrypted)
  - OpenStack identity endpoint
  - Data Center name
  - OpenStack username
4. Enter or edit the following **optional** parameters:
  - OpenStack project name
  - OpenStack user domain name
  - OpenStack project domain name
  - OpenStack region name
  - OpenStack perspective
  - Connection and read timeout (ms)
  - SSL Verification
5. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
6. Enter an **Observer job description** to explain the purpose of the job in more detail.

### Configure the OpenStack Observer rabbitmq job

7. Enter or edit the following parameters:
  - Unique ID
  - RabbitMQ username
  - RabbitMQ password (must be encrypted)
  - RabbitMQ hosts
  - Data center name
  - OpenStack username



8. Enter or edit the following **optional** parameters:

- OpenStack project name
- RabbitMQ virtual host name
- Use SSL?
- Nova v2 Oslo message queue
- Neutron v2 Oslo message queue
- Cinder v2 Oslo message queue
- Heat v2 Oslo message queue
- Number of consumer instances
- Observer job description

9. Click **Run job** to save your job and begin retrieving information.

## Configuring Rancher Observer jobs

Using the Rancher Observer, you can define a full load job that will read data from Rancher through the REST API and generate a topology.

### Before you begin

**Important:** The Rancher Observer supports the on-premise Rancher version 2.3.2 and API version 3.0.

The Rancher Observer is installed as part of the core installation procedure.

Ensure you have the Rancher details to hand, such as the Rancher host IP, accessKey, secretKey and clusterId, before running the observer job.

### About this task

The Rancher Observer supports a transient (one-off) Load job that loads all requested topology data via Rancher REST API to build the topology, and then exit.

You define and start the following job. You must edit the parameters in the configuration file before running this job.

#### Full Topology Upload job

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

Table 46. Rancher Observer parameters		
Parameter	Action	Details
Rancher Host IP	Enter the Rancher host IP address.	Required
Rancher Access Key ID	Supply the Rancher access key.	Required
Rancher Secret Access Key	Supply the Rancher secret key.	Required. Must be encrypted.
Rancher Cluster ID	Enter the cluster ID to be discovered	Required
Enable/Disable Secure Host Connection to Rancher's Host	Set to 'true' to secure the connection. Alternatively, set to 'false' to bypass secure host connection.	Required

Table 46. Rancher Observer parameters (continued)

Parameter	Action	Details
Certificate	If you have set the <b>enable_secure_host_connection</b> parameter to <b>true</b> , you must specify a certificate by name to load into the trustStore.	<b>On-prem</b> To create and store the certificate, see “Configuring password encryption, authentication certificates and keystores” on page 99  <b>OCP</b> To obtain the authentication certificate using OpenSSL and store them as secrets, see “Defining observer security” on page 99
SSL TrustStore File	If you have set the <b>enable_secure_host_connection</b> parameter to <b>true</b> , you must specify the TrustStore filename.	For <b>both on-prem and OCP</b> , provide the JKS name, and the observer will then create the JKS file accordingly.
Password SSL TrustStore File	If <b>enable_secure_host_connection</b> is set to 'true', set the password to decrypt the HTTPS trust store file. You <b>must</b> supply the SSL TrustStore Password in encrypted form.  <b>To encrypt the password on premise</b> Run the encrypt_password.sh script in the ASM_HOME/bin directory.  <b>To encrypt the password on OCP</b> Run the relevant <b>kubect1</b> command, as described <a href="#">here</a> .	Optional
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

**Encryption requirement:** The Load job requires the password and secret key in the configuration file in encrypted form. To encrypt the password, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

## Procedure

1. From the **Observer Configuration UI**, click **Configure** under the Rancher icon, or select an existing job to be edited.
2. Enter or edit the following parameters:

- Rancher Host IP
  - Rancher Access Key ID
  - Rancher Secret Access Key
  - Rancher Cluster ID
  - Enable/Disable Secure Host Connection to Rancher's Host
  - Certificate
  - SSL TrustStore File
  - Password SSL TrustStore File
3. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
  4. Enter an **Observer job description** to explain the purpose of the job in more detail.
  5. Click **Run job** to save your job and begin retrieving information.

## Results

The job gathers information and updates the topology.

## What to do next

See the related links for information on using rules for this observer.

### Related tasks

[“Defining rules” on page 309](#)

Rules help streamline topologies and conserve system resources, for example by merging different observer records of the same resource into a single composite resource, or by excluding specific changes from being recorded against a resource history.

## Configuring REST Observer jobs

Use the REST Observer for obtaining topology data via REST endpoints. This observer is a counterpart to the File Observer.

### Before you begin

The REST (or RESTful) Observer is installed as part of the core installation procedure.

### About this task

The REST Observer passes topology data to Agile Service Manager using a RESTful set of interfaces, which provide REST APIs that enable the following functionality:

- Management of Listen and bulk-replace job types.
- The insert-update (HTTP POST) of resources.
- The insert-update (HTTP POST) of relationships.
- The insert-replacement (HTTP PUT) of resources.
- The deletion (HTTP DELETE) of resources.
- A REST API that supports the deletion (HTTP DELETE) of all relationships of a given type from a specified resource.
- A REST API that supports the deletion (HTTP DELETE) of a specific relationship.

**Restriction:** Resources created via REST can have a provider, but not an observer.

## Benefits

Using the REST Observer rather than the File Observer or Topology Service APIs includes the following benefits:

- The ability to provide data to Agile Service Manager via HTTP REST Endpoints instead of files.
- The processing performed by all observers in their framework ensures that meta-data about observations from observers is managed correctly.
- A simple way of deleting all edges of a given type on a resource or a specific edge instance.

To use the REST Observer, a job request must be issued (HTTP POST) to the observer instance job management APIs before sending data to the Resource and Relationship APIs.

## Listen

A long-running listen job capable of consuming topology data over a long period of time.

A listen job is designed to support scenarios where the input data stream is unpredictable, or there is little or no consistency or versioning of resources within the data stream.

## Synchronize (bulk replace)

A long-running job with the same resource replace semantics as the File Observer.

Bulk-replace jobs are designed to support scenarios where a known set of resources are subject to updates or versioning, and a prior observation about resources is to be replaced with a new one.

This job can provide a new set of resources and relationships and synchronize them to Agile Service Manager, thereby causing any previous data provided by the Observer to be deleted and replaced with the new data.

**Note:** In both cases, a provider job parameter is required to identify the origin of the data being provided to the Observer job.

Once a job request has been successfully submitted, you can start to provide data to the Resource and Relationship APIs on behalf of a given job instance.

The Resource and Relationship APIs may respond with an HTTP 503 Service Unavailable response with a `Retry-After: 10` seconds in the header. This indicates that even though the request against those APIs is valid, the observer has not been able to ascertain that meta-data about the job is held in Agile Service Manager yet; this may be due to, for example, any prevailing conditions in the network that support the Agile Service Manager micro-services.

**Tip:** If such a response is received, try the request again later.

**Important:** Ensure that the body is structured correctly. When posting the body, information included in the body after the closing `}` that matches an opening `{` is ignored, and no error is recorded.

Table 47. REST Observer <b>listen</b> and <b>bulk replace</b> job parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job	Required
Provider	Specify the name of the program or system to provide data.	Required
Observer job description	Enter additional information to describe the job.	Optional

## Procedure

1. From the **Observer Configuration UI**, click **Configure** under the REST icon, or select an existing REST job to be edited.
2. Choose either **listen** or **bulk replace** from the job type drop-down.
3. Configure the following parameters for both **bulk replace** and **listen** jobs:
  - Unique ID

- Provider
4. Enter an **Observer job description** to explain the purpose of the job in more detail.
  5. Click **Run job** to save your job and begin retrieving information.

#### Related information

[Timestamps formats](#)

## Configuring ServiceNow Observer jobs

Using the ServiceNow Observer job, you can retrieve the configuration management database (CMDB) data from ServiceNow via REST API, using basic authentication credentials.

### Before you begin

**Important:** The ServiceNow Observer supports the cloud/SaaS ServiceNow version New York.

Ensure your user account has the `rest_api_explorer` and `web_service_admin` roles. These roles are required to access the resources from ServiceNow. Also, ensure you have the ServiceNow service details to hand, such as username, password, and URL.

The ServiceNow Observer is installed as part of the core installation procedure.

**Tip:** To run a load job using a non-admin account, assign the required roles to the account to provide read privileges for the list of API paths.

1. Sign in to the **ServiceNow** instance using an admin account.
2. Navigate to **User Administration** and select **Users** from the menu.
  - To create a new user, click **New**.
  - To edit an existing user, search and select the user.
3. From the user's information tab, select the **Roles** tab, then click **Edit**.
4. Assign the **cmdb\_read** and **service\_viewer** roles to the user.
5. Click **Save**, then **Update**.

For more detailed descriptions of available roles, navigate to **User Administration** and select the **Roles** page. You can also assign other roles with similar privileges to the non-admin account. To find other roles with similar privileges, navigate to **System Security** and select **Access Control (ACL)**.

### About this task

ServiceNow jobs retrieve configuration management database (CMDB) data from ServiceNow via REST API. The Observer loads and updates the resources and their relationships within the Agile Service Manager core topology service.

The observer discovers the following resources (based on the API path list):

- `/api/now/table/cmdb_ci`
- `/api/now/table/core_company`
- `/api/now/table/cmn_department`
- `/api/now/table/cmn_location`
- `/api/now/table/sys_user`

You define and start the following job.

#### ServiceNow job

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

Run this job whenever you want to refresh the ServiceNow topology data.

Table 48. ServiceNow Observer job parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job	Required
ServiceNow instance	Specify the ServiceNow instance.	Required
ServiceNow username	Specify the ServiceNow username.	Required
ServiceNow password	Specify the ServiceNow password.	Required. Must be encrypted.
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

#### Encryption requirement:

The ServiceNow job requires the password in the configuration file to be encrypted. To encrypt the password, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

#### Procedure

1. From the **Observer Configuration UI**, click **Configure** under the ServiceNow icon, or select an existing ServiceNow job to be edited.
2. Configure the following parameters for the ServiceNow job:
  - Unique ID
  - ServiceNow instance
  - ServiceNow username
  - ServiceNow password (must be encrypted)
3. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
4. Enter an **Observer job description** to explain the purpose of the job in more detail.
5. Click **Run job** to save your job and begin retrieving information.

#### What to do next

Run this job whenever you want to refresh the ServiceNow topology data.

## Configuring TADDM Observer jobs

IBM Tivoli Application Dependency Discovery Manager (TADDM) Observer jobs retrieve network topology data, including discovered applications, their components, configurations and dependencies, from the TADDM database server (running either a IBM DB2 or an Oracle database), and use this data to create topologies within the Agile Service Manager topology service.

## Before you begin

**Important:** The TADDM Observer supports the on-premise TADDM version.

To connect to a TADDM Oracle database, you must place the Oracle JDBC Driver into the \$ASM\_HOME/lib folder, and then restart the observer for it to take effect. You can download the driver from the Oracle website, or copy it from the Oracle server (**not** the Oracle client) from the following location: `../app/oracle/product/Oracle_version/dbhome/jdbc/lib/ojdbc6.jar`

Ensure you have the TADDM Rest API login access details in hand, such as the TADDM API URL, username and password.

The TADDM Observer is installed as part of the core installation procedure.

## About this task

TADDM Observer jobs retrieve topology data using the TADDM REST API. The observer loads and updates the resources and their relationships within the Netcool Agile Service Manager core topology service.

You define and start the following jobs.

### Load job

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

Table 49. TADDM Observer <b>load</b> job parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job	Required
TADDM API URL	Specify the TADDM endpoint to connect to.	Required.
TADDM username	Specify the TADDM user name.	Required
TADDM password	Specify the password for the TADDM user.	Required. Must be encrypted.
TADDM objects to observe	Select one or more options from the drop-down list.	Optional. If none are selected, all supported model objects are retrieved.
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

### Encryption requirement:

The Load job requires passwords in the configuration file to be encrypted. To encrypt the password, run the `encrypt_password.sh` script in the \$ASM\_HOME/bin directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

## Procedure

1. From the **Observer Configuration UI**, click **Configure** under the TADDM icon, or select an existing TADDM job to be edited.
2. Enter or edit the following parameters:

- Unique ID
  - TADDM API URL
  - TADDM username
  - TADDM password (must be encrypted)
  - TADDM objects to observe (optional)
3. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
  4. Enter an **Observer job description** to explain the purpose of the job in more detail.
  5. Click **Run job** to save your job and begin retrieving information.

## Configuring VMware NSX Observer jobs

You configure VMware NSX Observer jobs to dynamically load data from the VMware NSX REST interface.

### Before you begin

You can use the VMware NSX Observer when you have a VMware NSX appliance in your environment.

**Important:** The VMware NSX Observer supports the on-premise VMware NSX version 6.3.

Ensure you have the VMware NSX REST API details to hand, such as the VMware NSX URL, username, password, and SSL trustStore.

The VMware NSX Observer is installed as part of the core installation procedure.

### About this task

The VMware NSX Observer job extracts VMware NSX resource information via REST. The Observer loads and updates the resources and their relationships within the Netcool Agile Service Manager core topology service.

You define and start the following job.

#### VMware NSX Observer job (full topology load)

This job loads a baseline of topology data from an environment which contains a VMware NSX appliance.

Run this job whenever you need VMware NSX topology data refreshed.

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

The VMware NSX Observer loads the following resources and their relationship into the Netcool Agile Service Manager core topology service:

- NSX Appliance
- vCenter Appliance
- NSX Controller
- Edge Router - Logical (Distributed) Router, Edge Service Gateway
- Virtual Machines
- Host
- VNIC



Table 50. VMware NSX Observer job parameters

Parameter	Action	Details
Unique ID	Enter a unique name for the job.	Required
Network Virtualization and Security Platform password	Enter the password to authenticate with.	Required. Must be encrypted.
Network Virtualization and Security Platform API URL	Specify the API URL of the VMware NSX endpoint.	Required. Usually in the following format: <code>https://[hostname or IP address]/api</code>
SSL trustStore file	Specify the trustStore file name.	For <b>both on-prem and OCP</b> , provide the JKS name, and the observer will then create the JKS file accordingly.
SSL trustStore file password	Specify the trustStore password to decrypt the HTTPS trustStore file.	Required.
SSL Validation	Choose whether SSL validation is on or off. Turning SSL validation off will use HTTPS without host verification.	Optional
Connection and read timeout (ms)	Enter the time at which the connection and read actions time out.	Optional. Must be a value greater than 0 (zero), and the default is 5000 (5 seconds).
Data center name	Specify the data center in which the VMware NSX instance runs.	Required.
Network Virtualization and Security Platform username	Specify the username to connect as, or listen to.	Required
Network Virtualization and Security Platform tenant name	Specify the tenant.	Optional.
Network Virtualization and Security Platform certificate	Specify a certificate by name to load into the trustStore.	<p><b>On-prem</b> To create and store the certificate, see <a href="#">“Configuring password encryption, authentication certificates and keystores”</a> on page 99</p> <p><b>OCP</b> To obtain the authentication certificate using OpenSSL and store them as secrets, see <a href="#">“Defining observer security”</a> on page 99</p>
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

**Encryption requirement:** The job requires passwords in encrypted form. To encrypt the VMware NSX password and SSL trustStore file password (password\_ssl\_truststore\_file), run the `encrypt_password.sh` script in the `$ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

## Procedure

### To configure and run VMware NSX Observer jobs

1. From the **Observer Configuration UI**, click **Configure** under the VMware NSX icon, or select an existing VMware NSX job to be edited.
2. Enter or edit the following parameters:
  - Unique ID
  - Network Virtualization and Security Platform password
  - Network Virtualization and Security Platform API URL
  - SSL trustStore file
  - SSL trustStore file password (must be encrypted)
  - SSL Validation (optional)
  - Connection and read timeout (ms) (optional)
  - Data center name
  - Network Virtualization and Security Platform username
  - Network Virtualization and Security Platform tenant name (optional)
  - Network Virtualization and Security Platform certificate (optional)
3. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
4. Enter an **Observer job description** to explain the purpose of the job in more detail.
5. Click **Run job** to save your job and begin retrieving information.

### To acquire VMware NSX SSL certificate and build SSL truststore

6. For **OCP** Agile Service Manager deployments, use the relevant instructions in the following topic: [“Defining observer security” on page 99](#)
7. For **on-prem** Agile Service Manager deployments, use the relevant instructions in the following topic: [Defining VMware NSX Observer jobs \(on-prem\)](#)

## What to do next

Run this job whenever you need VMware NSX topology data refreshed.

## Configuring VMware vCenter Observer jobs

You configure VMware vCenter Observer jobs to dynamically load data from the VMware vCenter REST interface.

### Before you begin

**Important:** The VMware vCenter Observer supports the on-premise VMware vCenter versions 6.5 and 6.7.

Ensure you have the VMware vCenter service details to hand, such as username, password, SSL TrustStore and URL.

The VMware vCenter Observer is installed as part of the core installation procedure.

## About this task

The VMware vCenter Observer job extracts VMware vCenter resource information via REST. The Observer loads and updates the resources and their relationships within the Agile Service Manager core topology service.

You define and start the following job.

### VMware vCenter Observer job (full topology load)

Run this job whenever you need VMware vCenter topology data refreshed.

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

The VMware vCenter Observer loads the following resources and their relationship into the Agile Service Manager core topology service:

- ESXi / ESX Hosts
- Virtual Machines
- VNICS
- Storage

Table 51. VMware vCenter Observer job parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job.	Required
vCenter password	Enter the password to authenticate with.	Required. Must be encrypted.
vCenter API URL	Specify the API URL of the VMware vCenter endpoint (including port and version).	Required. Usually in the following format: <code>https://[hostname or IP address]/rest</code>
HTTPS trustStore file name	Specify the trustStore file name.	For <b>both on-prem and OCP</b> , provide the JKS name, and the observer will then create the JKS file accordingly.
trustStore file password	Specify the trustStore password to decrypt the HTTPS trustStore file.	Required.
Data center name	Specify the data center in which the VMware vCenter instance runs.	Required. If more than one, list them (comma-separated).
vCenter username	Specify the username to connect as, or listen to.	Required

Table 51. VMware vCenter Observer job parameters (continued)

Parameter	Action	Details
vCenter certificate	Specify a certificate by name to load into the trustStore.	<b>On-prem</b> To create and store the certificate, see <a href="#">“Configuring password encryption, authentication certificates and keystores”</a> on page 99  <b>OCP</b> To obtain the authentication certificate using OpenSSL and store them as secrets, see <a href="#">“Defining observer security”</a> on page 99
Certificate validation	Choose whether SSL validation is on or off. Turning SSL validation off will use HTTPS without host verification.	Optional
Connection and read timeout	Enter the time at which the connection and read actions time out.	Optional. Must be a value greater than 0 (zero), and the default is 5000 (5 seconds).
vCenter host's name regex to discover	Specify an exact match or a regular expression match for a host's name in order to discover all its virtual machines.	Optional
Connection retry times	Set the connection retry times.	Optional. The default value is 5.
Retry after delay (milliseconds)	Set the time delay before trying to reconnect (in milliseconds).	Optional. The default value is 1000.
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

**Encryption requirement:** The job requires passwords in encrypted form. To encrypt the VMware vCenter password (vcenter\_password) and SSL trustStore file password (password\_ssl\_truststore\_file), run the `encrypt_password.sh` script in the `$ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

## Procedure

### To configure and run VMware vCenter Observer jobs

1. From the **Observer Configuration UI**, click **Configure** under the VMware vCenter icon, or select an existing VMware vCenter job to be edited.
2. Enter or edit the following parameters:
  - Unique ID
  - vCenter password
  - vCenter API URL
  - HTTPS trustStore file name

- trustStore file password (must be encrypted)
  - Data center name
  - vCenter username
  - vCenter certificate (optional)
  - Certificate validation (optional)
  - Connection and read timeout (optional)
3. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
  4. Enter an **Observer job description** to explain the purpose of the job in more detail.
  5. Click **Run job** to save your job and begin retrieving information.
- To acquire VMware vCenter SSL certificate and build SSL truststore**
6. For **OCP** Agile Service Manager deployments, use the relevant instructions in the following topic:  
[“Defining observer security” on page 99](#)
  7. For **on-prem** Agile Service Manager deployments, use the relevant instructions in the following topic:  
[Defining VMware vCenter Observer jobs \(on-prem\)](#)

## What to do next

Run this job whenever you need VMware vCenter topology data refreshed.

## Configuring Zabbix Observer jobs

Using the Zabbix Observer functionality, you can load monitored servers and their associated network resources, and then visualize this data as a topology view in the Agile Service Manager UI. It is installed as part of the core installation procedure.

### Before you begin

The Zabbix Observer supports Zabbix Version 4.0.3.

Ensure you have the Zabbix server details to hand, such as the username, password and SSL TrustStore.

### About this task

A Zabbix Observer job extracts server information and its associated network resources from Zabbix via REST RPC. The Observer loads and updates the resources and their relationships within the Netcool Agile Service Manager core topology service.

You define and start the following job.

#### Full Topology Upload job

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

Table 52. Zabbix Observer parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job.	Required
The datacenter of zabbix server	Specify the data center in which the Zabbix instance runs.	Required. If more than one, list them (comma-separated).

Table 52. Zabbix Observer parameters (continued)

Parameter	Action	Details
The hostname of zabbix server	Enter the Zabbix virtual hostname.	Required
The username of zabbix server	Specify the Zabbix username.	Required
The password of zabbix server	Specify the Zabbix user password.	Required. Must be encrypted.
Zabbix ssl certificate	Specify a certificate file name.	<p><b>On-prem</b> To create and store the certificate, see <a href="#">“Configuring password encryption, authentication certificates and keystores”</a> on page 99</p> <p><b>OCP</b> To obtain the authentication certificate using OpenSSL and store them as secrets, see <a href="#">“Defining observer security”</a> on page 99</p>
SSL Validation	Choose whether to use SSL validation ('true' or 'false'). If false, HTTPS is used, but without hostname validation.	Optional
HTTPS trustStore file name	Specify the trustStore file name.	For <b>both on-prem and OCP</b> , provide the JKS name, and the observer will then create the JKS file accordingly.
trustStore file password	Specify the trustStore password to decrypt the HTTPS trustStore file.	Required. Must be encrypted.
Zabbix connection timeout (milliseconds) (optional)	Timeout, in ms, when querying the Zabbix REST API.	Optional. Default is 5000 (5s).
Job schedule	Specify when the job runs.	Optional. Load jobs only.
Observer job description	Enter additional information to describe the job.	Optional

## Procedure

- From the **Observer Configuration UI**, click **Configure** under the Zabbix icon, or select an existing Zabbix job to be edited.
- Configure the following parameters for the Zabbix job:
  - Unique ID
  - The datacenter of zabbix server
  - The hostname of zabbix server
  - The username of zabbix server
  - The password of zabbix server (must be encrypted)
  - Zabbix ssl certificate (optional)

- SSL Validation (optional)
  - HTTPS trustStore file name
  - trustStore file password (must be encrypted)
  - Zabbix connection timeout (milliseconds) (optional)
3. Define a **Job schedule** (for Load jobs only) by setting the time when the job should run, and whether it should run at regular intervals. By default, the job runs immediately, and only once.  
Optionally, you can specify a future date and time for the job to run, and then set it to run at regular intervals after that, if required. The run intervals must be at least 90 seconds apart, and if you set them at less than 15 minutes, a warning is displayed, as the frequency can impact system performance.
  4. Enter an **Observer job description** to explain the purpose of the job in more detail.
  5. Click **Run job** to save your job and begin retrieving information.

#### To acquire Zabbix SSL certificate and build SSL truststore

6. For **OCF** Agile Service Manager deployments, use the relevant instructions in the following topic: [“Defining observer security” on page 99](#)
7. For **on-prem** Agile Service Manager deployments, use the relevant instructions in the following topic: [Defining Zabbix Observer jobs \(on-prem\)](#)

## Defining observer jobs

Before observers can load data, you must first define and then run observer jobs. This section describes how to **manually** configure, schedule and run observers.

**Remember:** It is recommended that you use the Observer Configuration UI to create and run observer jobs, instead of editing job configuration files manually, as described here. However, to schedule jobs and configure trust stores and certificates, you can use the information in this section.

All prerequisites are deployed during the Agile Service Manager core installation. This includes the docker containers for the observers, which should be installed and running, as well as the required scripts to manage jobs.

You can verify that the observer docker containers are running using the following command:

```
/opt/ibm/netcool/asm/bin/asm_status.sh
```

The system will return text showing all running containers in a state of Up.

**SSL requirements:** To acquire SSL certificates and build SSL truststores, use the relevant instructions for both on-prem and OCP in the following section: [“Defining observer security” on page 99](#)

Observer jobs are configured and run from the Observer Configuration UI, and can be long-running or transient. For example, the Network Manager Observer topology 'load' job is a one-off, transient job (unless scheduled to run at set intervals), while the Network Manager 'listen' job is long-running and runs until explicitly stopped, or until the observer is stopped.

In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining ALM Observer jobs

Using the Agile Lifecycle Manager Observer, you can define jobs that dynamically load data associated with intent from the Agile Lifecycle Manager for analysis by Netcool Agile Service Manager.

### Before you begin

**Important:** The ALM Observer supports IBM Agile Lifecycle Manager Version 2.2.

Ensure you have the Agile Lifecycle Manager Kafka server host and topics to hand, such as the Agile Lifecycle Manager server, the Kafka port, and the topics used for lifecycle events.

**Important:** To access Agile Lifecycle Manager remotely, you must ensure that the Agile Lifecycle Manager installation has been configured with the **KAFKA\_ADVERTISED\_HOST\_NAME** so as to allow remote connections. For more information, see the Configuration reference topic in the Agile Lifecycle Manager Knowledge center at the following location: [https://www.ibm.com/support/knowledgecenter/SS8HQ3\\_2.2.0/GettingStarted/r\\_alm\\_quickreference.html](https://www.ibm.com/support/knowledgecenter/SS8HQ3_2.2.0/GettingStarted/r_alm_quickreference.html)

The Agile Lifecycle Manager Observer is installed as part of the core installation procedure.

### About this task

The Agile Lifecycle Manager Observer jobs listen to the Kafka 'state change' topics of Agile Lifecycle Manager, as well as the Agile Lifecycle Manager Resource Manager. Information is extracted from Agile Lifecycle Manager about Assemblies and Resources and a topology is created.

#### **alm\_observer\_common.sh**

The configuration file you use to customize the listening job for the Agile Lifecycle Manager lifecycle events topic.

The parameters defined here are then used by the `alm_observer_listen_start.sh` script to trigger the ALM Observer job.

#### **alm\_observer\_common\_rm.sh**

The configuration file you use to customize the listening job for the Agile Lifecycle Manager Resource Manager lifecycle events topic.

The parameters defined here are then used by the `alm_observer_listen_start_rm.sh` script to trigger the ALM Observer job.

After installation, you define and start the following two jobs. You must edit the parameters in the configuration file before running these jobs.

#### **Listener for Agile Lifecycle Manager lifecycle events**

A long-running job that monitors its source for updates and runs until it is explicitly stopped, or until the observer is stopped.

This job is started by the `alm_observer_listen_start.sh` script.

#### **Listener for Agile Lifecycle Manager Resource Manager lifecycle events**

A long-running job that monitors its source for updates and runs until it is explicitly stopped, or until the observer is stopped.

This job is started by the `alm_observer_listen_start_rm.sh` script.

**Remember:** Swagger documentation for the observer is available at the following default location: `https://<your host>/1.0/alm-observer/swagger`

### Procedure

1. Edit (at least) the following parameters in the `alm_observer_common.sh` configuration file:

#### **connection**

The host and port of the Agile Lifecycle Manager Kafka server.

2. Edit (at least) the following parameters in the `alm_observer_common_rm.sh` configuration file:



**topic**

The Kafka topic for the Agile Lifecycle Manager Resource Manager lifecycle events.

**connection**

The host and port of the Agile Lifecycle Manager Kafka server.

**Note:** The value of the **almInstall** parameter needs to be the same for both jobs to allow for the topology to be combined.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the configuration file settings.

3. To start the Agile Lifecycle Manager Observer **Listener for Agile Lifecycle Manager lifecycle events** job, use the following command:

```
$ASM_HOME/bin/alm_observer_load_start_rm.sh
```

The Listener for Agile Lifecycle Manager lifecycle events job monitors its source for updates and runs until it is stopped, or until the observer is stopped.

4. To start the ALM Observer **Listener for Agile Lifecycle Manager Resource Manager lifecycle events** job, use the following command:

```
$ASM_HOME/bin/itm_observer_listen_start_rm.sh
```

The Listener job monitors its source for updates and runs until it is stopped, or until the Observer is stopped.

## What to do next

You can also use the following scripts:

**alm\_observer\_listen\_stop.sh**

This script stops the Listener job for Agile Lifecycle Manager lifecycle events.

**alm\_observer\_listen\_stop\_rm.sh**

This script stops the Listener job for Agile Lifecycle Manager Resource Manager lifecycle events.

**alm\_observer\_job\_list.sh**

This script lists the current job status.

**alm\_observer\_log\_level.sh**

This script sets the log level.

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining Ansible AWX Observer jobs

The Ansible AWX observer supports load jobs. When a load job is run, it will gather and read data from an Ansible AWX instance of your choice. When a load job is run, it will load a baseline topology data view from Ansible AWX through its REST APIs.

### Before you begin

**Important:** The Ansible AWX Observer supports the on-premise Ansible AWX version 11.1.0 and Tower API version 3.6.1.

The Ansible AWX Observer is installed as part of the core installation procedure.

Ensure you have the Ansible AWX details to hand, such as the host, port, user and token before defining and running the observer job.

**Required:** Before defining the observer load job, you must obtain a personal access token (PAT) from AWX. The token authentication is recommended for the usage of Tower's API.

1. From the menu bar, navigate to the **Users** page that is under the **Access** section.
2. Select a username, then click the **Token** option.
3. Click on the **+** (plus) icon.
4. Leave the application field blank, and choose **read** for the token's access.

**Remember:** Swagger documentation for the observer is available at the following default location:  
`https://<your host>/1.0/ansibleawx-observer/swagger`

## About this task

The Ansible AWX Observer imports ITSM Resource Topology Service data to Agile Service Manager.

### **ansibleawx\_observer\_common.sh**

The configuration file you use to customize Ansible AWX Observer settings.

The parameters defined here are then used by `ansibleawx_observer_load_start.sh` to trigger the Ansible AWX Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the configuration file settings.

You define and start the following job. You must edit the parameters in the configuration file before running these jobs.

### **Full Topology Upload job**

A transient (one-off) job that loads all requested topology data.

This job is started by the `ansibleawx_observer_load_start.sh` script.

## Procedure

### **To edit the parameters in the configuration file**

1. Open the `ansibleawx_observer_common.sh` configuration file and edit at least the following parameters:

#### **ansibleawx\_host**

The hostname of the Ansible AWX instance

#### **ansibleawx\_port**

The port of the Ansible AWX instance

#### **ansibleawx\_user**

The owner of the token

#### **ansibleawx\_token**

The bearer token in encrypted format

#### **ansibleawx\_sslTruststoreFile**

The Ansible AWX truststore file in jks format

#### **ansibleawx\_passwordSslTruststoreFile**

The Ansible AWX truststore file's password in encrypted format

**Encryption requirement:** The Load job requires the password in the configuration file in encrypted form. To encrypt, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

**SSL requirements:** To acquire SSL certificates and build SSL truststores, use the relevant instructions for both on-prem and OCP in the following section: [“Defining observer security” on page 99](#)

### **To start the Load job**

2. To start the Ansible AWX Observer Full Topology Upload job, use the following command:

```
$ASM_HOME/bin/ansibleawx_observer_load_start.sh
```

## Results

The load job loads all requested topology data, and runs only once. Run this job whenever you need Ansible AWX topology data refreshed.

## What to do next

You can also use the following scripts:

**ansibleawx\_observer\_load\_stop.sh**

Stops the Load job

**ansibleawx\_observer\_job\_list.sh**

Lists the status of current jobs

**ansibleawx\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining AppDynamics Observer jobs

Using the AppDynamics Observer, you can define a full load job that will read data from the AppDynamics Controller via the REST API. This job can provide, for example, business applications, nodes, and tiers to Agile Service Manager.

### Before you begin

**Important:** The AppDynamics Observer supports the cloud/SaaS AppDynamics version 4.5.12 and API version 4.5.x.

Ensure you have the AppDynamics details to hand, such as the instance, account, username and password before running the observer job.

The AppDynamics Observer is installed as part of the core installation procedure.

**Remember:** Swagger documentation for the observer is available at the following default location: <https://<your host>/1.0/appdynamics-observer/swagger>

**Tip:** Before defining the observer load job, you must create an AppDynamics user with the correct permissions. This is required for REST API authentication.

1. On the AppDynamics **Administration** page, click the gear icon (top right).
2. Under the Admin group, select **Administration**, then select the **Users** tab.
3. Select the **AppDynamics** option from the **Display users** drop-down list.
4. Assign all available roles to the user by selecting **Add from Roles > Select All > Done**, then click **Save**.

### About this task

The AppDynamics Observer imports ITSM Resource Topology Service data to Agile Service Manager.

**appdynamics\_observer\_common.sh**

The configuration file you use to customize AppDynamics Observer settings.

The parameters defined here are then used by the `appdynamics_observer_load_start.sh` script to trigger the AppDynamics Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the configuration file settings.

You define and start the following job. You must edit the parameters in the configuration file before running this job.

### Full Topology Upload job

A transient (one-off) job that loads all requested topology data.

This job is started by the `appdynamics_observer_load_start.sh` script and loads all supported resources.

## Procedure

### To edit the parameters in the configuration file

1. Open the `appdynamics_observer_common.sh` configuration file and edit (at least) the following parameters:

**instance**

The instance of the AppDynamics

**account**

The tenant account of the instance

**username**

The user of the specified tenant

**password**

The password associated with the username in encrypted form

**Encryption requirement:**

The Load job requires the password in the configuration file in encrypted form. To encrypt the password, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return the encrypted password.

### To start the Load job

2. To start the AppDynamics Observer Full Topology Upload job, use the following command:

```
$ASM_HOME/bin/appdynamics_observer_load_start.sh
```

## Results

This job loads all requested topology data, and runs only once. Run this job whenever you need the AppDynamics topology data refreshed.

## What to do next

You can also use the following scripts:

**appdynamics\_observer\_load\_stop.sh**

Stops the Load job

**appdynamics\_observer\_job\_list.sh**

Lists the status of current jobs

**appdynamics\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its

logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining Application Discovery Observer jobs

You use the Application Discovery Observer to retrieve network topology data from the Application Discovery database via REST API, and then use this data to create topologies within the topology service.

### Before you begin

The Application Discovery services are installed as part of the core installation procedure. This includes the Application Discovery Observer, Application Discovery engines and Db2 docker containers, which have been installed and should be running, as well as the required scripts to manage jobs.

Verify the Application Discovery deployment before configuring the observer jobs.

1. Use the following command to verify that the docker containers are running:

```
ASM_HOME/bin/docker-compose ps
```

The system should return output indicating that the following containers have a state of Up and therefore are running:

```
...
asm_appdisco-observer_1
asm_app-disco-discovery_1
asm_app-disco-primarystorage_1
asm_db2_1
```

2. Check that storage and discovery services are running and have started:

```
$ docker exec -ti asm_app-disco-primarystorage_1 /opt/IBM/taddm/dist/bin/control status
```

**Tip:** You can also check the following logs:

```
ASM_HOME/logs/db2/create_database.log
```

```
ASM_HOME/logs/app-disco-pss/pss/wlp.log
```

**Remember:** Swagger documentation for the observer is available at the following default location:  
<https://<your host>/1.0/appdisco-observer/swagger>

### About this task

To use the Application Discovery component, you must complete two steps:

#### Configure the discovery information

Configure Application Discovery details. When successful, the configuration details will be saved in topology management artifacts.

This procedure is described here: [“Configuring application discovery service” on page 61](#)

#### Configure the full load job

Triggers Application Discovery, and discovers data.

This task is described in the following procedure.

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

The Application Discovery Observer retrieves application data from the database of Application Discovery via REST APIs. Once the discovery configuration has been successfully sent to the topology management artifact and Application Discovery, you can run a full load job to discover Application Discovery data.

**Note:** You can only run **one** job per profile.

## Procedure

1. Edit the configuration file as required.

See the following topic for more details: [“Configuring application discovery service” on page 61](#)

2. Open the `appdisco_observer_common.sh` configuration file and edit (at least) the following parameters:

### **unique\_id**

The unique ID of the job

### **config\_id**

Set the Application Discovery configuration ID.

The value must be the same value as defined for the `configId` property in the `commonconfig.json` file.

### **active\_discovery**

Set to 'true' to trigger Application Discovery **and** Application Discovery Observer discovery.

Set to 'false' to trigger Application Discovery Observer discovery only.

### **filter\_date**

Optional.

The date indicates the range of past resources to be discovered.

Leave empty for snapshot discovery.

The date must be in `dd-mm-yyyy HH:mm:ss` format.

3. To start the Application Discovery Observer **Load** job, use one of the following commands:

- This command triggers the job to load all requested topology data.

```
$ASM_HOME/bin/appdisco_observer_load_start.sh
```

- Alternatively, you can start the load job using the following cURL command:

```
curl -insecure -X POST --header 'Content-Type: application/json'
--header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-
a73a79c71255'
--header 'Authorization: Basic YXNtOmFzbQ==' -d '{
  "unique_id": "my job",
  "type": "string",
  "description": "job description",
  "parameters": {
    "config_id": "default-appdisco",
    "active_discovery": true,
    "filter_date": "dd-mm-yyyy HH:mm:ss"
  },
  "https://localhost/1.0/Application Discovery-observer/jobs/load"
```

## Results

The Application Discovery Observer job discovers and then uploads the application services data to the Agile Service Manager topology service.

## What to do next

You can also use the following scripts:

### **appdisco\_observer\_load\_stop.sh**

Stops the load job

**appdisco\_observer\_job\_list.sh**

Lists the status of current jobs

**appdisco\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

**Related tasks**

[“Configuring application discovery service” on page 61](#)

The Agile Service Manager application discovery service is installed together with the other core components. Once you have provided the discovery details, Agile Service Manager can discover application data.

[“Installing Db2 Community Edition on OCP” on page 70](#)

To deploy the Agile Service Manager application discovery service on OCP, you install Db2 Community Edition into an OpenShift cluster.

**Related reference**

[“Application discovery service configuration template” on page 65](#)

The following application discovery configuration template is provided (in json format).

[“Supported application discovery accessListTypes and properties” on page 68](#)

The following accessListTypes and their properties are supported by the Agile Service Manager application discovery service, and must be defined in the application discovery configuration file.

## Defining AWS Observer jobs

Using the AWS Observer, you can define jobs that read services data from the Amazon Web Services (AWS) through AWS SDK and then generate a topology. It is installed as part of the core installation procedure.

### Before you begin

**Important:** The AWS Observer supports the cloud/SaaS AWS version 1.11.

Ensure you have the AWS details to hand, such as AWS Region, Access Key ID and Access Secret Key.

**Remember:** Swagger documentation for the observer is available at the following default location:  
`https://<your host>/1.0/aws-observer/swagger`

### About this task

The AWS Observer supports multiple Amazon web services such as EC2 for its 'elastic compute' services.

**aws\_observer\_common.sh**

The configuration file you use to customize AWS Observer settings.

The parameters defined here are then used by the `aws_observer_load_start.sh` to trigger the AWS Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the configuration file settings.

You define and start the following job. You must edit the parameters in the configuration file before running this job.

**Full Topology Upload job**

A transient (one-off) job that loads all requested topology data.

This job is started by the `aws_observer_load_start.sh` script.

**Required:** In order for the AWS Observer to access the Amazon Web Services (AWS) account automatically, the **accessKey**, **secretKey** and **region** parameters are required.

- The access key and the secret access key are not the standard user name and password, but are special tokens that allow the services to communicate with the AWS account by making secure REST or Query protocol requests to the AWS service API.
- The region is the geographical location, for example **US East (Ohio)**, **Asia Pacific (Hong Kong)**, or **EU (London)**.

**Note:** The Full Topology Upload job also supports multi-region full loads, as well as properties filtering, but only via Swagger and UI.

- If you wish to discover more than one region, you will need to comma-separate each region, for example `["eu-east-1","ap-south-1"]`.
- If you wish to exclude more than one property, you will need to separate each property by comma, for example `"kernelId,state-code,keyName"`.

Multi-region full load and properties filtering are not supported in the `aws_observer_load_start.sh` script.

## Procedure

### To find your Access Key and Secret Access Key:

1. Log in to your AWS Management Console.
2. Click on your user name at the top right of the page.
3. Click on the **Security Credentials** link from the drop-down menu.
4. Find the **Access Credentials** section, and copy the latest Access Key ID.
5. Click on the **Show link** in the same row, and copy the Secret Access Key.

### To find the region

6. Check the region at the following location:  
<https://docs.aws.amazon.com/general/latest/gr/rande.html>

### To edit the parameters in the configuration file

7. Open the `aws_observer_common.sh` configuration file and edit (at least) the following Load parameters:

#### **accessKey**

AWS access key

#### **secretKey**

AWS secret key

#### **region**

AWS region to discover

#### **Proxy**

Enables proxy use

#### **Encryption requirement:**

The Load job requires the **secretKey** in the configuration file in encrypted form. To encrypt them, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the secret key. The encryption utility will return an encrypted **secretKey**.

### To start the Load job

8. To start the AWS Observer Full Topology Upload job, use the following command:



```
$ASM_HOME/bin/aws_observer_load_start.sh
```

## Results

This job loads all requested topology data, and runs only once. Run this job whenever you need AWS topology data refreshed.

## What to do next

You can also use the following scripts:

### **aws\_observer\_load\_stop.sh**

Stops the Load job

### **aws\_observer\_job\_list.sh**

Lists the status of current jobs

### **aws\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining Azure Observer jobs

Using the Azure Observer, you can define a full load job that will read data from Azure cloud services through its REST APIs and generate a topology.

### Before you begin

**Important:** The Azure Observer supports the cloud/SaaS Azure version.

Ensure you have the Azure details to hand, such as the Tenant ID, Client ID, and client password before running the observer job.

**Remember:** Swagger documentation for the observer is available at the following default location:  
[https://<your\\_host>/1.0/azure-observer/swagger](https://<your_host>/1.0/azure-observer/swagger)

The observer is installed as part of the core installation procedure.

### About this task

The Azure Observer retrieves topology data from Azure cloud services via REST APIs exposed by the Azure API server.

#### **azure\_observer\_common.sh**

The configuration file you use to customize Azure Observer settings.

The parameters defined here are then used by the `azure_observer_load_start.sh` script to trigger the Azure Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the configuration file settings.

You define and start the following job. You must edit the parameters in the configuration file before running this job.

#### **Full Topology Upload job**

A transient (one-off) job that loads all requested topology data.

This job is started by the `azure_observer_load_start.sh` script and loads all supported resources.

## Procedure

### To edit the parameters in the configuration file

1. Open the `azure_observer_common.sh` configuration file and edit (at least) the following parameters:

**data\_center**

The data center running the Azure server

**tenant\_id**

The tenant id associated with tenant

**client\_id**

The client/application id Azure generated when the application was registered

**client\_secret**

The client/application password associated with the client id in encrypted form

### Encryption requirement:

The Load job requires the `client_secret` in the configuration file in encrypted form. To encrypt the `client_secret`, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return the encrypted password.

### To start the Load job

2. To start the Azure Observer Full Topology Upload job, use the following command:

```
$ASM_HOME/bin/azure_observer_load_start.sh
```

## Results

This job loads all requested topology data, and runs only once. Run this job whenever you need the Azure topology data refreshed.

## What to do next

You can also use the following scripts:

**azure\_observer\_load\_stop.sh**

Stops the Load job

**azure\_observer\_job\_list.sh**

Lists the status of current jobs

**azure\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining BigFix Inventory Observer jobs

The Bigfix Inventory Observer is installed as part of the core installation procedure. Using the Bigfix Inventory Observer, you can define jobs that dynamically load Bigfix inventory data for analysis by Netcool Agile Service Manager.

### Before you begin

**Important:** The BigFix Inventory Observer supports the on-premise BigFix Inventory version 9.5.

Ensure you have the Bigfix Inventory service details to hand, such as API token, SSL TrustStore and URL.

**Remember:** Swagger documentation for the observer is available at the following default location:  
`https://<your host>/1.0/bigfixinventory-observer/swagger`

### About this task

The Bigfix Inventory Observer jobs extract Bigfix Inventory resources via REST. The Observer loads and updates the resources and their relationships within the Netcool Agile Service Manager core topology service.

#### **bigfixinventory\_observer\_common.sh**

The configuration file you use to customize Bigfix Inventory Observer settings.

The parameters defined here are then used by the `bigfixinventory_observer_load_start.sh` to trigger the Bigfix Inventory Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the configuration file settings.

You define and start the following job. You must edit the parameters in the configuration file before running this job.

#### **Full Topology Upload job**

A transient (one-off) job that loads all requested topology data.

This job is started by the `bigfixinventory_observer_load_start.sh` script.

## Procedure

### To edit the parameters in the configuration file

1. Open the `bigfixinventory_observer_common.sh` configuration file and edit (at least) the following Load parameters:

#### **instance\_url**

BigFix Inventory instance URL of the BigFix Inventory endpoint (including port)

Usually in the following format: `https://<hostname or IP address>:<port>`

#### **api\_token**

Bigfix Inventory API token

#### **resources**

Bigfix Inventory resources to discover.

Values are 'software', 'hardware', or '\*' (for both)

#### **truststore\_file**

Bigfix Inventory SSL trust store file for HTTPS authentication

#### **truststore\_password**

Password to decrypt and encrypt Bigfix Inventory SSL trust store file

#### **data\_center**

The data center(s) in which the Bigfix Inventory instance runs.

**Note:** The default data center specified in `bigfixinventory_observer_common.sh` is sufficient for a default deployment. If you require more than one custom data centers, enter them as a comma-separated list.

#### Encryption requirement:

The Load job requires the **api\_token** and **truststore\_password** in the configuration file in encrypted form. To encrypt them, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the token and password. The encryption utility will return an encrypted **api\_token** and **truststore\_password**.

#### To start the Load job

2. To start the Bigfix Inventory Observer Full Topology Upload job, use the following command:

```
$ASM_HOME/bin/bigfixinventory_observer_load_start.sh
```

## Results

This job loads all requested topology data, and runs only once. Run this job whenever you need Bigfix Inventory topology data refreshed.

## What to do next

You can also use the following scripts:

#### **bigfixinventory\_observer\_load\_stop.sh**

Stops the Load job

#### **bigfixinventory\_observer\_job\_list.sh**

Lists the status of current jobs

#### **bigfixinventory\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining Ciena Blue Planet Observer jobs

Using the Ciena Blue Planet Observer, you can define jobs that will gather and read all topology data from the Blue Planet MCP instance by REST API and generate a topology.

### Before you begin

**Important:** The Ciena Blue Planet Observer supports MCP 4.0.

The Ciena Blue Planet Observer is installed as part of the core installation procedure.

Ensure you have the Ciena Blue Planet details to hand, such as API username, API password, MCP URL, MCP certificate, truststore file and truststore password.

**Remember:** Swagger documentation for the observer is available at the following default location:  
`https://<your host>/1.0/cienablueplanet-observer/swagger`

## About this task

The Ciena Blue Planet Observer has two jobs, the restapi load and websocket listen jobs.

- When a load job is run, it loads baseline topology data through Blue Planet MCP APIs: Network Elements (constructs), EquipmentHolder, Equipment, TPE (Terminating Point Encapsulation), and FRE (Forwarding Relationship Encapsulation).
- When a listening job is run, the observer listens to changes in resource's status from the BluePlanet MCP instance through a websocket connection.

**Tip:** Defining observer jobs using the UI is the same for both on-premise and OCP.

### **cienablueplanet\_observer\_common.sh**

The configuration file you use to customize Ciena Blue Planet Observer settings.

The parameters defined here are then used by `cienablueplanet_observer_load_start.sh` and `cienablueplanet_observer_listen_start.sh` to trigger the Ciena Blue Planet Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the configuration file settings.

You define and start the following two jobs. You must edit the parameters in the configuration file before running these jobs.

### **Full Topology Upload job**

A transient (one-off) job that loads all requested topology data.

This job is started by the `cienablueplanet_observer_load_start.sh` script.

### **Listener job**

A long-running job that monitors its source for updates and runs until it is explicitly stopped, or until the Observer is stopped.

This job is started by the `cienablueplanet_observer_listen_start.sh` script.

## Procedure

### **To edit the parameters in the configuration file**

1. Open the `cienablueplanet_observer_common.sh` configuration file and edit the following **Load** parameters:

Table 53. Ciena Blue Planet Observer restapi <b>Load</b> parameters		
Parameter	Action	Details
Unique ID	Enter a unique name for the job.	Required
Server URL	The URL of the MCP server instance	Required
MCP Certificate	Certificate name to load into the trust store	Optional. If used, must be in the <code>/opt/ibm/netcool/asm/security</code> directory.
SSL truststore file	Exact HTTPS trust store file name	Required. The supported format is JKS and the file is relative to <code>\$ASM_HOME/security</code>
SSL truststore password	The password to decrypt HTTPS trust store file	Required. Must be encrypted.
Connection timeout	Sets the connection and read timeout in milliseconds	Optional. Must be a value greater than 0 (zero), and the default is 5000 (5 seconds).
Username	MCP API username	Required

Table 53. Ciena Blue Planet Observer restapi <b>Load</b> parameters (continued)		
Parameter	Action	Details
Password	MCP API password	Required
Data Center	The data center the MCP instance is running in	Required
Tenant Name	The tenant to use	Required

**Encryption requirement:** The jobs requires the **password** and the **ssl\_truststore\_password** in the configuration file in encrypted form. To encrypt them, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and confirm the password. The encryption utility returns passwords in encrypted format.

**SSL requirements:** To acquire SSL certificates and build SSL truststores, use the relevant instructions for both on-prem and OCP in the following section: [“Defining observer security” on page 99](#)

2. Open the `cienablueplanet_observer_listen_start.sh` script and edit the following **Listen** parameters:

Table 54. <b>Additional</b> Ciena Blue Planet Observer Websocket <b>Listen</b> parameters		
Parameter	Action	Details
Websocket URL	The MCP websocket URL to connect to	Required
Websocket Topics Subscriptions	The MCP topics to subscribe to	Required

#### To start the Load and Listen jobs

3. To start the Ciena Blue Planet Observer Full Topology Upload job, use the following command:

```
$ASM_HOME/bin/cienablueplanet_observer_load_start.sh
```

4. To start the Ciena Blue Planet Observer listen job, use the following command:

```
$ASM_HOME/bin/cienablueplanet_observer_listen_start.sh
```

## Results

The load job loads all requested topology data, and runs only once. Run this job whenever you need Ciena Blue Planet topology data refreshed.

The Listener job monitors its source for updates and runs until it is explicitly stopped, or until the observer is stopped.

## What to do next

You can also use the following scripts:

#### **cienablueplanet\_observer\_load\_stop.sh**

Stops the Load job

#### **cienablueplanet\_observer\_listen\_stop.sh**

Stops the Listener job

#### **cienablueplanet\_observer\_job\_list.sh**

Lists the status of current jobs

#### **cienablueplanet\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining Cisco ACI Observer jobs

The Cisco Application Centric Infrastructure (ACI) Observer is installed as part of the core installation procedure. You use the Cisco ACI Observer when you have a Cisco ACI environment with Cisco Application Policy Infrastructure Controller (APIC) in your environment. The Observer interfaces with Cisco APIC and makes active REST calls to Cisco APIC in the Cisco ACI environment. Using the Cisco ACI Observer, you can define jobs that dynamically load Cisco ACI data for analysis by Netcool Agile Service Manager.

### Before you begin

**Important:** The Cisco ACI Observer supports the on-premise Cisco ACI version 4.1.

Ensure you have the Cisco ACI service details to hand, such as the Cisco APIC username, Cisco APIC password, Cisco APIC SSL TrustStore and Cisco APIC URL.

**Remember:** Swagger documentation for the observer is available at the following default location:  
`https://<your host>/1.0/ciscoaci-observer/swagger`

### About this task

A Cisco ACI Observer job extracts Cisco ACI resources from Cisco APIC via REST. The Observer loads and updates the resources and their relationships within the Netcool Agile Service Manager core topology service.

#### **ciscoaci\_observer\_common.sh**

The configuration file you use to customize Cisco ACI Observer settings.

The parameters defined here are then used by the `ciscoaci_observer_query_start.sh` script to trigger the Cisco ACI Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the configuration file settings.

You define and start the following jobs. You must edit the parameters in the configuration file before running this job.

#### **Full Topology Upload job**

A transient (one-off) job that loads all requested topology data.

This job is started by the `ciscoaci_observer_query_start.sh` script.

#### **Listener**

A long-running job that monitors its source for updates and runs until it is explicitly stopped, or until the Observer is stopped.

This job is started by the `ciscoaci_observer_listen_start.sh` script.

The Cisco ACI Observer loads the following Cisco ACI objects and their relationships into the Netcool Agile Service Manager core topology service:

#### **Tenant Logical construct**

- (1) `fvTenant`
- (2) `fvAp`  
A policy owner in the virtual fabric
- (3) `fvAEPg`  
A set of requirements for the application-level EPG instance
- (4) `fvAEP`  
Abstract representation of an endpoint profile
- (5) `fvEpP`

An endpoint profile

(6) fvBD  
A bridge domain is a unique layer 2 forwarding domain that contains one or more subnets

(7) fvCtx  
The private layer 3 network context that belongs to a specific tenant or is shared

(8) vzBrCP  
A contract is a logical container for the subjects which relate to the filters that govern the rules for communication between endpoint groups (EPGs)

(9) vzOOBBBrCP  
An out-of-band binary contract profile can only be provided by an out-of-band endpoint group and can only be consumed by the external prefix set

(10) vzSubj  
A subject is a sub-application running behind an endpoint group (for example, an Exchange server). A subject is parented by the contract, which can encapsulate multiple subjects

(11) vzFilter  
A filter policy is a group of resolvable filter entries

(12) fvSubnet  
A subnet defines the IP address range that can be used within the bridge domain

(13) fvRsCons  
The Consumer contract profile information and on Cisco ACI gui the option to create this object is via Consumed Contract. Used to build relationship between fvAEPg and vzBrCP

(14) fvRsBD  
A source relation to the bridge domain associated to this endpoint group. Used to build relationship between fvBD and fvAEPg

(15) fvRsCtx  
A source relation to a private layer 3 network context that either belongs to a specific tenant or is shared. Used to build relationship between fvBD and fvCtx

(16) vzRsSubjFiltAtt  
The filter for the subject of a service contract. Used to build relationship between vzSubj and vzFilter

#### **Fabric Topology**

(1) fabricInst  
A container object for fabric policies

(2) fabricNode  
The root node for the APIC

(3) polUni  
Represents policy definition or resolution universe

(4) firmwareRunning  
Information about leaf or spine switch firmware running on a node

(5) firmwareCtrlrRunning  
Information about each controller firmware that is running

(6) eqptLCSlot  
The slot for the module card

(7) eqptLC  
A line card (IO card) contains IO ports

(8) eqptPsuSlot  
The power supply slot

(9) eqptPsu  
The power supply unit

(10) eqptFtSlot  
A fan tray slot

(11) eqptFan  
The fan in a fan tray

(12) topSystem  
Used to retrieve fabric node Operational State

(13) cnwPhysIf  
The physical interface assigned to the node cluster

(14) l1PhysIf  
The object that represents the Layer 1 physical Ethernet interface information object

(15) mgmtMgmtIf  
The management interface

(16) lldpAdjEp  
The LLDP neighbors, which contains the information regarding the neighbors

(17) eqptRsIoPhysConf  
A source relation to an L1 Ethernet interface. Used to build relationship between l1PhysIf and eqptLC

(18) mgmtRsOoBStNode  
An object which contains management ip address of fabric spine switches and fabric leaf switches

## **Procedure**

### **To edit the parameters in the configuration file**

1. Open the `ciscoaci_observer_common.sh` configuration file and edit (at least) the following parameters:



**ciscoapic\_api\_url**

Cisco APIC REST API endpoint

**ciscoapic\_username**

Cisco APIC user name for REST API

**ciscoapic\_password**

Cisco APIC user password for REST API.

Supply the Cisco APIC user password in encrypted text.

**ciscoapic\_tenant\_name**

Cisco APIC tenant

Set to 'admin' if there is no specific tenant

Set to '' to load Fabric Topology resources

**ssl\_truststore\_file**

Cisco APIC SSL trust store file for HTTPS authentication

JKS is the supported format and the file is relative to \$ASM\_HOME/security

**password\_ssl\_truststore\_file**

Password to decrypt and encrypt Cisco APIC SSL trust store file.

Supply Cisco APIC SSL trust store password in encrypted text.

**Encryption requirement:**

The Load and Listener jobs require passwords in encrypted form. To encrypt the `ciscoapic_password` and `password_ssl_truststore_file`, run the `encrypt_password.sh` script in the `$ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

**To acquire a Cisco APIC SSL certificate and build the SSL truststore**

2. Use the following command to use OpenSSL to connect to Cisco APIC over port 443, and extract a SSL Certificate from Cisco APIC to a `<certificate_file_name>.cert` file.

```
echo -n | openssl s_client -connect {Cisco APIC IPAddress}:443 | sed -ne
'/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > ./ {certificate_file_name}.cert
```

3. Use the following Java keytool command to import the Cisco APIC certificate file into a keystore and encrypt the keystore with a given password.

```
keytool -import -v -trustcacerts -alias {Cisco APIC Hostname}
-file {certificate_file_name}.cert -keystore {keystore file name}
-storepass {your plain text password to encrypt keystore}
```

**Tip:** You will need the following encryption information when editing `ciscoaci_observer_common.sh`

Table 55. Encryption parameters required for <code>ciscoaci_observer_common.sh</code>	
keystore parameter	<code>ciscoaci_observer_common.sh</code> parameter
keystore password	<code>password_ssl_truststore_file</code>
keystore file name	<code>ssl_truststore_file</code>

4. Copy the keystore file (`{keystore file name}`) to the `$ASM_HOME/security` directory to complete the SSL setup.

**To start the Load and Listener jobs**

5. To start the Cisco ACI Observer Full Topology Upload job, use the following command:

```
$ASM_HOME/bin/ciscoaci_observer_query_start.sh
```

This job loads all requested topology data. Run this job whenever you need Cisco ACI topology data refreshed.

6. To start the Cisco ACI Observer Listener job, use the following command:

```
$ASM_HOME/bin/ciscoaci_observer_listen_start.sh
```

This job monitors its source for updates and runs until it is explicitly stopped, or until the Observer is stopped.

## What to do next

You can also use the following scripts:

### **ciscoaci\_observer\_query\_stop.sh**

Stops the Full Topology Upload job

### **ciscoaci\_observer\_listen\_stop.sh**

Stops the Listener job

### **ciscoaci\_observer\_job\_list.sh**

Lists the status of current jobs

### **ciscoaci\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining Contrail Observer jobs

The Contrail observer is installed as part of the core installation procedure. Using the Contrail Observer, you can retrieve topology data from Juniper Network Contrail Release 4.1 via REST APIs exposed by the Contrail API server. This observer is developed against Juniper Network Contrail that integrates with OpenStack orchestration platform (Ubuntu 18.04 + Contrail Cloud - Ocata).

### Before you begin

**Important:** The Contrail Observer supports the on-premise Contrail version 4.1.0.

Ensure you have the Contrail service details to hand, such as the username, password, and URL.

**Remember:** Swagger documentation for the observer is available at the following default location:  
<https://<your host>/1.0/contrail-observer/swagger>

### About this task

Contrail Observer jobs retrieve topology data from Juniper Network Contrail Release 4.1 via REST APIs exposed by the Contrail API server. The observer loads and updates the resources and their relationships within the Netcool Agile Service Manager core topology service.

### **contrail\_observer\_common.sh**

The configuration file you use to customize Contrail Observer settings.

The parameters defined here are then used by the `contrail_observer_load_start.sh` and `contrail_observer_listen_start.sh` scripts to trigger the Contrail Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the configuration file settings.

You define and start the following jobs. You must edit the parameters in the configuration file before running these jobs.

### Load job

A transient (one-off) job that loads all requested topology data.

This job is started by the `contrail_observer_load_start.sh` script and loads all supported resources.

Run this job whenever you need the Contrail topology data refreshed.

### Listener

A long-running job that monitors its source for updates and runs until it is explicitly stopped, or until the Observer is stopped.

This job is started by the `contrail_observer_listen_start.sh` script and loads all supported resources during startup, and listens to RabbitMQ messages from 'vnc\_config.object-update' fanout exchange.

There is no need to run the Load job before running the Listen job, because the Listen job performs a Load job during initialization.

Table 56. Mapping of Contrail object types to Agile Service Manager entity types:

Contrail object types	Agile Service Manager entity types
domain	domain
project	project
bgp-as-a-service	service
bgpvpn	vpn
loadbalancer	loadbalancer
logical-router	router
network-ipam	ipam
service-instance	service
virtual-ip	ipaddress
virtual-machine-interface	networkinterface
virtual-network	network
virtual-router	router
physical-router	router
global-system-config	group
instance-ip	ipaddress
routing-instance	vrf
bgp-router	router
route-target	routetarget

## Procedure

### To edit the parameters in the configuration file

1. Open the `contrail_observer_common.sh` configuration file and edit (at least) the following parameters:

#### **api\_server\_url**

The Contrail API URL on which the Contrail API server is running

**os\_auth\_url**

The authentication URL for the identity service

**os\_user**

OpenStack username

**os\_password**

OpenStack password

**auth**

Optional

The supported authentication type types are 'keystone' and 'none'.

Use 'keystone' as default.

**os\_project\_domain\_name**

Optional

Openstack project domain name

**os\_user\_domain\_name**

Optional

Openstack user domain name

**os\_project\_name**

The OpenStack project name is required for version 3 authentication.

**os\_tenant\_name**

The OpenStack tenant name is required for version 2 authentication.

**os\_identity\_api\_version**

OpenStack Identity API version

Default to 3 = ['1', '2', '3'],

**connect\_read\_timeout\_ms**

Optional

Choose the time at which the connection and read action times out.

**Encryption requirement:**

The Load job require the os\_password in encrypted form. To encrypt the password, run the encrypt\_password.sh script in the ASM\_HOME/bin directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

2. Still in the contrail\_observer\_common.sh configuration file, edit (at least) the following Listen parameters:

**api\_server\_url**

The Contrail API URL on which the Contrail API server is running

**os\_auth\_url**

The authentication URL for the identity service

**os\_user**

OpenStack username

**os\_password**

OpenStack password

**rabbit\_server**

Hostname or IP address of RabbitMQ server

**rabbit\_port**

The port number to use for connection

**rabbit\_user**

The username to authenticate with RabbitMQ

**rabbit\_password**

The encrypted password to use authenticate with RabbitMQ

**auth**

Optional

The supported authentication type types are 'keystone' and 'none'.

Use 'keystone' as default.

**os\_project\_domain\_name**

Optional

Openstack project domain name

**os\_user\_domain\_name**

Optional

Openstack user domain name

**os\_project\_name**

The OpenStack project name is required for version 3 authentication.

**os\_tenant\_name**

The OpenStack tenant name is required for version 2 authentication.

**os\_identity\_api\_version**

OpenStack Identity API version

Default to 3 = ['1', '2', '3'],

**connect\_read\_timeout\_ms**

Optional

Choose the time at which the connection and read action times out.

**Encryption requirement:**

The Listener job requires the `os_password` and `rabbit_password` in encrypted form. To encrypt a password, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

**To start the Load and Listener jobs**

3. To start the Contrail Observer Load job, use the following command:

```
$ASM_HOME/bin/contrail_observer_load_start.sh
```

This job loads all requested topology data. This job runs only once.

4. To start the Contrail Observer Listener job, use the following command:

```
$ASM_HOME/bin/contrail_observer_listen_start.sh
```

This job monitors its source for updates and runs until it is explicitly stopped, or until the observer is stopped.

**What to do next**

You can also use the following scripts:

**contrail\_observer\_load\_stop.sh**

Stops the Load job

**contrail\_observer\_listen\_stop.sh**

Stops the Listener job

**contrail\_observer\_job\_list.sh**

Lists the status of current jobs

### **contrail\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## **Defining DNS Observer jobs**

The DNS Observer is installed as part of the core installation procedure. Using the DNS Observer, you can query internal DNS server performance, and use the returned information on response times and service addresses to create topologies within the topology service. The DNS Observer supports forward and reverse lookup calls, with **recurse** or **no recurse** options.

### **Before you begin**

Ensure you have the DNS access details to hand, such as domain, DNS server address and port number.

**Remember:** Swagger documentation for the observer is available at the following default location:

`https://<your host>/1.0/dns-observer/swagger`

### **About this task**

The DNS Observer (nasm-dns-observer) provides DNS query services and topological insight into how a specified DNS server is performing forward (name-to-IP address) or reverse (IP address-to-name) lookups. Query results include a list of addresses, information on how long it takes the DNS server to resolve a lookup, and, optionally (with the maximum number of recursive calls set at 200) how the DNS server is recursively resolving a given name or IP address.

Job data is automatically posted to the topology service, after which the job status expires, after a set amount of time. The Topology Viewer displays the results with color-coded lines representing the relationships between resources, and the lookup time in ms. A tabular view of the relationship details is also available.

**Tip:** The relationship types can be customized with line color, width and pattern functions. See the [“Creating custom relationship type styles” on page 283](#) topic for more information.

### **dns\_observer\_common.sh**

The configuration file you use to customize DNS Observer settings.

The parameters defined here are then used by the DNS forward and reverse lookup scripts (`dns_observer_forward_lookup_start.sh` and `dns_observer_reverse_lookup_start.sh`) to trigger the DNS Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the configuration file settings.

You define and start the following jobs. You must edit the parameters in the configuration file before running these jobs.

#### **Forward lookup job**

A transient (one-off) job that loads all requested DNS forward lookup topology data.

This job is started by the `dns_observer_forward_lookup_start.sh` script.

#### **Reverse lookup job**

A transient (one-off) job that loads all requested DNS reverse lookup topology data.

This job is started by the `dns_observer_reverse_lookup_start.sh` script.

## Procedure

1. Open the `dns_observer_common.sh` configuration file and edit the required parameters.

### **type**

Values can be either `forward` for name-to-IP address lookups, or `reverse` for IP address-to-name lookups.

### **address\_type**

IPV4 or IPV6

### **server**

The DNS server IP address

### **port**

The DNS server port number

### **recurse**

Values can be `false` to run the job without recursion, or `true` to initiate recursion, with the maximum number of calls set at 200.

### **domain\_name**

The domain name for the DNS forward lookup job

### **ip\_address**

The IP address for the DNS reverse lookup job

## Run the jobs

2. To start the DNS forward lookup job, use the following command:

```
$ASM_HOME/bin/dns_observer_forward_lookup_start.sh
```

3. To start the DNS reverse lookup job, use the following command:

```
$ASM_HOME/bin/dns_observer_reverse_lookup_start.sh
```

## Results

Data retrieved from the DNS query is displayed in the Agile Service Manager Topology Viewer.

## Example

**Example** of a **forward** DNS Observer job **with no** recursive lookup:

```
{
  "unique_id": "my job",
  "type": "forward",
  "parameters": {
    "address_types": "IPV4",
    "server": "8.8.8.8",
    "port": 53,
    "recurse": false,
    "domain_name": "yourdomain.com"
  }
}
```

**Example** of a **reverse** DNS Observer job **with** recursive lookup:

```
{
  "unique_id": "my job",
  "type": "reverse",
  "parameters": {
    "address_types": "IPV4",
    "server": "8.8.8.8",
    "port": 53,
    "recurse": true,
    "ip_address": "8.8.8.8"
  }
}
```

## What to do next

You can also use the following scripts:

### **dns\_observer\_lookup\_stop.sh**

Stops the DNS observer lookup job

### **dns\_observer\_job\_list.sh**

Lists the status of current jobs

### **dns\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining Docker Observer jobs

Using the Docker Observer functionality, you can discover Docker network resources, including Docker Swarm clusters, and then visualize this data as a topology view in the Agile Service Manager UI. In addition, it also discovers Docker clusters managed by Docker UCP.

### Before you begin

**Important:** The Docker Observer supports Docker version 3.1.0.

**Note:** Docker UCP v3.1.0 supports only TLS 1.2 for SSL negotiation and has removed support for TLS 1 and TLS 1.1.

The Docker Observer is installed as part of the core installation procedure.

**Remember:** Swagger documentation for the observer is available at the following default location:  
`https://<your host>/1.0/docker-observer/swagger`

**Update Notes:** If you have updated a previous version of Agile Service Manager with existing Docker Observer job data, you must run a data migration script (as documented in the release notes on-prem update topic) before running new observer jobs.

### About this task

The Docker Observer performs a single load job, which performs a one-off discovery of the Docker network.

The job definition indicates whether to connect to a local Docker on the same (UNIX) host as the observer using the **unix\_socket** parameter, or to a remote Docker using the **host** and **port** parameters.

#### Local docker

The default, if the job parameters are empty, is to try to connect to a UNIX socket at `/var/run/docker.sock`

If the location of the UNIX socket differs, the full path can be given in the **unix\_socket** parameter. The **host** and **port** parameters must **not** be supplied.

In either case, the socket must be accessible.

When the observer is running within the docker container to be monitored, `/var/run/docker.sock` must be available within the container. For example:

```
volumes:
  /var/run/docker.sock:/var/run/docker.sock
```



## Remote docker

The host and port parameters of the job can be used to identify the TCP port that Docker can be reached on. The **unix\_socket** parameter must **not** be supplied.

Docker is **not** accessible via TCP by default. To enable it, edit the `docker.service` file. On RedHat, this is available in `/usr/lib/systemd/system`. Amend the **ExecStart** option under the Service section to include a `-H` option. For example, to make it available externally on port 2375, you could add `-H tcp://0.0.0.0:2375`.

**Note:** If you want to continue to be able to access Docker via the default socket, for example if the Docker Observer container needs access, or if you want to be able to perform `docker ps -a` rather than `docker -H tcp://0.0.0.0:2375 ps -a`, then you need to also list it in the same line, as on the following example:

```
-H tcp://0.0.0.0:2375 -H unix:///var/run/docker.sock
```

You must reload the configuration:

```
sudo systemctl daemon-reload
sudo systemctl restart docker
```

**Tip:** If this fails to start Docker, and a Unix socket (or no socket at all) was specified, check that no directory with that name exists. If you start up docker with just a TCP socket and no Unix socket, this creates a `/var/run/docker.sock` directory, which you must delete after Docker is stopped, so that you can restart with access via that Unix socket.

## docker\_observer\_common.sh

The configuration file you use to customize Docker Observer settings.

The parameters defined here are then used by the `docker_observer_load_start.sh` script to trigger the Docker Observer job.

You can use the **view\_all** and **exclude\_containers** parameters to filter the scope of observations. These parameters are arrays or lists that can accept multiple values.

### view\_all

Use this parameter to force modeling of all containers, tasks and images.

By default, only running containers, running tasks, and images currently in use by modeled containers are modeled.

### exclude\_containers

Use this parameter to filter out containers that are not of interest, based on regular expression matches against the container name.

## Swagger UI usage examples

Using the Docker Observer, you can discover the following Docker resources:

- Remote Docker network resources via HTTP through TCP port exposure.  
Example:

```
{
  "unique_id": "my job",
  "type": "load",
  "parameters": {
    "host": "1.2.3.4",
    "port": 2375
  }
}
```

- Remote Docker network resources with HTTPS using a certificate.  
Example:

```
{
  "unique_id": "my job",
  "type": "load",
  "parameters": {
    "host": "1.2.3.4",

```

```

    "port": 2375,
    "username": "username",
    "password": "password",
    "docker_ssl_certificate": "certificate_file_name.crt",
    "docker_ssl_truststore_file": "truststore_file_name.jks",
    "password_ssl_truststore_file": "truststore_password"
  }
}

```

- Remote Docker network resources with HTTPS using certificate and truststore.

Example:

```

{
  "unique_id": "my job",
  "type": "load",
  "parameters": {
    "host": "1.2.3.4",
    "port": 2375,
    "username": "username",
    "password": "password",
    "docker_ssl_certificate": "certificate_file_name.crt",
    "docker_ssl_truststore_file": "truststore_file_name.jks",
    "password_ssl_truststore_file": "truststore_password"
  }
}

```

- Remote Docker network resources with HTTPS using truststore.

Example:

```

{
  "unique_id": "my job",
  "type": "load",
  "parameters": {
    "host": "1.2.3.4",
    "port": 2375,
    "username": "username",
    "password": "password",
    "docker_ssl_truststore_file": "truststore_file_name.jks",
    "password_ssl_truststore_file": "truststore_password"
  }
}

```

## Procedure

1. Edit the `docker_observer_common.sh` config file as required.

The Docker Observer supports multiple types of Docker configurations. Edit or populate the following parameters for different docker configurations:

### Local Docker

Empty parameters

### Remote Docker with HTTP

Populate **host** and **port**

### Remote Docker with HTTPS via certificate

The certificate will be added to the named truststore

Populate **host**, **port**, **username**, **password**, **docker\_ssl\_certificate**, **docker\_ssl\_truststore\_file** and **password\_ssl\_truststore\_file**

### Remote Docker with HTTPS via truststore

The truststore must contain the certificate

Populate **host**, **port**, **username**, **password**, **docker\_ssl\_truststore\_file** and **password\_ssl\_truststore\_file**

**Encryption requirement:** All jobs require passwords in encrypted form. To encrypt 'password' and 'password\_ssl\_truststore\_file', run the `encrypt_password.sh` script in the `$ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

**SSL requirements:** To acquire SSL certificates and build SSL truststores, use the relevant instructions for both on-prem and OCP in the following section: [“Defining observer security” on page 99](#)

2. To start a Docker Observer **Load** job, use the following command:

```
$ASM_HOME/bin/docker_observer_load_start.sh
```

### Usage examples for starting jobs:

#### Default job

```
$ASM_HOME/bin/docker_observer_load_start.sh
```

#### Local Docker

```
env unique_id=My job name $ASM_HOME/bin/docker_observer_load_start.sh
```

#### Remote Docker with HTTP

```
env unique_id=My job name host=1.2.3.4 port=2375  
$ASM_HOME/bin/docker_observer_load_start.sh
```

#### Remote Docker with HTTPS via certificate

```
env unique_id=My job name host=1.2.3.4 port=2375 username=username  
password=password docker_ssl_certificate=certificate_file_name.crt  
docker_ssl_truststore_file=truststore_file_name.jks  
password_ssl_truststore_file=truststore_password  
$ASM_HOME/bin/docker_observer_load_start.sh
```

#### Remote Docker with HTTPS via truststore

```
env unique_id=My job name host=1.2.3.4 port=2375 username=username  
password=password docker_ssl_truststore_file=truststore_file_name.jks  
password_ssl_truststore_file=truststore_password $ASM_HOME/bin/  
docker_observer_load_start.sh
```

## Results

The script triggers the Docker Observer Load job, which performs a one-off discovery of the Docker network you have specified.

## What to do next

You can also use the following scripts:

### **docker\_observer\_load\_stop.sh**

Stops the job

### **docker\_observer\_job\_list.sh**

Lists the current job status

### **docker\_observer\_log\_level.sh**

Sets the log level

**Tip:** For regular status updates, run the default Docker Observer job via a cron job.

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining Dynatrace Observer jobs

The Dynatrace Observer is installed as part of the core installation procedure. Using the Dynatrace Observer, you can define jobs that dynamically load Dynatrace data for analysis by Netcool Agile Service Manager.

### Before you begin

**Important:** The Dynatrace Observer supports Dynatrace cloud/SaaS and Dynatrace onPrem v1.182.102.

Ensure you have the Dynatrace service details to hand, such as API token and Base URL.

**Remember:** Swagger documentation for the observer is available at the following default location:  
`https://<your host>/1.0/dynatrace-observer/swagger`

### About this task

The Dynatrace Observer jobs extract Dynatrace resources via REST. The Observer loads and updates the resources and their relationships within the Netcool Agile Service Manager core topology service.

#### **`dynatrace_observer_common.sh`**

The configuration file you use to customize Dynatrace Observer settings.

The parameters defined here are then used by the `dynatrace_observer_load_start.sh` to trigger the Dynatrace Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the configuration file settings.

You define and start the following job. You must edit the parameters in the configuration file before running this job.

#### **Full Topology Upload job**

A transient (one-off) job that loads all requested topology data.

This job is started by the `dynatrace_observer_load_start.sh` script.

## Procedure

### To edit the parameters in the configuration file

1. Open the `dynatrace_observer_common.sh` configuration file and edit (at least) the following Load parameters:

#### **`api_token`**

Dynatrace API token

#### **`base_url`**

Dynatrace API base URL

Example of a Dynatrace cloud-based URL:

```
https://<your-environment-id>.live.dynatrace.com/api/v1/
```

Example of a Dynatrace managed or onPrem URL:

```
https://<your-domain>/e/<your-environment-id>/api/v1/
```

#### **Encryption requirement:**

The Load job requires the API token in the configuration file in encrypted form. To encrypt the **`api_token`**, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the API token. The encryption utility will return an encrypted **`api_token`**.

### To start the Load job

2. To start the Dynatrace Observer Full Topology Upload job, use the following command:

```
$ASM_HOME/bin/dynatrace_observer_load_start.sh
```

## Results

This job loads all requested topology data, and runs only once. Run this job whenever you need Dynatrace topology data refreshed.

## What to do next

You can also use the following scripts:

**`dynatrace_observer_load_stop.sh`**

Stops the Load job

**`dynatrace_observer_job_list.sh`**

Lists the status of current jobs

**`dynatrace_observer_log_level.sh`**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining File Observer jobs

Using the File Observer functionality, you can write bespoke data to a file in a specific format, upload this data to the topology service, and then visualize this data as a topology view in the Agile Service Manager UI. The File Observer is installed as part of the core installation procedure.

### Before you begin

**Remember:** Swagger documentation for the observer is available at the following default location:  
<https://<your host>/1.0/file-observer/swagger>

### About this task

File Observer jobs are HTTP POST requests that can be triggered via cURL or swagger, or via the example scripts provided in the `$ASM_HOME/bin` directory.

**`file_observer_common.sh`**

The config file you use to customize the File Observer job `unique_id` or service host.

The parameters defined here are then used by the `file_observer_load_start.sh` script to trigger the File Observer job.

The File Observer runs a 'loadFile' job that loads all requested topology data for each tenant. The loadFile job takes the name of the file to parse and load.

Lines starting with V: (vertex), E: (edge), D: (delete) or W: (wait) are treated as instruction lines to be processed. Other lines, for example lines that are empty or commented out, are ignored.

### Line format

**V:**

The JSON payload takes the format described in the swagger documentation of the POST /resources message body.

**E:**

The JSON payload takes the format described in the swagger documentation for the `_references` section of the `POST /resources` message body.

**W:**

Takes an integer period followed by a string specifying the units.

**D:**

Takes a single string which is the unique ID of the vertex to delete.

**Tip:**

- An example file is available in the `$ASM_HOME/data/file-observer` directory.
- See the related links for more information on available timestamps formats.

**Important:** Ensure that the file is structured correctly. For each line of the file, information included after the closing `}` that matches an opening `{` is ignored, and no error is recorded.

**Restriction:**

Files to be read by File Observer must be located in the following directory: `$ASM_HOME/data/file-observer`

A file name specified in a File Observer job must be relative to that directory (and not absolute).

**Procedure**

1. Edit the `file_observer_common.sh` config file.
2. Define your data file and copy the file to the following location: `/opt/ibm/netcool/asm/data/file-observer/`

For example:

```
cp dncim.file $ASM_HOME/data/file-observer/
```

3. To start the File Observer **Load** job, use one of the following commands:

**To define the data file via a command line argument**

```
./file_observer_load_start.sh --file dncim.file
```

**To define the data file via the environment**

```
env file=dncim.file $ASM_HOME/bin/file_observer_load_start.sh
```

The load job loads all requested topology data from the file specified. This job runs only once.

**Example**

The following cURL command example invokes the File Observer job:

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X POST --header 'Content-Type: application/json' --header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' -d '{ "unique_id": "dncim.file", "type": "load", "parameters": { "file": "dncim.file" } }' https://localhost/1.0/file-observer/jobs
```

**What to do next**

You can also use the following scripts:

**file\_observer\_load\_stop.sh**

Stops the job

**file\_observer\_job\_list.sh**

Lists the current job status

### **file\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

### **Related information**

[Timestamps formats](#)

## **Defining GoogleCloud Observer jobs**

Using the GoogleCloud Observer, you can define a full load job that will read services data from the Google Cloud Platform's Compute Services through Google's Compute Services SDK, and then generate a topology.

### **Before you begin**

**Important:** The Google Cloud Observer supports the cloud/SaaS Google Cloud version.

The GoogleCloud Observer is installed as part of the core installation procedure.

The GoogleCloud Observer supports GoogleCloud's compute services. Ensure you have the GoogleCloud details in hand, such as the Project ID, Service Account Key File and Zone, before running the observer job.

**Remember:** Swagger documentation for the observer is available at the following default location:  
<https://<your host>/1.0/googlecloud-observer/swagger>

### **About this task**

The GoogleCloud Observer supports a transient (one-off) Load job that loads all requested topology data via Google's Compute Services SDK to build the topology, and then exit.

#### **googlecloud\_observer\_common.sh**

The configuration file you use to customize GoogleCloud Observer settings.

The parameters defined here are then used by the `googlecloud_observer_load_start.sh` to trigger the GoogleCloud Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the configuration file settings.

You define and start the following job. You must edit the parameters in the configuration file before running this job.

#### **Full Topology Upload job**

A transient (one-off) job that loads all requested topology data.

This job is started by the `googlecloud_observer_load_start.sh` script.

**Note:** You must create a service account key file or use an existing one to allow the GoogleCloud Observer to discover resources from GoogleCloud.

## **Procedure**

### **To create a service account key file**

1. From the Google Cloud Platform dashboard, under your 'Project ID', go to **APIs and Services** and then choose **Credentials**.

The **Credentials** page is displayed listing a number of authentication methods.

2. Select the **Service account** authentication service
3. From **Create Credentials**, choose **Service account**.
4. Complete the service account details, then click **Create** to create a new service account.
5. Select the **Compute Engine > Compute Admin** role, then click **Continue** and **Done**.
6. To download the .json file from the **Credentials** page, click **Service account > Manage Service Accounts**.
7. Select the new service account, then choose **Create key** from the overflow menu (the three dots under the Actions column).
8. Choose a key type of **JSON**, then click **Create**. A .json file will be downloaded.
  - **For on-prem**, store the .json file under /opt/ibm/netcool/asm/security
  - **For ICP/OCF**, follow [these steps](#) to store the service account key file as a secret.

The filename will be used in the observer parameter (**service\_account\_key\_file**) for the full load job.

#### To edit the parameters in the configuration file

9. Open the `googlecloud_observer_common.sh` configuration file and edit (at least) the following Load parameters:

##### **project\_id**

Google Cloud Platform Compute Service's Project ID

##### **zone**

Google Cloud Platform Compute Service's zone to discover

##### **service\_account\_key\_file**

Google Cloud Platform Compute Service's Service Account Key File

(**For on-prem**) copy the json file to the `$ASM_HOME/security` directory.

#### To start the Load job

10. To start the GoogleCloud Observer Full Topology Upload job, use the following command:

```
$ASM_HOME/bin/googlecloud_observer_load_start.sh
```

## Results

This job loads all requested topology data, and runs only once. Run this job whenever you need GoogleCloud topology data refreshed.



**Trouble:** While the job is running, the status of discovered resources may appear as 'indeterminate' in the topology until the full upload is complete.

## What to do next

You can also use the following scripts:

##### **googlecloud\_observer\_load\_stop.sh**

Stops the Load job

##### **googlecloud\_observer\_job\_list.sh**

Lists the status of current jobs

##### **googlecloud\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.



## Defining IBM Cloud Observer jobs

The IBM Cloud Observer is installed as part of the core installation procedure. Use the IBM Cloud Observer when you have IBM Cloud installed in your environment to define jobs that perform REST calls to the IBM Cloud REST API. These jobs retrieve Cloud Foundry Apps information and services, and then dynamically load the retrieved data for analysis by Netcool Agile Service Manager.

### Before you begin

**Important:** The IBM Cloud Observer supports the cloud/SaaS IBM Cloud version.

**Important:** The IBM Cloud Observer supports Cloud Foundry API version 2.75.

Ensure you have the IBM Cloud access details to hand, such as username, password and region.

**Remember:** Swagger documentation for the observer is available at the following default location:  
`https://<your host>/1.0/ibmcloud-observer/swagger`

### About this task

In a typical IBM Cloud environment, you have access to four different region:

- US\_S (Dallas)
- UK (London)
- EU (Frankfurt)
- AP (Sydney & Tokyo)

You define which region is to be discovered, as IBM Cloud Observer supports all four regions.

Each region has its own URI, thus only a single region is discovered in a full load job. To discover different regions, a full load job needs to be triggered for each region. The prerequisites for a full load job are the IBM Cloud username, password and region.

**Note:** No listening job is supported at the moment.

**Tip:** You can configure IBM Cloud resources via the IBM Cloud GUI or the Cloud Foundry CLI.

#### **`ibmcloud_observer_common.sh`**

The config file you use to customize IBM Cloud Observer settings.

The parameters defined here are then used by the `ibmcloud_observer_load_start.sh` script to trigger the IBM Cloud Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the config file settings.

You define and start the following job. You must edit the parameters in the config file before running this job.

#### **Full Topology Upload job**

A transient (one-off) job that loads all requested topology data.

This job is started by the `ibmcloud_observer_load_start.sh` script.

Table 57. Mapping IBM Cloud model objects to Agile Service Manager entity types	
IBM Cloud resource object	Agile Service Manager entity types
stacks	operatingsystem
apps	application
routes	path
service bindings	hub
service instance	service

Table 57. Mapping IBM Cloud model objects to Agile Service Manager entity types (continued)

IBM Cloud resource object	Agile Service Manager entity types
user provided service instance	service
spaces	group
organization	organization
buildpacks	component

## Procedure

### To edit the parameters in the config file

1. Open the `ibmcloud_observer_common.sh` config file and edit (at least) the following parameters:

#### **username**

The user name for the IBM Cloud REST API

#### **password**

The user password for the IBM Cloud REST API

**Encryption requirement:** Jobs require the password in the configuration file to be encrypted. You encrypt the password using the `encrypt_password.sh` script in the `$ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

#### **region**

The IBM Cloud resource region (supported region codes are US\_S, UK, EU or AP)

### To start the Load job

2. To start the IBM Cloud Observer Full Topology Upload job, use the following command:

```
$ASM_HOME/bin/ibmcloud_observer_load_start.sh
```

This job loads all requested topology data. Run this job whenever you need the IBM Cloud topology data refreshed.

## What to do next

You can also use the following scripts:

#### **ibmcloud\_observer\_load\_stop.sh**

Stops the Full Topology Upload job

#### **ibmcloud\_observer\_job\_list.sh**

Lists the status of current jobs

#### **ibmcloud\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining ITNM Observer jobs

The ITNM Observer is installed as part of the core installation procedure. Using the ITNM Observer, you can define jobs that dynamically load data discovered by IBM Tivoli Network Manager for analysis by Netcool Agile Service Manager.

### Before you begin

**Important:** The ITNM Observer supports the on-premise ITNM version 4.2.

Ensure you have the ITNM service details to hand, such as the ITNM domain, host and port number.

**Remember:** Swagger documentation for the observer is available at the following default location:  
`https://<your host>/1.0/itnm-observer/swagger`

### About this task

The ITNM Observer jobs extract IBM Tivoli Network Manager resources using an Object Query Language JDBC driver. The Observer loads and updates the resources and their relationships within the Netcool Agile Service Manager core topology service.

#### **itnm\_observer\_common.sh**

The config file you use to customize ITNM Observer settings.

The parameters defined here are then used by the `itnm_observer_load_start.sh` and the `itnm_observer_listen_start.sh` scripts to trigger the ITNM Observer jobs.

After installation, you define and start the following two jobs. You must edit the parameters in the config file before running these jobs.

#### **Full Topology Upload**

A transient (one-off) job that loads all requested topology data.

This job is started by the `itnm_observer_load_start.sh` script.

#### **Listener**

A long-running job that monitors its source for updates and runs until it is explicitly stopped, or until the Observer is stopped.

This job is started by the `itnm_observer_listen_start.sh` script.

## Procedure

1. Edit (at least) the following parameters in the `itnm_observer_common.sh` config file:

#### **domain**

Network Manager domain name

#### **host**

Network Manager server

#### **port**

Port used to access the Network Manager `ncp_config` process

**Note:** The value of **port** will vary if multiple domains exist. To determine which port number to use for a Network Manager domain, look for the domain-specific `ncp_config` entry in the `$NCHOME/etc/precision/ServiceData.cfg` file.

#### **exclude\_no\_connection**

If true, only load entities that have connections including their dependencies are included.

#### **topology\_type\_edge\_type\_map**

Map of ITNM topology type to Agile Service Manager edge/relationship type

`{"topologyType": "edgeType"}` in JSON string format.

The default value is `{"ConvergedTopology": "connectedTo"}`.

This parameter accepts the default or custom Agile Service Manager edge types. For custom edge types, register the edge in the topology service swagger under the POST /types/edge API.

The value of topology type can be found in \$NCHOME/precision/disco/stitchers/DNCIM/PopulatedDNCIMTopologies.stch

Alternatively, run the following OQL statement against the model service to list the available topology type:

```
select ENTITYNAME from ncimCache.entityData where METAClass='Topology'
```

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the config file settings.

2. To start the ITNM Observer **Full Topology Upload** job, use the following command:

```
$ASM_HOME/bin/itnm_observer_load_start.sh
```

The Full Topology Upload job loads all requested topology data. This job runs only once.

3. To start the ITNM Observer **Listener** job, use the following command:

```
$ASM_HOME/bin/itnm_observer_listen_start.sh
```

The Listener job monitors its source for updates and runs until it is stopped, or until the Observer is stopped.

## What to do next

You can also use the following scripts:

### **itnm\_observer\_listen\_stop.sh**

Stops the Listener job

### **itnm\_observer\_load\_stop.sh**

Stops the Full Topology Upload job

### **itnm\_observer\_job\_list.sh**

Lists the current job status

### **itnm\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining Jenkins Observer jobs

Using the Jenkins Observer, you can define listen jobs that receive build information generated by the Agile Service Manager plugin for Jenkins.

### Before you begin

Ensure you have the Jenkins details to hand.

**Remember:** Swagger documentation for the observer is available at the following default location: <https://<your host>/1.0/jenkins-observer/swagger>

### About this task

You define and start the following job.

## Listen job

The standalone listen job receives Jenkins build notification data for a specified namespace and processes it as a topology.

The listen job is long-running, and runs until it is explicitly stopped or until the observer is stopped.

## Procedure

### To start the Listen job

1. To start the Jenkins Observer Listen job, use the following examples:

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X POST
--header 'Content-Type: application/json' --header 'Accept: application/json' --header
'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --header 'Authorization:
Basic YXNtOmFzbQ==' -d '{
  "unique_id": "my job",
  "type": "listen",
  "description": "job description",
  "parameters": {
    "jenkins_observation_namespace": "default"
  }
}' 'https://localhost/1.0/jenkins-observer/jobs/listen'
```

2. Verify that the job is running.

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X GET --header 'Accept:
application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255'
'https://localhost/1.0/jenkins-observer/jobs/my%20job'
```

## What to do next

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

### Related tasks

[“Configuring the Jenkins plugin” on page 53](#)

The Agile Service Manager software includes the Jenkins plugin, which you install on your Jenkins server using the Jenkins Plugin Manager Advanced installation wizard. From the Jenkins server, the plugin gathers and sends information to the Jenkins Observer.

[“Refining Jenkins integration and visualization” on page 57](#)

You can extend your Jenkins integration with rules to merge data from different sources, custom topology display conventions, and the use of templates for the automated generation of topologies.

[“Configuring Jenkins Observer jobs” on page 141](#)

Using the Jenkins Observer, you can define listen jobs that receive build information generated by the Agile Service Manager plugin for Jenkins.

### Related information

[“Jenkins Observer troubleshooting” on page 347](#)

## Defining Juniper CSO Observer jobs

Using the Juniper CSO Observer, you can define a full load job that will gather topology data from Juniper CSO. It is installed as part of the core installation procedure.

### Before you begin

**Important:** The Juniper CSO observer supports the on-premise Juniper CSO version 4.1.0.

Ensure you have the Juniper CSO details to hand, such as details of the Juniper CSO API server and its credentials.

**Remember:** Swagger documentation for the observer is available at the following default location:  
<https://<your host>/1.0/junipercso-observer/swagger>

## About this task

The Juniper CSO Observer retrieves topology data from Juniper CSO Release 4.1 via REST APIs exposed by CSO API server.

### **junipercso\_observer\_common.sh**

The configuration file you use to customize Juniper CSO Observer settings.

The parameters defined here are then used by the `junipercso_observer_load_start.sh` script to trigger the Juniper CSO Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the configuration file settings.

You define and start the following job. You must edit the parameters in the configuration file before running this job.

### **Load job**

A transient (one-off) job that loads all requested topology data.

This job is started by the `junipercso_observer_load_start.sh` script and loads all supported resources.

## Procedure

### **To edit the parameters in the configuration file**

1. Open the `junipercso_observer_common.sh` configuration file and edit (at least) the following parameters:

#### **cso\_central\_ms\_url**

CSO host

#### **auth\_url**

CSO authentication URL

#### **username**

CSO username

#### **password**

CSO password

#### **user\_domain\_name**

CSO's domain name

#### **domain\_project\_tenant\_name**

CSO domain or project or tenant to discover

#### **enable\_secure\_host\_connection**

Enable secure CSO host connection.

The default value is False.

#### **ssl\_truststore\_file**

SSL TrustStore file if you enable secure host connection

#### **password\_ssl\_truststore**

The SSL TrustStore password if you enable secure host connection

**Encryption requirement:** The Load job requires the password and password\_ssl\_truststore in the configuration file in encrypted form. To encrypt them, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the passwords. The encryption utility will return encrypted passwords.

**SSL requirements:** To acquire SSL certificates and build SSL truststores, use the relevant instructions for both on-prem and OCP in the following section: [“Defining observer security” on page 99](#)

### To start the Load job

2. To start the Juniper CSO Observer Full Topology Upload job, use the following command:

```
$ASM_HOME/bin/junipercso_observer_load_start.sh
```

This job loads all requested topology data, and runs only once. Run this job whenever you need the Juniper CSO topology data refreshed.

## What to do next

You can also use the following scripts:

### **junipercso\_observer\_load\_stop.sh**

Stops the Load job

### **junipercso\_observer\_job\_list.sh**

Lists the status of current jobs

### **junipercso\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining Kubernetes Observer jobs

The Kubernetes Observer is installed as part of the core installation procedure. Using this observer, you can define jobs that discover the services you run on Kubernetes, and display Kubernetes containers and the relationships between them.

### Before you begin

**Important:** The Kubernetes Observer supports Kubernetes version 1.18.

Ensure you have the Kubernetes service details to hand, such as the Kubernetes host IP and SSL Certificate details.

**Remember:** Swagger documentation for the observer is available at the following default location: `https://<your host>/1.0/kubernetes-observer/swagger`

### About this task

The Kubernetes Observer jobs query Kubernetes and extract information. The observer loads and updates the resources and their relationships within the Agile Service Manager core topology service.

### **kubernetes\_observer\_common.sh**

The configuration file you use to customize Kubernetes Observer settings.

The parameters defined here are then used by the `kubernetes_observer_query_start.sh`, the `kubernetes_observer_poll_start.sh` and the `kubernetes_observer_listen_start.sh` scripts to trigger the Kubernetes Observer jobs.

You must edit the parameters in the configuration file before running these jobs.

#### **Load**

A transient (one-off) job that loads all requested topology data.

This job is started by the `kubernetes_observer_query_start.sh` script.

#### **Poll**

A job that loads all requested topology data like the Load job, but repeated at set polling intervals.

This job is started by the `kubernetes_observer_poll_start.sh` script.

#### **Weave Scope Listen**

A standalone job that listens to the Weave Scope agent and continues to stream topology and state data to Agile Service Manager.

The listen job can maximally provide visibility of your Kubernetes services, pods, containers, deployments, stateful sets, Cron Jobs and processes for a specified namespace.

A long-running job that monitors its source for updates and runs until it is explicitly stopped, or until the Observer is stopped.

This job is started by the `kubernetes_observer_listen_start.sh` script.

### **What to do next**

**Tip:** You can start a job without editing the `kubernetes_observer_common.sh` script by providing a Kubernetes host IP or encrypted token directly, as in the following examples:

```
env kubernetes_token=<eyJhbGciOiJSUzI1NiIsInR5cCI6Ikp>
/opt/ibm/netcool/asm/bin/kubernetes_observer_query_start.sh
```

```
env kubernetes_master_ip=<host ip>
/opt/ibm/netcool/asm/bin/kubernetes_observer_query_start.sh
```

You can also use the following scripts:

#### **kubernetes\_observer\_query\_stop.sh**

Stops the Load job

#### **kubernetes\_observer\_poll\_stop.sh**

Stops the Poll job

#### **kubernetes\_observer\_listen\_stop.sh**

Stops the Weave Scope Listener job

#### **kubernetes\_observer\_job\_list.sh**

Lists the status of current jobs

#### **kubernetes\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.



## To define the full load and poll jobs

### Before you begin

**Required:** Before defining a Kubernetes Observer Load or Poll job, you must create a service account in the Kubernetes environment and obtain its token. The following example uses the default namespace. Edit the command based on the kubectl configuration.

1. Create a configuration file called `asm-k8s-observer.yaml` with the custom cluster role `asm:kubernetes-observer`

#### Use the following sample content

```
kind: ClusterRole
apiVersion: rbac.authorization.k8s.io/v1
metadata:
  annotations:
    rbac.authorization.kubernetes.io/autoupdate: "true"
  name: asm:kubernetes-observer
rules:
- apiGroups: [""]
  resources: ["pods", "namespaces", "nodes", "services"]
  verbs: ["get", "list"]
- apiGroups: ["apps"]
  resources: ["replicasets", "deployments"]
  verbs: ["get", "list"]
```

Run the following command to create the `asm:kubernetes-observer` custom cluster role with 'read' access to the resources that the observer discovers, for example pods, namespaces, and nodes.

```
kubectl create -f asm-k8s-observer.yaml
```

**Tip:** Verify that the cluster role `asm:kubernetes-observer` and its privileges exist using the following commands:

```
kubectl get clusterrole asm:kubernetes-observer
```

```
kubectl describe clusterrole asm:kubernetes-observer
```

**Version 1.1.10:** From Version 1.1.10, the Agile Service Manager Kubernetes Observer requires a new cluster role to discover target systems. Use the following sample content:

```
kind: ClusterRole
apiVersion: rbac.authorization.k8s.io/v1
metadata:
  annotations:
    rbac.authorization.kubernetes.io/autoupdate: "true"
  name: asm:kubernetes-observer
rules:
- apiGroups: [""]
  resources: ["pods", "namespaces", "nodes", "services", "endpoints"]
  verbs: ["get", "list"]
- apiGroups: ["apps"]
  resources: ["replicasets", "deployments", "statefulsets", "daemonsets"]
  verbs: ["get", "list"]
```

#### For a new (fresh) discovery

Create a configuration file called `asm-k8s-observer-v2.yaml` in the target environment using the yaml resource, as provided in the sample content.

Then run the following command:

```
kubectl create -f asm-k8s-observer-v2.yaml
```

Continue to the next step of creating a service account.

### For an existing discovery

Create a configuration file called `k8s-update-cr.yaml` using the yaml resource provided in the sample content.

Run the following command to patch the `clusterrole`:

```
kubectl patch clusterrole asm:kubernetes-observer --patch "$(cat k8s-update-cr.yaml)"
```

2. Create a service account. Assign the service account to the same namespace as `kubectl`.

```
kubectl create serviceaccount asm-k8s-account
```

**Tip:** Verify that the service account exists:

```
kubectl get serviceaccount
```

3. Bind the `asm:kubernetes-observer` role to the `asm-k8s-account` service account. The command binds the role to the service account that is created on the default namespace.

```
kubectl create clusterrolebinding asm-k8s --clusterrole=asm:kubernetes-observer  
--serviceaccount=default:asm-k8s-account
```

4. Obtain the Kubernetes service account token by completing the following steps:

- a. Get all secrets:

```
kubectl get secret
```

- b. Describe the `asm-k8s-account-token-*****` (which in this example is `ch47f`):

```
kubectl describe secret asm-k8s-account-token-ch47f
```

## Procedure

### To edit the parameters in the `kubernetes_observer_common.sh` configuration file

1. Open the `kubernetes_observer_common.sh` configuration file and edit the following parameters:

#### **data\_center**

Data centre running the Kubernetes instance, for example `dataCenter1`.

This parameter is used to ensure that observations of different Kubernetes clusters do not clash.

#### **master\_ip**

Kubernetes host IP

#### **token**

Kubernetes service account token, which must be encrypted.

**Encryption requirement:** The Load job requires the token in the configuration file to be encrypted. You encrypt the Kubernetes token using the `encrypt_password.sh` script in the `$ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

#### **namespace**

Optional parameter

**Without** the namespace parameter, the Kubernetes Observer uploads resources from all namespaces in the Kubernetes environment.

**With** the parameter defined, the Kubernetes observer uploads resources only from the given namespace in the Kubernetes environment.

Kubernetes namespaces are modeled as a group based on the resources discovered in an observation. A namespace group consists of server, pod, deployment and service resources that

are discovered in the namespace. Groups can be queried and viewed on the defined topologies section by searching for a particular namespace. The new **correlate** parameter enables the Event Analytics correlation of the namespace groups.

**Tip:** Run the following command in the Kubernetes environment to get a list of namespaces:

```
kubect1 get namespaces
```

### api\_port

The Kubernetes API Port

**Tip:** Get the Kubernetes master IP and its API port using the following command:

```
kubect1 cluster-info
```

The system returns the following information:

```
Kubernetes master is running at https://{master}:{port}
```

### trust\_all\_certificate

If you set this to **true**, the observer allows connection to the Kubernetes environment without a client certificate.

If you set this to **false**, the observer requires a valid certificate in \$ASM\_HOME/security (which must be provided to ssl\_certificate\_file).

### certificate

Kubernetes SSL certificate file name, which is the name of a file within \$ASM\_HOME/security. The files in the \$ASM\_HOME/security directory are made available in the observer container.

**Tip:** Obtain the SSL Certificate:

- a. Get the kubernetes master IP and its API port using:

```
kubect1 cluster-info
```

- b. Run the following OpenSSL command:

```
echo -n | openssl s_client -connect {master ip}:{api} | sed -ne  
'/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > ./certificate_file_name.crt
```

The certificate is saved as *certificate\_file\_name.crt*

- c. Copy the certificate file to the \$ASM\_HOME/security directory.

### connect\_read\_timeout\_ms

Connection timeout in milliseconds (ms), for example '5000'.

### hide\_terminated\_pods

The default value is 'false'.

- If you set this parameter to 'true', all pods with the phase Succeeded will be excluded from the Topology Viewer.
- If set to 'false', all pods regardless of any phase will be shown in Topology Viewer.

### To start the Load job

2. To start the Kubernetes Observer **Load** job, use the following command:

```
$ASM_HOME/bin/kubernetes_observer_query_start.sh
```

The Load job loads all requested topology data. This job runs only once.

### To start the Poll job

3. To start the Kubernetes Observer **Poll** job, use the following command:

```
$ASM_HOME/bin/kubernetes_observer_poll_start.sh
```

The Poll job loads all requested topology data at defined intervals.

## To define the Weave Scope Listen job

### Before you begin

**Required:** Before defining a Kubernetes Observer Weave Scope Listen job, you must install Weave Scope in your Kubernetes environment. For more information on Weave Scope, see the following location: <https://www.weave.works/docs/scope/latest/introducing/>

#### For OCP

1. Create Namespace 'weave' with 'ibm-privileged-psp'.

```
kubectl create namespace weave
kubectl -n weave create rolebinding weave-clusterrole-rolebinding --
clusterrole=ibm-privileged-clusterrole --group=system:serviceaccounts:
weave
```

2. Install Weave Scope using the following command:

```
kubectl apply -f "https://cloud.weave.works/k8s/scope.yaml?k8s-service-
type=NodePort&k8s-version=$(kubectl version | base64 | tr -d '\n')"
```

This will result in a port being opened that the Observer can use.

3. You can discover the NodePort using the following command:

```
kubectl -n weave describe service weave-scope-app
```

4. Launch the Weave Scope UI using the following URL:

```
https://<master ip>:<NodePort>
```

## Procedure

### To edit the parameters in the `kubernetes_observer_common.sh` configuration file

1. Open the `kubernetes_observer_common.sh` configuration file and edit the following parameters:

#### **cluster\_name**

Data centre running the Kubernetes instance, for example `mycluster`.

#### **host**

Kubernetes host IP

#### **port**

Weave Scope port, that is, NodePort.

**Tip:** The NodePort can be obtained using the following command:

```
kubectl -n weave describe service weave-scope-app
```

#### **namespace**

List of Kubernetes namespaces to listen for.

**Tip:** Run the following command in the Kubernetes environment to get a list of namespaces:

```
kubectl get namespaces
```

#### **topologies**

List of resources to include in the topology.

Available resources:

- containers
- hosts

- kube-controllers
- pods
- processes
- services

#### **exclude\_resources**

List of resources to exclude from the topology.

The default is to exclude containers named 'POD' and kube-system resources.

### **To start the Weave Scope Listen job**

2. To start the Kubernetes Observer **Weave Scope Listen** job, use the following command:

```
$ASM_HOME/bin/kubernetes_observer_listen_start.sh
```

The Listener job monitors its source for updates and runs until it is explicitly stopped, or until the Observer is stopped.

## **What to do next**

**Tip:** You can start a job without editing the `kubernetes_observer_common.sh` script by providing a Kubernetes host IP or encrypted token directly, as in the following examples::

```
env token=<eyJhbGciOiJIUzI1NiIsInR5cCI6Ikp>
$ASM_HOME/bin/kubernetes_observer_query_start.sh
env master_ip=<host ip>
$ASM_HOME/bin/kubernetes_observer_query_start.sh
```

However, when setting any of these parameters in the file or on the command line, all parameters must be valid.

## **Defining Network Discovery Observer jobs**

You use the Network Discovery Observer to retrieve network topology data from the Network Discovery database via REST API and use this data to create topologies within the topology service.

### **Before you begin**

The Network Discovery Services are installed as part of the core installation procedure. This includes the Network Discovery Observer, Network Discovery Collector and Network Discovery Engine, which have been installed and should be running, as well as the required scripts to manage jobs.

All observers have scripts to start and stop all available jobs, to list the status of a current job, to set its logging levels, and to configure its job parameters.

**Network Discovery Sizing:** The Agile Service Manager network discovery service supports a maximum number of discovered entities.

#### **For size0**

10,000 entities

#### **For size01**

1,000,000 entities

The number of entities created by a discovery depend on factors such as the number of devices in scope, and how many interfaces, cards, VLANs and so on are present on each device.

Verify the Network Discovery deployment before configuring the observer jobs.

1. Use the following command to verify that the docker containers are running:

```
ASM_HOME/bin/docker-compose ps
```

The system should return output indicating that the following containers have a state of Up and therefore are running:

asm_netdisco-observer_1	sh -c /opt/ibm/\$APPNAME/st	Up	9152/tcp, 9153/tcp
nasm-net-disco-collector	./entrypoint.sh	Up	8443/tcp, 8445/tcp
nasm-net-disco-config	./entrypoint.sh	Up	8443/tcp
nasm-net-disco-consul	/etc/init.sh agent -server	Up	0.0.0.0:8501->8501/tcp
nasm-net-disco-control	./entrypoint.sh	Up	8443/tcp
nasm-net-disco-schema	/etc/init.sh	Up	8081/tcp,
nasm-net-disco-security	./entrypoint.sh	Up	8443/tcp
nasm-net-disco-sftp	./entrypoint.sh	Up	0.0.0.0:9022->9022/tcp
nasm-net-disco-sidecar	./entrypoint.sh	Up	0.0.0.0:8765->8443/tcp
nasm-net-disco-status	./entrypoint.sh	Up	8443/tcp
nasm-net-disco-swagger	/etc/init.sh	Up	8443/tcp,
nasm-net-disco-swagger-ui	/etc/init.sh	Up	8443/tcp
nasm-net-disco-topogram	./entrypoint.sh	Up	8443/tcp
nasm-net-disco-worker	/bin/sh -c /opt/\$SERVICE_N	Up	

2. Check the status of the storage and discovery services. The status of the network discovery services can be checked in the Consul UI on port 8501, for example:

```
https://localhost:8501/ui/nasm_net_disco/services
```

**Known issue and workaround:** The Kafka service may not yet be fully initialized when the nasm-net-disco-schema-registry is started during an Agile Service Manager start up or restart, and it may subsequently fail a consul health check.

#### Workaround

1. Verify via consul that the nasm-net-disco-schema-registry has failed its healthcheck:

```
https://localhost:8501/ui/nasm_net_disco/services
```

2. Restart the nasm-net-disco-schema-registry:

```
docker restart nasm-net-disco-schema-registry
```

3. Wait until all Agile Service Manager services have registered with consul, and only then run the network discovery.

**Remember:** Swagger documentation for the observer is available at the following default location:  
<https://<your host>/1.0/netdisco-observer/swagger>

## About this task

To use the Network Discovery component, you must complete two steps:

### Configure the discovery process

Configure Network Discovery details (upon success, configuration details will be saved in the topology).

This procedure is described here: [“Configuring the network discovery services” on page 85](#)

### Configure the full load job

Triggers Network Discovery

Loads Network Discovery data

By default, Load jobs are one-off, transient jobs that perform a full upload of all requested topology data as soon as they are triggered.

You can also run these jobs (again) manually from the Observer UI, or schedule them to run at set times when configuring them.

The Network Discovery Observer receives network data from the network services. Once the discovery configuration has successfully been sent to the topology management artifact and Network Discovery engine, you can run a full load job to discover Network Discovery data.

## Procedure

### To edit the parameters in the configuration file

1. Open the `netdisco_observer_common.sh` configuration file and edit (at least) the following parameters:

#### **unique\_id**

The unique ID of the job.

#### **configId**

Set the Network Discovery configuration ID.

This is the same value as defined for the `configId` property in the `commonconfig.json` file.

2. To start the Network Discovery Observer **Load** job, use one of the following commands:

- This command triggers the job to load all requested topology data.

```
$ASM_HOME/bin/netdisco_observer_load_start.sh
```

- Alternatively, you can start the load job using the following cURL command:

```
curl -insecure -X POST --header 'Content-Type: application/json' --header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --header 'Authorization: Basic YXNtOmFzbQ==' -d '{
  "unique_id": "my job",
  "type": "string",
  "description": "job description",
  "parameters": {
    "configId": "default-netdisco"
  }
}' 'https://localhost/1.0/netdisco-observer/jobs/load'
```

## Results

The Network Discovery Observer job discovers and then uploads the network services data to the Agile Service Manager topology service.

## What to do next

You can also use the following scripts:

#### **netdisco\_observer\_load\_stop.sh**

Stops the load job

#### **netdisco\_observer\_job\_list.sh**

Lists the current job status

#### **netdisco\_observer\_log\_level.sh**

Sets the log level

#### **netdisco\_observer\_discovery\_status.sh**

Checks the discovery status

#### **netdisco\_observer\_discovery\_stop.sh**

Stops the discovery

**Important: Do not** run the `netdisco_observer_discovery_stop.sh` script as a standalone script. It is run by the `netdisco_observer_load_stop` script as part of a specific sequence to stop a network discovery before stopping an observer job.

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

**Tip:** Network discovery log are described [here](#).

## Related tasks

[“Configuring the network discovery services” on page 85](#)

The Agile Service Manager network discovery service is installed together with the other core components. Once you have configured it, Agile Service Manager can discover network data.

[“Configuring EMS and SSH File Transfer Protocol \(SFTP\) Network Connectivity for Collector Discovery on OCP” on page 92](#)

You set up SSH File Transfer Protocol for use by the network discovery service.

[“Configuring Network Discovery Observer jobs” on page 148](#)

The Network Discovery Observer job retrieves network topology data, including discovered devices and relationships, from the Network Discovery database via REST API, and uses this data to create topologies within the Agile Service Manager topology service.

## Related reference

[“Network discovery service configuration template” on page 94](#)

The following network discovery configuration template is provided (in json format).

## Defining New Relic Observer jobs

The New Relic Observer is installed as part of the core installation procedure. Use New Relic Observer when you have a New Relic account with a New Relic Infrastructure subscription. Using New Relic Observer, you can define jobs that dynamically load New Relic Infrastructure resource data via New Relic for analysis by Netcool Agile Service Manager.

## Before you begin

**Important:** The New Relic Observer supports the cloud/SaaS New Relic version.

Ensure you have the New Relic account and New Relic Infrastructure subscription details to hand, such as the account name, account ID, and New Relic Insights API query key.

**Remember:** Swagger documentation for the observer is available at the following default location:  
`https://<your host>/1.0/newrelic-observer/swagger`

**Restriction:** New Relic applies a 1000 results limit on all New Relic Query Language (NRQL) queries. To accommodate this limit when retrieving data from the SystemSample, StorageSample, ProcessSample and NetworkSample event tables, the New Relic Observer uses the following NRQL query time clause:

```
"SINCE 4 hours ago LIMIT 1000"
```

## About this task

The Observer uses the New Relic Infrastructure subscription and makes active New Relic Query Language (NRQL) calls over REST to New Relic Insights to download New Relic Infrastructure resource data.

The New Relic Observer loads the following New Relic Infrastructure resources and their relationships to the Agile Service Manager core topology service:

- Host
- Storage
- OS
- Network Interfaces
- Processes

The New Relic Observer job extracts New Relic Infrastructure resources from New Relic using New Relic Query Language (NRQL) over REST. The observer loads and updates the resources and their relationships within the Agile Service Manager core topology service.

### **newrelic\_observer\_common.sh**

The configuration file you use to customize New Relic Observer settings.



The parameters defined here are then used by the `newrelic_observer_load_start.sh` script to trigger the New Relic Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the configuration file settings.

You define and start the following job. You must edit the parameters in the configuration file before running this job.

### Full Topology Upload job

A transient (one-off) job that loads all requested topology data.

This job is started by the `newrelic_observer_load_start.sh` script.

## Procedure

### To edit the parameters in the configuration file

1. Open the `newrelic_observer_common.sh` configuration file and edit (at least) the following Load parameters:

#### **accountName**

New Relic account name or tenant name

#### **accountId**

New Relic account ID.

To obtain the account ID, first log into the New Relic login page:

<https://login.newrelic.com/login> and then obtain the account ID from this URL:

<https://rpm.newrelic.com/accounts/<accountId>>

#### **insightsQueryAPIKey**

New Relic Insights Query API Key in encoded format.

A new Relic user with a new Relic Infrastructure subscription is required to generate a new Relic Insights query API Key as outlined here: <https://docs.newrelic.com/docs/insights/insights-api/get-data/query-insights-event-data-api>

Use the Agile Service Manager encryption tool to encode the New Relic Insights query API key before using it in job parameter.

#### **Encryption requirement:**

The Load job requires the `insightsQueryAPIKey` in encrypted form. To encrypt the `insightsQueryAPIKey`, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

#### **filterCriteria**

Extend the result set returned to Agile Service Manager.

The default is 'SINCE 4 hours ago LIMIT 1000'.

For more information, see the documentation for New Relic Query Language.

### To start the Load job

2. To start the New Relic Observer Full Topology Upload job, use the following command:

```
$ASM_HOME/bin/newrelic_observer_load_start.sh
```

## Results

This job loads all requested topology data. Run this job whenever you need New Relic topology data refreshed.

## What to do next

You can also use the following scripts:

### **newrelic\_observer\_load\_stop.sh**

Stops the Load job

### **newrelic\_observer\_job\_list.sh**

Lists the status of current jobs

### **newrelic\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining OpenStack Observer jobs

The OpenStack Observer is installed as part of the core installation procedure. Using the OpenStack Observer, you can define jobs that dynamically load OpenStack data for analysis by Netcool Agile Service Manager.

### Before you begin

**Important:** The OpenStack Observer supports the on-premise OpenStack version Stein.

Ensure you have the OpenStack service details to hand, such as username, password, and URL.

**Remember:** Swagger documentation for the observer is available at the following default location:  
[https://<your\\_host>/1.0/openstack-observer/swagger](https://<your_host>/1.0/openstack-observer/swagger)

**Note:** OpenStack uses RBAC-based protection of its API by defining policy rules based on an RBAC approach. Availability of resources retrieved by the observer is also governed by the same policy. For example, a VM created in project A by users with the admin role may only be available to other users with the same admin role. This can be configured or modified according to user requirements in the OpenStack's policy configuration.

### About this task

The OpenStack Observer jobs extract OpenStack resources via REST or RabbitMQ. The Observer loads and updates the resources and their relationships within the Netcool Agile Service Manager core topology service.

#### **openstack\_observer\_common.sh**

The config file you use to customize OpenStack Observer settings.

The parameters defined here are then used by the `openstack_observer_query_start.sh` and the `openstack_observer_listen_start.sh` scripts to trigger the OpenStack Observer jobs.

You define and start the following two jobs. You must edit the parameters in the config file before running these jobs.

#### **Full Topology Upload**

A transient (one-off) job that loads all requested topology data.

This job is started by the `openstack_observer_query_start.sh` script.

**Restriction:** An OpenStack environment that has a list of endpoints whereby the heat-cfn service comes first before the heat service, will encounter a JSON parsing error in the logs due to a known issue in the `openstack4j` library. When this happens, the full load for the heat service will be skipped entirely. The other service will run as normal.

## Listener

A long-running job that monitors its source for updates and runs until it is explicitly stopped, or until the Observer is stopped.

This job is started by the `openstack_observer_listen_start.sh` script.

**Restriction:** Only one listening job should be listening to one queue (or sets of queues) at any one time. If you need to listen to multiple projects, then separate queues must be set up in OpenStack, with appropriate names, before separate listening jobs are submitted for each. For example, for Nova via the `rmq_nova_notify` attribute, for Neutron via the `rmq_neutron_notify` attribute.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the config file settings.

## Procedure

### To edit the parameters in the `openstack_observer_common.sh` config file

1. Open the `openstack_observer_common.sh` config file and edit (at least) the following **Load** parameters:

**os\_auth\_url**

OpenStack identity endpoint

**os\_username**

OpenStack user name

**os\_password**

OpenStack user password

**Encryption requirement:**

The Load and Listener jobs require passwords in the configuration file to be encrypted. To encrypt the `os_password`, run the `encrypt_password.sh` script in the `$ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password, for example:

```
2IuExvqz5SGnGgR0YGLAQg==
```

**os\_tenant\_name**

OpenStack tenant

**Restriction:** The `os_tenant_name` parameter should only be specified for jobs of version 2 authentication (and **not** version 3). When using authentication version 3, specify the `os_project_name` parameter in place of the `os_tenant_name` parameter.

**os\_perspective**

OpenStack network perspective

**ssl\_verification\_disabled**

OpenStack connection SSL

**os\_certificate**

The certificate file of the Openstack host server

**ssl\_truststore\_file**

The SSL truststore to use to authenticate to Openstack host

**ssl\_truststore\_password**

Password to the truststore

**SSL requirements:** To acquire SSL certificates and build SSL truststores, use the relevant instructions for both on-prem and OCP in the following section: [“Defining observer security” on page 99](#)

2. Still in the `openstack_observer_common.sh` config file, edit (at least) the following **Listen** parameters:

**rmq\_hosts**

RMQ connection details

**rmq\_username**

RMQ user name

**rmq\_password**

RMQ user password

**Encryption requirement:**

The Load and Listener jobs require passwords in the configuration file to be encrypted. To encrypt the `rmq_password`, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

**os\_project\_name**

OpenStack project

**Remember:** The `os_project_name` parameter should be specified in place of `os_tenant_name` when using authentication version 3.

To configure the OpenStack installation method

3. Do one of the following depending on whether you have used, or are planning to use, DevStack or another method to install OpenStack.

- **DevStack installation**

**If you have already installed OpenStack using DevStack**

Add the following code to the end of the `local.conf` file, and then reinstall OpenStack.

**If you are planning to install OpenStack using DevStack**

Add the following code to the end of the `local.conf` file before installation.

```
[[post-config|$NOVA_CONF]]
[DEFAULT]
notification_topics = notifications,com.ibm.asm.obs.nova.notify
notification_driver=messagingv2
notify_on_state_change=vm_and_task_state
notify_on_any_change=True
```

- **Other installation**

**For standard (or any other) OpenStack installations**

Add the following code under the `[DEFAULT]` section of the `nova.conf` file, and then restart the nova compute service.

```
notification_topics = notifications,com.ibm.asm.obs.nova.notify
notification_driver=messagingv2
notify_on_state_change=vm_and_task_state
notify_on_any_change=True
```

To start the Load and Listener jobs

4. To start the OpenStack Observer **Full Topology Upload** job, use the following command:

```
$ASM_HOME/bin/openstack_observer_query_start.sh
```

The Full Topology Upload job loads all requested topology data. This job runs only once.

5. To start the OpenStack Observer listener job, use the following command:

```
$ASM_HOME/bin/openstack_observer_listen_start.sh
```

The Listener job monitors its source for updates and runs until it is explicitly stopped, or until the Observer is stopped.

## What to do next

You can also use the following scripts:

**openstack\_observer\_query\_stop.sh**

Stops the Full Topology Upload job

**openstack\_observer\_listen\_stop.sh**

Stops the Listener job

**openstack\_observer\_job\_list.sh**

Lists the status of current jobs

**openstack\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining Rancher Observer jobs

Using the Rancher Observer, you can define a full load job that will read data from Rancher through the REST API and generate a topology.

### Before you begin

**Important:** The Rancher Observer supports the on-premise Rancher version 2.3.2 and API version 3.0.

The Rancher Observer is installed as part of the core installation procedure.

Ensure you have the Rancher details to hand, such as the Rancher host IP, accessKey, secretKey and clusterId, before running the observer job.

### About this task

**Remember:** Swagger documentation for the observer is available at the following default location:

<https://<your host>/1.0/rancher-observer/swagger>

The Rancher Observer supports a transient (one-off) Load job that loads all requested topology data via Rancher REST API to build the topology, and then exit.

#### **rancher\_observer\_common.sh**

The configuration file you use to customize Rancher Observer settings.

The parameters defined here are then used by the `rancher_observer_load_start.sh` to trigger the Rancher Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the configuration file settings.

You define and start the following job. You must edit the parameters in the configuration file before running this job.

#### **Full Topology Upload job**

A transient (one-off) job that loads all requested topology data.

This job is started by the `rancher_observer_load_start.sh` script and loads all supported resources.

## Procedure

### To edit the parameters in the configuration file

1. Open the `rancher_observer_common.sh` configuration file and edit (at least) the following Load parameters:

**Rancher Host IP**

The API base URL

**accessKey**

The Rancher access key

**secretKey**

The Rancher secret key (must be encrypted)

**clusterId**

The Rancher cluster ID

**enableSecuredConnection**

Set to 'true' to secure the connection.

Set to 'false' to bypass secure host connection.

**certificate**

Optional.

If **enable\_secure\_host\_connection** is set to 'true':

**On Premise**

Enter the certificate name (in the `/opt/ibm/netcool/asm/security` directory) to load into the trust store.

**On OCP**

Update the global secret with encrypted certificate information.

See the relevant step [here](#).

**ssl\_truststore\_file**

The SSL truststore file

**password\_ssl\_truststore\_file**

The SSL truststore password (must be encrypted)

**Encryption requirement:** The Load job requires the password and secret key in the configuration file in encrypted form. To encrypt the password, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

**To start the Load job**

2. To start the Rancher Observer Full Topology Upload job, use the following command:

```
$ASM_HOME/bin/rancher_observer_load_start.sh
```

## Results

This job loads all requested topology data, and runs only once. Run this job whenever you need Rancher topology data refreshed.

## What to do next

See the related links for information on using rules for this observer.

You can also use the following scripts:

**rancher\_observer\_load\_stop.sh**

Stops the Load job

**rancher\_observer\_job\_list.sh**

Lists the status of current jobs

### **rancher\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

### **Related tasks**

[“Defining rules” on page 309](#)

Rules help streamline topologies and conserve system resources, for example by merging different observer records of the same resource into a single composite resource, or by excluding specific changes from being recorded against a resource history.

## **Defining REST Observer jobs**

The REST (or RESTful) Observer is installed as part of the core installation procedure. Use the REST Observer for obtaining topology data via REST endpoints. This observer is a counterpart to the File Observer.

### **Before you begin**

**Remember:** Swagger documentation for the observer is available at the following default location:  
`https://<your host>/1.0/rest-observer/swagger`

### **About this task**

The REST Observer passes topology data to Agile Service Manager using a RESTful set of interfaces, which provide REST APIs that enable the following functionality:

- Management of Listen and bulk-replace job types.
- The insert-update (HTTP POST) of resources.
- The insert-update (HTTP POST) of relationships.
- The insert-replacement (HTTP PUT) of resources.
- The deletion (HTTP DELETE) of resources.
- A REST API that supports the deletion (HTTP DELETE) of all relationships of a given type from a specified resource.
- A REST API that supports the deletion (HTTP DELETE) of a specific relationship.

**Restriction:** Resources created via REST can have a provider, but not an observer.

### **Benefits**

Using the REST Observer rather than the File Observer or Topology Service APIs includes the following benefits:

- The ability to provide data to Agile Service Manager via HTTP REST Endpoints instead of files.
- The processing performed by all observers in their framework ensures that meta-data about observations from observers is managed correctly.
- A simple way of deleting all edges of a given type on a resource or a specific edge instance.

To use the REST Observer, a job request must be issued (HTTP POST) to the Observer instance job management APIs before sending data to the Resource and Relationship APIs.

### **Listen**

A long-running listen job capable of consuming topology data over a long period of time.

A listen job is designed to support scenarios where the input data stream is unpredictable, or there is little or no consistency or versioning of resources within the data stream.

**Note:** These examples assume that the environment variables have been set in `rest_observer_common.sh`

#### start

```
$ASM_HOME/bin/rest_observer_listen_start.sh
```

#### stop

Default job

```
./bin/rest_observer_listen_stop.sh
```

Named job

```
env unique_id='My job name' $ASM_HOME/bin/rest_observer_listen_stop.sh
```

### Bulk replace

A long-running job with the same resource replace semantics as the File Observer.

Bulk-replace jobs are designed to support scenarios where a known set of resources are subject to updates or versioning, and a prior observation about resources is to be replaced with a new one.

This job can provide a new set of resources and relationships and synchronize them to Agile Service Manager, thereby causing any previous data provided by the Observer to be deleted and replaced with the new data.

**Note:** These examples assume that the environment variables have been set in `rest_observer_common.sh`

#### start

Default job:

```
./bin/rest_observer_bulk_replace_start.sh
```

Job with `bulk_replace_unique_id` and provider given:

```
env bulk_replace_unique_id=manageDataCenter provider=MyJavaProgram  
$ASM_HOME/bin/rest_observer_bulk_replace_start.sh
```

#### synchronize

Default job

```
./bin/rest_observer_bulk_replace_synchronize.sh
```

Named job

```
env unique_id='My job name'  
$ASM_HOME/bin/rest_observer_bulk_replace_synchronize.sh
```

#### stop

Default job

```
$ASM_HOME/bin/rest_observer_bulk_replace_stop.sh
```

Named job

```
env unique_id='My job name'  
$ASM_HOME/bin/rest_observer_bulk_replace_stop.sh
```

Once a job request has been successfully submitted, you can start to provide data to the Resource and Relationship APIs on behalf of a given job instance.

The Resource and Relationship APIs may respond with an `HTTPS 503 Service Unavailable` response with a `Retry-After: 10` seconds in the header. This indicates that even though the request against those APIs is valid, the observer has not been able to ascertain that meta-data about the job is



held in Agile Service Manager yet; this may be due to, for example, any prevailing conditions in the network that support the Agile Service Manager micro-services.

**Tip:** If such a response is received, try the request again later.

**Important:** Ensure that the body is structured correctly. When posting the body, information included in the body after the closing `}` that matches an opening `{` is ignored, and no error is recorded.

## Procedure

### Listen job process and examples

The following procedure (steps one to ten) includes examples that show how to use the REST Observer listen job to create and adjust a small topology.

1. Start the Listen job.

Use the following example as a model.

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X POST --header 'Content-Type: application/json' --header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' -d '{
  "unique_id": "my job",
  "type": "listen",
  "parameters": {
    "provider": "MyListenJob"
  }
}' 'https://localhost/1.0/rest-observer/jobs/listen'
```

2. Verify that the job is running.

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X GET --header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' 'https://localhost/1.0/rest-observer/jobs/my%20job'
```

3. Create a 'person' resource.

This example creates a person resource called 'Thomas Watson'.

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X POST --header 'Content-Type: application/json' --header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --header 'JobId: my job' -d '{
  "name": "Thomas Watson",
  "uniqueId": "Thomas Watson",
  "entityTypes": [
    "person"
  ]
}' 'https://localhost/1.0/rest-observer/rest/resources'
```

4. Create an 'organization' resource.

This example creates an 'organization' resource of 'IBM'.

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X POST --header 'Content-Type: application/json' --header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --header 'JobId: my job' -d '{
  "name": "IBM",
  "uniqueId": "IBM",
  "entityTypes": [
    "organization"
  ]
}' 'https://localhost/1.0/rest-observer/rest/resources'
```

5. Create a 'manages' relationship between Thomas Watson and IBM.

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X POST --header 'Content-Type: application/json' --header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --header 'JobId: my job' -d '{
  "_fromUniqueid": "Thomas Watson",
  "_edgeType": "manages",
  "_toUniqueid": "IBM"
}' 'https://localhost/1.0/rest-observer/rest/references'
```

6. Create a new 'location' resource and relate it to Thomas Watson.

This example creates a 'location' resource of 'Campbell, New York' for Thomas Watson, and a location relationship (an edgeType of locatedAt).

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X POST --header 'Content-Type: application/json' --header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --header 'JobId: my job' -d '{
  "name": "Campbell, New York",
  "uniqueId": "Campbell, New York",
  "entityTypes": [
    "location"
  ],
  "_references": [
    {
      "_fromUniqueId": "Thomas Watson",
      "_edgeType": "locatedAt"
    }
  ]
}' 'https://localhost/1.0/rest-observer/rest/resources'
```

7. Replace the location resource with one having latitude and longitude properties.

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X PUT --header 'Content-Type: application/json' --header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --header 'JobId: my job' -d '{
  "name": "Campbell, New York",
  "uniqueId": "Campbell, New York",
  "entityTypes": [
    "location"
  ],
  "latitude": 42.232909,
  "longitude": -77.196918
}' 'https://localhost/1.0/rest-observer/rest/resources'
```

8. Delete all locatedAt relationships from Thomas Watson.

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X DELETE --header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --header 'JobId: my job' 'https://localhost/1.0/rest-observer/rest/resources/Thomas%20Watson/references/both/locatedAt'
```

9. Delete the Campbell, New York location.

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X DELETE --header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --header 'JobId: my job' 'https://localhost/1.0/rest-observer/rest/resources/Campbell%2C%20New%20York'
```

10. Delete the manages relationship between Thomas Watson and IBM.

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X DELETE --header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --header 'JobId: my job' 'https://localhost/1.0/rest-observer/rest/resources/Thomas%20Watson/references/both/manages/IBM'
```

## Bulk Replace job process and examples

The following procedure (steps 11 - 21) includes examples that show how to use the REST Observer bulk-replace job to create and adjust a small topology.

**Note:** These examples use a mixture of the provided scripts in \$ASM\_HOME/bin and the cURL command.

11. Submit a bulk-replace job request with the unique ID of 'my bulk replace'.

```
[root@asm-backend asm]# env bulk_replace_unique_id="my bulk replace" provider="Me" bin/rest_observer_bulk_replace_start.sh
```

12. Verify that the job is running.

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X GET --header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' 'https://localhost/1.0/rest-observer/jobs/my%20job'
```

13. Submit a location resource for the city of Coventry.

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X POST --header 'Content-Type: application/json'
```

```
--header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255'
--header 'JobId: my bulk replace' -d '{
  "name": "Coventry",
  "uniqueId": "Coventry",
  "entityTypes": [
    "location"
  ]
}' 'https://localhost/1.0/rest-observer/rest/resources'
```

#### 14. Submit a location resource for the town of Rugby

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X POST --header 'Content-Type: application/json'
--header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255'
--header 'JobId: my bulk replace' -d '{
  "name": "Rugby",
  "uniqueId": "Rugby",
  "entityTypes": [
    "location"
  ]
}' 'https://localhost/1.0/rest-observer/rest/resources'
```

#### 15. Submit a location resource for the Town of Leamington Spa with relationships to the existing resources of Coventry and Rugby.

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X POST --header 'Content-Type: application/json'
--header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --header 'JobId: my bulk replace' -d '{
  "name": "Leamington Spa",
  "uniqueId": "Leamington Spa",
  "entityTypes": [
    "location"
  ],
  "_references": [
    {
      "_toUniqueId": "Coventry",
      "_edgeType": "routesVia"
    },
    {
      "_toUniqueId": "Rugby",
      "_edgeType": "routesVia"
    }
  ]
}' 'https://localhost/1.0/rest-observer/rest/resources'
```

#### 16. Check your progress by rendering the topology.

#### 17. Having completed this set of observations, initiate a synchronize request for 'my bulk replace' job.

```
[root@asm-backend asm]# env bulk_replace_unique_id="my bulk replace" ./bin/rest_observer_bulk_replace_synchronize.sh
```

#### 18. Provide a new topology for the 'my bulk replace' job.

This example submits a location resource for the town of Leamington Spa with relationships to the towns of Warwick and Stratford-Upon-Avon (resource placeholders).

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X POST --header 'Content-Type: application/json'
--header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255'
--header 'JobId: my bulk replace' -d '{
  "name": "Leamington Spa",
  "uniqueId": "Leamington Spa",
  "entityTypes": [
    "location"
  ],
  "_references": [
    {
      "_toUniqueId": "Warwick",
      "_edgeType": "routesVia"
    },
    {
      "_toUniqueId": "Stratford-upon-Avon",
      "_edgeType": "routesVia"
    }
  ]
}' 'https://localhost/1.0/rest-observer/rest/resources'
```

19. Provide the resource data for the Town of Warwick resource placeholder, thus fully creating the resource.

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X POST --header 'Content-Type: application/json' --header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --header 'JobId: my bulk replace' -d '{
  "name": "Warwick",
  "uniqueId": "Warwick",
  "entityTypes": [
    "location"
  ]
}' 'https://localhost/1.0/rest-observer/rest/resources'
```

20. Provide the resource data for the town of Stratford-upon-Avon resource placeholder, thus fully creating the resource.

```
curl -u PROXY_USER[:PROXY_PASSWORD] -X POST --header 'Content-Type: application/json' --header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' --header 'JobId: my bulk replace' -d '{
  "name": "Stratford-upon-Avon",
  "uniqueId": "Stratford-upon-Avon",
  "entityTypes": [
    "location"
  ]
}' 'https://asm-backend.rtp.raleigh.ibm.com/1.0/rest-observer/rest/resources'
```

21. Initiate a synchronize request for the 'my bulk replace' job. This signifies to the Observer that it should instruct ASM to replace the previous set of observations with the new ones.

**Note:** The new data is available immediately as it is provided to ASM. The synchronize request simply deletes any resources previously observed that were not observed this time. In the current example, Coventry and Rugby were not observed, and therefore they are deleted.

```
[root@asm-backend asm]# env bulk_replace_unique_id="my bulk replace" ./bin/rest_observer_bulk_replace_synchronize.sh
```

## What to do next

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

### Related information

[Timestamps formats](#)

## Defining ServiceNow Observer jobs

The ServiceNow Observer is installed as part of the core installation procedure. Using the ServiceNow Observer, you can retrieve the configuration management database (CMDB) data from ServiceNow via REST API. Currently, the observer only supports load job. The load job queries the configuration items (CI) from CMDB via ServiceNow REST API using basic authentication credentials.

### Before you begin

**Important:** The ServiceNow Observer supports the cloud/SaaS ServiceNow version New York.

Ensure your user account has the `rest_api_explorer` and `web_service_admin` roles. These roles are required to access the resources from ServiceNow. Also, ensure you have the ServiceNow service details to hand, such as username, password, and URL.

**Remember:** Swagger documentation for the observer is available at the following default location:

`https://<your host>/1.0/servicenow-observer/swagger`

**Tip:** To run a load job using a non-admin account, assign the required roles to the account to provide read privileges for the list of API paths.

1. Sign in to the **ServiceNow** instance using an admin account.
2. Navigate to **User Administration** and select **Users** from the menu.
  - To create a new user, click **New**.
  - To edit an existing user, search and select the user.
3. From the user's information tab, select the **Roles** tab, then click **Edit**.
4. Assign the **cmdb\_read** and **service\_viewer** roles to the user.
5. Click **Save**, then **Update**.

## About this task

ServiceNow jobs retrieve the configuration management database (CMDB) data from ServiceNow via REST API. The Observer loads and updates the resources and their relationships within the Netcool Agile Service Manager core topology service.

The observer discovers the following resources (based on the API path list):

- /api/now/table/cmdb\_ci
- /api/now/table/core\_company
- /api/now/table/cmn\_department
- /api/now/table/cmn\_location
- /api/now/table/sys\_user

### **servicenow\_observer\_common.sh**

The configuration file you use to customize ServiceNow Observer settings.

The parameters defined here are then used by the `servicenow_observer_load_start.sh` script to trigger the ServiceNow Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the configuration file settings.

You define and start the following job. You must edit the parameters in the configuration file before running this job.

### **Load job**

A transient (one-off) job that loads all requested topology data.

This job is started by the `servicenow_observer_load_start.sh` script.

Run this job whenever you need the ServiceNow topology data refreshed. .

Table 58. Mapping of ServiceNow object types to Agile Service Manager entity types:	
ServiceNow object types	Agile Service Manager entity types
cmdb_ci	based on sys_class_name attribute
cmn_department	department
cmn_location	location
core_company	company
sys_user	person

## Procedure

### To edit the parameters in the configuration file

1. Open the `servicenow_observer_common.sh` configuration file and edit (at least) the following parameters:

**ServiceNow username**

The username of the ServiceNow instance

**ServiceNow password**

The password of the ServiceNow instance

**Encryption requirement:**

The Load job requires the password in the configuration file to be encrypted. To encrypt the password, run the `encrypt_password.sh` script in the `ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

**Instance URL**

The URL on which the ServiceNow instance is running

**To start the Load job**

2. To start the ServiceNow Observer Load job, use the following command:

```
$ASM_HOME/bin/servicenow_observer_load_start.sh
```

This job loads all requested topology data. Run this job whenever you need the ServiceNow topology data refreshed.

## Results

You can now use data retrieved from the ServiceNow configuration management database (CMDB) to create topologies in the Agile Service Manager topology service.

## What to do next

You can also use the following scripts:

**`servicenow_observer_load_stop.sh`**

Stops the Load job

**`servicenow_observer_job_list.sh`**

Lists the status of current jobs

**`servicenow_observer_log_level.sh`**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining TADDM Observer jobs

The TADDM Observer is installed as part of the core installation procedure. Using the TADDM Observer, you can retrieve network topology data from the TADDM database and use this data to create topologies within the topology service.

### Before you begin

**Important:** The TADDM Observer supports the on-premise TADDM version.

Ensure you have the TADDM Rest API login access details in hand, such as the TADDM API URL, username and password.

All prerequisites are deployed during the Agile Service Manager core installation. This includes the TADDM Observer docker container, which has been installed and should be running, as well as the required scripts to manage jobs. All observers have scripts to start and stop all available jobs, to list the status of a current job, to set its logging levels, and to configure its job parameters.

You can verify that the TADDM Observer's docker container is running using the following command:

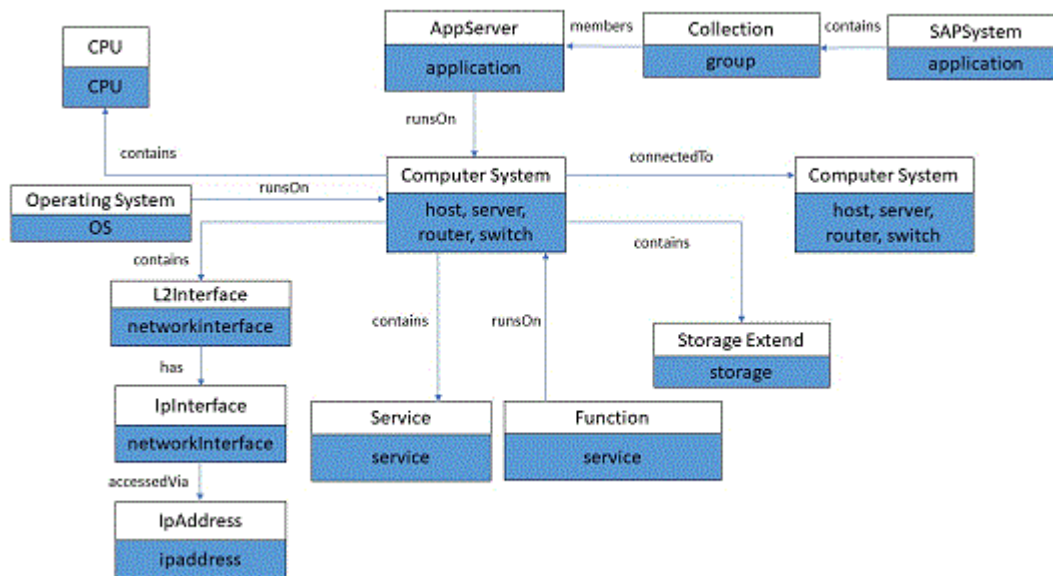
```
$ASM_HOME/bin/docker-compose ps
```

The system should return text indicating that `asm_taddm-observer_1` has a state of Up and therefore is running.

**Remember:** Swagger documentation for the observer is available at the following default location:  
<https://<your host>/1.0/taddm-observer/swagger>

## About this task

The TADDM Observer is built on the Observer framework:



- A computer system can be a host, server, router or switch
- A computer system contains CPU, L2Interface and storage
- Operating system, application server and service run on a computer system
- A computer system can connect to another computer system
- A SAPSystem contains collection
- An application server can be a member of a collection

Table 59. Mapping TADDM model objects to Agile Service Manager entity types	
TADDM model object	Agile Service Manager entity types
AppServer	application
ComputerSystem	host, server, router, switch
CPU	cpu
L2Interface	networkinterface

Table 59. Mapping TADDM model objects to Agile Service Manager entity types (continued)

TADDM model object	Agile Service Manager entity types
IpInterface	networkinterface
IpAddress	ipaddress
OperatingSystem	os
Service	service
StorageExtent	storage
Function	service
SAPSystem	application
Collection	group

The TADDM Observer job retrieves topology data using the TADDM REST API. The observer loads and updates the resources and their relationships within the Netcool Agile Service Manager core topology service.

#### **taddm\_observer\_common.sh**

The config file you use to customize TADDM Observer settings.

The parameters defined here are then used by the `taddm_observer_load_start.sh` and `taddm_observer_poll_start.sh` scripts to trigger the TADDM Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the configuration file settings.

You define and start the following job. You must edit the parameters in the config file before running a job.

#### **Load job**

A transient (one-off) job that loads all requested topology data.

This job is started by the `taddm_observer_load_start.sh` script.

### **Procedure**

#### **To edit the parameters in the config file**

1. Open the `$ASM_HOME/bin/taddm_observer_common.sh` configuration file and edit (at least) the following parameters:

##### **username**

TADDM user

##### **password**

TADDM password, in encrypted form. Use the `$ASM_HOME/bin/encrypt_password.sh` utility to generate an encrypted password.

##### **api\_url**

TADDM API URL

##### **model\_objects**

Optional

List of supported TADDM model object names to be observed.

Keep the default to let the observer fetch all the supported model objects. Supported model objects are ["AppServer", "ComputerSystem", "CPU", "StorageExtent", "L2Interface", "IpInterface", "IpAddress", "OperatingSystem", "Function", "SAPSystem", "Collection"]

#### **To start the Load job**

2. To start the TADDM Observer load job, use the following command:



```
$ASM_HOME/bin/taddm_observer_load_start.sh
```

This job loads all requested topology data. Run this job whenever you need the TADDM topology data refreshed.

## Results

You can now use data retrieved from the TADDM database to create topologies in the Agile Service Manager topology service.

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining VMware NSX Observer jobs

The VMware NSX Observer is installed as part of the core installation procedure. Use the VMware NSX Observer when you have VMware NSX installed in your environment to define jobs that dynamically load VMware NSX data for analysis by Netcool Agile Service Manager.

### Before you begin

You can use the VMware NSX Observer when you have a VMware NSX appliance in your environment.

**Important:** The VMware NSX Observer supports the on-premise VMware NSX version 6.3.

Ensure you have the VMware NSX service details to hand, such as username, password, SSL TrustStore and URL.

**Remember:** Swagger documentation for the observer is available at the following default location: <https://<your host>/1.0/vmwarensx-observer/swagger>

### About this task

The VMware NSX Observer job extracts VMware NSX resource information via REST. The Observer loads and updates the resources and their relationships within the Netcool Agile Service Manager core topology service.

#### **vmwarensx\_observer\_common.sh**

The config file you use to customize VMware NSX Observer settings.

The parameters defined here are then used by the `vmwarensx_observer_query_start.sh` script to trigger the VMware NSX Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the config file settings.

You define and start the following job. You must edit the parameters in the config file before running this job.

#### **Full Topology Upload job**

A transient (one-off) job that loads all requested topology data.

This job is started by the `vmwarensx_observer_query_start.sh` script.

The VMware NSX Observer loads the following resources and their relationship into the Netcool Agile Service Manager core topology service:

- NSX Appliance
- vCenter Appliance
- NSX Controller

- Edge Router - Logical (Distributed) Router, Edge Service Gateway
- Virtual Machines
- Host
- VNIC

## Procedure

### To edit the parameters in the config file

1. Open the `vmwarensx_observer_common.sh` config file and edit (at least) the following parameters:

**api\_url**

VMware NSX REST API endpoint

**username**

VMware NSX user name for REST API

**password**

VMware NSX user password for REST API

Supply the VMware NSX user password in encrypted text.

**tenant\_name**

VMware NSX tenant

Set to 'default' if there is no specific tenant.

**ssl\_truststore\_file**

VMware NSX SSL trust store file for HTTPS authentication

JKS is the supported format and the file is relative to the `$ASM_HOME/security` directory.

**password\_ssl\_truststore\_file**

Password to decrypt an encrypted VMware NSX SSL trust store file

Supply the VMware NSX SSL trust store password in encrypted format.

**Encryption requirement:**

The Load job requires passwords in encrypted form. To encrypt the `nsx_password` and `password_ssl_truststore_file`, run the `encrypt_password.sh` script in the `$ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

### To acquire VMware NSX SSL certificate and build SSL truststore

2. Use the following command to use OpenSSL to connect to VMware NSX over port 443, and extract a SSL Certificate from VMware NSX to a `<certificate_file_name>.crt` file.

```
echo -n | openssl s_client -connect {VMware NSX IpAddress}:443 | sed -ne
'/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > ./ {certificate_file_name}.crt
```

3. Use the following Java keytool command to import the VMware NSX certificate file into a keystore and encrypt the keystore with a given password.

```
keytool -import -v -trustcacerts -alias {VMware NSX Hostname}
-file {certificate_file_name}.crt -keystore {keystore file name}
-storepass {your password to encrypt keystore}
```

**Tip:** You will need the following encryption information when editing `vmwarensx_observer_common.sh`

Table 60. Encryption parameters required for <code>vmwarensx_observer_common.sh</code>	
keystore parameter	vmwarensx_observer_common.sh parameter
keystore password	password_ssl_truststore_file

Table 60. Encryption parameters required for <code>vmwarensx_observer_common.sh</code> (continued)	
keystore parameter	vmwarensx_observer_common.sh parameter
keystore file name	ssl_truststore_file

4. Copy the keystore file (`{keystore file name}`) to the `$ASM_HOME/security` directory to complete the SSL setup.

### To start the Load job

5. To start the VMware NSX Observer Full Topology Upload job, use the following command:

```
$ASM_HOME/bin/vmwarensx_observer_query_start.sh
```

This job loads all requested topology data. Run this job whenever you need VMware NSX topology data refreshed.

## What to do next

You can also use the following scripts:

### **vmwarensx\_observer\_query\_stop.sh**

Stops the Full Topology Upload job

### **vmwarensx\_observer\_job\_list.sh**

Lists the status of current jobs

### **vmwarensx\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining VMware vCenter Observer jobs

The VMware vCenter Observer is installed as part of the core installation procedure. Use the VMware vCenter Observer when you have VMware vCenter installed in your environment to define jobs that dynamically load VMware vCenter data for analysis by Netcool Agile Service Manager.

### Before you begin

**Important:** The VMware vCenter Observer supports the on-premise VMware vCenter versions 6.5 and 6.7.

Ensure you have the VMware vCenter service details to hand, such as username, password, SSL TrustStore and URL.

**Remember:** Swagger documentation for the observer is available at the following default location: `https://<your host>/1.0/vmcenter-observer/swagger`

### About this task

The VMware vCenter Observer job extracts VMware vCenter resource information via REST. The Observer loads and updates the resources and their relationships within the Netcool Agile Service Manager core topology service.

### **vmcenter\_observer\_common.sh**

The config file you use to customize VMware vCenter Observer settings.

The parameters defined here are then used by the `vmvcenter_observer_query_start.sh` script to trigger the VMware vCenter Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the config file settings.

You define and start the following job. You must edit the parameters in the config file before running this job.

### Full Topology Upload job

A transient (one-off) job that loads all requested topology data.

This job is started by the `vmvcenter_observer_query_start.sh` script.

The VMware vCenter Observer loads the following resources and their relationship into the Netcool Agile Service Manager core topology service:

- ESXi / ESX Hosts
- Virtual Machines
- VNICS
- Storage

## Procedure

### To edit the parameters in the config file

1. Run the following command to check if a specific username has access to obtain the session-id token.

```
curl -kX POST -u 'username':'password' -H "Content-Type: application/json"
'https://<host>/rest/com/vmware/cis/session'
```

If the output is a value, then the username can be used to obtain a session-id token.

2. Open the `vmvcenter_observer_common.sh` config file and edit (at least) the following parameters:

#### **vcenter\_api\_url**

VMware vCenter REST API endpoint

#### **vcenter\_username**

VMware vCenter user name for REST API

#### **vcenter\_password**

VMware vCenter user password for REST API

#### **ssl\_truststore\_file**

VMware vCenter SSL trust store file for HTTPS authentication

JKS is the supported format and the file is relative to `$ASM_HOME/data/vmcenter-observer`

#### **password\_ssl\_truststore\_file**

Password to decrypt and encrypt VMware vCenter SSL trust store file

#### **Encryption requirement:**

The Load job requires passwords in encrypted form. To encrypt the `vcenter_password` and `password_ssl_truststore_file`, run the `encrypt_password.sh` script in the `$ASM_HOME/bin` directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

#### **certificate**

The certificate name (in the `/opt/ibm/netcool/asm/security` directory)

3. Optional: Edit the following optional parameters:

#### **connect\_read\_timeout\_ms**

Set the connection and read timeout value (in milliseconds).

The default value is 5000.

### **include**

The VMware vCenter host's name regex to discover.

Specify an exact match or a regular expression match for a host's name in order to discover all its virtual machines.

By default, it discovers all.

### **connect\_retry**

Set the connection retry times.

The default value is 5.

### **connect\_retry\_delay**

Set the time delay before trying to reconnect (in milliseconds).

The default value is 1000.

## **To acquire VMware vCenter SSL certificate and build SSL truststore**

4. Use the following command to use OpenSSL to connect to VMware vCenter over port 443, and extract a SSL Certificate from VMware vCenter to a `<certificate_file_name>.cert` file.

```
echo -n | openssl s_client -connect {VMware vCenter IPAddress}:443 | sed -ne
'/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > ./ {certificate_file_name}.cert
```

5. Use the following Java keytool command to import the VMware vCenter certificate file into a keystore and encrypt the keystore with a given password.

```
keytool -import -v -trustcacerts -alias {VMware vCenter Hostname} -file
{certificate_file_name}.cert -keystore {keystore file name}
-storepass {your password to encrypt keystore}
```

**Tip:** You will need the following encryption information when editing `vmvcenter_observer_common.sh`

Table 61. Encryption parameters required for <code>vmvcenter_observer_common.sh</code>	
keystore parameter	vmvcenter_observer_common.sh parameter
keystore password	password_ssl_truststore_file
keystore file name	ssl_truststore_file

6. Copy the keystore file (`{keystore file name}`) to the `$ASM_HOME/security` directory to complete the SSL setup.

## **To start the Load job**

7. To start the VMware vCenter Observer Full Topology Upload job, use the following command:

```
$ASM_HOME/bin/vmvcenter_observer_query_start.sh
```

This job loads all requested topology data. Run this job whenever you need VMware vCenter topology data refreshed.

## **What to do next**

You can also use the following scripts:

### **vmcenter\_observer\_query\_stop.sh**

Stops the Full Topology Upload job

### **vmcenter\_observer\_job\_list.sh**

Lists the status of current jobs

### **vmcenter\_observer\_log\_level.sh**

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.

## Defining Zabbix Observer jobs

Using the Zabbix Observer functionality, you can load monitored servers and their associated network resources, and then visualize this data as a topology view in the Agile Service Manager UI. It is installed as part of the core installation procedure.

### Before you begin

The Zabbix Observer supports Zabbix Version 4.0.3.

Ensure you have the Zabbix server details to hand, such as the username, password and SSL TrustStore.

**Remember:** Swagger documentation for the observer is available at the following default location:

`https://<your host>/1.0/zabbix-observer/swagger`

### About this task

A Zabbix Observer job extracts server information and its associated network resources from Zabbix via REST RPC. The Observer loads and updates the resources and their relationships within the Netcool Agile Service Manager core topology service.

#### **zabbix\_observer\_common.sh**

The configuration file you use to customize Zabbix Observer settings.

The parameters defined here are then used by the `zabbix_observer_load_start.sh` script to trigger the Zabbix Observer jobs.

**Tip:** Alternatively, you can set the appropriate environment variables. If an environment variable is set, it takes precedence over the config file settings.

You define and start the following job. You must edit the parameters in the config file before running this job.

#### **Full Topology Upload job**

A transient (one-off) job that loads all requested topology data.

This job is started by the `zabbix_observer_load_start.sh` script.

## Procedure

### To edit the parameters in the config file

1. Open the `zabbix_observer_common.sh` config file and edit (at least) the following parameters:

#### **hostname**

Zabbix hostname or ipaddress

#### **username**

Zabbix user name

#### **password**

Zabbix user password

Must be supplied in encrypted format

#### **certificate**

Optional certificate name. If provided, then a certificate file with the same name must exist in the `$ASM/security` directory.

### ssl\_truststore\_file

Zabbix SSL trust store file for HTTPS authentication

JKS is the supported format and the file is relative to \$ASM\_HOME/security

### truststore\_password

Password to decrypt and encrypt Zabbix SSL trust store file

Must be encrypted

#### Encryption requirement:

The Load job requires passwords in encrypted form. To encrypt the password and truststore\_password, run the encrypt\_password.sh script in the \$ASM\_HOME/bin directory:

```
./bin/encrypt_password.sh
```

Enter and then confirm the password. The encryption utility will return an encrypted password.

### connect\_read\_timeout\_ms

Connection timeout in milliseconds (ms), for example '5000'.

#### To acquire Zabbix SSL certificate and build SSL truststore

2. Use the following command to use OpenSSL to connect to Zabbix, and extract an SSL Certificate from Zabbix to a <certificate\_file\_name>.crt file.

```
echo -n | openssl s_client -connect {Zabbix IpAddress}:{SSL port} | sed -ne  
'/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > ./ {certificate_file_name}.crt
```

3. Use the following Java keytool command to import the Zabbix certificate file into a keystore and encrypt the keystore with a given password.

```
keytool -import -v -trustcacerts -alias {Zabbix Hostname} -file  
{certificate_file_name}.crt -keystore {keystore file name}  
-storepass {your plain text password to encrypt keystore}
```

**Tip:** You will need the following encryption information when editing zabbix\_observer\_common.sh

Table 62. Encryption parameters required for zabbix_observer_common.sh	
keystore parameter	zabbix_observer_common.sh parameter
keystore password	truststore_password
keystore file name	ssl_truststore_file

4. Copy the keystore file ({keystore file name}) to the \$ASM\_HOME/security directory to complete the SSL setup.

#### To start the Load job

5. To start the Zabbix Observer Full Topology Upload job, use the following command:

```
$ASM_HOME/bin/zabbix_observer_load_start.sh
```

This job loads all requested topology data. Run this job whenever you need Zabbix topology data refreshed.

## What to do next

You can also use the following scripts:

### zabbix\_observer\_query\_stop.sh

Stops the Full Topology Upload job

### zabbix\_observer\_job\_list.sh

Lists the status of current jobs

### zabbix\_observer\_log\_level.sh

Sets the log level

**Remember:** In addition to being configurable from the Observer Configuration UI, all on-prem observer jobs also have scripts to start and stop all available jobs, to list the status of a current job, and to set its logging levels. Scripts can be run with **-h** or **--help** to display help information, and with **-v** or **--verbose** to print out the details of the actions performed by the script, including the full cURL command. For the on-prem version of Agile Service Manager, observer scripts are configured for specific jobs by editing the script configuration files.



---

## Chapter 6. Using Agile Service Manager

You use the Agile Service Manager UI to visualize your topology data. First you define a seed resource on which to base your view, then choose the levels of networked resources around the seed that you wish to display, before rendering the view. You can then further expand or analyze the displayed topology in real time, or compare it to previous versions within a historical time window.

The Netcool Agile Service Manager Topology Viewer has four toolbars and a visualization display area.

### Navigation toolbar

You use the navigation toolbar to select the seed resource, define the number of relationship hops to visualize from the seed resource, and specify the type of relationship hop traversal to make (either host-to-host, or element-to-element).

### Resource filter toolbar

You use the resource filter toolbar to apply entity- or relationship-type filters to the resources displayed in the topology.

### Visualization toolbar

You use the Visualization toolbar to customize the topology view, for example by zooming in and panning.

### History toolbar

You use the History toolbar to compare and contrast a current topology with historical versions.

### Topology visualization panel

You use the Topology visualization panel to view the topology, and access the resource nodes for further analysis performed via the context menu.

---

## Logging into the UI (OCP)

You construct the Agile Service Manager OCP logon URL from the Netcool Operations Insight Helm release name.

### Before you begin

Also see the following Netcool Operations Insight topic for more information: [https://www.ibm.com/support/knowledgecenter/SSTPTP\\_1.6.3/com.ibm.netcool\\_ops.doc/soc/start/task/start\\_getting-started-rhocp.html](https://www.ibm.com/support/knowledgecenter/SSTPTP_1.6.3/com.ibm.netcool_ops.doc/soc/start/task/start_getting-started-rhocp.html)

**Tip:** To prevent UI timeout errors, you can increase the timeout values for the topology, layout and search services. See the following troubleshooting topic for more details: [“User interface timeout errors” on page 351](#)

### About this task

You login to the Agile Service Manager OCP installation using a URL of the following format (example):

```
http://netcool.noi.apps.<your-ocp-cluster>/ibm/console
```

Where *noi* is the Netcool Operations Insight Helm release name. Use the following command to retrieve the DASH URL:

```
helm status NOI helm release name --tls
```

**Tip:** When first logging on, an Information prompt may appear asking you to accept the Agile Service Manager UI certificate. To do so, click on any one of the provided links and accept the certificate when prompted. You may then close the information prompt.

**Tip:** When opening an Agile Service Manager UI page in DASH using Internet Explorer or Edge, the following certificate failure can occur: Content was blocked because it was not signed by a

valid security certificate. To fix this, your administrator must obtain and install the Agile Service Manager root certificate manually, as described in this [troubleshooting](#) topic.

## Logging into the UI (Watson AIOps)

---

You access the standalone OCP version of Agile Service Manager deployed with IBM Watson AIOps from the OCP Administration console.

### Before you begin

**Tip:** To prevent UI timeout errors, you can increase the timeout values for the topology, layout and search services. See the following troubleshooting topic for more details: [“User interface timeout errors”](#) on page 351

### About this task

The URL to access the UI of the OCP version of Agile Service Manager deployed with IBM Watson AIOps is constructed from the OCP-cluster-host and the Agile Service Manager tenant ID. For example:

```
https://<OCP-cluster-host>/aiops/cfd95b7e-3bc7-4006-a4a8-a73a79c71255
```

### Procedure

1. Go to **Network > Routes** in the OCP Administration console and identify the **<operator-instance-name>-ibm-hdm-common-ui** route.
2. Click the route to access the Agile Service Manager UI.

### What to do next

#### Administrator access

To access the Agile Service Manager UI as an administrator user, you require administrator privileges. To do so, you pass an authorization code to an Agile Service Manager UI URL as an access GET parameter to the endpoint.

1. Go to **Workloads > Secrets** in the OCP Administration console and search for **<operator-instance-name>-ibm-hdm-common-ui-watson-authorization-code-secret**.
2. Open the secret, and under **Data** reveal or copy the **code** entry.
3. Using a browser, use this string to pass an authorization code to an Agile Service Manager UI URL. For example:

```
https://<OCP-cluster-host>/aiops/cfd95b7e-3bc7-4006-a4a8-a73a79c71255?  
access=<authorization-code>
```

You will remain logged on as administrator until your session is ended by inactivity.

#### Related information

[IBM Watson AIOps: Installing extensions](#)

## Accessing the Topology Viewer in DASH (on-prem)

---

The Netcool Agile Service Manager UI consists of the topology viewer, which you access through DASH.

### Before you begin

Ensure you have the correct user roles and WAS security settings to access the topology viewer in DASH.

#### WebSphere Application Server (WAS) security setting

Set the SSL Quality of Protection setting in the WAS Admin Console to either **TLSv1.2** or **SSL\_TLS**, then restart your instance of WAS.

1. In the WAS Admin Console, go to **SSL certificate and key management > SSL configuration > NodeDefaultSSLSetting > Quality of protection (QoP) settings**.
2. Select one of the following from the **Protocol** drop-down list:
  - **TLSv1.2**
  - **SSL\_TLS**
3. Restart WAS.

#### User roles

Ensure you have the appropriate user roles.

### About this task

The Topology Viewer is accessed through an existing DASH deployment, giving you access to all its functionality. Typically, this would be part of an integration deployment that also gives you the use of other NOI applications.



**Attention:** During startup, topology services may try to connect to the Cassandra datastore before it has fully started, thereby causing an error, as also described in the related [troubleshooting](#) section. It will try again until the datastore is ready, and the error becomes void.

### Procedure

1. Using a compatible browser, open DASH using the DASH URL.

For example:

```
https://<DASH HOST>:<DASH PORT>/ibm/console/
```

2. Log in using your user credentials.

DASH is displayed.

**Tip:** When first logging on, an Information prompt may appear asking you to accept the Agile Service Manager UI certificate. To do so, click on any one of the provided links and accept the certificate when prompted. You may then close the information prompt.

3. In DASH, open the **Incident** menu.
4. Click **Topology Viewer** (under the Agile Service Management heading).

The **Topology Viewer** is displayed.

**Tip:** When opening an Agile Service Manager UI page in DASH using Internet Explorer or Edge, the following certificate failure can occur: Content was blocked because it was not signed by a valid security certificate. To fix this, your administrator must obtain and install the Agile Service Manager root certificate manually, as described in this [troubleshooting](#) topic.

### Results

The Netcool Agile Service Manager UI connects to the topology service, which provides the data needed to render the topology visualization. By default it refreshes the view dynamically every thirty seconds.

### What to do next

Once you have accessed the Topology Viewer, you define the seed resource on which you want to base your topology view, and then choose the level of connected resources that you wish to render.

## Accessing topologies via direct-launch URL string

You view specific topologies using a direct-launch URL by setting the topology navigation settings in the URL parameters to directly load a specific topology configuration within DASH. This gives you and others

quick access to specific views. You can obtain a direct-launch URL string from a defined topology in Topology Viewer, or create it manually.

## Before you begin

To obtain a direct-launch URL from the Topology Viewer accessed through DASH, you must have the appropriate DASH user roles. You must also complete the process of visualizing a topology. To share a direct-launch URL with others, they must be DASH users with the appropriate user roles.

**Restriction:** When launching a direct Netcool Agile Service Manager window in your browser, you **must** log onto DASH and keep the DASH window open as a separate tab in your browser. If you close the DASH browser tab, your DASH session will expire after the session timeout period, and you will be logged out of Agile Service Manager.

## About this task

**From version 1.1.9:** The URL has changed. From Version 1.1.9 onwards a typical URL starts with the Agile Service Manager core host URL and the tenant ID:

```
https://<ASM CORE HOST>/aiops/<TENANT ID>/federated/topology-viewer
```

Example:

```
https://asm.example.host.com/aiops/cfd95b7e-3bc7-4006-a4a8-a73a79c71255/federated/topology-viewer?resourceId=s11RQ5NCR_ykgciQCTkFGA&hops=1&layoutType=5&refreshTime=30000
```

This example includes the **resourceId** to identify the seed resource, **hop** for the number of hops from the seed to be included in the topology, **layoutType** to specify the type of display, and **refreshTime** to specify how often the view is refreshed. In this case the hop number is set at 'one' (note that the hop type is 'element-to-element', which does not have to be specified as it is the default). The layout type is five, and the refresh time is set at 30000 ms (30 seconds).

### Tips:

- For a list of supported parameters, see the following topic: [“Launch-in-context parameters” on page 308](#).
- By specifying 'type', you can include **Advanced** resource or relationship filters in your URL (not available for **Simple** filters). See the following topic for more [information on filtering](#).

## Procedure

### To obtain a topology URL from Topology Viewer in DASH

1. Once you have rendered a topology in the Topology Viewer, use the **Additional Actions** drop-down on the Navigation bar to obtain a direct launch URL string.

### To define a URL manually

2. You can define a URL by editing the parameters, as shown in the following examples.

#### Empty topology URL

To open the Topology Viewer page directly with no configuration, use this type of URL.

```
https://<ASM CORE HOST>/aiops/<TENANT ID>/federated/topology-viewer
```

#### Specify seed resource by ID in the URL

To open the page and create a visualization from a specific seed resource by ID, use this format.

```
https://<ASM CORE HOST>/aiops/<TENANT ID>/federated/topology-viewer?resourceId=<resource_id>
```

#### Specify seed resource by name in the URL

To open the page and create a visualization from a specific seed resource by name, use this format.

```
https://<ASM CORE HOST>/aiops/<TENANT ID>/federated/topology-viewer?resourceName=<name>
```

### Specify seed resource and hops URL

To open the page and create a visualization from a specific seed resource with a specific number of hops and hop type, use this format.

The **hopType** parameter is only required if you wish to use the host-to-host (that is, host) hop type (element-to-element is the default).

```
https://<ASM CORE HOST>/aiops/<TENANT ID>/federated/topology-viewer?
resourceName=<name>&hops=<hops>&hopType=<host>
```

### Specify focus areas for large topologies in the URL

When loaded, this topology view zooms in on the area specified, thus preventing very large topologies from being displayed with potentially too much information on view.

The **focusId** parameter is required for the topology view to zoom into the area of a given resource.

```
https://<ASM CORE HOST>/aiops/<TENANT ID>/federated/topology-viewer?
resourceId=<resource_id>&hops=<hop number>&focusId=<resource_id>
```

### Specify additional resources to be highlighted in the URL

Specify an array of resource IDs to highlight additional resource in the topology view. These selections override any already selected seed resource.

The **selectedIds** parameter is required for the topology view to highlight non-seed resources when the topology is loaded.

```
https://<ASM CORE HOST>/aiops/<TENANT ID>/federated/topology-viewer?
resourceId=<resource_id>&hops=<number>&selectedIds=["<resource_id>" " " ...]
```

### Specify seed resource, hops and get neighbor expansions URL

To open the page and create a visualization from a specific seed resource with a specific number of hops, as well as resource neighbors, use this format.

This example sets the hop type as element-to-element, and it uses the parameter **neighbourRequests**, which expects an array of space separated id strings.

```
https://<ASM CORE HOST>/aiops/<TENANT ID>/federated/topology-viewer?
resourceName=<name>&hops=<hops>&neighbourRequests=["<resource_id>" "<resource id>" ...]
```

### Specify search parameter (partial resource name) in the URL

To open the page and create a visualization from a specific resource search value, use this format.

- If a single resource is found that matches the search parameter, the resource is used as a seed and the topology is drawn.
- If more than one resource is found that matches the parameter, the **Search Results** page is displayed listing possible results, including their names, types and other resource properties.
- If no matches are found, an error message is displayed.

This example searches for a resource by name, and will return results that match the name either fully or partially.

```
https://<ASM CORE HOST>/aiops/<TENANT ID>/federated/topology-viewer?search=<name>
```

### Load topology using unique ID to set seed

```
https://<ASM CORE HOST>/aiops/<TENANT ID>/federated/topology-viewer?
resourceUniqueId=<unique id>
```

### Load topology using advanced resource filters

```
https://<ASM CORE HOST>/aiops/<TENANT ID>/federated/topology-viewer?
resourceFilter=["Type1" "Type2"]
```

### Load topology using advanced relationship filters

```
https://<ASM CORE HOST>/aiops/<TENANT ID>/federated/topology-viewer?  
relationFilter=["Type1" "Type2"]
```

### Show topology at a given time point

```
https://<ASM CORE HOST>/aiops/<TENANT ID>/federated/topology-viewer?  
resourceName=<name>..<any other configuration>..&time=<unixTimeMilliseconds>
```

### Show topology in delta history mode, with 'time' as a reference point and deltaTime as the point to delta against

```
https://<ASM CORE HOST>/aiops/<TENANT ID>/federated/topology-viewer?  
resourceName=<name>..<any other  
configuration>..&time=<unixTimeMilliseconds>&deltaTime=<unixTimeMilliseconds>
```

### Load topology with side toolbar hidden

```
https://<ASM CORE HOST>/aiops/<TENANT ID>/federated/topology-viewer?hideToolbar=<true>
```

### Load topology with top search bar hidden

```
https://<ASM CORE HOST>/aiops/<TENANT ID>/federated/topology-viewer?hideSearch=<true>
```

## What to do next

Having created direct-launch URL topology visualizations, you can save them for quick access to specific views, or share them with others to provide them with specific, instant information.

## Rendering static topologies via UI-API

You can use the UI-API to render a static topology view via cURL command, which you then download as an image. You use the **resourceId** parameter to seed the topology, and then define a number of additional parameters to generate a static server-side topology image.

### Before you begin

To define a cURL string that renders a topology via the UI-API, you must know the URL for your core host, the resource ID for your seed resource, as well as the Agile Service Manager basic access authentication credentials.

### About this task

The Agile Service Manager UI-API is available on the following endpoints:

```
https://<ASM CORE HOST>/1.0/ui-api/render/svg
```

```
https://<ASM CORE HOST>/1.0/ui-api/render/png
```

You construct your cURL command as described in this procedure, then use it to download a static topology view as an image (either PNG or SVG).

In addition to the **resourceId** parameter, you use a **Basic** authentication header when calling the UI-API. This lets you call the API without having to log into the Agile Service Manager UI first.

### PNG example

```
curl -X GET "https://asm.example.host.com/1.0/ui-api/render/png?  
resourceId=s11RQ5NCR_ykgciQCTkFGA" -H "accept: image/png" -H "authorization: Basic  
YXNtOmFzbQ=="
```

## SVG example

```
curl -X GET "https://asm.example.host.com/1.0/ui-api/render/svg?resourceId=s11RQ5NCR_ykgciQCTkFGA" -H "accept: image/svg+xml" -H "authorization: Basic YXNtOmFzbQ=="
```

Where:

- The Agile Service Manager core host is `asm.example.host.com`
- The `resourceId` that identifies the seed resource is `s11RQ5NCR_ykgciQCTkFGA`
- The output format is specified by either `png` or `svg`
- The authorization method is specified as `Basic`
- The encrypted user credentials (username and password) are `YXNtOmFzbQ==`

### Tips:

- You can use the **width** and **height** URL parameters (exclusively for use with the UI-API) to set the size of the image produced; the default size is 1000 x 700.
- The **hideToolbar** and **hideSearch** parameters are **not** available when constructing an UI-API call. Parameters are listed [here](#).
- You can include advanced resource or relationship filters in your URL by specifying 'type'. Topology filtering is described [here](#).
- You can change the label length for your resource types, should the default label legibility be less than optimal in your topology image. Resource labels are described [here](#).

## Procedure

1. Obtain your resource ID, authentication credentials, and core Agile Service Manager host URL.
2. Define a cURL call by editing the parameters, as shown in the following examples.

### Specify a seed resource by ID

The following cURL call example renders a simple Agile Service Manager topology view with a host URL of `asm.example.host.com` and a port number of 3080. The view is based on the **resourceId** parameter with the value of `s11RQ5NCR_ykgciQCTkFGA`. It returns a **png** image, and performs **Basic** access authentication using the encrypted Agile Service Manager username and password credentials (`YXNtOmFzbQ==`).

```
curl -X GET "https://asm.example.host.com:3080/1.0/ui-api/render/png?resourceId=s11RQ5NCR_ykgciQCTkFGA" -H "accept: image/png" -H "authorization: Basic YXNtOmFzbQ=="
```

### Specify the hop numbers and types

The image returned by this cURL call has the same seed resource as before, but now with the number of **hops** set at 3, and a **hopType** of `host`. The **hopType** parameter is only required if you wish to use the host-to-host hop type (element-to-element is the default).

```
curl -X GET "https://asm.example.host.com:3080/1.0/ui-api/render/png?resourceId=s11RQ5NCR_ykgciQCTkFGA&hops=3&hopType=host" -H "accept: image/png" -H "authorization: Basic YXNtOmFzbQ=="
```

### Specify focus areas for large topologies

The image returned by this cURL call has the same seed resource as before (now with a hop count of three), but now zooms in on the area specified by the **focusId** parameter, thereby preventing very large topologies from being displayed with potentially too much information on view.

```
curl -X GET "https://asm.example.host.com:3080/1.0/ui-api/render/png?resourceId=s11RQ5NCR_ykgciQCTkFGA&hops=3&focusId=<resource_id>" -H "accept: image/png" -H "authorization: Basic YXNtOmFzbQ=="
```

### Specify additional resources to be highlighted

The image returned by this cURL call has the same seed resource as before with a hop count of three, but now uses the **selectedIds** parameter to specify additional non-seed resources (via their IDs) to highlight when the image is rendered.

```
curl -X GET "https://asm.example.host.com:3080/1.0/ui-api/render/png?resourceId=s11RQ5NCR_ykgciQCTkFGA&hops=3&selectedIds=<resource_id> <resource_id>" -H "accept: image/png" -H "authorization: Basic YXNtOmFzbQ=="
```

### Specify 'get neighbor' expansion

The image returned by this cURL now includes a **neighbourRequests** parameter that gathers a number of additional neighboring resources for display, each identified by its resource ID.

The **neighbourRequests** parameter expects an array of space-separated id strings.

```
curl -X GET "https://asm.example.host.com:3080/1.0/ui-api/render/png?resourceId=s11RQ5NCR_ykgciQCTkFGA&hops=3&neighbourRequests=<resource_id> <resource_id>" -H "accept: image/png" -H "authorization: Basic YXNtOmFzbQ=="
```

### Specify advanced resource filters

The image returned now includes a **resourceFilter** parameter that filters out all resources of a specified type.

The **resourceFilter** parameter expects an array of space-separated, previously defined resource types.

```
curl -X GET "https://asm.example.host.com:3080/1.0/ui-api/render/png?resourceId=s11RQ5NCR_ykgciQCTkFGA&hops=3&resourceFilter=<Type1> <Type2>" -H "accept: image/png" -H "authorization: Basic YXNtOmFzbQ=="
```

### Specify advanced relationship filters

The image returned includes the relation filter parameter (**relationFilter**) to filter out all relationships of a specified type.

The **relationFilter** parameter expects an array of space-separated, previously defined resource types.

```
curl -X GET "https://asm.example.host.com:3080/1.0/ui-api/render/png?resourceId=s11RQ5NCR_ykgciQCTkFGA&hops=3&relationFilter=<Type1> <Type2>" -H "accept: image/png" -H "authorization: Basic YXNtOmFzbQ=="
```

### Specify time point

The image returned reflects the state of the topology at a specific point in history (in unixtime).

The **time** parameter expects the time entered in unixtime.

```
curl -X GET "https://asm.example.host.com:3080/1.0/ui-api/render/png?resourceId=s11RQ5NCR_ykgciQCTkFGA&hops=3&time=<unixTimeMilliseconds>" -H "accept: image/png" -H "authorization: Basic YXNtOmFzbQ=="
```

### Specify delta history mode

The image returned shows the topology in delta history mode, with the specified **time** as a reference point and the **deltaTime** parameter as the point to delta against.

```
curl -X GET "https://asm.example.host.com:3080/1.0/ui-api/render/png?resourceId=s11RQ5NCR_ykgciQCTkFGA&hops=3&time=<unixTimeMilliseconds>&deltaTime=<unixTimeMilliseconds>" -H "accept: image/png" -H "authorization: Basic YXNtOmFzbQ=="
```

## What to do next

Having created topology images, you can save them for quick access to specific views, or share them with others to provide them with specific, instant information.



## Rendering (visualizing) a topology

---

You define the scope of the topology that you want to render by specifying a seed resource, the number of relationship hops surrounding that resource, as well as the types of hops. The topology service then supplies the data required to visualize the topology.

### Before you begin

To visualize a topology, your topology service must be running, and your Observer jobs must be active.

### About this task

You use this task to render a topology based on a specified seed resource.

**Note:** The UI has a default timeout set at 30 seconds. If service requests are not received in that time, a timeout message is shown, as in the following example:

A time-out has occurred. No response was received from the Proxy Service within 30 seconds.

See [Topology render timeout](#) for more information on addressing this issue.

### Procedure

1. Access the Topology Viewer.

The **Search** page is displayed immediately. From here, you search for a seed resource to build your topology.

2. Find a resource.

#### Search for a resource

The seed resource of the topology visualization.

You define the seed resource around which a topology view is rendered using the **Search for a resource** field. As you type in a search term related to the resource that you wish to find, such as name or server, a drop-down list is displayed with suggested search terms that exist in the topology service.

If the resource that you wish to find is unique and you are confident that it is the first result in the list of search results, then instead of selecting a result from the suggested search terms, you can choose to click the shortcut in the **Suggest** drop-down, which will render and display the topology for the closest matching resource.

If you select one of the suggested results, the **Search Results** page is displayed listing possible resource results.

The Results are listed under separate **Resources** and **Topologies** tabs.

#### Note:

##### Deleted resources

You can select the **Include deleted resources** toggle in order to include deleted resources in the Results page (unless an administrator has disabled that feature).

Deleted resources are marked as such, and the date and time that they were deleted is displayed.

If you click **View topology**, the resource is displayed in the Topology Viewer with the history timeline open, and set to 1ms before deletion. When closed, the topology is cleared.

When you search for resources with a previous point in time set in the history timeline, deleted resources are **not** returned and the toggle to include them is not available.

Deleted templates or defined topologies are **not** included in any search results.

##### Defined topologies

Defined topologies **must** be defined by an administrator user before they are listed.

If you are an administrator defining topology templates in the **Topology template builder**, search results are listed under separate **Resources** and **Templates** tabs.

To add a defined topology search result to the collection of topologies accessible in the Topology Dashboard, tag it as a favorite by selecting the **star** icon next to it.

For each result, the name, type and other properties stored in the Elasticsearch engine are displayed.

If a status other than clear exists for a search result, the maximum severity is displayed in the information returned, and a color-coded information bar above each result displays all non-clear statuses (in proportion).

You can expand a result in order to query the resource or defined topology further and display more detailed, time-stamped information, such as its state and any associated severity levels, or when the resource was previously updated or replaced (or deleted).

You can click the **View Topology** button next to a result to render the topology.

#### Defined topology restriction:

- When you load a predefined topology, it is displayed in a 'defined topology' version of the Topology Viewer, which has restricted functionality. You are unable to follow its neighbors, or change its hops, or make use of its advanced filters.
- You can recenter the defined topology from the context menu, which loads it in the Topology Viewer with all its standard functionality.

#### From the Navigation toolbar, perform the following actions:

3. Select a number between one and four to define the number of relationship hops to be visualized.  
See the [“Defining advanced topology settings” on page 285](#) topic for more information on customizing the maximum hop count.
4. Choose one of the following hop types:
  - The **Element to Element** hop type performs the traversal using all element types in the graph.
  - The **Host to Host** hop type uses an aggregate traversal across elements with the entity type 'host'.
  - The **Element to Host** hop type provides an aggregated hop view like the 'Host to Host' type, but also includes the elements that are used to connect the hosts.
5. Filter the topology before rendering it.  
Open the Filter toolbar using the **Filter** toggle, and apply the filters required. For more information on using filters, see the [Filter the topology](#) section in the 'Viewing a topology' topic.
6. Click **Render** to render the topology.

## Results

The Agile Service Manager topology viewer connects to the topology service and renders the topology. By default the view is refreshed every thirty seconds, unless specified otherwise (by an administrator user).



**Trouble: Topology render timeout:** If you receive a timeout message, this may be due to a number of reasons:

- Large amounts of data being retrieved for complex topologies
- Too many hop counts specified
- Issues with the back-end services

#### Workarounds

- Check that all services are running smoothly. You can verify that the docker containers are running using the following command:

```
$ASM_HOME/bin/docker-compose ps
```

The system should return text indicating that all containers have a state of Up.

- Lower the hop count to reduce the service load. See the [“Defining advanced topology settings” on page 285](#) topic for more information on customizing the maximum hop count.

- **An administrator user** can increase the default 60 seconds timeout limit by changing the following setting in the \$ASM\_HOME/.env file:

```
LAYOUT_SERVICE_TIMEOUT=60
```

If increasing this to two minutes or more, you should also increase the timeouts used by the proxy. Change `proxy_read_timeout 120` in the following files:

- /opt/ibm/netcool/asm/etc/nginx/conf.d/nasm-ui-api.rules
- /opt/ibm/netcool/asm/etc/nginx/conf.d/nasm-layout.rules

You must restart Agile Service Manager for the changes to take effect:

- To stop Agile Service Manager, run:

```
<ASM_HOME>/bin/asm_stop.sh
```

- Once stopped, re-start Agile Service Manager:

```
<ASM_HOME>/bin/asm_start.sh
```

## What to do next

Next, you can refine and manipulate the view for further analysis.

## Viewing a topology

---

Once you have rendered a topology, you can refine and manipulate the view.

### Before you begin

To refine a topology, you must have previously defined a topology, as described in the [“Rendering \(visualizing\) a topology”](#) on page 253 topic.

**Note:** You can change a topology if and as required while viewing or refining an existing topology.

### About this task

You can perform the following actions once you have rendered a topology:

#### View the topology

You can zoom in and out of the specific areas of the topology, and pan across it in various ways.

You can also auto-fit the topology into the available display window, draw a mini map, or redraw the entire topology.

#### Use the Update Manager

With auto-updates turned off, you can work with your current topology until you are ready to integrate the new resources into the view.

#### Filter resources

You can filter the types of resources displayed, or the types of relationships rendered.

## Procedure

### View a topology (created earlier)

1. From the Visualization toolbar below the Navigation toolbar, you can manipulate the topology using a number of visualization tools.

#### Select tool submenu

When you hover over the Select tool icon, a submenu is displayed from which you can choose the **Select**, **Pan** or **Zoom Select** tool.

**Select tool**

Use this icon to select individual resources using a mouse click, or to select groups of resources by creating a selection area (using click-and-drag).

**Pan tool**

Use this icon to pan across the topology using click-and-drag on a blank area of the visualization panel.

**Zoom Select tool**

Use this icon to zoom in on an area of the topology using click-and-drag.

**Zoom In**

Use this icon to zoom in on the displayed topology.

**Zoom Out**

Use this icon to zoom out of the displayed topology.

**Zoom Fit**

Use this icon to fit the entire topology in the current view panel.

**Overview Toggle**

Use this icon to create the overview mini map in the bottom right corner.

The mini map provides an overview of the entire topology while you zoom in or out of the main topology. The mini map displays a red rectangle to represent the current topology view.

**Layout**

Use this icon to recalculate, and then render the topology layout again.

You can choose from a number of layout types and orientations.

**Layout 1**

A layout that simply displays all resources in a topology without applying a specific layout structure.

**Layout 2**

A circular layout that is useful when you want to arrange a number of entities by type in a circular pattern.

**Layout 3**

A grouped layout is useful when you have many linked entities, as it helps you visualize the entities to which a number of other entities are linked. This layout helps to identify groups of interconnected entities and the relationships between them.

**Layout 4**

A hierarchical layout that is useful for topologies that contain hierarchical structures, as it shows how key vertices relate to others with peers in the topology being aligned.

**Layout 5**

A peacock layout is useful when you have many interlinked vertices, which group the other linked vertices.

**Layout 6**

A planar rank layout is useful when you want to view how the topology relates to a given vertex in terms of its rank, and also how vertices are layered relative to one another.

**Layout 7**

A rank layout is useful when you want to see how a selected vertex and the vertices immediately related to it rank relative to the remainder of the topology (up to the specified amount of hops). The root selection is automatic.

For example, vertices with high degrees of connectivity outrank lower degrees of connectivity. This layout ranks the topology automatically around the specified seed vertex.

**Layout 8**

A root rank layout similar to layout 7, except that it treats the selected vertex as the root. This layout is useful when you want to treat a selected vertex as the root of the tree, with others being ranked below it.

Ranks the topology using the selected vertex as the root (root selection: Selection)

### Layout orientation

For layouts 4, 6, 7 and 8, you can set the following layout orientations:

- Top to bottom
- Bottom to top
- Left to right
- Right to left

### History toggle

Use this to open and close the Topology History toolbar. The topology is displayed in history mode by default.

### Configure Refresh Rate

When you hover over the **Refresh Rate** icon, a submenu is displayed from which you can configure the auto-update refresh rate.

You can pause the topology data refresh, or specify the following values: 10 seconds, thirty seconds (default), one minute, or five minutes.

### Resource display conventions

**Deleted:** A minus icon shows that a resource has been deleted since last rendered.

Displayed when a topology is updated, and in the history views.

**Added:** A purple plus (+) icon shows that a resource has been added since last rendered.

Displayed when a topology is updated, and in the history views.

**Added (neighbors):** A blue asterisk icon shows that a resource has been added using the 'get neighbors' function.

### Use the Update Manager

2. If auto-updates have been turned off, the Update Manager informs you if new resources have been detected. It allows you to continue working with your current topology until you are ready to integrate the new resources into the view.

The Update Manager is displayed in the bottom right of the screen.

### Show details

Displays additional resource information.

### Render

Integrates the new resources into the topology.

Choosing this option will recalculate the topology layout based on your current display settings, and may therefore adjust the displayed topology significantly.

### Cogwheel icon

When clicked, provides you with quick access to change your user preferences:

- **Enable auto-refresh:** Switches auto-refresh back on, and disables the Update Manager.
- **Remove deleted resources:** Removes the deleted resources from your topology view when the next topology update occurs.

### Hide

Reduces the Update Manager to a small purple icon that does not obstruct your current topology view.

When you are ready to deal with the new resources, click on the icon to display the Update Manager again.

### Modify a topology

3. The displayed topology consists of resource nodes and the relationship links connecting the resources. You can interact with these nodes and links using the mouse functionality.

### Dragging a node

Click and drag a node to move it.

### Selecting a node

Selection of a node highlights the node, and emphasizes its first-order connections by fading all other resources.

### Context menu (right-click)

You open the context menu using the right-click function. The context menu provides access to the resource-specific actions you can perform.

For resource entities, you can perform the following:

#### Resource Details

When selected, displays a dialog that shows all the current stored properties for the specified resource in tabular and raw format.

When selected while viewing a topology history with Delta mode **On**, the properties of the resource at both the reference time and at the delta time are displayed.








#### Resource Status

If statuses related to a specific resource are available, the resource will be marked with an icon depicting the status severity level, and the Resource Status option will appear in the resource context menu.

When selected, Resource Status displays a dialog that shows the time-stamped statuses related to the specified resource in table format. The Severity and Time columns can be sorted, and the moment that Resource Status was selected is also time-stamped.

In addition, if any status tools have been defined, the status tool selector (three dots) is displayed next to the resource's statuses. Click the status tool selector to display a list of any status tools that have been defined, and then click the specific tool to run it. Status tools are only displayed for the states that were specified when the tools were defined.

The **severity** of a status ranges from 'clear' (white tick on a green square) to 'critical' (white cross on a red circle).

Table 63. Severity levels	
Icon	Severity
	clear
	indeterminate
	information
	warning
	minor
	major
	critical

#### Comments

When selected, this displays any comments recorded against the resource.

By default, resource comments are displayed by date in ascending order. You can sort them in the following way:

- Oldest first
- Newest first
- User Id (A to Z)
- User Id (Z to A)

Users with the `inasm_operator` role can view comments, but not add any. Users with `inasm_editor` or `inasm_admin` roles can also add new comments. See the [“Configuring DASH user roles” on page 32](#) topic for more information on assigning user roles.

To add a new comment, enter text into the New Comment field, and then click **Add Comment** to save.

### Get Neighbors

When selected, opens a menu that displays the resource types of all the neighboring resources. Each resource type lists the number of resources of that type, as well as the maximum severity associated with each type.

You can choose to get all neighbors of the selected resource by adding a **Filter tag**, or only the neighbors of a specific type by adding a **Filter neighbor tag**. This lets you expand the topology in controlled, incremental steps.

Selecting **Get Neighbors** overrides any existing filters.

You can **Undo** the last neighbor request made.

### Follow Relationship

When selected, opens a menu that displays all adjacent relationship types.

Each relationship type lists the number of relationships of that type, as well as the maximum severity associated with each type.

You can choose to follow all relationships, or only the neighbors of a specific type.

### Show last change in timeline

When selected, will display the history timeline depicting the most recent change made to the resource.

### Show first change in timeline

When selected, will display the history timeline depicting the first change made to the resource.

### Find path

When selected, opens the **Path tool** dialog, from where you can define the criteria for the system to find the shortest path between two resources in your topology. See the [“Using the Path tool” on page 267](#) topic for more information.

### Recenter View

When selected, this updates the displayed topology with the specified resource as seed.

## Filter the topology

4. Open the Resource Filter toolbar using the **Filter** toggle in the Topology Visualization toolbar. From here, you can apply filters to the topology in order to refine the types of resources or relationships displayed.

The Filter toolbar is displayed as a panel on the right-hand side of the page, and consists of a **Simple** and an **Advanced** tab. If selected, each tab provides you with access to lists of resource types and relationship types. Any types that have been excluded from view in order to optimize performance when rendering topologies are listed here. The seed resource of a topology **can not** be excluded.

**Note:** You can override the exclusion using the context menu functionality.

- **If you are filtering a topology before rendering it:** All resource types and relationship types are displayed, except those that have been excluded. After rendering the topology, you can toggle the **Show all types** switch so that only types relevant to your topology are displayed.
- **If you are filtering a topology already displayed in the viewer:** Only types relevant to your topology are displayed, for example **host**, **ipaddress**, or **operatingsystem**. You can toggle the **Show all types** switch so that all types are listed (apart from the excluded types).

### Simple tab

When you use the Simple tab to filter out resource or relationship types, all specified types are removed from view, including the seed resource.

It **only** removes the resources matching that type, leaving the resources below, or further out from that type, based on topology traversals.

By default, all types are **On**. Use the **Off** toggle to remove specific types from your view.

#### Advanced tab

The Advanced tab performs a server-side topology-based filter action.

It can exclude or include types:

- **Exclude** removes the resources matching that type, **as well as** all resources below that type. However, the seed resource is **not** removed from view, even if it is of a type selected for removal.
- **Include** displays the selected types in the topology.

#### Tips

**Reset or invert all filters:** Click **Reset** to switch all types back on, or click **Invert** to invert your selection of types filtered.

**Hover to highlight:** When a topology is displayed, hover over one of the filtering type options to highlight them in the topology.

## Viewing topology history

---

You can view a topology dynamically, or use the history timeline function to compare and contrast the current topology with historical versions.

### Before you begin

To refine a topology, you must have previously defined a topology, as described in the [“Rendering \(visualizing\) a topology” on page 253](#) topic.

**Note:** You can change a topology if and as required while viewing or refining an existing topology.

### About this task

**Tip:** The topology is displayed in history mode by default.

### Procedure

1. Open the Topology History toolbar by clicking the **History** toggle in the Topology Visualization toolbar (on the left).
2. You can display and refine topology history in a number of ways.

#### Update mode

The topology is displayed in update mode by default with Delta mode set to **Off**.

While viewing the timeline in update mode with Delta mode set to **On**, any changes to the topology history are displayed on the right hand side of the timeline, with the time pins moving apart at set intervals. By clicking **Render**, you reset the endpoint to 'now' and the pins form a single line again.

While viewing the timeline in update mode with Delta mode set to **Off**, only a single pin is displayed.

#### Delta mode

You toggle between delta mode **On** and **Off** using the Delta switch above the topology.

When Delta mode is **On** with Update mode also **On**, differences in topology are displayed via purple plus or minus symbols next to the affected resource.

When Delta mode is **On** with History mode **On** (that is, Update mode set to **Off**), you can compare two time points to view differences in topology. Historical change indicators (blue dots) are displayed next to each affected resource.



**Note:** For efficiency reasons, historical change indicators are only displayed for topologies with fifty or fewer resources. You can reduce (but not increase) this default by changing the Historical Change Threshold as described in [“Defining advanced topology settings”](#) on page 285.

### **Lock time pin**

Click the **Lock** icon on a time pin's head to lock a time point in place as a reference point, and then use the second time slider to view topology changes.

### **Compare resource properties**

Click **Resource Properties** on a resource's context menu to compare the resource's data at the two selected time points. You can view and compare the resource's property names and values in table format, or raw JSON format.

## **History timeline**

You open the Topology History toolbar using the **History** toggle in the Topology Visualization toolbar (on the left).

**Note:** Resources with a large amount of history may take longer to display. To avoid timeout issues, you can [adjust the default timeout values](#).

You use the time pins to control the topology shown. When you move the pins, the topology updates to show the topology representation at that time.

While in delta mode you can move both pins to show a comparison between the earliest pin and the latest. The timeline shows the historic changes for a single selected resource, which is indicated in the timeline title. You can lock one of the time pins in place to be a reference point.

When you first display the history timeline, coach marks (or tooltips) are displayed, which contain helpful information about the timeline functionality. You can scroll through these, or switch them off (or on again) as required.

To view the timeline for a different resource, you click on it, and the heading above the timeline changes to display the name of the selected resource. If you click on the heading, the topology centers (and zooms into) the selected resource.

The history timeline is displayed above a secondary time bar, which displays a larger time segment and indicates how much of it is depicted in the main timeline. You can use the jump buttons to move back and forth along the timeline, or jump to the current time.

You can use the time picker, which opens a calendar and clock, to move to a specific second in time.

To view changes made during a specific time period, use the two time sliders to set the time period. You can zoom in and out to increase or decrease the granularity using the + and - buttons on the right, or by double-clicking within a time frame. The most granular level you can display is an interval of one second. The granularity is depicted with time indicators and parallel bars, which form 'buckets' that contain the recorded resource change event details.

The timeline displays changes to a resource's state, properties, and its relationships with other resources. These changes are displayed through color-coded bars and dash lines, and are elaborated on in a tooltip displayed when you hover over the change. You can exclude one or more of these from display.

### **Resource state changes**

The timeline displays the number of state changes a resource has undergone.

### **Resource property changes**

The timeline displays the number of times that resource properties were changed.

Each time that property changes were made is displayed as one property change event regardless of whether one or more properties were changed at the time.

### **Resource relationship changes**

The number of relationships with neighboring resources are displayed, and whether these were changed.

The timeline displays when relationships with other resources were changed, and also whether these changes were the removal or addition of a relationship, or the modification of an existing relationship.

## Rebuilding a topology

---

Once you have rendered a topology, you can search for (or define) a new seed resource and build a topology around it, change the number of hops rendered, and switch between element-to-element, host-to-host and element-to-host hop types.

### Before you begin

To refine a topology, you must have previously defined a topology, as described in the [“Rendering \(visualizing\) a topology”](#) on page 253 topic.

**Note:** You can change a topology if and as required while viewing or refining an existing topology.

### About this task

**Tip:** For information on re-indexing the Search service, see the 'Re-indexing Search' information in the task troubleshooting section of this topic.

### Procedure

From the Navigation toolbar, you can again search for a resource around which to build a topology, change the number of hops and the type of hop, and re-render the topology.

#### Topology Search

If you conduct a resource search from the navigation toolbar with a topology already loaded, the search functionality searches the loaded topology as well as the topology database.

As you type in a search term, a drop-down list is displayed that includes suggested search results from the displayed topology listed under the **In current view** heading.

If you hover over a search result in this section, the resource is highlighted in the topology window.

If you click on a search result, the topology view zooms in on that resource and closes the search.

#### No. Hops

The number of relationship hops to visualize from the seed resource, with the default set at 'one'.

You define the number of relationship hops to be performed, which can be from zero to four, unless this setting has been customized. See the [“Defining advanced topology settings”](#) on page 285 topic for more information on customizing the maximum hop count.

#### Type of Hop

The type of graph traversal used.

The options are:

##### Element to Element hop type

This type performs the traversal using all element types in the graph.

##### Host to Host hop type

This type generates a view showing host to host connections.

##### Element to Host hop type

This type provides an aggregated hop view like the Host to Host type, but also includes the elements that are used to connect the hosts.

**Tip:** The URL captures the hopType as 'e2h'. When launching a view using a direct URL, you can use the hopType=e2h URL parameter.

#### Filter toggle

Use this icon to display or hide the filter toolbar. You can filter resources that are displayed in the topology, or set filters before rendering a topology to prevent a large, resource-intensive topology from being loaded.

If a filter has been applied to a displayed topology, the text 'Filtering applied' is displayed in the status bar at the bottom of the topology.

### Render

This performs the topology visualization action, rendering the topology based on the settings in the navigation toolbar.

Once rendered, the topology will refresh on a 30 second interval by default. You can pause the auto-update refresh, or select a custom interval.

**Tip:** The UI can time out if a large amount of data is being received. See the [timeout troubleshooting](#) section in the following topic for information on how to address this issue, if a timeout message is displayed: “Rendering (visualizing) a topology” on page 253

## Performing topology administration

---

From the Topology Viewer, you can obtain direct-launch URLs, perform a system health check, and set user preferences.

### Before you begin

Access the Topology Viewer.

### About this task

You can perform the following admin actions:

#### Share direct launch URL

You can copy and save a URL to quickly access a currently defined topology view.

#### Export a topology snapshot

You can share a snapshot of a topology in either PNG or SVG format.

#### View system health

You can view your system's health.

#### Set user preferences

You can set user preferences that define the default settings for rendering your topology.

### Procedure

You perform the following actions from the **Navigation bar > Additional actions** or **Navigation bar > Sharing options** menus.

#### Sharing options

You can share a topology either by obtaining a direct URL linking to the topology view, or by exporting a view of the topology as an image.

##### Obtain Direct URL

Open the **Sharing options** drop-down menu, and then use the **Obtain Direct URL** option to display the **Direct Topology URL** dialog.

The displayed URL captures the current topology configuration, including layout type (layout orientation is not tracked).

Click **Copy** to obtain a direct-launch URL string, then click **Close** to return to the previous screen.

Use the direct-launch URL for quick access to a given topology view within DASH.

**Tip:** You can share this URL with all DASH users with the required permissions.

##### Export as PNG / SVG

You can share a snapshot of a topology in either PNG or SVG format, for example with someone who does not have DASH access.

Open the **Sharing options** drop-down menu, and then use either the **Export as PNG** or the **Export as SVG** option.

Specify a name and location, then click **Save** to create a snapshot of your topology view.

You can now share the image as required.

#### **Additional actions > View System Health**

Open the **Additional actions** drop-down menu, and then use the **View System Health** option to access your Netcool Agile Service Manager deployment's system health information.

#### **Additional actions > Edit User Preferences**

Open the **Additional actions** drop-down menu, and then use the **Edit User Preferences** option to access the **User Preferences** window. Click **Save**, then **Close** when done.

You can customize the following user preferences to suit your requirements:

#### **Updates**

##### **Default auto refresh rate (seconds)**

The rate at which the topology will be updated.

The default value is 30.

You must reopen the page before any changes to this user preference take effect.

##### **Maximum number of resources to load with auto refresh enabled**

When the resource limit set here is reached, auto-refresh is turned off.

The maximum value is 2000, and the default is set to 500.

**Tip:** If you find that the default value is too high and negatively impacts your topology viewer's performance, reduce this value.

#### **Auto render new resources**

Enable this option to display new resources at the next scheduled or ad-hoc refresh as soon as they are detected.

#### **Remove deleted topology resources**

Enable this option to remove deleted resources at the next scheduled or ad-hoc refresh.

#### **Layout**

Set **Default layout type** including the layout orientation for some of the layout types. You can also configure a default layout in User Preferences.

You can choose from a number of layout types, and also set the orientation for layouts 4, 6, 7 and 8.

**Tip:** A change to a layout type is tracked in the URL (layout orientation is not tracked). You can manually edit your URL to change the layout type display settings.

The following numbered layout types are available:

##### **Layout 1**

A layout that simply displays all resources in a topology without applying a specific layout structure.

##### **Layout 2**

A circular layout that is useful when you want to arrange a number of entities by type in a circular pattern.

##### **Layout 3**

A grouped layout is useful when you have many linked entities, as it helps you visualize the entities to which a number of other entities are linked. This layout helps to identify groups of interconnected entities and the relationships between them.

##### **Layout 4**

A hierarchical layout that is useful for topologies that contain hierarchical structures, as it shows how key vertices relate to others with peers in the topology being aligned.

##### **Layout 5**

A force-directed (or 'peacock') layout is useful when you have many interlinked vertices, which group the other linked vertices.

### Layout 6

A planar rank layout is useful when you want to view how the topology relates to a given vertex in terms of its rank, and also how vertices are layered relative to one another.

### Layout 7

A rank layout is useful when you want to see how a selected vertex and the vertices immediately related to it rank relative to the remainder of the topology (up to the specified amount of hops). The root selection is automatic.

For example, vertices with high degrees of connectivity outrank lower degrees of connectivity. This layout ranks the topology automatically around the specified seed vertex.

### Layout 8

A root rank layout similar to layout 7, except that it treats the selected vertex as the root. This layout is useful when you want to treat a selected vertex as the root of the tree, with others being ranked below it.

Ranks the topology using the selected vertex as the root (root selection: Selection)

### Layout orientation

**For layouts 4, 6, 7 and 8**, you can set the following layout orientations:

- Top to bottom
- Bottom to top
- Left to right
- Right to left

### Misc

#### Information message auto hide timeout (seconds)

The number of seconds that information messages are shown for in the UI.

The default value is 3.

**Tip:** If you are using a screen reader, it may be helpful to increase this value to ensure that you do not miss the message.

#### Screen reader support for graphical topology

You can enable the display of additional Help text on screen elements, which can improve the usability of screen readers.

You must reopen the page before any changes to this user preference take effect.

#### Enhanced client side logging, for problem diagnosis

If enabled, additional debug output is generated, which you can use for defect isolation.

**Tip:** Use this for specific defect hunting tasks, and then disable it again. If left enabled, it can reduce the topology viewer's performance.

You must reopen the page before any changes to this user preference take effect.

## Using the topology dashboard

---

You can use the Topology Dashboard to tag, view and access your most commonly used 'favorite' defined topologies.

### About this task

The Topology Dashboard presents a single view of all defined topologies that have been tagged as favorites. They are displayed as a collection of circles, each surrounded by a color-coded band that indicates the states of all the constituent resources in proportional segments. In addition, each defined topology may display an icon, if assigned by the owner. From here, you can access each defined topology for further investigation or action.

## Procedure

### Tagging a defined topology as a favorite

1. As the admin user, log into your DASH web application, then select **Administration** from the DASH menu.
2. Select **Topology Dashboard** under the Agile Service Management heading.  
All defined topologies tagged as 'favorites' are displayed.
3. To add defined topologies as favorites, enter a search term in the **Search for a defined topology** field.  
The **Search Results** page is displayed listing possible results.
4. To save a topology as a favorite, click the star icon next to it on the **Search Results** page.  
You can remove the favorite tag by deselecting the star again.

### Using the Topology Dashboard to view favorites

5. View defined topology information on the **Topology Dashboard**.

Option	Description
<b>At a glance</b>	If you hover over a specific defined topology, an extended tooltip is displayed listing the number of resources against their states, which is also proportionally represented by the color-coded band.
<b>More details</b>	<p>If you click a defined topology, it is displayed in the bottom half of the screen, and the window with the defined topology favorites is moved to the top half of the screen.</p> <p>For the displayed topology:</p> <ul style="list-style-type: none"><li>• The state of each resource or relationship is displayed in color, and you can use the context (right-click) menu to obtain further information.</li><li>• In the upper 'favorites' display window, all defined topologies that intersect with the selected topology remain in focus, while the others are grayed out.</li><li>• If you select a specific resource in your displayed topology, only the displayed favorites that contain the resource remain in focus. You can remove the favorite tag by deselecting the star displayed in the top right corner of the displayed topology.</li><li>• A breakdown of all states for the displayed topology during a historical period is displayed in the Status severity breakdown chart. The bars in the chart are hoverable, displaying additional information, such as the number of relationships added or removed, or the number of property changes. You can adjust (via toggle) the historical period from one hour to seven days.</li><li>• A visualisation of the churn of the displayed topology is shown in the Topology Churn chart. The churn represents changes in the properties or relationships for a historical period. The lines on the chart are hoverable, displaying additional information, and you can adjust the historical period.</li></ul>

Searching for and viewing defined topologies (or resources) is described in more detail in the [Search](#) section of the Topology Viewer reference topic, and in the [Chapter 6, “Using Agile Service Manager,”](#) on [page 245](#) topics.

**Remember:** Defined topologies are updated automatically when the topology database is updated, and generated dynamically as they are displayed in the Topology Viewer.

### Related tasks

[“Using templates to generate defined topologies” on page 315](#)

You can create topology templates to generate defined topologies, which search your topology database for instances that match its conditions. These defined topologies can then be searched for in the Topology Viewer and displayed.

## Using the Path tool

---

You can use the Path tool to discover the shortest path between two specified resources.

### About this task

To discover the shortest path between two resources, you select the 'start' resource in the Topology Viewer, then open the Path tool from the context menu, where you select the 'end' resource and specify additional criteria to discover the shortest path.

### Procedure

1. From the Topology Viewer, identify your start resource.
2. Select the start resource, then choose **Find path** from the context (right-click) menu.  
The **Path tool** dialog is displayed.
3. Use the **Search for the end path resource** field to search for and select the end resource.
4. Choose to **Exclude** or **Include** the resource or relationship types from the path.  
Identify the types to be included/excluded from the **Resource types to filter** and **Relationships types to filter** drop-down lists (which include all defined types).
5. Select the **Maximum number of hops** permitted when retrieving the path. The values range from one hop to thirty.
6. Select **Include shortcut relationships** to consider the relationships from a resource to a parent, when available.
7. Click **Find path**.  
The time taken to calculate the shortest path depends on the complexity of the links between the selected resources.  
A separate Topology Viewer tab labeled **Path topology** opens, which will display the shortest path or paths between the start and end resources, up to a maximum of five paths.  
**Tip:** You can toggle back to, and work with, your initial topology view while the system calculates the path in the background.
8. View or refine the displayed path.
  - When the **Path topology** tab displays the shortest paths, the start and end resources are highlighted (via blue circle).
  - When you hover over a link between resources, all the paths that include these resources are highlighted.
  - You can further filter the displayed paths using the [Filter](#) functionality.

### Related tasks

[“Rendering \(visualizing\) a topology” on page 253](#)

You define the scope of the topology that you want to render by specifying a seed resource, the number of relationship hops surrounding that resource, as well as the types of hops. The topology service then supplies the data required to visualize the topology.

[“Viewing a topology” on page 255](#)

Once you have rendered a topology, you can refine and manipulate the view.





---

## Chapter 7. Administration

Use the following topics to understand administration tasks, such as monitoring system health and logging.

### Configuring authentication

---

You can customize secure access to, and between, core services.

#### Configuring the authentication method when connecting to core services

You can change the authentication method used to access core services.

##### Before you begin

Agile Service Manager must be installed and running.

To configure the authentication method when connecting to core services, you use the username and password supplied during the Agile Service Manager installation, or use the default asm username and password.

##### About this task

Agile Service Manager services have access to two mechanisms of authentication:

###### Basic authentication (default)

Environment variable based using the **ASM\_USER** and **ASM\_PASS** variables configured in `$ASM_HOME/.env`

###### LDAP

LDAP authentication configured using environment variables in `$ASM_HOME/.env`

##### Procedure

###### Switch authentication method

1. The values of the **ASM\_AUTHENTICATOR** parameter in `.env` determine the authentication method:
  - `com.ibm.itsm.topology.service.auth.BasicAuthenticator`
  - `com.ibm.itsm.topology.service.auth.LdapAuthenticator`

###### Configure authentication access details

2. Once you have defined the authentication method, configure access details:
  - If using **basic authentication**, edit `$ASM_HOME/.env` and change the values for **ASM\_USER** and **ASM\_PASS**
  - If using **LDAP authentication**, edit `$ASM_HOME/.env` and change the following values:

###### Service name

`LDAP_SERVICE_NAME: ${LDAP_SERVICE_NAME:-localhost}`

###### Service port

`LDAP_SERVICE_PORT: ${LDAP_SERVICE_PORT:-389}`

###### SSL

`LDAP_USE_SSL: ${LDAP_USE_SSL:-0}`

###### LDAP base

`LDAP_BASE_DN: ${LDAP_BASE_DN:-dc=example,dc=com}`

###### Supply authentication credentials to Swagger

3. To run REST requests via Swagger, you must use the configured username and password.

**Tip:** When changing user credentials in Swagger, the interface sometimes prevents you from logging out.

**Workaround:** Restart your browser.

#### Related information

[apache.org httpd - Manage user files for basic authentication](http://apache.org/httpd/docs/auth/htpasswd.html)

## Changing default passwords

Passwords are encrypted by default and stored in a number of environment (.env) files.

### About this task

Passwords are stored in the .env files located in the \$ASM\_HOME directory.

Using the \$ASM\_HOME/bin/encrypt\_password.sh script, you can change the following passwords:

- ASM\_PASS
- KAFKA\_ADMIN\_PASSWORD
- KAFKA\_CLIENT\_PASSWORD
- ZK\_KAFKA\_CLIENT\_PASSWORD
- CASSANDRA\_PASS
- COUCHDB\_PASSWORD
- REDIS\_PASSWORD
- DASH\_PASS

**Tip:** These passwords are encrypted (by default), which is indicated by the **\_ENCRYPTED** postfix appended to them.

### Procedure

#### To change a password

1. Use the \$ASM\_HOME/bin/encrypt\_password.sh script to encrypt the password.  
The encryption tool will prompt you to confirm the new password, and then display an encrypted version of the password.
2. Update the appropriate .env file with the encrypted password.
3. Restart Agile Service Manager using the \$ASM\_HOME/bin/asm\_start.sh script.

#### To configure the use of unencrypted passwords

**Note:** Not recommended for production systems.

4. Remove the **\_ENCRYPTED** postfix from the environment variable.
5. Change the password value to plain text.

### Related concepts

[“Installing Agile Service Manager on-prem” on page 13](#)

To install Netcool Agile Service Manager, you complete a number of prerequisites tasks. You then install the Netcool Agile Service Manager core components, UI and observers, before deploying the Netcool Hybrid Deployment Option Integration Kit using the IBM Installation Manager.

## Customizing UI elements

---

You can customize a number of Agile Service Manager UI elements for your deployment, such as tooltips, link styles and icons. You can also create custom tools which users can access through a topology's context menu.

**Tip:** The Topology Viewer loads all UI configuration settings into memory when it opens. If you make changes to your UI configurations after opening the Topology Viewer, you must reopen it before these changes take effect.

## Configuring custom tools

As an administrator or advanced user, you can create custom topology tools, which users can then access from within a topology's context menu. This functionality allows you to access properties of a selected item (such as a resource or relationship) and execute some custom functionality within the context of that item.

### Before you begin

- To access the Topology Tools page, you must have the admin role **inasm\_admin** assigned to you. See the [“Configuring DASH user roles” on page 32](#) topic for more information.
- You must be proficient in JavaScript to define custom tools.

### About this task

Custom tools are written in JavaScript, and accessed through the right-click (context) menu in the UI. All **inasm\_operator** users can access the tools, but only **inasm\_admin** users can customize them.



#### Trouble:

##### Some topology tools fail to load external content

New security measures implemented in recent versions of Firefox and Chrome to prevent click-jacking and similar risks may also prevent the loading of third party content over HTTPS into an iframe.

This mechanism may be used by custom Agile Service Manager topology tools to display contextual information, and as such may be affected.

#### Workaround:

##### Firefox

The browser presents the option to load the content in a separate window.

##### Chrome

None

**Tip:** The Refresh icon reloads the tools list. This can be useful if other users are customizing the tools.

## Procedure

1. As the admin user, log into your DASH web application.
2. Select **Administration** from the DASH menu.
3. Select **Topology Tools** under the Agile Service Management heading.
4. Use the following information to complete the **Topology Tools - Details** page.

#### Name

Unique name used as an internal reference.

Required.

#### Menu label

The menu label is the text displayed in the context menu.

This can be the same name as used by other tools, which is why the unique name is required.

Required

#### Description

A description to help administrator users record the tool's purpose.

Not displayed in the context menu.

Optional.

#### Menu priority

The menu priority slider defines where in the context menu the tool is displayed.

For example, tools with a priority of two will be displayed higher in the menu than tools that have a priority of four.

Available values are one to ten.

Optional.

#### Navigation

You can move to the next page by using the page selector.

The minimum requirement to save the tool and open the **Topology Tools - Implementation** page is the name and label.

5. Use the following information to complete the **Topology Tools - Implementation** page.

Here you create the JavaScript implementation for the tool, which defines the action that will occur when a user selects this option from the menu. JavaScript examples are included after these steps. To help you create custom Agile Service Manager tools, you have access to the following custom JavaScript helper functions:

#### asmProperties

The tool implementation has access to the properties of the relevant **resource**, **relationship** or **status** via the `asmProperties` JavaScript object, which contains all the properties.

You can access the properties using standard JavaScript, but you must protect against a value not being present.

For example if you intend to use the property 'latitude', you must verify that it is present before using it. To do so, use the following check command:

```
asmProperties.hasOwnProperty('latitude')
```

If the property is present, the Boolean value `true` will be returned.

#### Status tools properties

When creating **status** tools, you use JavaScript that is similar to the script that you use when creating **resource** or **relationship** tools. However, the properties you use in your status tool scripts, such as `asmProperties`, reference the properties for the **status** item; unlike the properties you use in your resource or relationship tool scripts, which reference the properties for the resources or relationships. For example, if you use `asmProperties.location` in a status tool script, there must be a corresponding 'location' property in the status record.

When creating status tools, the `asmProperties` object has a property that takes the form of an array called **resources**, which represents the resources in the topology with which this status is associated. Each item in the resources array is an object with properties that represent the properties of that resource. For example, if a status is associated with two resources, the **uniqueId** property of the first of those two resources could be referenced in the script by using `asmProperties.resources[0].uniqueId`

In addition, you can access the properties of a resource against which you are running a status tool by using the **asmSourceProperties** object when scripting the status tool.

## asmSourceProperties

You can access information about the source properties of any **relationships** or **status** the custom tool is acting on via the asmSourceProperties JavaScript object.

Example of using the source resource properties in a custom relationship stroke definition:

```
if (asmSourceProperties.myProp === 'high') {  
    return 'blue';  
} else {  
    return 'black';  
}
```

**Remember:** The arrows indicating a relationship point from the source to the target.

## asmTargetProperties

You can access information about the target properties of **relationships** the custom tool is acting on via the asmTargetProperties JavaScript object.

## asmFunctions

You can use a number of other helper functions, which are accessed from the asmFunctions object, which includes the following:

### showConfirmationPopup(title, message, onOk)

Creates a popup confirmation allowing the tool to confirm an action.

Takes a title and message, which is displayed on the popup, and a function definition, which is run if the user clicks the OK button on the popup.

### showToasterMessage(status, message)

Shows a popup toaster with the appropriate status coloring and message.

### showPopup(title, text)

Shows a popup with a given title and text body (including markdown), which can be generated based on the properties of the resource or relationship.

**Tip:** The **asmFunctions.showPopup** helper function lets you use markdown to create more sophisticated HTML popups. For more information on markdown syntax, consult a reputable markdown reference site.

### showIframe(url)

Displays a popup filling most of the page which wraps an iframe showing the page of the given URL.

Allows you to embed additional pages.

### sendPortletEvent(event)

Allows you to send DASH portlet events from the Topology Viewer that can be used to manipulate other DASH portlets, such as the Event Viewer within IBM Tivoli Netcool/OMNIBus Web GUI.

**Note:** You can send events to other DASH portlets only if you are running Agile Service Manager within DASH (rather than in a direct-launch browser window), and if the receiving DASH portlets subscribe to the types of events being sent. See the [“sendPortletEvent examples” on page 277](#) topic for more information.

### getResourceStatus(<resource\_id>, <callback\_function>, [<time\_stamp>])

Allows you to request status information from a tool definition for a given resource using its **\_id** parameter.

#### resource\_id

Required

Can be obtained from a resource via `asmProperties._id` and from a relationship using `asmSourceProperties._id` or `asmTargetProperties._id`

#### callback\_function

Required

Is called once the status data has been collected from the topology service, with a single argument containing an array of status objects

#### **time\_stamp**

Optional

Unix millisecond timestamp to get the status from a given point in history

The following example prints the status information of a source resource from a relationship context to the browser console log:

```
let printStatusCallback = function(statuses) {
    statuses.forEach(function(status) {
        console.log('status:', status.status,
                    'state:', status.state,
                    'severity:', status.severity,
                    'time:', new Date(status.time));
    })
}
asmFunctions.getResourceStatus(asmSourceProperties._id,
printStatusCallback);
```

#### **sendHttpRequest(url, options)**

Lets you send an HTTP or HTTPS request to a remote web server using the Agile Service Manager backend server rather than the browser, thereby avoiding any browser domain-blocking.

##### **url**

Required

The full URL of the remote site to be accessed.

For example:

```
https://data-svr-01.uk.com/inv?id=1892&offset=0
```

**Restriction:** You must add any websites referenced by the url parameter to a list of trusted sites as described in the [“Defining advanced topology settings” on page 285](#) topic (in this example data-svr-01.uk.com).

##### **options**

Optional

##### **method**

HTTP method used:

- GET
- POST
- PUT
- DELETE

The default is GET

##### **headers**

An object defining any special request headers needed

##### **body**

A string containing the body data for the request

POST and PUT requests only

##### **autoTrust**

A flag to indicate if the remote web server can be automatically trusted

This flag is **required** if the web site uses a self-signed SSL certificate with no CA, or a CA that is unknown to the Agile Service Manager server.

True or false, with a default of false

##### **onSuccess**

A callback function to run if the HTTP request is successful

This function will be passed the following three parameters:

- Response text
- HTTP status code
- Response headers

#### **onError**

A callback function to run if the HTTP request fails

This function will be passed the following three parameters:

- Response text
- HTTP status code
- Response headers

#### **Options parameter script sample:**

```
{
  method: 'GET',
  headers: {
    Content-Type: 'application/json',
    X-Locale: 'en'
  },
  body: '{ "itemName": "myData1" }',
  autoTrust: true,
  onSuccess: _onSuccessCallback,
  onError: _onErrorCallback
}
```

6. Use the following information to complete the **Topology Tools - Conditions** page.

Here you define the resource or relationship conditions under which this tool is available.

#### **Applicable item type for tool definition**

From this drop-down, select the types to which the tool is applicable: **Resource**, **Relationship**, **Resource and Relationship**, or **Status**.

Depending on your selection, a number of check boxes are displayed, which you use to configure which resources, relationships or states are included.

#### **All types / All states**

Select this option if you want the tool to be displayed for all resource and relationship types, or all states (for Status).

The tool will also be displayed for any specific types not listed here.

#### **Resource types**

Select one or more resource types from the list displayed.

#### **Relationship types**

Select one or more relationship types from the list displayed.

#### **Status**

Select from the following possible states for which the tool will be available:

- **Open**
- **Clear**
- **Closed**

**Remember:** When creating status tools, the properties you use in your status tool scripts reference the properties for the status item, while the properties you use in your resource or relationship tools reference the properties for the resources or relationships.

## **Results**

Once you have saved your changes, you must close the Topology Viewer and then reopen it to make the new tools available (tools will be available for use depending on the conditions set).

## Example

### Example of a lookup tool based on a 'person' resource with the properties 'name' and 'email'.

You first set the **Topology Tools - Conditions** page to use **Resource** as the applicable item type, and then select the **person** resource type only to ensure that the custom tool is only displayed for 'person' resources.

This example demonstrates the JavaScript for a tool that checks for properties, makes an HTTP request, and handles the response either by creating a popup with the response, or by showing an error message.

**Note:** The function `asmHelperFunctions.showToasterMessage` accepts the following values for the status:

- information
- normal
- warning
- escalated
- critical

```
// Check that the resource has the properties email and name
if(asmProperties.hasOwnProperty('email') && asmProperties.hasOwnProperty('name')){
    var emailAddress = asmProperties.email;
    var personName = asmProperties.name;

    // Build a http request to collect information from another service using the email
    var request = new XMLHttpRequest();
    var url = 'https://my-lookup-service?email=' + emailAddress;

    request.open('GET', url, true);

    request.onreadystatechange = function() {
        if (request.readyState === 4) {
            if (request.status === 200) {
                // On successful request show popup in UI window with custom title and
                the json response from the request
                var title = personName + ' Details: ';
                var content = JSON.stringify(JSON.parse(this.responseText), null, 4);
                asmFunctions.showPopup(title, content);
            } else {
                // On request error show popup in UI with custom message.
                var messageStatus = 'critical';
                var message = 'API error, ' + this.responseText;
                asmFunctions.showToasterMessage(messageStatus, message);
            }
        }
    };

    request.send();
} else {

    // If resource doesn't have the expected properties, show toaster message with warning.
    var messageStatus = 'warning';
    var message = 'Unable to load service information, email address not provided.';
    asmFunctions.showToasterMessage(messageStatus, message);
}
```

**Example of a tool using the `asmFunctions.sendHttpRequest(url, options)` function:** This example sends an HTTP GET request for external REST data, and then displays the retrieved information in a popup dialog.

```
// Get the unique id from the resource and use it to send a REST request
var extId = asmProperties.uniqueId;
if (!extId || extId === '') {
    asmFunctions.showToasterMessage('warning', 'The resource does not have a unique identifier. ');
} else {
    var options = {
        method: 'GET',
        onSuccess: _onSuccess,
        onError: _onError
    };
    asmFunctions.sendHttpRequest('https://api.data.uk/resource/' + extId, options);
}
```



```

}

function _onSuccess(response, status, headers) {
    // Start to build text for the info dialog
    var output = 'Here are details for ' + asmProperties.name + ':\n\n';

    var data;
    try {
        data = JSON.parse(response);
    } catch (e) {
        data = {};
    }
    var props = data.additionalProperties;
    if (!props || props.length === 0) {
        asmFunctions.showToastMessage('information', 'No information exists for this
resource.');
```

```

    } else {
        for (var idx in props) {
            var prop = props[idx].key;
            var val = props[idx].value;
            if (!prop.endsWith('Id')) {
                output += props[idx].key + ' = ' + props[idx].value + '\n';
            }
        }

        // Show the output in an ASM popup dialog
        asmFunctions.showPopup('External Details', output);
    }
}

function _onError(response, status, headers) {
    if (!response || response === '') {
        asmFunctions.showToastMessage('critical', 'HTTP request failed with status code ' +
status);
    } else {
        asmFunctions.showToastMessage('critical', 'HTTP request failed: ' + response);
    }
}

```

## sendPortletEvent examples

Using the Agile Service Manager custom tool functionality, you can send Agile Service Manager DASH portlet events to other DASH portlets, provided these subscribe to the types of events being sent. You use the `sendPortletEvent` helper function to define these portlet events. You can only send events to other DASH portlets if you are running Agile Service Manager within DASH, rather than in a direct-launch browser window.

### Example 1: Topology Viewer and DASH Web Widget on same page

This event opens an internal personnel directory in the web widget for the clicked-on 'person' resource.

```

var address = 'http://adminsyst:8080/staff?name=' + asmProperties.name;

var eventPayload = {
    'name': 'http://ibm.com/TIP#DisplayURL',
    'URL': address
};

asmFunctions.sendPortletEvent(eventPayload);

```

### Example 2: Topology Viewer and Event Viewer on same page

This event updates the Netcool/OMNIBus Web GUI Event Viewer to display events where 'Node' matches the clicked resource name. It displays the Agile Service Manager Topology Viewer and Netcool/OMNIBus Event Viewer on the same page.

```

var whereClause = 'Node = \'' + asmProperties.name + '\'';

var eventPayload = {
    name: "http://ibm.com/tip#NodeClickedOn",
    payload: {
        product: {

```

```

        OMNIBusWebGUI: {
            displaycontext: {
                "filterName": "HostEvents",
                "filterType": "user_transient",
                "registerFilter": "true",
                "sql": whereClause,
                "forceOverwrite": "true",
                "viewName": "Default",
                "viewType": "global",
                "dataSource": "OMNIBUS"
            }
        }
    }
};

asmFunctions.sendPortletEvent(eventPayload);

```

### Example 3: Event Viewer on it own page

This event updates the Netcool/OMNIBus Web GUI Event Viewer to display events where 'Node' matches the clicked resource name. It displays the Netcool/OMNIBus Event Viewer on its own page.

```

var whereClause = 'Severity > 2 and Node = \'' + asmProperties.name + '\'';

var eventPayload = {
    name: "http://ibm.com/isclite#launchPage",
    NavigationNode: "item.desktop.navigationElement.EventViewer",
    switchPage: true,
    payload: {
        product: {
            OMNIBusWebGUI: {
                displaycontext: {
                    "filterName": "HostEvents",
                    "filterType": "user_transient",
                    "registerFilter": "true",
                    "sql": whereClause,
                    "forceOverwrite": "true",
                    "viewName": "Default",
                    "viewType": "global",
                    "dataSource": "OMNIBUS"
                }
            }
        }
    }
};

asmFunctions.sendPortletEvent(eventPayload);

```

### Related tasks

[“Configuring custom tools” on page 271](#)

As an administrator or advanced user, you can create custom topology tools, which users can then access from within a topology's context menu. This functionality allows you to access properties of a selected item (such as a resource or relationship) and execute some custom functionality within the context of that item.

### Splunk integration example

Using the Agile Service Manager custom tool functionality, you can create a custom tool integration with Splunk. The example in this topic lets you search a Splunk repository via the right-click menu of a resource in the topology viewer.

### Splunk custom tool: 'Search' queries

**Tip:** You must have Splunk Enterprise and Splunk Forwarder installed before you can define a custom tool. The tool sample provided was tested with Splunk Version 8.1.0. See the [Splunk website](#) for more information.

The following sample code creates a custom (right-click) tool that lets you use the properties of the resource from which you launch the tool to search a specific Splunk repository.

This specific tool searches Splunk for logs matching the name of an Agile Service Manager resource. This lets you find out whether Splunk has registered any log information about the resource from which you are launching the tool.

To customize this sample code, change the values for the following properties:

**splunkUrl**

The base URL of your Splunk instance

**splunkRepository**

The repository you want to search

**searchValue**

The dynamically created searchTerm used in Splunk

**Note:** You can use any Agile Service Manager properties associated with the resource to create your tool.

You can also define the tool's response when its search is unsuccessful.

**Sample code:**

```
// Set the Splunk instance url
var splunkUrl = 'host:port';
// This will store the value that you wish to search in Splunk
var searchValue = '';
// This will search results for defined period
// 0 = all time
// -15m = last 15 mins
// -24h = last 24 hours
// -7d = last 7 days
var searchEarliest = "0";
// This flag is set if you find a valid search value
var foundValidSearchValue = false;
if (asmProperties && asmProperties.name) {
    searchValue = "\"" + asmProperties.name + "\"";
    foundValidSearchValue = true;
} else {
    var status = 'critical';
    var message = 'Unable to find name property for Splunk search request';
    asmFunctions.showToastMessage(status, message);
}
// DO NOT MODIFY THE LINES BELOW
if (foundValidSearchValue) {
    var searchLatest = '';
    if(searchEarliest!="0") {
        searchLatest="now"
    }
    var encodedSearchValue = encodeURIComponent(searchValue);
    var searchQuery = '/en-US/app/search/search?q=search%20index%3D*%20'+encodedSearchValue
    +"&earliest="+searchEarliest+"&latest="+searchLatest;
    window.open(splunkUrl+searchQuery);
}
```

**Related information**

[Splunk website](#)

## Defining custom icons

You can add custom icons for resources that are displayed in the Agile Service Manager UI using the **Custom Icons** page accessed through DASH.

### Before you begin

To access the **Custom Icons** page, you must have the admin role **inasm\_admin** assigned to you. See the [“Configuring DASH user roles” on page 32](#) topic for more information.

### About this task

**Tip:** Instead of using the following procedure, you can first select a specific resource type from the **Resource Types** page, and then define a custom icon for that specific resource type only. This is described in step four [“Editing resource type styles” on page 281](#). However, if you intend to create a

number of icons without assigning them, or simply want to edit or delete icons, use the following procedure.

### Icon properties

Each custom icon must have a name that uniquely identifies the icon when assigning it to a type.

**Remember:** You cannot change the name of an existing icon. If you want an icon to have a different name, create a new icon, then delete the old one.

Icons are defined inside an SVG editor, which performs an XML validation check.

Each icon definition must be valid svg xml with a given viewBox, which is important to ensure scaling of the image.

The svg definition must include inline styling of the image, such as stroke color and fill color. Use of style classes is not advised, as it can cause visualization issues on some browsers. If style classes must be used, naming must be unique for each svg image to prevent class definitions from being overwritten.

Optionally, each icon can be assigned to a category, which allows you to group icons of the same type or function together when displaying them in a list.

### Procedure

1. As a user with the `inasm_admin` role, log into your DASH web application.
2. Select **Administration** from the DASH menu.
3. Select **Custom Icons** under the Agile Service Management heading.
4. Click **New** to create a new icon, and type a unique name. Alternatively, click the 'edit' symbol to edit an existing icon.

**Remember:** You cannot change the name of an existing icon. If you want an icon to have a different name, create a new icon, then delete the old one.

5. Enter the SVG XML to define the icon.

Use the editor to enter valid XML. The XML editor includes a **Preview** area where the results of your edits are displayed.

**Example:** Use the following definition for the 'disk' icon as guidance:

```
<svg xmlns="http://www.w3.org/2000/svg" viewBox="0 0 64 64">
  <ellipse style="fill-opacity:0;stroke:currentColor;stroke-width:12.12270069;"
    id="path4139" cx="33.627117" cy="32.949142" rx="16.803904" ry="17.210684"/>
  <circle cx="33.827423" cy="33.055576" r="3.3037829"/>
</svg>
```

6. Enter a category name, if required.
7. Click **Save**.

### What to do next

Next, you assign these icons to particular resource types using the Resource Types page accessed through DASH. There, you can also apply further style edits to the resource types.

### Related tasks

[“Editing resource type styles” on page 281](#)

You assign existing or new custom icons to particular resource types using the **Resource Types** page accessed through DASH. Here you can also apply further resource type style edits, such as adding custom labels to a resource type, and changing its shape, size, border and background.

## Editing resource type styles

You assign existing or new custom icons to particular resource types using the **Resource Types** page accessed through DASH. Here you can also apply further resource type style edits, such as adding custom labels to a resource type, and changing its shape, size, border and background.

### Before you begin

To access the **Resource Types** page, you must have the admin role **inasm\_admin** assigned to you. See the [“Configuring DASH user roles” on page 32](#) topic for more information.

### About this task

**Tip:** If you intend to create a number of icons without assigning them to specific resource types, or simply want to edit or delete icons, use the procedure described in the [“Defining custom icons” on page 279](#) topic. However, it may be more convenient to define custom icons as described from step four of the following procedure, as the custom icon is then immediately assigned to the previously selected resource type.

### Procedure

1. As a user with the **inasm\_admin** role, log into your DASH web application.
2. Select **Administration** from the DASH menu.
3. Select **Resource Types** under the Agile Service Management heading.

The **Resource Types** page is displayed, which lists all existing resource types in table format in sortable columns. The table also displays the icons for the resource types and their names, categories (if defined), whether they are system icons or custom icons, and whether they have custom labels, styles or shapes. From here, you edit existing resource types, and then associate existing icons with them, and apply resource styling.

4. Click **Edit** to edit an existing resource type.

The **Edit Resource Type** page is displayed. You can toggle between the **Identification** and the **Styling** tabs.

**Restriction:** You cannot change the name of an existing resource type.

5. On the **Identification** tab, define the selected resource type's icon, label and shape.
  - a) Choose an icon to associate with the resource type using one of the following methods:

#### From the Icon drop-down list

If you open the **Icon** drop-down list, all icons are listed by name in alphabetical order.

#### From the View all icons button

If you click the **View all icons** button, all icons are displayed in alphabetical order.

Click an icon to associate it with the resource type.

#### From the Quick assign button

If an icon exists with the same name as the resource type, click the **Quick assign** button to select it without having to sort through all existing icons.

This function is useful if you have added a custom icon, and are now assigning it to a resource type with the same name.

#### From the Define new custom icon button

From here you can define a custom icon, which is automatically associated with the resource type when done.

Click the **Define new custom icon** button to display the **Custom Icons** page.

Click **New** to create a new icon. Alternatively, click the 'edit' symbol to edit an existing icon. After you have selected or created an icon, the **Configure Resource Type** page is displayed. Use the following information to define the icon:

### Icon properties

Each custom icon must have a name that uniquely identifies the icon when assigning it to a type.

**Remember:** You cannot change the name of an existing icon. If you want an icon to have a different name, create a new icon, then delete the old one.

Icons are defined inside an SVG editor, which performs an XML validation check.

Each icon definition must be valid svg xml with a given viewBox, which is important to ensure scaling of the image.

The svg definition must include inline styling of the image, such as stroke color and fill color. Use of style classes is not advised, as it can cause visualization issues on some browsers. If style classes must be used, naming must be unique for each svg image to prevent class definitions from being overwritten.

Optionally, each icon can be assigned to a category, which allows you to group icons of the same type or function together when displaying them in a list.

**Remember:** You can also create custom icons from the **Custom Icons** page accessed through DASH, which is described in the [“Defining custom icons” on page 279](#) topic.

- b) Define the label for the resource type by editing the following fields:

#### Label

By default the name of a property (`asmProperties.name`) is used as the label that is displayed in the topology.

You can replace this by typing in a custom label for the resource type.

#### Label Maximum Length

You can override the default label length of 15 characters to avoid truncating the displayed label.



**Warning:** Avoid labels that are too long, as long labels may overlap in the topology. The maximum suggested label length is 20 characters.

- c) Choose a shape for the resource type from the **Resource Shape** drop-down list.

The default shape for a resource in the Topology Viewer is a circle, or a square for a host server. You can change the shape of the resource type to one of the following shapes:

- Circle
- Square
- Hexagon
- Vertical hexagon

6. On the **Styling** tab, define the selected resource type's border color, border pattern, background color, and resource display size.

a) Change the border color by entering a hex definition. The default border color is `#171717`.

b) Change the border pattern. The default pattern is ' '.

c) Change the background color for the resource. The default is `#F3F3F3` (a very light grey).

d) Select a size for the resource type. The default is medium.

You can further refine the size function by specifying how certain resource properties effect the size of the resource type displayed. For example, resource types can appear larger depending on the number of connections they have with other resources.

7. Click **Save** to return to the **Resource Types** page.

## Results

The changes you have made to the resource type and its associated icon are now displayed in the Resource Types table.

## Creating custom relationship type styles

You can customize the styles and labels for specified relationship types using the **Relationship Types** page accessed through DASH. You can also delete existing, or create new relationship type styles.

### Before you begin

To access the **Relationship Types**, you must have the admin role **inasm\_admin** assigned to you. See the [“Configuring DASH user roles” on page 32](#) topic for more information.

### About this task

**Note:** Any resource or relationship properties used in custom relationship styles must be added as required properties to the Agile Service Manager 'Advanced topology settings' (**Administration > Advanced topology settings**). However, if they do not exist in the Advanced topology settings, you will be prompted to add them when you save your new relationship styles at the end of the following procedure.

You can set a relationship's line color, thickness, and style, as well as the label.

You customize styles in two steps. First, you identify the relationship type to be customized, or create a new one. Then you customize the style elements for the relationship type.

#### Default relationship style

When you customize relationship styles, you change the following default settings:

Line Color Function:

```
return '#171717';
```

Line Width Function:

```
return '0.5px';
```

Line Pattern Function:

```
return null;
```

#### Style restrictions

The line color (strokeWidth) property can have a maximum value of 9px, or it will cause styling problems.

The label (linkLabel) property must return a valid string, which will be displayed alongside the relationship.

See the following external sites for more detailed SVG style definition information:

#### Line color (stroke)

<https://developer.mozilla.org/en-US/docs/Web/SVG/Attribute/stroke>

#### Line width (stroke-width)

<https://developer.mozilla.org/en-US/docs/Web/SVG/Attribute/stroke-width>

#### Line pattern (stroke-dasharray)

<https://developer.mozilla.org/en/docs/Web/SVG/Attribute/stroke-dasharray>

#### Accessing properties for styling functions

When defining styles you can access dynamic properties of the relationship or connected resources.

To access the properties of resources, use the asmProperties JavaScript object.

To access the properties of relationships, use the asmSourceProperties or asmTargetProperties JavaScript objects.

**Remember:** The arrows indicating a relationship point from the source to the target.

To access the highest open severity status from the source or target resource, use the `asmSourceProperties._hasStatus` or `asmTargetProperties._hasStatus` JavaScript objects. The following example uses the **\_hasStatus** parameter to modify the relationship label:

```
if (asmSourceProperties._hasStatus || asmTargetProperties._hasStatus) {
  // object of all ASM's status severities
  let severityRank = {
    clear: 0,
    indeterminate: 1,
    information: 2,
    warning: 3,
    minor: 4,
    major: 5,
    critical: 6
  };
  let sourceSeverityRank = severityRank[asmSourceProperties._hasStatus] || 0;
  let targetSeverityRank = severityRank[asmTargetProperties._hasStatus] || 0;
  let labelString = asmProperties._edgeType;
  // Show highest Status on relationship label
  if (sourceSeverityRank > targetSeverityRank) {
    labelString += ': Status = ' + asmSourceProperties._hasStatus;
  } else {
    labelString += ': Status = ' + asmTargetProperties._hasStatus;
  }
  return labelString;
} else {
  // No status for source or target resources use plain label
  return asmProperties._edgeType;
}
```

## Procedure

1. As a user with the `inasm_admin` role, log into your DASH web application.
2. Select **Administration** from the DASH menu.
3. Select **Relationship Types** under the Agile Service Management heading.

The **Relationship Types** page is displayed, which lists all your existing customized relationship types in table format, also displaying a Last Updated time stamp, and whether a relationship type has a custom label, custom color or custom width defined. From here, you can delete or edit configurations. You can also create a new relationship type configuration.

4. Do one of the following:
  - To delete a relationship type, click the **Delete** icon.
  - To edit an existing relationship type, click the **Edit** icon. The **Configure Relationship Type** page is displayed.
  - To create a new relationship type, click **New**. The **Configure Relationship Type** page is displayed.

### Configure relationship types

5. Select the **Identification** tab on the **Configure Relationship Type** page.

#### Relationship Type

Choose the relationship type that you want to configure from the dropdown list.

**Note:** Only relationship types that exist in your topology database are listed.

#### Label Function

Example of a JavaScript function to define a label for the relationship using the resource properties:

```
return asmProperties.labelText;
```

6. Select the **Style** tab on the **Configure Relationship Type** page.

#### Line Color Function

Example of a JavaScript function to define the ink color:



```
return 'blue';
```

### Line Width Function

Example of a JavaScript function to define the line width for the relationship using the source resource properties:

```
return asmSourceProperties.myProp > 10 ? '9px' : '1.5px';
```

### Line Pattern Function

Example of a JavaScript function to define the line pattern:

```
return '3 4 5';
```

7. Click **Save**.

**Note:** If any of the resource or link properties used have not yet been defined in 'Advanced topology settings', you will be prompted to save them now.

## Results

The relationship style and label for the specified relationship type has been customized.

## Defining advanced topology settings

As a system administrator, you can define advanced topology settings, such as the URLs for trusted sites, the required properties for tooltips or relationship type styles, or the maximum hop numbers a user can choose in the Agile Service Manager topology viewer. You do this from the **Advanced topology settings** page accessed through DASH.

### Before you begin

To access the **Advanced topology settings** page, you must have the admin role **inasm\_admin** assigned to you. See the [“Configuring DASH user roles” on page 32](#) topic for more information.

**Important:** Ensure that you understand your system's data capacity. If you allow users to set a high hop count, this will place greater demands on your network, with more information being sent from the topology service to the topology viewer, and a greater workload on the topology service itself, and in the browser when rendering the topology.

### About this task

#### Hop count

Choose a number between zero and thirty.

The default maximum hop count is four.

#### Historical change threshold

The maximum size of a topology as defined by its constituent resources for which the historical changes indicators will be enabled when in Delta mode.

#### Required properties

To improve system performance only specific resource and relationship properties are loaded into the Topology Viewer. You can specify additional properties to be fetched [here](#).

These properties may then be used in tooltips, custom status tools or to customize UI elements (via custom JavaScript code).

You can exclude these required properties from being displayed in a tooltip in the Topology Viewer.

#### Excluded types

In order to optimize performance when rendering complex topologies, you can exclude specific resource or relationship types. The seed resource of a topology **can not** be excluded.

Excluded resource or relationship types are not displayed in the Topology Viewer, though operators can see a list of excluded types on the Filtering toolbar.

## Topology search

You can include deleted resources in the search results.

You can toggle between **Disabled** and **Enabled** (the default).

## Trusted sites

If you need to link to an HTTP address instead of HTTPS while writing a custom tool definition, and need to request data from a site that uses HTTP instead of HTTPS, then you can use the Agile Service Manager UI server as a proxy. The URL to the HTTP proxy takes the actual target HTTP site URL as a parameter. The proxy then sends the request itself, and returns the response back to the browser.

For security reasons, the HTTP proxy can only access the URLs that have been defined by an administrator user as 'trusted sites'.

The trustedSites property values are a comma-separated list of the trusted sites to which the proxy server can link, and from which it can retrieve information.

Trusted sites operate under a 'starts with' condition.

## Procedure

1. As a user with the `inasm_admin` role, log into your DASH web application.
2. Select **Administration** from the DASH menu.
3. Select **Advanced topology settings** under the Agile Service Management heading.

The **Advanced topology settings** page is displayed consisting of the expandable Topology Rendering, Topology Search, and Topology Tools sections.

### Under Topology Rendering

4. To set the maximum number of hops, select a value between zero and thirty from the **Maximum number of hops allowed** drop-down list.



**CAUTION:** Ensure that you set a maximum hop count that, if selected by an operator, will not overload your system's data capacity.

5. To set the maximum topology size, enter a value up to, but not exceeding fifty, in the **Historical changes threshold** field.

This setting defines the maximum size of a topology (based on the number of resources) for which the historical changes indicators will be enabled when in delta mode.

6. To retrieve all relationships between resources in a topology, turn on the **Include adjacent relationships on topological hops** toggle.
7. To add a required property, enter the property into the **Properties required for tooltips or relationship style** text box, then click **Add +**.

For example: `location`

#### To exclude a required property from being displayed in a tooltip

Clear the **Show in tooltip?** check box next to the required property.

#### To delete a property

Click **Delete** (the bin icon) next to the property.

8. To exclude a resource type from being displayed in the Topology Viewer, enter the type into the **Resource to exclude from topologies** field, then click **Add +**.

**Note:** The seed resource of a topology **can not** be excluded.

The excluded resource types are not displayed in the Topology Viewer, though operators will see a list of excluded types on the Filter toolbar.

9. To exclude a relationship type from being displayed in the Topology Viewer, enter the type into the **Relationship to exclude from topologies** field, then click **Add +**.

The excluded relationship types are not displayed in the Topology Viewer, though operators will see a list of excluded types on the Filter toolbar.

### Under Topology Search

10. To prevent deleted resources from being returned in the topology search, set the **Include deleted resources in search results** toggle to 'Disabled'.

#### Under Topology Tools

11. To add a trusted website, click **Add +**, and enter the required URLs into the **Trusted websites that can be accessed via HTTP** text box.

For example, you could add the following trusted sites:

```
www.ibm.com
data-server.intranet.com:8080/network/statistics
```

In this example, the HTTP proxy will allow requests for all target URLs that start with 'www.ibm.com' or with 'data-server.intranet.com:8080/network/statistics', but no others, as illustrated in the following examples.

#### Allowed http proxy targets

```
www.ibm.com/cloud
data-server.intranet.com:8080/network/statistics/1893713/info
```

#### Not allowed http proxy targets

```
data-server.intranet.com:8080
www.ibm.co.uk
www.news-page.com
```

## Configuring retention period for resource history

Agile Service Manager retains historical topology data for a default period of 30 days. You can increase this to a maximum of 90 days by increasing the 'time to live' (TTL) value.

### Before you begin



**CAUTION:** The performance and scalability of Agile Service Manager is affected by the number of resources managed, the amount of history present for each resource, and the ingestion rate. If you increase the retention period for historical resource data from the default of 30 days (up to a maximum of 90 days), your system performance (when rendering views) may deteriorate, if as a result of this increase resources have in excess of 25,000 historical entries.

### About this task

The retention period for historical resource data is configured via the HISTORY\_TTL configuration property.

When a resource or edge in Agile Service Manager is deleted, it will no longer appear in the UI. Historic representations of the resource or edge, however, are retained and can be accessed in the UI history timeline until the history TTL limit has been reached, after which the data is deleted.

**Note:** Unless explicitly deleted, live resources remain current and the TTL limit does not apply to them.

For an illustration, see the example section.

### Procedure

#### Edit the configuration file

1. Open the /opt/ibm/netcool/asm/.env file using an appropriate editor.
2. Edit the HISTORY\_TTL setting. Enter the values in hours.  
Change the TTL value from the default of 720 (30 days) to any value up to 2160 (90 days).
3. Restart Agile Service Manager:

```
$ASM_HOME/bin/asm_start.sh
```

## Example

In the following example scenario, the default TTL value of 30 days (720 hours) applies.

Table 64. TTL example for the 'sprocket' resource			
Date	Action	Topology view	Historical data
01-January-2019	<b>sprocket</b> resource created	<b>sprocket</b> resource visible in current and history	History is current resource
01-March-2019	<b>sprocket</b> resource modified	<b>sprocket</b> resource visible in current and history	History has both current and previous resource
31-March-2019	TTL expires for historic resource created on 01-March	<b>sprocket</b> resource visible in current and history	History is current (modified) resource only
01-May-2019	<b>sprocket</b> resource deleted	<b>sprocket</b> resource visible in history only	History still contains deleted resource
31-May-2019	TTL expires for historic resource created on 01-May	<b>sprocket</b> resource not visible in current or history	No history remains

## Increasing default timeouts

To give your deployment of Agile Service Manager enough time to process very large data sets, you can increase the default system timeout values.

### About this task

To change the default timeouts, you edit the \$ASM\_HOME/.env file.

### Procedure

1. Open the \$ASM\_HOME/.env file using an appropriate editor.
2. Edit the following settings for the listed Agile Service Manager services.  
Enter the values in seconds, for example 120.

```
TOPOLOGY_SERVICE_TIMEOUT:  
LAYOUT_SERVICE_TIMEOUT:  
SEARCH_SERVICE_TIMEOUT:  
MERGE_SERVICE_TIMEOUT:  
OBSERVER_SERVICE_TIMEOUT:
```

3. Restart Agile Service Manager.

- To stop Agile Service Manager, run:

```
<ASM_HOME>/bin/asm_stop.sh
```

- Once stopped, re-start Agile Service Manager:

```
<ASM_HOME>/bin/asm_start.sh
```

## Configuring the Helm chart to use alternate storage (OCP)

Agile Service Manager by default supports local volumes for storage on Kubernetes. To configure an alternate storage backend, you must set the storage class. This requires a manual command line installation, and **not** an installation via Helm.

### Before you begin



**CAUTION:** You cannot configure storage class for installation from OCP. This is an advanced command line installation that you should only perform if you have the required expertise.

### About this task

**Assumption:** A suitable **storageclass** has been defined and configured to provision storage volumes

### Procedure

1. Find the storage class:

```
# kubectl get storageclass
NAME                                PROVISIONER                AGE
vsphere                            kubernetes.io/vsphere-volume 7d18h
```

2. Set the storage class in the persistent volume claim:

**Important:** The following configuration sample contains the relevant settings for storage class and volume capacity only. You must merge these settings with the other installation parameters.

```
global:
  persistence:
    enabled: true
    useDynamicProvisioning: true
    storageClassName: "vsphere"      # to match value from 'kubectl get storageclass'
    storageClassOption:
      cassandradata: "default"
      zookeeperdata: "default"
      kafkadata: "default"
    storageSize:
      cassandradata: 50Gi
      kafkadata: 15Gi
      zookeeperdata: 5Gi
      elasticdata: 75Gi
```

### Example

After you have changed the settings, the Agile Service Manager PersistentVolumeClaims will now include a storage class. On a system with an appropriate provisioner in place, the PersistentVolumes should be generated automatically, as in the following example:

NAME	READY	STATUS	RESTARTS	AGE
pod/asm-cassandra-0	1/1	Running	0	10m
pod/asm-elasticsearch-0	1/1	Running	0	10m
pod/asm-kafka-0	2/2	Running	0	10m
pod/asm-kubernetes-observer-698cfd746b-q7z9l	1/1	Running	0	10m
pod/asm-layout-84474476bc-96cz6	1/1	Running	0	10m
pod/asm-merge-c4cd8f8b7-nq589	1/1	Running	0	10m
pod/asm-search-778b9f9574-hhsxf	1/1	Running	0	10m
pod/asm-system-health-cronjob-1560349200-qgkrr	0/1	Completed	0	9m10s
pod/asm-system-health-cronjob-1560349500-qqdnn	0/1	Completed	0	19s
pod/asm-topology-6b6c4b4b54-pp92j	1/1	Running	0	10m
pod/asm-ui-api-7849c55b4c-5qq2m	1/1	Running	0	10m
pod/asm-zookeeper-0	1/1	Running	0	10m

NAME	CAPACITY	ACCESS MODES	STORAGECLASS	AGE	STATUS	VOLUME
persistentvolumeclaim/data-asm-cassandra-0	50Gi	RWO	vsphere	10m	Bound	pvc-06784b19-8d1d-11e9-a948-005056b47772
persistentvolumeclaim/data-asm-elasticsearch-0	75Gi	RWO	vsphere	10m	Bound	pvc-06f05e48-8d1d-11e9-a948-005056b47772
persistentvolumeclaim/data-asm-kafka-0	15Gi	RWO	vsphere	10m	Bound	pvc-07e53665-8d1d-11e9-a948-005056b47772
persistentvolumeclaim/data-asm-zookeeper-0	5Gi	RWO	vsphere	10m	Bound	pvc-085f9547-8d1d-11e9-a948-005056b47772

## Porting data for testing, backup and recovery

You can create backups of your Agile Service Manager UI configuration data in order to run a test configuration, or simply to safeguard your custom settings. You can also back up and restore your topology data.

### Backing up and restoring database data (on-prem)

You can back up and restore existing topology data using scripts included in the Agile Service Manager **on-prem** installation. This can be helpful when updating your system, or for maintenance reasons.

#### About this task

Agile Service Manager stores its information in the Cassandra database, using in turn Elasticsearch for indexing and fast-searching a subset of the data (listed [here](#)). The backup and restore procedures are therefore based on backing up the Cassandra database (where all the information is stored) and, upon its restoration, triggering a topology re-index to re-populate Elasticsearch to keep the consistency of the indexed information.

The `backup_cassandra.sh` and `restore_cassandra.sh` scripts are stored in the `ASM_HOME/bin` directory.

**Tip:** The backup utility runs from inside the Cassandra Docker Container, and thus filesystem paths are relative from inside that container. Typically these file paths can be accessed from the host system via

```
${ASM_HOME}/logs/cassandra
```

and

```
${ASM_HOME}/data/cassandra/backup
```

To restore topology data on a new system, you start with a clean instance of Cassandra, then start the Topology Service to create keys and indexes in Cassandra, and when the topology is complete, you run the restore script.

**Note:** To restore Cassandra to a previous state, for example in scenarios where the Docker container has not been re-created, you run the Restore script without cleaning Cassandra first.

## Procedure

### Backing up data

1. From the `ASM_HOME` directory, run the backup script:

#### Example

```
# ./bin/backup_cassandra.sh

# Description : The backup script will complete the backup in multiple phases -
# 1. Take statistics of the keyspace(s) before backup
# 2. Clear existed snapshots
# 3. Take backup of keyspace(s) SCHEMA in temporary BACKUP_TEMP_DIR
# 4. Take snapshot of keyspace(s)
# 5. Copy snapshot to temporary BACKUP_TEMP_DIR
# 6. Compact the temporary BACKUP_TEMP_DIR in one tar file and send it to BACKUP_DIR

USAGE: backup_cassandra.sh
[ -k keyspace to backup ] # default is ALL keyspaces
[ -b temporary backup dir ] # default is /opt/ibm/cassandra/data/backup/../../backup_temp
[ -d datadir ] # default is /opt/ibm/cassandra/data/data
[ -s storagedir ] # default is /opt/ibm/cassandra/data/backup
[ -u Cassandra username ]
[ -p Cassandra password ]
[ -log logdir ] # default is /opt/ibm/cassandra/logs
[ -speed tardiskspeed ] # default is 17M
[ -f ] # for non interactive mode
```

```

***** START CONFIGURATION *****
KEYSPACE_TO_BACKUP=ALL
BACKUP_TEMP_DIR=/opt/ibm/cassandra/data/backup/./backup_temp
BACKUP_DIR=/opt/ibm/cassandra/data/backup
CASSANDRA_DATA=/opt/ibm/cassandra/data/data
LOG_PATH=/opt/ibm/cassandra/logs
TAR_SPEED_LIMIT=17M
FORCE=N
USER=cassandra
PASS=XXXX
***** END CONFIGURATION *****

Tue Mar  5 13:15:26 UTC 2019 *****
Tue Mar  5 13:15:26 UTC 2019 Do you want to continue (y/n) ?

```

## Restore topology data (on an existing system)

2. From the ASM\_HOME directory, run the restore script:

### Example

```

# ./bin/restore_cassandra.sh

# Description : The restore script will complete the restore in multiple phases -
1. Take statistics of the cassandra node before restore
2. Check if the keyspace exists and if it does not exist, create it using the schema cql file
   saved in the backup file
3. Truncate all tables in keyspace
4. Clear all files in commitlog directory
5. Copy contents of desired snapshot to active keyspace.
6. Refresh all tables in that keyspace
7. Take statistics of the cassandra node after restore and compare with statistics taken
   before
backup, making sure number of keys per table is the same

```

```

USAGE: restore_cassandra.sh
      -k keyspaceName # compulsory parameter
      [-h backup hostname] # if backup was done on a different hostname than 96c6953586a3
      [-b temporary backup dir ] # default is /opt/ibm/cassandra/data/backup/./backup_temp
      [-d dataDir ] # default is /opt/ibm/cassandra/data/data
      [-t snapshotTimestamp ] # timestamp of type date YYYY-MM-DD-HHMM-SS - default is latest
      [-s storageDir ] # default is /opt/ibm/cassandra/data/backup
      [-u Cassandra username ]
      [-p Cassandra password ]
      [-log logDir ] # default is /opt/ibm/cassandra/logs
      [-f ] # for non interactive mode

```

```

***** START CONFIGURATION *****
BACKUP_TEMP_DIR=/opt/ibm/cassandra/data/backup/./backup_temp
BACKUP_DIR=/opt/ibm/cassandra/data/backup
DATA_DIR=/opt/ibm/cassandra/data/data
LOG_PATH=/opt/ibm/cassandra/logs
LOCAL_HOSTNAME=96c6953586a3
BACKUP_HOSTNAME=96c6953586a3
SNAPSHOT_DATE_TO_RESTORE=latest
KEYSPACE_TO_RESTORE=janusgraph
FORCE=N
USER=cassandra
PASS=XXXX
***** END CONFIGURATION *****

Tue Mar  5 13:17:22 UTC 2019 *****
Tue Mar  5 13:17:22 UTC 2019 Do you want to continue (y/n) ?

```

3. Rebroadcast data to ElasticSearch (that is, re-index Elasticsearch).

If data in Elasticsearch is out of sync with data in the Cassandra database, resynchronize it by calling the rebroadcast API of the topology service. This triggers the rebroadcast of all known resources on Kafka, and the Search service will then index those resources in Elasticsearch.

### Workaround

Call the rebroadcast API of the Topology service, specifying a tenantId:

```
https://master_fqdn/1.0/topology/swagger#!/Crawlers/rebroadcastTopology
```

## Restore topology data (on a new system)

4. Clean the Cassandra database.

- a) Ensure the topology service and database are not running.
- b) Remove all database content:

```
rm -rf /opt/ibm/netcool/asm/data/cassandra/data/*
```

c) Start Cassandra.

```
docker-compose up -d cassandra
```

5. Start the topology service.

This will create keys and indexes in Cassandra.

6. Run the restore script from the Cassandra container.

**Note:** For a new host, you must use the -h switch to identify the backup hostname, as it will be different to the hostname into which you are restoring.

```
./restore_cassandra.sh -h <backup_hostname>
```

**For example:**

```
./restore_cassandra.sh -h 96c6953586a3
```

## Backing up UI configuration data (on-prem)

Agile Service Manager includes a backup facility, which lets you backup UI configuration settings such as user preferences, topology tools, custom icons, relationship types, and advanced topology settings. This topic describes how to **export** these settings.

### Before you begin

**Remember:** Backing up and restoring your UI configuration is a two step process. The first step, exporting data, is described in this topic. The second step, importing data to restore previous configurations, is described in the [“Restoring UI configuration data \(on-prem\)”](#) on page 294 topic.

The tool detects the host, port, and tenant id for the Topology Service from the following environment variables:

- TOPOLOGY\_SERVICE\_HOST
- TOPOLOGY\_SERVICE\_PORT
- TENANT\_ID

**Important:** If you have a standard installation of Agile Service Manager core, and none of the default settings have been changed, the tool will work without you having to reset any of these environment variables. However, if you do have a non-standard installation, you need to reset these before running the tool.

### About this task

#### When to export configuration data

Agile Service Manager UI configuration settings are stored in the topology service database. If that database is deleted, the configuration settings are also deleted. You may therefore want to create a backup of configuration data if you intend to conduct testing that may involve changes to your database. After installing or rebuilding a new database, you can then restore the configuration data.

You can export configuration data as part of your data protection strategy, to provide a backup in case of data corruption or accidental data deletion.

You can export configuration data from a staging system in order to then import it into a live system.



## Syntax

This command **must** be run from the ASM\_HOME directory.

```
docker-compose -f tools.yml run --rm backup_ui_config [-config <config_type>]
[-out <output_filename>] [-force] [-verbose]
```

**Tip:** For help syntax:

```
docker-compose -f tools.yml run --rm backup_ui_config -help
```

## Parameters

All parameters are optional.

**Note:** You can run the backup\_ui\_config command without setting any parameters. If you do, all Agile Service Manager UI configuration settings will be exported to the following default file:  
\$ASM\_HOME/data/tools/asm\_ui\_config.txt

### config

The -config flag allows the type of configuration you want to export to be specified.

By default all UI configuration settings are backed up.

#### Settings for -config <config\_type>

Backs up the following Agile Service Manager UI configuration settings

##### all

All UI configurations (default)

##### tools

Topology tools definitions

##### icons

Custom icon definitions

##### types

Entity type definitions

##### links

Relationship type definitions

##### preferences

User preferences

##### settings

Advanced topology settings

### out

The -out flag is the name of the backup file name to create.

The name must be a file name only, with no directory paths.

The default output file name is asm\_ui\_config.txt, and the output location is fixed as \$ASM\_HOME/data/tools.

**Note:** If the file already exists, the tool will indicate this and quit. For the existing file to be overwritten with new output, use the -force parameter.

### force

If you set the -force parameter, the tool overwrites an existing output file with new content.

### verbose

The -verbose flag runs the tool in verbose mode, whereby extra log messages are printed to the shell during execution.

This parameter is useful if a problem occurs while running the tool, and you want to re-run it with extra information made available.

## Procedure

1. Using the syntax information provided, determine the export (backup) <options> you need to set, if any.

**Remember:** Run the `backup_ui_config` command without any options set to backup all UI configuration settings to `$ASM_HOME/data/tools/asm_ui_config.txt`.

2. Run the `backup_ui_config` command from your `ASM_HOME` directory, as in the following example. This runs the tool inside a docker container, and the `--rm` flag then causes the exited container to be deleted once the tool has completed.

```
docker-compose -f tools.yml run --rm backup_ui_config <options>
```

## Results

The Agile Service Manager UI configuration data is exported to the specified file in the `$ASM_HOME/data/tools` directory. If the `-force` flag has been set, an existing backup file is overwritten.

From the backup file, you can restore the settings using the [“Restoring UI configuration data \(on-prem\)”](#) on page 294 topic.

## Restoring UI configuration data (on-prem)

Agile Service Manager includes a backup facility, which lets you backup UI configuration settings such as user preferences, topology tools, custom icons, relationship types, and advanced topology settings. This topic describes how to **import** previously exported (backed up) settings in order to restore your previous configurations.

### Before you begin

You can only import and restore configuration settings that have been previously exported, as described in the [“Backing up UI configuration data \(on-prem\)”](#) on page 292 topic

The tool detects the host, port, and tenant id for the Topology Service from the following environment variables:

- `TOPOLOGY_SERVICE_HOST`
- `TOPOLOGY_SERVICE_PORT`
- `TENANT_ID`

**Important:** If you have a standard installation of Agile Service Manager core, and none of the default settings have been changed, the tool will work without you having to reset any of these environment variables. However, if you do have a non-standard installation, you need to reset these before running the tool.

### About this task

#### Syntax

This command **must** be run from the `ASM_HOME` directory.

```
docker-compose -f tools.yml run --rm import_ui_config -file <input_file>
[-overwrite] [-verbose]
```

**Tip:** For help syntax:

```
docker-compose -f tools.yml run --rm import_ui_config -help
```

## Parameters

### file

The `-file` parameter is the name of the backup file from which to import definitions. It **must** be a file name only with no directory paths included, and it **must** exist in the tools data directory (`$ASM_HOME/data/tools`).

### overwrite

By default, as the import tool reads the backup file it looks up each item in the topology service to see if it already exists. Any configuration definitions which already exist are **not** updated.

However, if you set the `-overwrite` flag, the existing definitions are overwritten with the values from the backup file.

### verbose

The `-verbose` flag runs the tool in verbose mode, whereby extra log messages are printed to the shell during execution.

Useful if a problem occurs running the tool and you want to re-run it with extra information made available.

## Procedure

1. Using the syntax information provided, enter the file name of the previously exported backup file.
2. Determine if any other `<options>` you need to set, such as the `-overwrite` flag.
3. Run the `import_ui_config` command from your `ASM_HOME` directory, as in the following example.  
This runs the tool inside a docker container, and the `--rm` flag then causes the exited container to be deleted once the tool has completed.

```
docker-compose -f tools.yml run --rm import_ui_config <options>
```

## Results

The Agile Service Manager UI configuration data is imported from the specified file in the `ASM_HOME/data/tools` directory. If the `-overwrite` flag has been set, existing configuration data will be overwritten.

## Backing up database data (OCP)

You can back up (and later restore) existing topology data for the Agile Service Manager **OCP** installation. This can be helpful when updating your system, as part of your company's data management best-practice, or for maintenance reasons.

### About this task

Agile Service Manager stores its information in the Cassandra database, using in turn Elasticsearch for indexing and fast-searching a subset of the data (listed [here](#)). The backup and restore procedures are therefore based on backing up the Cassandra database (where all the information is stored) and, upon its restoration, triggering a topology re-index to re-populate Elasticsearch to keep the consistency of the indexed information.

Agile Service Manager can be installed on RedHat OpenShift Container Platform (OCP) following different deployment models:

- As part of a Netcool Operations Insight solution
- Or as a stand-alone product (that is, without NOI) as part of a Watson AIOps deployment

### Naming conventions for Agile Service Manager on OCP (as part of NOI or AIOps Event Manager)

When Agile Service Manager is deployed as part of NOI, the naming convention the Agile Service Manager pods follow is: `{releaseName}-topology-{resource}-{suffix}`, where 'suffix' is a number for the Stateful Sets (such as a Cassandra database) that represents the pod replica number, or a hash-like uid for the Deployment pods.

For example, Cassandra pods might be {releaseName}-topology-cassandra-0 for the first replica of the set, whereas the topology service pod could be {releaseName}-topology-topology-57477f4978-qp9gz.

In cases where Cassandra is shared, the shared instance pods Agile Service Manager uses would be {releaseName}-cassandra-{suffix} instead.

### Naming conventions for Agile Service Manager on OCP standalone (pre-reqs for AIOps AI Manager)

When Agile Service Manager is installed as a standalone product, the naming convention for its pods is {releaseName}-{resource}-{suffix}.

When deployed as part of Netcool Operations Insight, and depending on whether you are installing for the first time or upgrading, you can share the Cassandra database instances between all the NOI products, or deploy separate instances for Agile Service Manager and NOI. When sharing the Cassandra instances, the database will contain data from the rest of the components in NOI that make use of Cassandra, **as well as** the Agile Service Manager data.

**Assumption:** The backup and restore procedures in these topics assume a standard Agile Service Manager NOI deployment (and **not** an standalone deployment), a shared use of Cassandra, and that the release name used is 'noi'. Adjust the samples provided to your circumstances.

### Backup

The backup procedure documented here performs a backup of **all** the keyspaces in the Cassandra database, including those not specific to Agile Service Manager.

### Restore

The restore procedures focus on restoring **only** the keyspace that is relevant to Agile Service Manager (that is, 'janusgraph').

## Procedure

### Preparing your system for backup

1. Authenticate into the Kubernetes namespace where Agile Service Manager is deployed as part of your solution.
2. Deploy the following kPodLoop bash shell function.

**kPodLoop** is a bash shell function that allows a command to be run against all the matching Kubernetes containers. You can copy it into the shell.

```
kPodLoop() {
  __podPattern=$1
  __podCommand=$2
  __podList=$( kubectl get pods --field-selector=status.phase=Running --no-headers=true --
output=custom-columns=NAME:.metadata.name | grep ${__podPattern} )
  printf "Pods found: $(echo -n ${__podList})\n"
  for pod in ${__podList}; do
    printf "\n==== EXECUTING COMMAND in pod: %-42s =====\n" ${pod}
    kubectl exec ${pod} -- bash -c "${__podCommand}"
    printf '%.0s' {1..80}
    printf "\n"
  done;
}
```

This kPodLoop bash shell function filters the pods to run the commands against only those that are in a 'Running' phase. This filter ensures that the configuration pods that are only run as part of your installation, like the secret generator pod, are skipped.

3. Make a note of the scaling of Agile Service Manager pods.

```
kubectl get pods --field-selector=status.phase=Running --no-headers=true --output=custom-
columns=CNAME:.metadata.ownerReferences[0].name | grep noi-topology | uniq --count
```

### Example output:

```
1 noi-topology-dns-observer
1 noi-topology-docker-observer
3 noi-topology-elasticsearch
1 noi-topology-file-observer
1 noi-topology-kubernetes-observer
```

```
1 noi-topology-layout
1 noi-topology-merge
1 noi-topology-noi-gateway
1 noi-topology-noi-probe
1 noi-topology-observer-service
1 noi-topology-search
1 noi-topology-status
1 noi-topology-topology
1 noi-topology-ui-api
```

4. Verify access to each Cassandra database (this command will return a list of keyspace from each Cassandra node). Adjust the Cassandra pod names based on your deployment model [naming convention](#).

```
kPodLoop noi-cassandra "cqlsh -u \${CASSANDRA_USER} -p \${CASSANDRA_PASS} -e \"DESC
KEYSPACES;\""
```

5. Scale down Agile Service Manager pods.

```
kubectl scale deployment --replicas=0 noi-topology-dns-observer
kubectl scale deployment --replicas=0 noi-topology-file-observer
kubectl scale deployment --replicas=0 noi-topology-kubernetes-observer
kubectl scale deployment --replicas=0 noi-topology-observer-service
kubectl scale deployment --replicas=0 noi-topology-noi-gateway
kubectl scale deployment --replicas=0 noi-topology-noi-probe
kubectl scale deployment --replicas=0 noi-topology-layout
kubectl scale deployment --replicas=0 noi-topology-merge
kubectl scale deployment --replicas=0 noi-topology-status
kubectl scale deployment --replicas=0 noi-topology-search
kubectl scale deployment --replicas=0 noi-topology-ui-api
kubectl scale deployment --replicas=0 noi-topology-topology
```

The Cassandra and Elasticsearch pods (noi-cassandra and noi-topology-elasticsearch) are left active. Cassandra pods need to be running in order to execute the backup of their data, whereas the Elasticsearch pods have no interaction with nor influence on the Cassandra contents, so can be kept running.

**Important:** Include in this scale down any additional observers in your deployment.

6. Verify that scaling down was successful.

```
kubectl get pods --field-selector=status.phase=Running | grep noi-topology
```

The Agile Service Manager services have now been scaled down, and the Cassandra database contents will not be modified anymore.

## Backing up data

7. Deploy the pbkc bash shell function.

The **pbkc** function attempts to backup the Cassandra database on all nodes as close to simultaneously as possible. You can copy it into the shell. Modify the following example script to conform to your installation's Cassandra pod names based on your deployment model [naming conventions](#).

```
pbkc() {
  ## Parallel Backup of Kubernetes Cassandra
  DATE=$( date +"%F-%H-%M-%S" )
  LOGFILEBASE=/tmp/clusteredCassandraBackup-${DATE}-
  declare -A PIDWAIT
  declare -A LOG

  ## get the current list of cassandra pods.
  podlist=$( kubectl get pods --field-selector=status.phase=Running --no-headers=true --
output=custom-columns=NAME:.metadata.name | grep noi-cassandra )
  for pod in ${podlist}; do
    LOG[$pod]=${LOGFILEBASE}${pod}.log
    echo -e "BACKING UP CASSANDRA IN POD $pod (logged to ${LOG[$pod]})"
    kubectl exec $pod -- bash -c "/opt/ibm/backup_scripts/backup_cassandra.sh -u \${
CASSANDRA_USER} -p \${CASSANDRA_PASS} -f" > ${LOG[$pod]} & PIDWAIT[$pod]=$!
  done

  echo -e "${#PIDWAIT[@]} Backups Active ..."

  for pod in ${podlist}; do
```

```

    wait ${PIDWAIT[$pod]}
    echo -e "Backup of ${pod} completed, please verify via log file (${LOG[$pod]})"
done
}

```

8. Run a clean-up on all keyspaces in all Cassandra instances.

Example Cassandra keyspaces cleanup:

```

kPodLoop noi-cassandra "nodetool cleanup system_schema"
kPodLoop noi-cassandra "nodetool cleanup system"
kPodLoop noi-cassandra "nodetool cleanup system_distributed"
kPodLoop noi-cassandra "nodetool cleanup system_auth"
kPodLoop noi-cassandra "nodetool cleanup janusgraph"
kPodLoop noi-cassandra "nodetool cleanup system_traces"

```

9. Run backup on all Cassandra instances (using the pbkc shell function just deployed).

```
pbkc
```

10. Check the final output in the log file for each backup.

Adjust the date in the grep command as appropriate.

```

grep "BACKUP DONE SUCCESSFULLY" /tmp/clusteredCassandraBackup-2019-06-14-14-09-50*
/tmp/clusteredCassandraBackup-2019-06-14-14-09-50-noi-cassandra-0.log:Fri Jun 14 14:11:04
UTC 2019 BACKUP DONE SUCCESSFULLY !!!
/tmp/clusteredCassandraBackup-2019-06-14-14-09-50-noi-cassandra-1.log:Fri Jun 14 14:11:16
UTC 2019 BACKUP DONE SUCCESSFULLY !!!
/tmp/clusteredCassandraBackup-2019-06-14-14-09-50-noi-cassandra-2.log:Fri Jun 14 14:11:16
UTC 2019 BACKUP DONE SUCCESSFULLY !!!

```

**Tip:** For additional information about backing up secrets and the system\_auth keyspace, see the [Secrets and system\\_auth keyspace](#) note in the 'Restoring database data (OCP)' topic.

When backup has successfully completed, you restore your Agile Service Manager services to normal operation.

## Restore services

11. Scale up the services to the original level.

The original level was obtained in a [previous step](#).

```

kubectl scale deployment --replicas=1 noi-topology-topology
kubectl scale deployment --replicas=1 noi-topology-layout
kubectl scale deployment --replicas=1 noi-topology-merge
kubectl scale deployment --replicas=1 noi-topology-status
kubectl scale deployment --replicas=1 noi-topology-search
kubectl scale deployment --replicas=1 noi-topology-observer-service
kubectl scale deployment --replicas=1 noi-topology-noi-gateway
kubectl scale deployment --replicas=1 noi-topology-noi-probe
kubectl scale deployment --replicas=1 noi-topology-ui-api
kubectl scale deployment --replicas=1 noi-topology-dns-observer
kubectl scale deployment --replicas=1 noi-topology-file-observer
kubectl scale deployment --replicas=1 noi-topology-rest-observer
kubectl scale deployment --replicas=1 noi-topology-kubernetes-observer

```

## Results

The backup procedure stores the backup generated files inside the Agile Service Manager Cassandra pods inside the /opt/ibm/cassandra/data/backup\_tar/ directory.

## What to do next

You can restore your backed up data as and when required.

### Related tasks

[“Restoring database data \(OCP\)” on page 299](#)

You can restore existing topology data for the Agile Service Manager **OCP** installation, if backed up earlier. This can be helpful when updating your system, or for maintenance reasons.

[“Backing up and restoring UI configuration data \(OCP\)” on page 305](#)

Agile Service Manager on OCP includes a backup facility, which lets you backup UI configuration settings such as user preferences, topology tools, custom icons, relationship types, and advanced topology settings.

## Restoring database data (OCP)

You can restore existing topology data for the Agile Service Manager **OCP** installation, if backed up earlier. This can be helpful when updating your system, or for maintenance reasons.

### Before you begin

Before restoring a Cassandra database, you must back it up, as described in the following topic :[“Backing up database data \(OCP\)”](#) on page 295

### About this task

**Assumption:** The backup and restore procedures in these topics assume a standard Agile Service Manager NOI deployment (and **not** an standalone deployment), a shared use of Cassandra, and that the release name used is 'noi'. Adjust the samples provided to your circumstances.

#### Backup

The backup procedure documented here performs a backup of **all** the keyspaces in the Cassandra database, including those not specific to Agile Service Manager.

#### Restore

The restore procedures focus on restoring **only** the keyspace that is relevant to Agile Service Manager (that is, 'janusgraph').

#### Secrets and system\_auth keyspace:

During Agile Service Manager deployment, a secret called {release}-topology-cassandra-auth-secret is generated, if none already exists with that name. Cassandra is protected with the user and password of that secret, which will be used by the Agile Service Manager services to connect to the database.

In the restore scenarios described, it is assumed that Agile Service Manager is deployed in a standard way, meaning that the connection to Cassandra is set with the described secret. If you were to restore the system\_auth keyspace (instead of just the janusgraph keyspace), you would have to make sure the user and password in the mentioned secret matches the credentials contained in the keyspace for the version being restored.

The following command 'exports' the secret from your Kubernetes cluster. However, the process of backing up and restoring secrets is **not** described here, as it is contingent on your company security policies, which you should follow.

```
kubect1 get secret noi-cassandra-auth -o yaml
```

**Restore scenarios:** You may encounter a variety of typical data restore scenarios while administering Agile Service Manager on OCP.

#### Same cluster data restore (rollback)

This scenario covers the restore of a database backup to the same cluster from where it was taken, as documented in the [procedure](#); essentially rolling back your deployment to a previous state.

You typically perform such a data 'rollback' due to data corruption, or the need to revert some changes made to your Agile Service Manager deployment.

#### Restoring to a different cluster

In this scenario you recreate your Agile Service Manager deployment data in a different cluster from the one from which you have taken a backup.

You must have Agile Service Manager successfully deployed on your new target cluster.

Typically, you would restore data to a different cluster in a disaster recovery situation, where your primary cluster is not accessible; or in a situation where you want to clone a production system to a test environment.

The backup procedure stores the backup generated files inside the Agile Service Manager Cassandra pods inside the `/opt/ibm/cassandra/data/backup_tar/` directory. Ensure that these backup files are present in the target cluster Cassandra pods before attempting this scenario; either copy them to the new location, or mount the external storage to that location. Once the backup files are in the correct target location, you restore the backed up data as documented in the [procedure](#).

**Tip:** The `backup_tar` directory may not exist if you did not set a `storageClass` for `cassandrak` during the Agile Service Manager installation.

### Losing a Cassandra node in your OCP cluster

This scenario describes the steps to perform should you lose a worker node in your OCP cluster where one of the Agile Service Manager Cassandra pods is running, thereby effectively losing one of your Cassandra replica nodes.

This might happen for a variety of reasons and will leave your Cassandra cluster with two remaining functioning nodes, to which you then add a new node to restore your three-node configuration.

See the [“Replacing a Cassandra node in your OCP cluster”](#) on page 303 steps for more details.

## Procedure

### Preparing your system for data restoration

**Restoring to a different cluster:** The backup procedure stores the backup-generated files inside the Agile Service Manager Cassandra pods inside the `/opt/ibm/cassandra/data/backup_tar/` directory. **If you are restoring your data to a different cluster**, ensure that these backup files are present in the target cluster Cassandra pods: either copy them to the new location, or mount the external storage to that location.

1. Authenticate into the Agile Service Manager Kubernetes namespace.
2. Deploy the `kPodLoop` bash shell function.

**kPodLoop** is a bash shell function that allows a command to be run against matching Kubernetes containers. You can copy it into the shell.

```
kPodLoop() {
  __podPattern=$1
  __podCommand=$2
  __podList=$( kubectl get pods --field-selector=status.phase=Running --no-headers=true --
output=custom-columns=NAME:.metadata.name | grep ${__podPattern} )
  printf "Pods found: $(echo -n ${__podList})\n"
  for pod in ${__podList}; do
    printf "\n==== EXECUTING COMMAND in pod: %-42s =====\n" ${pod}
    kubectl exec ${pod} -- bash -c "${__podCommand}"
    printf '%.0s' {1..80}
    printf "\n"
  done;
}
```

This `kPodLoop` bash shell function filters the pods to run the commands against only those that are in a 'Running' phase. This filter ensures that the configuration pods that are only run as part of your installation, like the secret generator pod, are skipped.

3. Make a note of the scaling of Agile Service Manager pods.

```
kubectl get pods --field-selector=status.phase=Running --no-headers=true --output=custom-
columns=CNAME:.metadata.ownerReferences[0].name | grep noi-topology | uniq --count
```

### Example output:

```
1 noi-topology-dns-observer
1 noi-topology-docker-observer
3 noi-topology-elasticsearch
1 noi-topology-file-observer
1 noi-topology-kubernetes-observer
1 noi-topology-layout
```



```
1 noi-topology-merge
1 noi-topology-noi-gateway
1 noi-topology-noi-probe
1 noi-topology-observer-service
1 noi-topology-search
1 noi-topology-status
1 noi-topology-topology
1 noi-topology-ui-api
```

4. Verify access to each Cassandra database (this command will return a list of keyspaces from each Cassandra node).

```
kPodLoop noi-cassandra "cqlsh -u ${CASSANDRA_USER} -p ${CASSANDRA_PASS} -e \"DESC
KEYSPACES;\""
```

5. Scale down Agile Service Manager pods.

```
kubect1 scale deployment --replicas=0 noi-topology-dns-observer
kubect1 scale deployment --replicas=0 noi-topology-file-observer
kubect1 scale deployment --replicas=0 noi-topology-kubernetes-observer
kubect1 scale deployment --replicas=0 noi-topology-observer-service
kubect1 scale deployment --replicas=0 noi-topology-noi-gateway
kubect1 scale deployment --replicas=0 noi-topology-noi-probe
kubect1 scale deployment --replicas=0 noi-topology-layout
kubect1 scale deployment --replicas=0 noi-topology-merge
kubect1 scale deployment --replicas=0 noi-topology-status
kubect1 scale deployment --replicas=0 noi-topology-search
kubect1 scale deployment --replicas=0 noi-topology-ui-api
kubect1 scale deployment --replicas=0 noi-topology-topology
```

The Cassandra and ElasticSearch pods (noi-cassandra and noi-topology-elasticsearch) are left active. Cassandra pods need to be running in order to execute the backup of their data, whereas the ElasticSearch pods have no interaction with nor influence on the Cassandra contents, so can be kept running.

**Important:** Include in this scale down any additional observers you have installed in your deployment.

6. Verify that scaling down was successful.

```
kubect1 get pods --field-selector=status.phase=Running | grep noi-topology
```

The Agile Service Manager services have now been scaled down, and the Cassandra database contents will not be modified anymore.

## Restore data

7. Update the Cassandra restore script to suppress the truncation of restored data.

**Note:** The `restore_cassandra.sh` tool truncates all data in the target table each time it is used, and despite the restore being targeted at one Cassandra node only, the truncate is propagated to all nodes. In order to suppress the truncate step, you must update the restore script on all but the first node.

- a) Copy `cassandra_functions.sh` out of one of the asm-cassandra nodes.

```
kubect1 cp noi-cassandra-0:/opt/ibm/backup_scripts/cassandra_functions.sh /tmp/
cassandra_functions.sh
```

- b) Edit `cassandra_functions.sh`

```
vi /tmp/cassandra_functions.sh
```

Locate the call to `truncate_all_tables` within the `restore()` function and comment out the appropriate lines, as in the following example:

```
Printf "`date` Starting Restore \n"

#### truncate_all_tables
#### testResult $? "truncate tables"

repair_keyspace
```

c) Save the file, then copy the file back to all nodes, except the first Cassandra node.

```
kubectl cp /tmp/cassandra_functions.sh noi-cassandra-1:/opt/ibm/backup_scripts/
cassandra_functions.sh

kubectl cp /tmp/cassandra_functions.sh noi-cassandra-2:/opt/ibm/backup_scripts/
cassandra_functions.sh
```

8. Locate the timestamps of the backups from each Cassandra node to restore.

Each node's backup was started at a similar time, so the timestamps may differ by a few seconds. In the following example a backup was performed at about 2019-06-11 09:36, and grep is then used to filter to these backup archives.

**Tip:** You can ignore this step if you are about to apply the most recent backup. If you do, the **-t** parameter can be omitted during all subsequent steps.

```
kPodLoop noi-cassandra "ls -larth \${CASSANDRA_DATA}/../backup_tar | grep 2019-06-11-09"

Pods found: noi-cassandra-0 noi-cassandra-1 noi-cassandra-2

===== EXECUTING COMMAND in pod: noi-cassandra-0 =====
-rwxrwxr-x 1 cassandra cassandra 524M Jun 11 09:37 cassandra_noi-topology-
cassandra-0_KS_system_schema_KS_system_KS_system_distributed_KS_system_auth_KS_janusgraph_KS
_system_traces_date_2019-06-11-0936-04.tar
-----

===== EXECUTING COMMAND in pod: noi-cassandra-1 =====
-rwxrwxr-x 1 cassandra cassandra 565M Jun 11 09:37 cassandra_noi-topology-
cassandra-1_KS_system_schema_KS_system_KS_system_distributed_KS_system_auth_KS_janusgraph_KS
_system_traces_date_2019-06-11-0936-07.tar
-----

===== EXECUTING COMMAND in pod: noi-cassandra-2 =====
-rwxrwxr-x 1 cassandra cassandra 567M Jun 11 09:37 cassandra_noi-topology -
cassandra-2_KS_system_schema_KS_system_KS_system_distributed_KS_system_auth_KS_janusgraph_KS
_system_traces_date_2019-06-11-0936-07.tar
-----
```

9. Working across each Cassandra node, restore the relevant backup of the janusgraph keyspace.

**Note:** For information about the system\_auth keyspace see the [Secrets and system\\_auth keyspace](#) note.

The specific node pod names can vary depending on your [installation and release names used](#)

a) noi-cassandra-0

**Remember:** This will cause the existing data in the janusgraph keyspace tables to be truncated.

```
kPodLoop noi-cassandra-0 "/opt/ibm/backup_scripts/restore_cassandra.sh -k janusgraph -t
2019-06-11-0936-04 -u \${CASSANDRA_USER} -p \${CASSANDRA_PASS} -f"

kPodLoop noi-cassandra-0 "nodetool repair --full janusgraph"
```

b) noi-cassandra-1

```
kPodLoop noi-cassandra-1 "/opt/ibm/backup_scripts/restore_cassandra.sh -k janusgraph -t
2019-06-11-0936-07 -u \${CASSANDRA_USER} -p \${CASSANDRA_PASS} -f"

kPodLoop noi-cassandra-1 "nodetool repair --full janusgraph"
```

c) noi-cassandra-2

```
kPodLoop noi-cassandra-2 "/opt/ibm/backup_scripts/restore_cassandra.sh -k janusgraph -t
2019-06-11-0936-07 -u \${CASSANDRA_USER} -p \${CASSANDRA_PASS} -f"

kPodLoop noi-cassandra-2 "nodetool repair --full janusgraph"
```

## Restore services

10. Scale up the services to the original level.

The original level was obtained in a [previous step](#).

```
kubectl scale deployment --replicas=1 noi-topology-topology
kubectl scale deployment --replicas=1 noi-topology-layout
kubectl scale deployment --replicas=1 noi-topology-merge
kubectl scale deployment --replicas=1 noi-topology-status
kubectl scale deployment --replicas=1 noi-topology-search
kubectl scale deployment --replicas=1 noi-topology-observer-service
kubectl scale deployment --replicas=1 noi-topology-noi-gateway
kubectl scale deployment --replicas=1 noi-topology-noi-probe
kubectl scale deployment --replicas=1 noi-topology-ui-api
kubectl scale deployment --replicas=1 noi-topology-dns-observer
kubectl scale deployment --replicas=1 noi-topology-file-observer
kubectl scale deployment --replicas=1 noi-topology-rest-observer
kubectl scale deployment --replicas=1 noi-topology-kubernetes-observer
```

#### 11. Rebroadcast data to ElasticSearch (that is, re-index Elasticsearch).

If data in Elasticsearch is out of sync with data in the Cassandra database, resynchronize it by calling the `rebroadcast` API of the topology service. This triggers the rebroadcast of all known resources on Kafka, and the Search service will then index those resources in Elasticsearch. Call the `rebroadcast` API of the Topology service, specifying a `tenantId`:

```
https://master_fqdn/1.0/topology/swagger#!/Crawlers/rebroadcastTopology
```

### Related tasks

[“Backing up database data \(OCP\)” on page 295](#)

You can back up (and later restore) existing topology data for the Agile Service Manager **OCP** installation. This can be helpful when updating your system, as part of your company's data management best-practice, or for maintenance reasons.

[“Backing up and restoring UI configuration data \(OCP\)” on page 305](#)

Agile Service Manager on OCP includes a backup facility, which lets you backup UI configuration settings such as user preferences, topology tools, custom icons, relationship types, and advanced topology settings.

## Replacing a Cassandra node in your OCP cluster

Agile Service Manager configuration for Cassandra sets a three-node cluster with a replication factor of three, which means your deployment would still be fully functional should you lose one node. To mitigate the risk of a second node failing, however, you would perform the steps documented here to restore the level of service to a three node Cassandra cluster.

### About this task

**Important:** These steps provide basic instructions on how to diagnose and recover from the described situation. However, it is likely that in your specific situation you may encounter differences in how your OCP cluster or Cassandra behave under such a failure, in which case you should engage your company Cassandra administrators.

## Procedure

### Verify the state of your Cassandra cluster

1. Authenticate into the Kubernetes namespace where Agile Service Manager is deployed as part of your solution.
2. Check the status of your Cassandra pods.

The following example system output shows that a pod is being terminated because of a node on the cluster failing:

NAME	READY	STATUS	RESTARTS	AGE	
noi-cassandra-0	1/1	Running	1	4d2h	
noi-cassandra-1	1/1	Running	1	4d2h	
noi-cassandra-2	1/1	Terminating	1	4d2h	

3. Verify that the Cassandra node is down.

For the pods that are still running, use a command like the following example:

```
kubectl exec noi-cassandra-0 -- bash -c "/opt/ibm/cassandra/bin/nodetool status -r"
```

Sample system output, where the line starting with 'D' indicates a node that is down.:

```
Datacenter: datacenter1
=====
Status=Up/Down
-- State=Normal/Leaving/Joining/Moving
-- Address      Load      Tokens     Owns (effective)  Host ID
ID            Address      Load      Tokens     Owns (effective)  Host ID
UN noi-cassandra-0.noi-cassandra.<project>.svc.cluster.local 64.34 MiB 256
100.0%        f6e6f151-ca7b-4117-be87-97245e61d7e9 rack1
UN noi-cassandra-1.noi-cassandra.<project>.svc.cluster.local 64.34 MiB 256
100.0%        989027b6-896b-4622-b282-9aa1dc2d9e39 rack1
DN 10.254.4.4    64.31 MiB 256          100.0%          ce054185-d72a-4b48-9c34-e8199b6e1559 rack1
```

### Restore a three-node Cassandra cluster

4. Authenticate into the Kubernetes namespace where Agile Service Manager is deployed as part of your solution.
5. Remove the Cassandra node that is down from the Cassandra cluster before the new one comes online.

Use a command like the following example (using **your node ID**).

This command runs Cassandra's **nodetool** utility in one of the nodes still up (in this case noi-cassandra-0) to remove the Cassandra node that is marked as down.

```
kubectl exec noi-cassandra-0 -- bash -c "/opt/ibm/cassandra/bin/nodetool removenode ce054185-d72a-4b48-9c34-e8199b6e1559"
```

6. Confirm the deletion using a command like the following example against one of the running pods:

```
kubectl exec noi-cassandra-0 -- bash -c "/opt/ibm/cassandra/bin/nodetool status -r"
```

7. Delete the pod that was running in the lost node reporting Terminating state:

```
kubectl delete pod noi-cassandra-2 --grace-period=0 --force
```

**Tip:** You may not need to delete the pod, depending on your cluster configuration and the outcome of your investigation. It will, however, be necessary to delete it if the pod is permanently stuck in the 'Terminating' state, as in this example.

8. Bring the new node online in your OCP cluster.

The container is initialized to join the Cassandra cluster, replacing the removed node.

**Troubleshooting:** Check whether the newly added Cassandra node lists itself as a seed. Run the following command to check the seeds configured. In this example, the added node is noi-cassandra-2 pod.

```
kubectl exec noi-cassandra-2 -- bash -c "grep seeds: /opt/ibm/cassandra/conf/cassandra.yaml"
```

System output example listing the seeds configured for the node:

```
- seeds: "10.254.12.2,10.254.8.7"
```

If the newly added node is listing itself as a seed, it can report inconsistent information with unexpected results, should it be the node queried by the Agile Service Manager services. To limit the potential impact, follow the [Preparing your system for data restoration](#) steps to stop your Agile Service Manager services from accessing your data until you have stabilized the new node.

9. Perform a full repair.

The following command instructs Cassandra to perform a repair of the data:

```
kubectl exec noi-cassandra-0 -- bash -c "/opt/ibm/cassandra/bin/nodetool repair --full"
```

**Tip:** For large data sets, it is preferred that you run the previous repair command several times for a limited range of tokens each time. You can get a list of tokens with the following command:

```
kubectl exec noi-cassandra-0 -- bash -c "/opt/ibm/cassandra/bin/nodetool info --tokens"
```

Example system output:

```
ID                : e2494466-0cc9-4268-a5f8-5d5fa363faaa
...
Token             : -9026954462746495840
Token             : -8998340199710379626
...
Token             : 9099714334544743528
Token             : 9120502118133589206
```

Run the repair command several times, each time specifying a different range of tokens, for example:

```
kubectl exec noi-cassandra-0 -- bash -c "/opt/ibm/cassandra/bin/nodetool repair --start-token -9026954462746495840 --end-token -8998340199710379626"
```

You can check the progress of the repair in the Cassandra pods logs.

#### 10. Perform a checksum verification of your data.

For example:

```
kubectl exec noi-cassandra-0 -- bash -c "/opt/ibm/cassandra/bin/nodetool verify -e"
```

If any errors are returned, repeat the [full repair step](#) until the checksum verification no longer returns any errors.

## Backing up and restoring UI configuration data (OCP)

Agile Service Manager on OCP includes a backup facility, which lets you backup UI configuration settings such as user preferences, topology tools, custom icons, relationship types, and advanced topology settings.

### About this task

The UI configuration data is stored in the Cassandra database and therefore backed up with the rest of the data, as described in the [backup](#) and [restore](#) topics. However, using the following procedure, you can test UI customizations in a staging environment, before exporting them back into the production environment.

### Procedure

1. Find the name of the topology pod, as in the following example:

```
$ kubectl get pod --namespace default --selector app=topology
NAME                                READY    STATUS    RESTARTS   AGE
noi-topology-topology-577dc5497b-2wbxk  1/1      Running   0           12h
```

2. Run the backup tool using `kubectl exec`, as in the following examples:

#### Example A

```
$ kubectl exec -ti noi-topology-topology-577dc5497b-2wbxk -- /opt/ibm/graph.tools/bin/
backup_ui_config -help

usage: backup_ui_config [-config <config_type>] [-out <output_filename>] [-force] [-verbose]
where 'config-type' can be set to one of the following:

all                - backup all ASM UI configuration (default)
```

```

tools      - backup topology tools definitions
icons      - backup custom icon definitions
types      - backup entity type definitions
links      - backup relationship type definitions
preferences - backup user preferences
settings   - backup advanced topology settings

```

### Example B

```

$ kubectl exec -ti noi-topology-topology-577dc5497b-2wbxk --
/opt/ibm/graph.tools/bin/backup_ui_config -out backup-20180908.json
INFO : Topology Service REST host detected: localhost:8080
INFO : Topology Service tenant ID detected: cfd95b7e-3bc7-4006-a4a8-a73a79c71255
WARNING : No topology tool definitions were found
WARNING : No custom icon definitions were found
INFO : Backing up entity type: container
INFO : Backing up entity type: cpu
INFO : Backing up entity type: deployment
INFO : Backing up entity type: image
INFO : Backing up entity type: namespace
INFO : Backing up entity type: namespace
INFO : Backing up entity type: networkinterface
INFO : Backing up entity type: operatingsystem
INFO : Backing up entity type: pod
INFO : Backing up entity type: server
INFO : Backing up entity type: service
INFO : Backing up entity type: volume
WARNING : No relationship type definitions were found
WARNING : No user preferences definitions were found
WARNING : No advanced topology settings definitions were found
INFO : Output file has been created: /opt/ibm/netcool/asm/data/tools/
backup-20180908.json

Program complete.

```

### 3. Run the import tool, as in the following example:

```

$ kubectl exec -ti noi-topology-topology-577dc5497b-2wbxk --
/opt/ibm/graph.tools/bin/import_ui_config -file backup-20180908.json -overwrite
INFO : Topology Service REST host detected: localhost:8080
INFO : Topology Service tenant ID detected: cfd95b7e-3bc7-4006-a4a8-a73a79c71255
INFO : Skipping import of entity type because it matches the existing definition:
container
INFO : Skipping import of entity type because it matches the existing definition: cpu
INFO : Skipping import of entity type because it matches the existing definition:
deployment
INFO : Skipping import of entity type because it matches the existing definition: image
INFO : Skipping import of entity type because it matches the existing definition:
networkinterface
INFO : Skipping import of entity type because it matches the existing definition: psu
INFO : Skipping import of entity type because it matches the existing definition: router
INFO : Skipping import of entity type because it matches the existing definition: sensor
INFO : Skipping import of entity type because it matches the existing definition: server
INFO : Skipping import of entity type because it matches the existing definition: service
INFO : Skipping import of entity type because it matches the existing definition: subnet
INFO : Skipping import of entity type because it matches the existing definition: switch
INFO : Skipping import of entity type because it matches the existing definition: vlan
INFO : Skipping import of entity type because it matches the existing definition: vpn

Program complete.

```

### 4. To save a copy of your backup, copy the file out of the topology container using the `kubectl cp` command.

For example:

```

$ kubectl cp noi-topology-topology-577dc5497b-2wbxk:/opt/ibm/netcool/asm/data/tools/
backup-20180908.json
/tmp/backup-20180809.json
$ find /tmp/backup*
/tmp/backup-20180809.json

```

### 5. To import files, copy them into the `/opt/ibm/netcool/asm/data/tools` location inside the container:

```

$ kubectl cp /tmp/backup-20180809.json noi-topology-topology-577dc5497b-2wbxk:/opt/ibm/
netcool/asm/data/

```

```
tools/backup-20180909.json
$ kubect1 exec -ti noi-topology-topology-577dc5497b-2wbxk -- find /opt/ibm/netcool/asm/data/
tools/
/opt/ibm/netcool/asm/data/tools/backup-20180908.json
/opt/ibm/netcool/asm/data/tools/backup-20180909.json
```

## Launching in context from OMNIbus Event Viewer

You can set up launch-in-context functionality in DASH from a Netcool/OMNIbus Web GUI event list (Event Viewer) to an Agile Service Manager Topology Viewer portlet, using a DASH **NodeClickedOn** action event.

Using the NodeClickedOn DASH event to launch the topology viewer with a set of parameters is similar to using a direct-launch URL (as described in the following topic: [“Accessing topologies via direct-launch URL string”](#) on page 247).

### Advantages of using a NodeClickedOn event for launch-in-context (rather than a direct URL)

Removes the need to re-load the entire Topology Viewer when rendering a new topology.

Allows the use of the `asmFunctions.sendPortletEvent` function to the opened Topology Viewer.

## Updating a topology on the same DASH page

You can add a right-click menu item to the Netcool/OMNIbus Web GUI event list, which you can use to update an already-open Agile Service Manager Topology Viewer that is on the **same** page in DASH as the event list.

### Procedure

1. As a user with the Web GUI admin role, log into your DASH web application.
2. Open the Netcool/OMNIbus Web GUI tool configuration page: **Administration > Event Management Tools > Tool Configuration**
3. Create a new Script tool, with a script command similar to the following example:

In this example the topology seed is a vertex name whose value is derived from the Node field in a Netcool/OMNIbus event.

```
var eventPayload = {
  "name": "http://ibm.com/tip#NodeClickedOn",
  "payload": {
    "product": {
      "AgileServiceManager": {
        "resourceName": "{@Node}",
        "hops": 2,
        "hopType": "e2h"
      }
    }
  }
};
{$param.portletNamespace}sendPortletEvent(eventPayload);
```

**Tip:** Use the following list of supported parameters: [Table 65 on page 308](#)

4. Open the Netcool/OMNIbus Web GUI menu configuration page: **Administration > Event Management Tools > Menu Configuration**
5. Add the tool to the **Alerts** menu.
6. In DASH, click **Console Settings > General > Pages** and create a new page.
7. Add an Event Viewer portlet and a Topology Viewer portlet to the page, and arrange them as required.

### Results

When you select an event from your Event Viewer portlet and launch your new tool, the Topology Viewer on the same DASH page will be updated, and render the topology for the seed whose resource name property matches the Node field value from the selected event.

## Updating a topology on a different DASH page

You can add a right-click menu item to the Netcool/OMNIbus Web GUI event list, which you can use to update a Topology Viewer that is on a **different** DASH page from the event list.

### Procedure

1. As a user with the Web GUI admin role, log into your DASH web application.
2. Open the Netcool/OMNIbus Web GUI tool configuration page: **Administration > Event Management Tools > Tool Configuration**
3. Create a new Script tool, with a script command similar to the following example:

The 'NavigationNode' value must be the Page Unique Name for the DASH page that you want to launch. In this example it is the unique name of the out-of-the-box Topology Viewer page.

```
var eventPayload = {
  "name": "http://ibm.com/isclite#launchPage",
  "NavigationNode": "web.netcool.asm.topologyViewer.page",
  "switchPage": "true",
  "payload": {
    "product": {
      "AgileServiceManager": {
        "resourceName": "{@Node}",
        "hops": "3",
        "layoutType": "4",
        "layoutOrientation": "TopToBottom"
      }
    }
  }
};
{$param.portletNamespace}sendPortletEvent(eventPayload);
```

**Tip:** Use the following list of supported parameters: [Table 65 on page 308](#)

4. Open the Netcool/OMNIbus Web GUI menu configuration page: **Administration > Event Management Tools > Menu Configuration**
5. Add the tool to the **Alerts** menu.

### Results

When you select an event from your Event Viewer portlet and launch your new tool, the Topology Viewer page will be opened or updated, and render the topology for the seed whose resource name property matches the Node field value from the selected event.

## Launch-in-context parameters

The event payload of a DASH event used to update a Topology Viewer can have any of the parameters listed in this topic.

### Parameters

**Tip:** The following parameters are supported by the [direct-launch URL](#) and the [UI-API render](#) facilities.

Table 65. Launch-in-context parameters		
Parameter	Type	Purpose
deltaTime	Integer	A point in history to compare to, as a unixtime
focusId	String	The 'id' of a resource or the '_id' of a resource status, used for zooming into a resource area
hideSearch	String	Hides the top search bar if set to 'true'
hideToolbar	String	Hides the side toolbar if set to 'true'



Table 65. Launch-in-context parameters (continued)		
Parameter	Type	Purpose
hops	Integer	The number of hops from the seed to display
hopType	String	The type of hops to display, for example 'host', 'e2h'
layoutOrientation	String	The layout orientation, for example 'TopToBottom'
layoutType	Integer	The topology <a href="#">layout type</a> , as a number
neighbourRequests	String	The 'id' of a neighboring resource or the '_id' of a neighboring resource status included with the seed resource
refreshTime	Integer	The refresh rate in milliseconds (when <b>not</b> in historical mode).  You can specify a value of 0 (zero) for a paused view.
relationFilter	String	One or more advanced, predefined relationship filters
resourceFilter	String	One or more advanced, predefined resource filters
resourceId	String	The 'id' of the seed resource or the '_id' of a resource status
resourceName	String	The 'name' of the seed resource
resourceUniqueId	String	The 'uniqueId' of the seed resource
search	String	The partial name of a resource used to search for and return either topologies, or the Search Page (if more than one result)
selectedIds	String	The 'id' of a resource or the '_id' of a resource status, used for highlighting one or more additional (non-seed) resources
time	Integer	The point in history to view the topology, as a unixtime

## Defining rules

Rules help streamline topologies and conserve system resources, for example by merging different observer records of the same resource into a single composite resource, or by excluding specific changes from being recorded against a resource history.

### Before you begin

- You must know your tenant ID.
- You also need to know specific details about resources for which you intend to develop rules. For example, to create merge rules you must know which resources exist as duplicate records before you can merge these into composites.

### Version 1.1.5 Notice:

Existing Agile Service Manager 1.1.4 merge rules will work without the need of any migration, but any scripts that contain rules need to be changed to the new Agile Service Manager 1.1.5 format if they are to be run on Version 1.1.5 (**or later**).

## About this task

### Rule Types:

You can use a number of different types of rules for different purposes. The rule type (**ruleType**) can be one the following:

#### **mergeRule**

A merge rule populates the tokens of resources matched by the rule to prevent duplicate records of the same resource from being displayed in the topology.

See [About Merge Rules](#) for more information.

#### **tagsRule**

The tags rule allows the value of another field to be copied to the tags field so that it becomes searchable in the UI.

It populates the **tags** of resources matched by the rule.

See [About Tag Rules](#) for more information.

**Tip:** Any field that isn't indexed and can therefore not normally be searched for becomes searchable if copied to the **tags** property. For a list of indexed properties, see [Table 67 on page 354](#).

#### **matchTokensRule**

A match tokens rule allows the value of another field to be copied to the matchTokens field so that it can be used to match events to resources.

It populates the **matchTokens** of resources matched by the rule.

See [About Match Token Rules](#) for more information.

#### **historyRule**

A history rule identifies properties that change every observation, but that don't indicate an important change, for example a host's sysUpTime property.

A history rule allows you to exclude properties from being retained in history, thereby saving resources by not maintaining detailed historical records of changes to these properties.

See [About History Rules](#) for more information.

### Custom Observer Rules:

Use the **observers** parameter to list the observers to which a rule applies. You can find more information for specific observers here:

#### **Jenkins Observer**

[“Jenkins Observer rules” on page 314](#)

## Procedure

- Using the live Swagger interface or cURL commands, write a rule for each record, and POST it to the Rules API.

Use the following information as guidance when defining each rule:

#### **Name**

The name of the rule, which must be unique within the context of the tenant.

#### **Rule type (ruleType)**

The rule type specifies the type of rule, and can be one the following:

- **mergeRule**
- **tagsRule**
- **matchTokensRule**

## – historyRule

### Status (ruleStatus)

The rule status can be either enabled or disabled, and the observers will only apply rules which are in an enabled state.

### Merge tokens (tokens)

The tokens set in a merge rule contains the list of resource parameter names which will become merge tokens for those resources to which the rule is applied.

Merge tokens can be constructed using variable substitutions, which allows you to combine more than one property value in a token, and also combine them with literal strings, as shown in this [example](#).

**Important:** The tokens are the shared elements that the duplicate records to be merged have in common.

### Entity types (entityTypes)

The entity types set in a rule contain the list of resource entity types for which this rule is valid.

If omitted, the rule will apply to all entity types.

### Observers (observers)

The observers set contains the list of names of the observers to which this rule applies.

If omitted, or set to include the value '\*', the rule will apply to all observers.

### Providers (providers)

The providers set contains the list of names of the providers so which this rule applies.

If omitted, or set to include the value '\*', the rule will apply to all providers.

You can use the mutually exclusive **excludeTokens** and **includeTokens** properties to filter providers.

### Exclude tokens (excludeTokens)

These properties discard any values that match the regular expression.

### Include tokens (includeTokens)

These properties apply a token only if the value matches the regular expression.

## Results

After you have created and posted your rules, these are applied before each observer job is sent to the topology service.

## Example

### Sample rules for merging resources

A composite resource has its own unique identifier, and the merged resources continue to exist separately. The following example shows the individual and composite resources when retrieved from the topology **Resources** API.

#### Resource one

`https://<your_NASM_host>/1.0/topology/resources/ABC`

```
{
  "_id": "ABC",
  "name": "host1",
  "propertyAbc": "This property only exists on ABC",
  "entityTypes": [ "host" ],
  "tokens": [ "host1MergeToken" ],
  "_compositeId": "XYZ"
}
```

The resource has an id of ABC, and the value of compositeId denotes the single, merged resource, which is XYZ.

#### Resource two

`https://<your host>/1.0/topology/resources/DEF`

```
{
  "_id": "DEF",
  "name": "host1",
  "propertyDef": "This property only exists on DEF",
  "entityTypes": [ "host" ],
  "tokens": [ "host1MergeToken" ],
  "_compositeId": "XYZ"
}
```

The resource has an id of DEF, and the value of compositeId denotes the single, merged resource, which is XYZ.

### Composite resource

[https://<your\\_NASM\\_host>/1.0/topology/resources/XYZ](https://<your_NASM_host>/1.0/topology/resources/XYZ)

```
{
  "_id": "XYZ",
  "name": "host1",
  "propertyAbc": "This property only exists on ABC",
  "propertyDef": "This property only exists on DEF",
  "entityTypes": [ "host" ],
  "tokens": [ "host1MergeToken" ],
  "_compositeOfIds": [ "ABC", "DEF " ]
}
```

The resource has an id of XYZ, and the value of compositeOfIds lists the ids of the merged resources, in this case ABC and DEF. The XYZ composite resource includes the properties from both of the merged resources.

### Resource with variable substitutions and exclude list

```
{
  "name": "sysNameMatching",
  "tokens": [ "sysName", "${name}/${customField}" ],
  "ruleStatus": "enabled",
  "entityTypes": [ "host", "server" ],
  "observers": [ "ITNM", "TADDM" ],
  "providers": [ "*" ],
  "customField": "string",
  "excludeTokens": [ "^asm-default.*" ]
}
```

The `^asm-default.*` value set for `excludeTokens` ensures that any values that match the regular expressions are excluded.

The merge token with the value of `${name}/${customField}` combine the `${name}` and `${customField}` properties using the `{}{}` syntax, and demonstrate how variable substitutions work.

- Literal values are entered as they are in the merge token, which in this case is the `/` character.
- To be backwards compatible, tokens consisting of a single value, as in the `sysName` example, are treated as variable substitutions, that is, as if they are `${sysName}`.

**Tip:** You can also view composite resources in the live Swagger Composite API in the Merge Service, which returns the properties of the composite itself, and provides methods to get all composite vertices: [https://<your\\_host>/1.0/merge/swagger/#/Composites](https://<your_host>/1.0/merge/swagger/#/Composites)

### Related information

[Blog: Using Merge Rules in Agile Service Manager](#)

## Rules examples and samples

Use the following information to define rules.

### Rule types

#### Merge Rules:

Different observers deployed as part of the Agile Service Manager solution may record and then display the same resource as two (or more) resources. To prevent this, you create a merge rule that ensures that

the separate records of the same resource share values in their tokens set, which then triggers the Merge Service to create a single composite resource vertex. Composite, merged resources are displayed in the Topology Viewer as a single resource, which includes the properties of all merged resources.

Merge rules are applied to a resource in a observer job before it is sent to the topology service. Rules can be managed using the Rules REST API in the Merge Service. For each Agile Service Manager observer, merge rules control which tokens are considered merge tokens. Live Swagger documentation for merge rules is here:

```
https://<your host>/1.0/merge/swagger/#/Rules
```

The following example is the default merge rule defined for the Docker Observer:

```
rules:
- name: dockerId
  ruleType: mergeRule
  ruleStatus: enabled
  tokens: [ dockerId ]
  entityTypes: null
  observers: [ docker-observer ]
  providers: null
- name: almExternalId
  ruleType: mergeRule
  ruleStatus: enabled
  tokens: [ externalId ]
  entityTypes: null
  observers: [ alm-observer ]
  providers: null
```

Notice that the rules name in this example is **dockerId**, and that it applies only to instances of observers named **docker-observer**. The ruleType property here specifies the **mergeRule** rule type. This merge rule applies to all entity types and all providers and will copy the **dockerId** property into the merge tokens for all resources that have this property set.

### Tags Rules:

The tags rule allows the value of another field to be copied to the tags field so that it becomes searchable in the UI. It populates the tags of resources matched by the rule.

**Example 1:** Simple example of a tagsRule:

```
- name: tagRuleCustomProp
  ruleType: tagsRule
  ruleStatus: enabled
  tokens: [ name ]
  entityTypes: null
  observers: null
  providers: null
```

**Example 2:** The following example creates tags rules for the Docker Observer and will add **com.docker.compose.service** and **com.docker.compose.project** properties to the tags list, for entities of the type **container**:

```
curl -k -s -X POST --user $PROXY_USER:$PROXY_PASSWORD --header 'Content-Type: application/json'
--header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255'
-w "Response_code: %{http_code}\n" -d '{
  "tokens": [
    "com.docker.compose.service",
    "com.docker.compose.project"
  ],
  "name": "tag-docker-container-service-project",
  "ruleStatus": "enabled",
  "ruleType": "tagsRule",
  "observers": [
    "docker_observer"
  ],
  "entityTypes": [
    "container"
  ]
}' "https://$PROXY_HOST/1.0/merge/rules"
```

To verify:

```
{
  "_createdAt": "2019-09-10T02:24:03.652Z",
  "_executionTime": 10,
  "_id": "YEaGjuCJT5-wCH-SAfCsaA",
  "_modifiedAt": "2019-09-10T15:29:33.829Z",
  "_observedAt": "2019-09-10T15:29:33.829Z",
  "_startedAt": "2019-09-10T15:29:33.829Z",
  "architecture": "x86_64",
  "changeTime": 1568129373829,
  "com.docker.compose.config-hash":
    "6afc3ec12fd17d2b815df61e4ab276af6a27324a1487567f1bbba29a4a4bca1e",
  "com.docker.compose.container-number": "1",
  "com.docker.compose.oneoff": "False",
  "com.docker.compose.project": "asm",
  "com.docker.compose.service": "search",
  "com.docker.compose.version": "1.11.2",
  "createTime": 1568082243652,
  "dockerId": "dc39b5ea19cb79755e48a525bfbd2d479777464b1e80c708cb851da1b22fa952",
  "entityTypes": [
    "container"
  ],
  "matchTokens": [
    "asm_search_1",
    "dc39b5ea19cb79755e48a525bfbd2d479777464b1e80c708cb851da1b22fa952"
  ],
  "mergeTokens": [
    "dc39b5ea19cb79755e48a525bfbd2d479777464b1e80c708cb851da1b22fa952"
  ],
  "name": "asm_search_1",
  "observedTime": 1568129373829,
  "tags": [
    "asm_search_1",
    "asm",
    "search"
  ],
  "uniqueId": "dc39b5ea19cb79755e48a525bfbd2d479777464b1e80c708cb851da1b22fa952",
  "vcs-type": "git",
  "vendor": "IBM",
  "version": "1.0.10-SNAPSHOT"
}
```

### Match Token Rules:

A match tokens rule allows the value of another field to be copied to the **matchTokens** field so that it can be used to match events to resources. It populates the **matchTokens** of resources matched by the rule.

Simple example of a matchTokensRule:

```
- name: matchRule
  ruleType: matchTokensRule
  ruleStatus: enabled
  tokens: [ name ]
  entityTypes: null
  observers: null
  providers: null
```

### History Rules:

A history rule populates the **historyExcludeTokens** field with field names.

When a resource is updated, the topology service checks if **all** the updates are on fields that are listed in **historyExcludeTokens**, and if they are, it updates the resource **without** creating history.

### Jenkins Observer rules

A **mergeRule** and an **eventsFilterRule** are included out-of-the-box for the Jenkins Observer.

A merge rule to merge Artifactory's published Docker modules information with the information provided by the Docker Observer is provided here: [“Defining Jenkins data merge rules” on page 58.](#)

## Using templates to generate defined topologies

---

You can create topology templates to generate defined topologies, which search your topology database for instances that match its conditions. These defined topologies can then be searched for in the Topology Viewer and displayed.

### Before you begin

You must have the admin role **inasm\_admin** assigned to you. See the [“Configuring DASH user roles”](#) on page 32 topic for more information.

### About this task

You can create a new template or edit an existing template.

To create a new template, you construct a topology view centered around a seed resource. This lets you search your database for topologies with resources and relationships that match the conditions defined in the template.

These defined topologies are dynamically generated, maintained and updated, and can be searched for and accessed in the Topology Viewer. Defined topologies are dynamically refreshed when accessed.

#### Dynamic changes to defined topologies

If edges or resources are removed that render an existing defined topology that has already been indexed incomplete, it will remain available, and can therefore still be found and displayed via the topology search, or via a previously created direct URL.

If edges or resources are added or changed that render a defined topology that was not indexed previously complete, the topology will be indexed automatically, thereby becoming accessible via the search service.

**Tip:** Searching for and then rendering resources is described in more detail in the [Search](#) section of the Topology Viewer reference topic.

### Procedure

1. As the admin user, log into your DASH web application.
2. Select **Administration** from the DASH menu.
3. Select **Topology Templates** under the Agile Service Management heading.
4. Enter a resource or template name.

As you type, a drop-down list is displayed with suggested search terms that already exist in the topology service. If you select one of these, the **Search Results** page is displayed, which lists the results under separate **Resources** and **Templates** tabs.

- For each resource or template, the name and other properties are displayed.
- Click **View Topology** next to a result to create a new template, or **View Template** to open an existing one.

The **Topology template builder** page is displayed, consisting of a topology view on the right, and a number of template definition fields on the left.

#### Edit an existing template

5. If you have selected an existing template from the Search results, you edit the template conditions as you would when creating a new template.
  - You can also delete an existing template from here.

#### Create a new template

If you have selected a resource rather than an existing template, a basic topology view centered around the resource as seed is displayed in the **Topology Visualization** panel on the right.

6. Enter a name for the template.

This name **must** be unique.

7. Define a defined topology naming pattern.

The seed name of the defined topology is substituted when a defined topology is generated, but you can define either a prefix, suffix, or both to create an appropriate name for each defined topology.

8. Add a tag to be included for each defined topology.

Tags allow you to group the individually generated defined topologies by matching them to resources

**Restriction:** The following special characters are **not** supported when specifying tags:

```
. ? + * | { } [ ] ( ) " \ # @ & < > ~
```

If existing resources have tags with these characters, you must change them before you can use them to define defined topologies.

9. Define the type of template used.

- **Exact:** You can define a template that uses the selected resource only. This allows you to create a single defined topology, which operators can access from the Topology Viewer.

Although an **exact** template type focuses on a single resource, it remains dynamic, which means as long as the seed resource itself remains unchanged, the defined topology can change through the removal or addition of connected resources or relationships.

- **Dynamic:** You can define a template for resources with types and tags similar to the selected resource. This allows you to search for defined topologies that follow the same relationship patterns.

10. Choose an icon to associate with the template using one of the following methods:

**From the Icon drop-down list**

If you open the **Icon** drop-down list, all icons are listed by name in alphabetical order.

**From the View all icons button**

If you click the **View all icons** button, all icons are displayed in alphabetical order.

Click an icon to associate it with the template.

**From the Quick assign button**

If an icon exists with the same name as the template, click the **Quick assign** button to select it without having to sort through all existing icons.

This function is useful if you have added a custom icon, and are now assigning it to a template with the same name.

**From the Define new custom icon button**

From here you can define a custom icon, which is automatically associated with the template when done.

Click the **Define new custom icon** button to display the **Custom Icons** page.

Click **New** to create a new icon. Alternatively, click the 'edit' symbol to edit an existing icon. Use the following information to define the icon:

**Icon properties**

Each custom icon must have a name that uniquely identifies the icon when assigning it to a type.

**Remember:** You cannot change the name of an existing icon. If you want an icon to have a different name, create a new icon, then delete the old one.

Icons are defined inside an SVG editor, which performs an XML validation check.

Each icon definition must be valid svg xml with a given viewBox, which is important to ensure scaling of the image.

The svg definition must include inline styling of the image, such as stroke color and fill color. Use of style classes is not advised, as it can cause visualization issues on some browsers. If style classes must be used, naming must be unique for each svg image to prevent class definitions from being overwritten.



Optionally, each icon can be assigned to a category, which allows you to group icons of the same type or function together when displaying them in a list.

**Remember:** You can also create custom icons from the **Custom Icons** page accessed through DASH, which is described in the [“Defining custom icons”](#) on page 279 topic.

All defined topologies based on this template will display the icon.

11. Define the topology conditions.

Use the context (right-click) menu to perform a number of resource-specific actions. You can view the resource's details, get its neighbors, or follow adjacent relationships.

As you define the template conditions, a preview function retrieves information from the topology database, indicating the number of matches. If these exceed 500, you cannot generate the defined topologies. Refine your conditions until the number of defined topologies are within the limit.

12. Click **Save template and generate defined topologies**.

A list of defined topologies is displayed. From here, you can view individual defined topologies to verify that they meet your requirement. If they do not, refine your template.

13. Define a default layout for the defined topologies.

**Layout types**

You can choose from a number of layout types and orientations.

**Layout 1**

A layout that simply displays all resources in a topology without applying a specific layout structure.

**Layout 2**

A circular layout that is useful when you want to arrange a number of entities by type in a circular pattern.

**Layout 3**

A grouped layout is useful when you have many linked entities, as it helps you visualize the entities to which a number of other entities are linked. This layout helps to identify groups of interconnected entities and the relationships between them.

**Layout 4**

A hierarchical layout that is useful for topologies that contain hierarchical structures, as it shows how key vertices relate to others with peers in the topology being aligned.

**Layout 5**

A peacock layout is useful when you have many interlinked vertices, which group the other linked vertices.

**Layout 6**

A planar rank layout is useful when you want to view how the topology relates to a given vertex in terms of its rank, and also how vertices are layered relative to one another.

**Layout 7**

A rank layout is useful when you want to see how a selected vertex and the vertices immediately related to it rank relative to the remainder of the topology (up to the specified amount of hops). The root selection is automatic.

For example, vertices with high degrees of connectivity outrank lower degrees of connectivity. This layout ranks the topology automatically around the specified seed vertex.

**Layout 8**

A root rank layout similar to layout 7, except that it treats the selected vertex as the root. This layout is useful when you want to treat a selected vertex as the root of the tree, with others being ranked below it.

Ranks the topology using the selected vertex as the root (root selection: Selection)

**Layout orientation**

**For layouts 4, 6, 7 and 8**, you can set the following layout orientations:

- Top to bottom

- Bottom to top
- Left to right
- Right to left

When a defined topology based on this template is selected in the Topology Dashboard, the default layout is used.

## Results

The defined topologies are saved in the topology database and can now be searched for by operators using the standard Search functionality, and then accessed in the Topology Viewer.

**Remember:** Defined topologies are updated automatically when the topology database is updated, and generated dynamically as they are displayed in the Topology Viewer.

## What to do next

Searching for and then rendering resources is described in more detail in the [Search](#) section of the Topology Viewer reference topic, and in the [Chapter 6, “Using Agile Service Manager,”](#) on page 245 topics.

## Improving database performance

You can improve the performance of Agile Service Manager, such as fine-tuning Cassandra database cluster operations.

### Changing the Cassandra `gc_grace_seconds` value (OCP)

The Cassandra database cluster performance may be impacted by tombstone occurrences which results in slower query response times. When this performance degrades the Agile Service Manager installation, the following procedure to set `gc_grace_seconds` can be used to mitigate this degradation.

#### About this task

For an environment with a Cassandra cluster, `gc_grace_seconds` can be reduced from the default value of 864000 seconds (10 days).

This parameter impacts the ability of the Cassandra cluster to repair itself after a node has been offline.

Define a value for `gc_grace_seconds` that is greater than the duration of any anticipated Cassandra node outage.

#### Procedure

1. Find the name of a Cassandra pod. The change can be carried out on any node as the change will be replicated across the nodes.

```
kubectl get pods | grep cass
```

##### For example:

```
$ kubectl get pods | grep cass
asm-cassandra-0      1/1      Running    0      9d
```

The pod is identified as `asm-cassandra-0`

2. Run the following command to exec into the pod and start `cqlsh`:

```
kubectl exec -ti {pod name} -- cqlsh -u cassandra -p cassandra
```

##### For example:

```
$ kubectl exec -ti asm-cassandra-0 -- cqlsh -u cassandra -p cassandra
Connected to apm_cassandra at asm-cassandra-0:9042.
```

```
[cqlsh 5.0.1 | Cassandra 3.11.3 | CQL spec 3.4.4 | Native protocol v4]
Use HELP for help.
cassandra@cqlsh>
```

Repeat steps 3 to 5 for all tables within the janusgraph key space.

3. Verify the current setting of gc\_grace\_seconds.

```
SELECT table_name,gc_grace_seconds FROM system_schema.tables
WHERE keyspace_name='janusgraph';
```

For example:

```
cassandra@cqlsh> SELECT table_name,gc_grace_seconds
FROM system_schema.tables WHERE keyspace_name='janusgraph';
```

table_name	gc_grace_seconds
edgestore	864000
edgestore_lock_	864000
graphindex	864000
graphindex_lock_	864000
janusgraph_ids	864000
system_properties	864000
system_properties_lock_	864000
systemlog	864000
txlog	864000

4. Update the values using the **ALTER TABLE** command:

```
ALTER TABLE janusgraph.{table name} WITH gc_grace_seconds = {gc_grace_seconds
value};
```

For example:

```
cassandra@cqlsh> ALTER TABLE janusgraph.edgestore WITH gc_grace_seconds = 345600;
```

5. Verify the settings have worked.

```
SELECT table_name,gc_grace_seconds FROM system_schema.tables
WHERE keyspace_name='janusgraph';
```

For example:

```
cassandra@cqlsh> SELECT table_name,gc_grace_seconds
FROM system_schema.tables WHERE keyspace_name='janusgraph';
```

table_name	gc_grace_seconds
edgestore	345600
edgestore_lock_	864000
graphindex	864000
graphindex_lock_	864000
janusgraph_ids	864000
system_properties	864000
system_properties_lock_	864000
systemlog	864000
txlog	864000

(9 rows)

**Remember:** Repeat steps 3 to 5 for all tables within the janusgraph key space.

6. Exit **cqlsh**:

Example:

```
cassandra@cqlsh> exit
$
```

## Changing the Cassandra gc\_grace\_seconds value (on-prem)

The Cassandra database cluster performance may be impacted by tombstone occurrences which results in slower query response times. When this performance degrades the Agile Service Manager installation, the following procedure to set gc\_grace\_seconds can be used to mitigate this degradation.

### About this task

For an on-prem environment, gc\_grace\_seconds can be safely set to 0.

### Procedure

1. Log into a server where Agile Service Manager is running.
2. Find the name of a Cassandra container.

```
docker ps | grep cassandra
```

#### For example:

```
$ docker ps | grep cassandra
000000000000 nasm-cassandra:3.11.3.62 "/opt/ibm/start-ca..."
5 hours ago Up 5 hours asm_cassandra_1
```

The pod is identified as asm-cassandra-1

3. Run the following command to exec into the container and start **cqlsh**:

```
docker exec -ti {container name} cqlsh -u cassandra -p cassandra
```

#### For example:

```
$ docker exec -ti asm_cassandra_1 cqlsh -u cassandra -p
Connected to topology_cassandra at asm_cassandra_1:9042.
[cqlsh 5.0.1 | Cassandra 3.11.3 | CQL spec 3.4.4 | Native protocol v4]
Use HELP for help.
cassandra@cqlsh>
```

Repeat steps 4 to 6 for all tables within the janusgraph key space.

4. Verify the current setting of gc\_grace\_seconds.

```
SELECT table_name,gc_grace_seconds FROM system_schema.tables
WHERE keyspace_name='janusgraph';
```

#### For example:

```
cassandra@cqlsh> SELECT table_name,gc_grace_seconds
FROM system_schema.tables WHERE keyspace_name='janusgraph';
```

table_name	gc_grace_seconds
edgestore	864000
edgestore_lock_	864000
graphindex	864000
graphindex_lock_	864000
janusgraph_ids	864000
system_properties	864000
system_properties_lock_	864000
systemlog	864000
txlog	864000

5. Change the value to 0 (zero) using the **ALTER TABLE** command:

```
ALTER TABLE janusgraph.{table name} WITH gc_grace_seconds = {gc_grace_seconds
value};
```

#### For example:

```
cassandra@cqlsh> ALTER TABLE janusgraph.edgestore WITH gc_grace_seconds = 0;
```

6. Verify the settings have worked.

```
SELECT table_name,gc_grace_seconds FROM system_schema.tables WHERE
keyspace_name='janusgraph';
```

**For example:**

```
cassandra@cqlsh> SELECT table_name,gc_grace_seconds FROM system_schema.tables
WHERE keyspace_name='janusgraph';
```

table_name	gc_grace_seconds
edgestore	0
edgestore_lock_	864000
graphindex	864000
graphindex_lock_	864000
janusgraph_ids	864000
system_properties	864000
system_properties_lock_	864000
systemlog	864000
txlog	864000

(9 rows)

**Remember:** Repeat steps 4 to 6 for all tables within the janusgraph key space.

7. Exit **cqlsh**:

**Example:**

```
cassandra@cqlsh> exit
$
```

## Changing the Cassandra `dclocal_read_repair_chance` value (OCP)

In OCP environments, the Cassandra cluster performance can be improved by setting `dclocal_read_repair_chance` to 0, thereby removing this Cassandra functionality. This functionality is not required as consistency issues are resolved by all Agile Service Manager 'read' and 'write' activities using a consistency level of QUORUM.

### Procedure

1. Find the name of a Cassandra pod. The change can be carried out on any node as the change will be replicated across the nodes.

```
kubectl get pods | grep cass
```

**For example:**

```
$ kubectl get pods | grep cass
asm-cassandra-0      1/1      Running    0      9d
```

The pod is identified as `asm-cassandra-0`

2. Run the following command to exec into the pod and start **cqlsh**:

```
kubectl exec -ti {pod name} -- cqlsh -u cassandra -p cassandra
```

**For example:**

```
$ kubectl exec -ti asm-cassandra-0 -- cqlsh -u cassandra -p cassandra
Connected to apm_cassandra at asm-cassandra-0:9042.
[cqlsh 5.0.1 | Cassandra 3.11.3 | CQL spec 3.4.4 | Native protocol v4]
Use HELP for help.
cassandra@cqlsh>
```

**Repeat steps 4 to 6 for all tables within the janusgraph key space.**

3. Verify the current setting of `dclocal_read_repair_chance`.

```
SELECT table_name,dclocal_read_repair_chance FROM system_schema.tables
WHERE keyspace_name='janusgraph';
```

**For example:**

```
cassandra@cqlsh> SELECT table_name,dclocal_read_repair_chance
FROM system_schema.tables WHERE keyspace_name='janusgraph';
```

table_name	dclocal_read_repair_chance
edgestore	0.1
edgestore_lock_	0.1
graphindex	0.1
graphindex_lock_	0.1
janusgraph_ids	0.1
system_properties	0.1
system_properties_lock_	0.1
systemlog	0.1
txlog	0.1

(9 rows)

4. Update the value to 0 (zero) using the **ALTER TABLE** command:

```
cassandra@cqlsh> ALTER TABLE janusgraph.edgestore WITH dclocal_read_repair_chance
= 0;
```

5. Verify the change has worked.

```
SELECT table_name,gc_grace_seconds,dclocal_read_repair_chance FROM
system_schema.tables WHERE keyspace_name='janusgraph';
```

**For example:**

```
cassandra@cqlsh> SELECT table_name,gc_grace_seconds,dclocal_read_repair_chance
FROM system_schema.tables WHERE keyspace_name='janusgraph';
```

table_name	dclocal_read_repair_chance
edgestore	0
edgestore_lock_	0.1
graphindex	0.1
graphindex_lock_	0.1
janusgraph_ids	0.1
system_properties	0.1
system_properties_lock_	0.1
systemlog	0.1
txlog	0.1

(9 rows)

**Remember:** Repeat steps 4 to 6 for all tables within the janusgraph key space.

6. Exit **cqlsh**:

**Example:**

```
cassandra@cqlsh> exit
$
```

## Configuring scaling for OCP

You can scale your Agile Service Manager deployment vertically or horizontally, and you can scale out as well as in.

**Remember:**

### scaling

You can scale up you system horizontally or vertically.

**Horizontal scaling** means increasing the replication factor of a particular service, and may also require adding additional hardware.

**Vertical scaling** means that you add more power (CPU or RAM) to an existing machine.

**Important:** To avoid data loss, enable persistence before scaling up your system.

**Tip:** The redistribution process places additional load on the cluster, so should be performed at quiet times.

**Restriction:** Scaling for Agile Service Manager observers is **not** supported, and neither is scaling from a single to multiple instances.

## Scaling vertically

To scale up Agile Service Manager vertically, you increase available CPU or RAM resources.

### About this task

The amount of CPU and RAM requested by Agile Service Manager services is controlled by the configuration option **global.environmentSize** with valid values being:

#### size0

Specifies the least amount of resources required to run Agile Service Manager.

Recommended for testing or proof-of-concept deployments only, and not suitable for high availability (HA) mode.

#### size1

Specifies the resource requirements that allow Agile Service Manager to run production workloads.

Suitable for HA mode.

**Network Discovery Sizing:** The Agile Service Manager network discovery service supports a maximum number of discovered entities.

#### For size0

10,000 entities

#### For size01

1,000,000 entities

The number of entities created by a discovery depend on factors such as the number of devices in scope, and how many interfaces, cards, VLANs and so on are present on each device.

## Scaling horizontally

To scale up Agile Service Manager horizontally, you first provision additional persistent volumes, then scale the component, and then redistribute, reassign or repair data across all nodes.

The following Agile Service Manager components can be scaled horizontally by adding additional machines, or pods, to your deployment:

- **Cassandra** database
- **ElasticSearch** search and analytics engine
- **Kafka** message bus
- **Zookeeper** synchronization service

**Assumption:** This task and the examples provided assume that you have deployed a standard production environment with three instances of each of the components that are to be scaled horizontally.

## Scaling horizontally: Persistence

Before you scale any of the components, you provision new storage volumes for the additional replicas to use.

### Before you begin

#### Source the kubhelper.sh helper function

The kubhelper.sh is provided in the Agile Service Manager pak\_extensions directory.

Source the helper script as follows:

```
$ source pak_extensions/common/kubhelper.sh
```

### About this task

The helper function has the following parameters:

- **Kubernetes worker node** for the volume
- **Helm release name** containing the application that needs a volume
- **Kubernetes namespace** where application is to be installed
- **Claim name** that needs storage
- **Storage capacity** required
- **Path** on the worker that will be used for storage

### Procedure

**Example:** The following example adds storage for an additional three Kafka brokers.

1. Three brokers have already been installed on the first three worker nodes.

```
$ kubectl get pod -l app=kafka,release=asm -o wide
NAME          READY   STATUS    RESTARTS   AGE   IP            NODE          NOMINATED NODE
asm-kafka-0    2/2     Running   0           108m  10.1.205.80   172.16.188.122 <none>
asm-kafka-1    2/2     Running   0           108m  10.1.34.201   172.16.154.233 <none>
asm-kafka-2    2/2     Running   0           108m  10.1.91.94    172.16.183.205 <none>
```

2. Considering the three brokers on the first three worker nodes, create volumes on other hosts so that the Kafka cluster remains resilient to node failures.

These are the other nodes:

```
$ kubectl get nodes
NAME                                STATUS    ROLES                                AGE   VERSION
172.16.153.121                     Ready    etcd,management,master,proxy,va    35h   v1.13.5+ocp-ee
172.16.154.233                     Ready    worker                               35h   v1.13.5+ocp-ee
172.16.183.205                     Ready    worker                               35h   v1.13.5+ocp-ee
172.16.188.122                     Ready    worker                               35h   v1.13.5+ocp-ee
172.16.190.218                     Ready    worker                               35h   v1.13.5+ocp-ee
172.16.191.196                     Ready    worker                               35h   v1.13.5+ocp-ee
172.16.192.211                     Ready    worker                               35h   v1.13.5+ocp-ee
172.16.192.70                      Ready    worker                               35h   v1.13.5+ocp-ee
```

3. Choose the next three workers in the list (that is, 172.16.190.218, 172.16.191.196 and 172.16.192.211) and then create the volumes as follows:

```
$ createPersistentVolume 172.16.190.218 asm netcool data-asm-kafka-3 15 /opt/ibm/netcool/asm/data/kafka
Wed Jun 12 13:35:26 PDT 2019 INFO: Checking if '172.16.190.218' is a valid worker node - OK
Wed Jun 12 13:35:26 PDT 2019 INFO: Checking if 'netcool' is a valid namespace - OK
Wed Jun 12 13:35:26 PDT 2019 INFO: Creating volume for pvc 'netcool/data-asm-kafka-3' with capacity
'15Gi'
at path '/opt/ibm/netcool/asm/data/kafka' on node '172.16.190.218'
persistentvolume/172.16.190.218-data-asm-kafka-3 created
```

```
$ createPersistentVolume 172.16.191.196 asm netcool data-asm-kafka-4 15 /opt/ibm/netcool/asm/data/kafka
Wed Jun 12 13:35:56 PDT 2019 INFO: Checking if '172.16.191.196' is a valid worker node - OK
Wed Jun 12 13:35:57 PDT 2019 INFO: Checking if 'netcool' is a valid namespace - OK
Wed Jun 12 13:35:57 PDT 2019 INFO: Creating volume for pvc 'netcool/data-asm-kafka-4' with capacity
'15Gi'
```



```
at path '/opt/ibm/netcool/asm/data/kafka' on node '172.16.191.196'
persistentvolume/172.16.191.196-data-asm-kafka-4 created
```

```
$ createPersistentVolume 172.16.192.211 asm netcool data-asm-kafka-5 15 /opt/ibm/netcool/asm/data/kafka
Wed Jun 12 13:36:38 PDT 2019 INFO: Checking if '172.16.192.211' is a valid worker node - OK
Wed Jun 12 13:36:38 PDT 2019 INFO: Checking if 'netcool' is a valid namespace - OK
Wed Jun 12 13:36:38 PDT 2019 INFO: Creating volume for pvc 'netcool/data-asm-kafka-5' with capacity
'15Gi'
at path '/opt/ibm/netcool/asm/data/kafka' on node '172.16.192.211'
persistentvolume/172.16.192.211-data-asm-kafka-5 created
```

4. Finally, create the paths on these workers.

```
$ ssh root@172.16.190.218 mkdir -p /opt/ibm/netcool/asm/data/kafka
$ ssh root@172.16.191.196 mkdir -p /opt/ibm/netcool/asm/data/kafka
$ ssh root@172.16.192.211 mkdir -p /opt/ibm/netcool/asm/data/kafka
```

## Scaling horizontally: Cassandra database

This task described how to scale the Cassandra database (both out and in).

### Before you begin

A default production deployment will have a replication factor of three. That means that with three nodes, each one should contain a complete copy of the data, as shown in this example:

```
nodetool status for asm-cassandra-0
Datacenter: datacenter1
=====
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
-- Address      Load        Tokens      Owns (effective)  Host ID                               Rack
UN  10.1.91.112   126.21 MiB   256         100.0%            dec6be10-4dcc-493e-8c26-330c19a32da2 rack1
UN  10.1.34.216   126.18 MiB   256         100.0%            e79c86a8-2105-4fad-b2ab-d10cdcd8354d rack1
UN  10.1.205.91   126.12 MiB   256         100.0%            77851686-5715-4237-86e4-31b62603ac2b rack1
```

### About this task

#### Scale out

To scale out and spread the data load, you add additional nodes to the Cassandra cluster.

#### Scale in

To scale in a Cassandra cluster, you must first decommission nodes, starting with the highest numbered pod.

During this process data must be moved to the remaining cluster nodes. The decommission process instructs the node being decommissioned to move its data elsewhere (which is essentially the opposite of bootstrapping).

**Remember:** This process places additional load on the cluster, so ideally needs to be performed at quiet times.

### Procedure

#### Scale out

1. Provision extra storage.

Before you scale Cassandra, you will need to provision extra storage, as also described in more detail in the [“Scaling horizontally: Persistence”](#) on page 324 topic. You must provision the additional storage on worker nodes other than the current ones being used, so that the Cassandra cluster remains resilient to node failures, as in the following example:

```
$ source pak_extensions/common/kubhelper.sh
$ createPersistentVolume 172.16.190.218 asm netcool data-asm-cassandra-3 50 /opt/ibm/netcool/asm/data/
cassandra
$ createPersistentVolume 172.16.191.196 asm netcool data-asm-cassandra-4 50 /opt/ibm/netcool/asm/data/
cassandra
$ ssh root@172.16.190.218 mkdir -p /opt/ibm/netcool/asm/data/cassandra
$ ssh root@172.16.191.196 mkdir -p /opt/ibm/netcool/asm/data/cassandra
```

2. Update deployment configuration.

Update the installation configuration with the desired number of Cassandra nodes in the cluster. For clarity, only the Cassandra cluster size is shown here.

```
global:
  cassandraNodeReplicas: 5
```

3. Perform a helm upgrade using the updated configuration:

```
helm upgrade asm ocp-local/ibm-netcool-asm-prod --values=asm-config.yaml --tls
```

4. Once the upgrade completes, check that the additional Cassandra pods are ready.

```
$ watch kubectl get pod -lapp=cassandra,release=asm
NAME          READY   STATUS    RESTARTS   AGE
asm-cassandra-0 1/1     Running   1          4d21h
asm-cassandra-1 1/1     Running   2          4d21h
asm-cassandra-2 1/1     Running   1          4d21h
asm-cassandra-3 1/1     Running   0          3m36s
asm-cassandra-4 1/1     Running   0          3m36s
```

5. Verify that the data is now distributed across all nodes.

Check the cluster 'nodetool' status according to each node:

```
for pod in `kubectl get pod -l app=cassandra,release=asm | grep -v NAME | awk '{print $1}'`; do
  echo "nodetool status for $pod"
  kubectl exec $pod /opt/ibm/cassandra/bin/nodetool status
done
```

Scaled out to to five nodes, the three copies should be distributed across five nodes with each having about 60% of the data:

```
nodetool status for asm-cassandra-4
Datacenter: datacenter1
=====
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
-- Address      Load          Tokens         Owns (effective)  Host ID                               Rack
UN  10.1.87.160   72.4 MiB      256            59.8%            d89e123f-3e98-401f-8442-d5521ad50daf rack1
UN  10.1.91.112   132.22 MiB    256            56.4%            dec6be10-4dcc-493e-8c26-330c19a32da2 rack1
UN  10.1.34.216   132.33 MiB    256            63.0%            e79c86a8-2105-4fad-b2ab-d10cdcd8354d rack1
UN  10.1.205.91   132.21 MiB    256            59.1%            77851686-5715-4237-86e4-31b62603ac2b rack1
UN  10.1.54.122   103.34 MiB    256            61.7%            0a99bbb8-3383-4cd0-b3c6-87d9f3909981 rack1
```

## Scale in

In the following example the use of five nodes is scaled back to three nodes.

6. Decommission the highest numbered node.

You go into the fifth container, check the current status of the cluster and then start the decommission process, which moves the data elsewhere.

```
$ kubectl exec -ti asm-cassandra-4 bash
[cassandra@asm-cassandra-4 /]$ /opt/ibm/cassandra/bin/nodetool status
Datacenter: datacenter1
=====
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
-- Address      Load          Tokens         Owns (effective)  Host ID                               Rack
UN  10.1.34.242   155.64 MiB    256            63.0%            e79c86a8-2105-4fad-b2ab-d10cdcd8354d rack1
UN  10.1.87.166   100.22 MiB    256            59.8%            d89e123f-3e98-401f-8442-d5521ad50daf rack1
UN  10.1.91.70    155.57 MiB    256            56.4%            dec6be10-4dcc-493e-8c26-330c19a32da2 rack1
UN  10.1.205.93   147.69 MiB    256            59.1%            77851686-5715-4237-86e4-31b62603ac2b rack1
UN  10.1.54.124   128.39 MiB    256            61.7%            0a99bbb8-3383-4cd0-b3c6-87d9f3909981 rack1

[cassandra@asm-cassandra-4 /]$ /opt/ibm/cassandra/bin/nodetool decommission
[cassandra@asm-cassandra-4 /]$
```

7. After the decommission process completes, check the cluster status.

This should show that five nodes have been reduced to four, each now having approximately 75% of the data.

```
$ /opt/ibm/cassandra/bin/nodetool status
Datacenter: datacenter1
=====
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
-- Address      Load          Tokens         Owns (effective)  Host ID                               Rack
UN  10.1.34.242   173.04 MiB    256            80.1%            e79c86a8-2105-4fad-b2ab-d10cdcd8354d rack1
```

UN	10.1.91.70	155.54 MiB	256	71.4%	dec6be10-4dcc-493e-8c26-330c19a32da2	rack1
UN	10.1.205.93	147.69 MiB	256	73.1%	77851686-5715-4237-86e4-31b62603ac2b	rack1
UN	10.1.54.124	147.99 MiB	256	75.4%	0a99bbb8-3383-4cd0-b3c6-87d9f3909981	rack1

## 8. Repeat the process for the fourth Cassandra node.

After completion, there should be three nodes, each with 100% of the data:

```
[cassandra@asm-cassandra-3 ~]$ /opt/ibm/cassandra/bin/nodetool status
Datacenter: datacenter1
=====
Status=Up/Down
-- State=Normal/Leaving/Joining/Moving
-- Address      Load      Tokens     Owns (effective)  Host ID                               Rack
UN 10.1.34.242   203.76 MiB 256        100.0%            e79c86a8-2105-4fad-b2ab-d10cdcd8354d rack1
UN 10.1.91.70    167.36 MiB 256        100.0%            dec6be10-4dcc-493e-8c26-330c19a32da2 rack1
UN 10.1.205.93   161.14 MiB 256        100.0%            77851686-5715-4237-86e4-31b62603ac2b rack1
```

## 9. Update the deployment configuration to remove the additional pods:

```
global:
  cassandraNodeReplicas: 3
```

## 10. Perform a helm upgrade using the updated configuration:

```
helm upgrade asm ocp-local/ibm-netcool-asm-prod --values=asm-config.yaml --tls
```

## 11. Check that the additional pods are stopped, and the cluster status of the remaining nodes.

## 12. Deprovision the additional storage.

- Clean up the persistent volumes, starting with the claims first:

```
$ kubectl delete pvc data-asm-cassandra-3 data-asm-cassandra-4
persistentvolumeclaim "data-asm-cassandra-3" deleted
persistentvolumeclaim "data-asm-cassandra-4" deleted
```

- Check that the persistent volumes are released:

```
ibmadmin@asm-prod-master:~$ kubectl get pv | grep cass
NAME             CAPACITY  ACCESS MODES  RECLAIM POLICY  STATUS  CLAIM              STORAGECLASS
REASON  AGE
172.16.154.233-data-asm-cassandra-0  50Gi      RWO           Retain          Bound   netcool/data-asm-7d2h
172.16.183.205-data-asm-cassandra-1  50Gi      RWO           Retain          Bound   netcool/data-asm-7d2h
172.16.188.122-data-asm-cassandra-2  50Gi      RWO           Retain          Bound   netcool/data-asm-7d2h
172.16.190.218-data-asm-cassandra-3  50Gi      RWO           Retain          Released netcool/data-asm-45h
172.16.191.196-data-asm-cassandra-4  50Gi      RWO           Retain          Released netcool/data-asm-45h
cassandra-4
```

- Delete the now redundant persistent volumes:

```
$ kubectl delete pv 172.16.190.218-data-asm-cassandra-3 172.16.191.196-data-asm-cassandra-4
persistentvolume "172.16.190.218-data-asm-cassandra-3" deleted
persistentvolume "172.16.191.196-data-asm-cassandra-4" deleted
```

- Clean the actual volumes on the worker nodes:

```
$ ssh root@172.16.190.218 rm -rf /opt/ibm/netcool/asm/data/cassandra
$ ssh root@172.16.191.196 rm -rf /opt/ibm/netcool/asm/data/cassandra
```



**Warning:** Always make absolutely sure that you are cleaning the correct nodes.

## Scaling horizontally: Elasticsearch search and analytics engine

Elasticsearch data is stored in an index split into a number of shards, which distribute data around a cluster. To achieve high availability, these shards are replicated and distributed across the cluster.

### Before you begin

The Agile Service Manager Search service by default creates a number of indices, which essentially are current resources and historical resources. The default number of shards is five, which you can customize using the **ELASTICSEARCH\_SHARDS** variable.

## About this task

The number of replica shards depends on how many Elasticsearch nodes there are, but there would never be more than two replicas. When there are three Elasticsearch nodes, an index has five shards and there are two replicas of those shards, meaning there are three copies of the data spread over three nodes. One of those copies is elected the Primary.

Replicas are used to provide redundant copies of your data to protect against hardware failure, and to serve read requests, like searching or retrieving a document.

### Scale out

To scale out Elasticsearch, you provision extra storage, update the deployment configuration, and then perform a helm upgrade.

You can verify that the scale out was successful by checking cluster status and health, as well as node and shard health.

### Scale in

To avoid data loss, you scale in Elasticsearch one node at a time.

Elasticsearch automatically redistributes shards, but you must wait for the scale in of each node to succeed before proceeding to scale in the next node.

## Procedure

### Scale out

1. Provision extra storage.

Before you scale Elasticsearch, you will need to provision extra storage, as also described in more detail in the [“Scaling horizontally: Persistence”](#) on page 324 topic. You must provision the additional storage on worker nodes other than the current ones being used, so that the Elasticsearch cluster remains resilient to node failures, as in the following example:

```
$ source pak_extensions/common/kubhelper.sh
$ createPersistentVolume 172.16.190.218 asm netcool data-asm-elasticsearch-3 75 /opt/ibm/netcool/asm/data/elasticsearch
$ createPersistentVolume 172.16.191.196 asm netcool data-asm-elasticsearch-4 75 /opt/ibm/netcool/asm/data/elasticsearch
$ ssh root@172.16.190.218 mkdir -p /opt/ibm/netcool/asm/data/elasticsearch
$ ssh root@172.16.191.196 mkdir -p /opt/ibm/netcool/asm/data/elasticsearch
```

2. Update deployment configuration.

Update the installation configuration with the desired number of Elasticsearch nodes in the cluster. For clarity, only the Elasticsearch cluster size is shown here.

```
global:
  elasticsearch:
    replicaCount: 5
```

3. Perform a helm upgrade using the updated configuration:

```
helm upgrade asm ocp-local/ibm-netcool-asm-prod --values=asm-config.yaml --tls
```

4. Once the upgrade completes, check that the additional Elasticsearch pods are ready.

**Note:** The new pods should become ready, and then the existing pods will be updated with the new configuration via a rolling-update, each pod updated in turn while waiting for each to become ready before updating the next. During this process some pods may restart waiting for a quorum of master nodes. The elected master will often change a few times, and this can sometimes happen just prior to a node being updated.

```
$ watch kubectl get pods -l release=asm,app=elasticsearch --namespace=netcool
NAME                                READY   STATUS    RESTARTS   AGE
asm-elasticsearch-0                 1/1     Running   0           4m4s
asm-elasticsearch-1                 1/1     Running   1          5m52s
asm-elasticsearch-2                 1/1     Running   1          7m35s
asm-elasticsearch-3                 1/1     Running   1          9m43s
asm-elasticsearch-4                 1/1     Running   1          9m43s
```

5. Monitor the state of the Search pods.

Search may restart with the new configuration.

```
$ watch kubectl get pod -lapp=search,release=asm
NAME                                READY    STATUS    RESTARTS   AGE
asm-search-5485cf6579-xsdsp        1/1     Running   0           10m
```

## 6. Verify that the shards are now distributed across all nodes.

Check the cluster status according to each node:

```
for pod in `kubectl get pod -l app=elasticsearch | grep -v NAME | awk '{print $1}'`; do
  echo -n "$pod cluster status = "
  kubectl exec $pod -- curl -s localhost:9200/_cluster/health | jq .status
done
```

The system output should indicate a cluster status of 'green'. Yellow would be functioning, although perhaps without the required number of replicas during a node outage, while a status of Red would indicate a problem:

```
asm-elasticsearch-0 cluster status = "green"
asm-elasticsearch-1 cluster status = "green"
asm-elasticsearch-2 cluster status = "green"
asm-elasticsearch-3 cluster status = "green"
asm-elasticsearch-4 cluster status = "green"
```

## 7. Check the cluster health.

View the unfiltered cluster health as in the following example:

```
for pod in `kubectl get pod -l app=elasticsearch | grep -v NAME | awk '{print $1}'`; do
  echo -n "$pod cluster status = "
  kubectl exec $pod -- curl -s localhost:9200/_cluster/health | jq
done
```

Example system output for one node providing full cluster health details:

```
asm-elasticsearch-0 cluster status = {
  "cluster_name": "elastic_production",
  "status": "green",
  "timed_out": false,
  "number_of_nodes": 5,
  "number_of_data_nodes": 5,
  "active_primary_shards": 20,
  "active_shards": 60,
  "relocating_shards": 0,
  "initializing_shards": 0,
  "unassigned_shards": 0,
  "delayed_unassigned_shards": 0,
  "number_of_pending_tasks": 0,
  "number_of_in_flight_fetch": 0,
  "task_max_waiting_in_queue_millis": 0,
  "active_shards_percent_as_number": 100
}
```

## 8. Check the node health.

View a summary of node health as in the following example:

```
for pod in `kubectl get pod -l app=elasticsearch | grep -v NAME | awk '{print $1}'`; do
  echo "$pod node health"
  kubectl exec $pod -- curl -s localhost:9200/_cat/nodes?v
  echo
done
```

Example system output for two nodes (all nodes should agree who the master node is):

```
asm-elasticsearch-0 node health
ip          heap.percent ram.percent cpu load_1m load_5m load_15m node.role master name
10.1.87.147 12          -756      3    0.38    0.36    0.42 mdi      -    10th8s1
10.1.54.116 8           -675      3    1.37    1.39    0.96 mdi      -    ax1XbEG
10.1.205.83 14          -716      3    0.27    0.37    0.61 mdi      *    1A5OuMp
10.1.91.100 14          -170      5    0.47    0.71    0.75 mdi      -    C4tNw5s
10.1.34.229 7           79       14    2.28    3.30    4.17 mdi      -    TlsN1kz

asm-elasticsearch-1 node health
ip          heap.percent ram.percent cpu load_1m load_5m load_15m node.role master name
10.1.205.83 14          -716      3    0.27    0.37    0.61 mdi      *    1A5OuMp
10.1.34.229 7           79       14    2.28    3.30    4.17 mdi      -    TlsN1kz
10.1.91.100 14          -170      5    0.47    0.71    0.75 mdi      -    C4tNw5s
10.1.87.147 12          -756      3    0.38    0.36    0.42 mdi      -    10th8s1
10.1.54.116 8           -675      3    1.37    1.39    0.96 mdi      -    ax1XbEG
```

## 9. Check the shard health.

View the shard status as in the following example:

```
for pod in `kubectl get pod -l app=elasticsearch | grep -v NAME | awk '{print $1}'`; do
  echo "$pod shard status"
  kubectl exec $pod -- curl -s localhost:9200/_cat/shards?v
  echo
done
```

Example system output for one node (all nodes should report the same status), and for one index (for clarity):

```
asm-elasticsearch-4 shard status
index          shard prirep state docs store ip          node
searchservice_v8 1      p     STARTED 102 2.7mb 10.1.54.116 ax1XbEG
searchservice_v8 1      r     STARTED 102 2.7mb 10.1.34.229 TlsN1kz
searchservice_v8 1      r     STARTED 102 2.9mb 10.1.87.147 10th8s1
searchservice_v8 4      r     STARTED 115 2.9mb 10.1.91.100 C4tNw5s
searchservice_v8 4      r     STARTED 115 2.9mb 10.1.205.83 1A50uMp
searchservice_v8 4      p     STARTED 115 2.9mb 10.1.87.147 10th8s1
searchservice_v8 3      r     STARTED 110 2.7mb 10.1.91.100 C4tNw5s
searchservice_v8 3      r     STARTED 110 2.7mb 10.1.54.116 ax1XbEG
searchservice_v8 3      p     STARTED 110 2.7mb 10.1.87.147 10th8s1
searchservice_v8 2      r     STARTED 102 2.8mb 10.1.91.100 C4tNw5s
searchservice_v8 2      p     STARTED 102 2.8mb 10.1.205.83 1A50uMp
searchservice_v8 2      r     STARTED 102 2.8mb 10.1.34.229 TlsN1kz
searchservice_v8 0      r     STARTED 97 2.2mb 10.1.54.116 ax1XbEG
searchservice_v8 0      p     STARTED 97 2.2mb 10.1.205.83 1A50uMp
searchservice_v8 0      r     STARTED 97 2.2mb 10.1.34.229 TlsN1kz
```

Note the following:

- There are five shards per index.
- Each shard has one primary and two replicas.
- Primary and replica shards are all started (not UNASSIGNED).
- Primary and replica shards are spread over different nodes.
- Primary and replica shards should contain the same number of docs.

## Scale in

In the following example the use of five nodes is scaled back to three nodes.

10. Update the deployment configuration to remove one Elasticsearch node:

In the following example, the node count (**replicaCount**) in the installation configuration is reduced from five nodes to four:

```
global:
  elasticsearch:
    replicaCount: 4
```

11. Perform a helm upgrade using the updated configuration:

```
helm upgrade asm ocp-local/ibm-netcool-asm-prod --values=asm-config.yaml --tls
```

12. Once the upgrade completes, check the status of the ElasticSearch pods.

- The number of pods will reduce, and a rolling update will happen to apply the new expected number of nodes.

```
$ watch kubectl get pods -l release=asm,app=elasticsearch --namespace=netcool
NAME          READY STATUS RESTARTS AGE
asm-elasticsearch-0 1/1 Running 0 2m6s
asm-elasticsearch-1 1/1 Running 0 4m9s
asm-elasticsearch-2 1/1 Running 1 5m24s
asm-elasticsearch-3 1/1 Running 0 6m49s
```

- During this process some the cluster health will often report as **yellow**, meaning operational, but not with the desired number of replicas. After the scale in, cluster health should return to green:

```
asm-elasticsearch-0 cluster status = "green"
asm-elasticsearch-1 cluster status = "green"
asm-elasticsearch-2 cluster status = "green"
asm-elasticsearch-3 cluster status = "green"
```

13. To further reduce the number of Elasticsearch nodes from four to three, repeat steps 10 to 12.
14. Verify cluster status, node health and shard health by following steps 6 to 9.
15. Deprovision the additional storage (or persistent volumes).

See the related “[Scaling horizontally: Cassandra database](#)” on page 325 topic for an example of storage deprovisioning.



**Warning:** Always make absolutely sure that you are cleaning the correct nodes.

### Related information

[Elasticsearch horizontal scaling](#)

[Elasticsearch cluster health](#)

[Elasticsearch failover](#)

## Scaling horizontally: Kafka message bus

All messages on Kafka are organized into topics. Agile Service Manager has several different topics for different purposes. These topics are further divided into partitions, and these partitions are spread over the available brokers (instance of Kafka in the cluster). Partitions have replicas, but all reads and writes happen on the leader for a partition.

### About this task

You can **list the topics in Kafka** as in the following example:

```
$ ./bin/kafka-topics.sh --zookeeper $ZOOKEEPER_URL --list
__consumer_offsets
itsm.monitor.json
itsm.nodes.json
itsm.rebroadcast.json
itsm.resources.json
kafka.topic.notification.json
providers.json
resources.json
status.json
```

You can **describe the topics in Kafka** as in the following example, which shows the number of partitions, how those partitions are spread over the brokers, who is the leader, and where the partition replicas are:

```
$ ./bin/kafka-topics.sh --zookeeper $ZOOKEEPER_URL --describe --topic resources.json
Topic:resources.json    PartitionCount:24      ReplicationFactor:3    Configs:
Topic: resources.json   Partition: 0           Leader: 1               Replicas: 1,0,2        Isr: 1,0,2
Topic: resources.json   Partition: 1           Leader: 2               Replicas: 2,1,0        Isr: 2,1,0
Topic: resources.json   Partition: 2           Leader: 0               Replicas: 0,2,1        Isr: 0,2,1
Topic: resources.json   Partition: 3           Leader: 1               Replicas: 1,2,0        Isr: 1,2,0
Topic: resources.json   Partition: 4           Leader: 2               Replicas: 2,0,1        Isr: 2,0,1
Topic: resources.json   Partition: 5           Leader: 0               Replicas: 0,1,2        Isr: 0,1,2
Topic: resources.json   Partition: 6           Leader: 1               Replicas: 1,0,2        Isr: 1,0,2
Topic: resources.json   Partition: 7           Leader: 2               Replicas: 2,1,0        Isr: 2,1,0
Topic: resources.json   Partition: 8           Leader: 0               Replicas: 0,2,1        Isr: 0,2,1
Topic: resources.json   Partition: 9           Leader: 1               Replicas: 1,2,0        Isr: 1,2,0
Topic: resources.json   Partition: 10          Leader: 2               Replicas: 2,0,1        Isr: 2,0,1
Topic: resources.json   Partition: 11          Leader: 0               Replicas: 0,1,2        Isr: 0,1,2
Topic: resources.json   Partition: 12          Leader: 1               Replicas: 1,0,2        Isr: 1,0,2
Topic: resources.json   Partition: 13          Leader: 2               Replicas: 2,1,0        Isr: 2,1,0
Topic: resources.json   Partition: 14          Leader: 0               Replicas: 0,2,1        Isr: 0,2,1
Topic: resources.json   Partition: 15          Leader: 1               Replicas: 1,2,0        Isr: 1,2,0
Topic: resources.json   Partition: 16          Leader: 2               Replicas: 2,0,1        Isr: 2,0,1
Topic: resources.json   Partition: 17          Leader: 0               Replicas: 0,1,2        Isr: 0,1,2
Topic: resources.json   Partition: 18          Leader: 1               Replicas: 1,0,2        Isr: 1,0,2
Topic: resources.json   Partition: 19          Leader: 2               Replicas: 2,1,0        Isr: 2,1,0
Topic: resources.json   Partition: 20          Leader: 0               Replicas: 0,2,1        Isr: 0,2,1
Topic: resources.json   Partition: 21          Leader: 1               Replicas: 1,2,0        Isr: 1,2,0
Topic: resources.json   Partition: 22          Leader: 2               Replicas: 2,0,1        Isr: 2,0,1
Topic: resources.json   Partition: 23          Leader: 0               Replicas: 0,1,2        Isr: 0,1,2
```

### Scale out

To scale out Kafka, you provision extra storage, update the deployment configuration, and then perform a helm upgrade.

To redistribute the topic partitions over all available brokers, you then reassign topic partitions.

### Scale in

To scale in Kafka, you reassign topic partitions, update the deployment configuration, perform a helm upgrade, and deprovision storage.

## Procedure

### Scale out

#### 1. Provision extra storage.

Before you scale Kafka, you will need to provision extra storage, as also described in more detail in the [“Scaling horizontally: Persistence”](#) on page 324 topic. You must provision the additional storage on worker nodes other than the current ones being used, so that the Kafka cluster remains resilient to node failures, as in the following example:

```
$$ source pak_extensions/common/kubhelper.sh
$ createPersistentVolume 172.16.190.218 asm netcool data-asm-kafka-3 15 /opt/ibm/netcool/asm/data/kafka
$ createPersistentVolume 172.16.191.196 asm netcool data-asm-kafka-4 15 /opt/ibm/netcool/asm/data/kafka
$ createPersistentVolume 172.16.192.211 asm netcool data-asm-kafka-5 15 /opt/ibm/netcool/asm/data/kafka
$ ssh root@172.16.190.218 mkdir -p /opt/ibm/netcool/asm/data/kafka
$ ssh root@172.16.191.196 mkdir -p /opt/ibm/netcool/asm/data/kafka
$ ssh root@172.16.192.211 mkdir -p /opt/ibm/netcool/asm/data/kafka
```

#### 2. Update deployment configuration.

Update the installation configuration with the desired number of Kafka nodes in the cluster. For clarity, only the Kafka cluster size is shown here.

```
global:
  kafka:
    clusterSize: 6
```

#### 3. Perform a helm upgrade using the updated configuration:

```
helm upgrade asm ocp-local/ibm-netcool-asm-prod --values=asm-config.yaml --tls
```

#### 4. Once the upgrade completes, check that the additional Kafka pods are ready.

```
$ kubectl get pod -lapp=kafka,release=asm
NAME          READY   STATUS    RESTARTS   AGE
asm-kafka-0   2/2     Running   0           14h
asm-kafka-1   2/2     Running   0           14h
asm-kafka-2   2/2     Running   0           14h
asm-kafka-3   2/2     Running   0           12m
asm-kafka-4   2/2     Running   0           12m
asm-kafka-5   2/2     Running   0           12m
```

### Reassign topic partitions

**Note:** In this example, six Kafka brokers are running. If you describe a Kafka topic as described above, it will show that the topics are only distributed over the original three brokers. The topic partitions must be distributed over all available brokers.

#### 5. Define topic to move or reassign.

Run the following commands inside one of the existing Kafka pods. The following topics are the Agile Service Manager topics:

```
cat <<EOF | tee /tmp/topics-to-move.json
{
  "topics": [
    {
      "topic": "itsm.monitor.json"
    },
    {
      "topic": "itsm.nodes.json"
    },
    {
      "topic": "itsm.rebroadcast.json"
    },
    {
      "topic": "itsm.resources.json"
    },
    {
      "topic": "providers.json"
    },
    {
      "topic": "resources.json"
    },
    {
      "topic": "status.json"
    },
    {
      "topic": "__consumer_offsets"
    }
  ]
}
```



```

    {
      "topic": "kafka.topic.notification.json"
    }
  ],
  "version": 1
}
EOF

```

## 6. Generate a topic reassignment plan.

Based on the topic list from the previous step, run the following command to generate a plan:

```
kafka-reassign-partitions.sh --generate
```

The list of available brokers, six in this example, is displayed by `--broker-list` and can be checked by running the following command:

```

$ /opt/kafka/bin/zookeeper-shell.sh $ZOOKEEPER_URL <<< "ls /brokers/ids"
Connecting to asm-zookeeper:2181 ...
[0, 1, 2, 3, 4, 5]

```

To save the proposed partition reassignment configuration to file:

```

/opt/kafka/bin/kafka-reassign-partitions.sh --generate --zookeeper asm-zookeeper:2181 --topics-to-move-
json-file
/tmp/topics-to-move.json --broker-list 0,1,2,3,4,5 | grep version | tail -1 > /tmp/reassignment-plan.json

```

## 7. Execute the reassignment plan:

```

/opt/kafka/bin/kafka-reassign-partitions.sh --execute --zookeeper $ZOOKEEPER_URL --reassignment-json-file
/tmp/reassignment-plan.json

```

## 8. Verify the topic reassignment.

Kafka will redistribute the leaders and replicas amongst all the specified brokers, which may take some time:

```

$ /opt/kafka/bin/kafka-topics.sh --zookeeper $ZOOKEEPER_URL --describe --topic resources.json
Topic:resources.json PartitionCount:24 ReplicationFactor:3 Configs:
Topic: resources.json Partition: 0 Leader: 1 Replicas: 1,3,4 Isr: 1,3,4
Topic: resources.json Partition: 1 Leader: 2 Replicas: 2,4,5 Isr: 5,2,4
Topic: resources.json Partition: 2 Leader: 3 Replicas: 3,5,0 Isr: 0,5,3
Topic: resources.json Partition: 3 Leader: 4 Replicas: 4,0,1 Isr: 1,0,4
Topic: resources.json Partition: 4 Leader: 5 Replicas: 5,1,2 Isr: 2,1,5
Topic: resources.json Partition: 5 Leader: 0 Replicas: 0,2,3 Isr: 0,2,3
Topic: resources.json Partition: 6 Leader: 1 Replicas: 1,4,5 Isr: 5,1,4
Topic: resources.json Partition: 7 Leader: 2 Replicas: 2,5,0 Isr: 2,0,5
Topic: resources.json Partition: 8 Leader: 3 Replicas: 3,0,1 Isr: 0,1,3
Topic: resources.json Partition: 9 Leader: 4 Replicas: 4,1,2 Isr: 1,2,4
Topic: resources.json Partition: 10 Leader: 5 Replicas: 5,2,3 Isr: 5,2,3
Topic: resources.json Partition: 11 Leader: 0 Replicas: 0,3,4 Isr: 0,3,4
Topic: resources.json Partition: 12 Leader: 1 Replicas: 1,5,0 Isr: 1,5,0
Topic: resources.json Partition: 13 Leader: 2 Replicas: 2,0,1 Isr: 2,1,0
Topic: resources.json Partition: 14 Leader: 3 Replicas: 3,1,2 Isr: 2,1,3
Topic: resources.json Partition: 15 Leader: 4 Replicas: 4,2,3 Isr: 2,3,4
Topic: resources.json Partition: 16 Leader: 5 Replicas: 5,3,4 Isr: 5,3,4
Topic: resources.json Partition: 17 Leader: 0 Replicas: 0,4,5 Isr: 0,5,4
Topic: resources.json Partition: 18 Leader: 1 Replicas: 1,0,2 Isr: 1,2,0
Topic: resources.json Partition: 19 Leader: 2 Replicas: 2,1,3 Isr: 2,1,3
Topic: resources.json Partition: 20 Leader: 3 Replicas: 3,2,4 Isr: 2,3,4
Topic: resources.json Partition: 21 Leader: 4 Replicas: 4,3,5 Isr: 5,3,4
Topic: resources.json Partition: 22 Leader: 5 Replicas: 5,4,0 Isr: 0,5,4
Topic: resources.json Partition: 23 Leader: 0 Replicas: 0,5,1 Isr: 0,1,5

```

## Scale in

Scaling in is essentially the same operation, where the reassignment plan it to move topics to the brokers that will remain after the scale in. In this example, six brokers are scaled in to three brokers.

## 9. Reassign topics:

Reassign the same topics as for the scale out example as described [\(from\) here](#). Adjust the reassignment plan to use the brokers that remain after the scale in, in this case:

```
--broker-list 0,1,2
```

Verify the topics are reassigned to the first three brokers, as such:

```

Topic:resources.json PartitionCount:24 ReplicationFactor:3 Configs:
Topic: resources.json Partition: 0 Leader: 1 Replicas: 1,2,0 Isr: 0,1,2

```

Topic: resources.json	Partition: 1	Leader: 2	Replicas: 2,0,1	Isr: 0,1,2
Topic: resources.json	Partition: 2	Leader: 0	Replicas: 0,1,2	Isr: 0,1,2
Topic: resources.json	Partition: 3	Leader: 1	Replicas: 1,0,2	Isr: 1,0,2
Topic: resources.json	Partition: 4	Leader: 2	Replicas: 2,1,0	Isr: 2,1,0
Topic: resources.json	Partition: 5	Leader: 0	Replicas: 0,2,1	Isr: 2,0,1
Topic: resources.json	Partition: 6	Leader: 1	Replicas: 1,2,0	Isr: 0,1,2
Topic: resources.json	Partition: 7	Leader: 2	Replicas: 2,0,1	Isr: 2,0,1
Topic: resources.json	Partition: 8	Leader: 0	Replicas: 0,1,2	Isr: 1,0,2
Topic: resources.json	Partition: 9	Leader: 1	Replicas: 1,0,2	Isr: 2,1,0
Topic: resources.json	Partition: 10	Leader: 2	Replicas: 2,1,0	Isr: 0,1,2
Topic: resources.json	Partition: 11	Leader: 0	Replicas: 0,2,1	Isr: 0,1,2
Topic: resources.json	Partition: 12	Leader: 1	Replicas: 1,2,0	Isr: 1,0,2
Topic: resources.json	Partition: 13	Leader: 2	Replicas: 2,0,1	Isr: 2,1,0
Topic: resources.json	Partition: 14	Leader: 0	Replicas: 0,1,2	Isr: 2,1,0
Topic: resources.json	Partition: 15	Leader: 1	Replicas: 1,0,2	Isr: 0,1,2
Topic: resources.json	Partition: 16	Leader: 2	Replicas: 2,1,0	Isr: 0,1,2
Topic: resources.json	Partition: 17	Leader: 0	Replicas: 0,2,1	Isr: 0,1,2
Topic: resources.json	Partition: 18	Leader: 1	Replicas: 1,2,0	Isr: 2,1,0
Topic: resources.json	Partition: 19	Leader: 2	Replicas: 2,0,1	Isr: 2,1,0
Topic: resources.json	Partition: 20	Leader: 0	Replicas: 0,1,2	Isr: 0,1,2
Topic: resources.json	Partition: 21	Leader: 1	Replicas: 1,0,2	Isr: 0,1,2
Topic: resources.json	Partition: 22	Leader: 2	Replicas: 2,1,0	Isr: 0,1,2
Topic: resources.json	Partition: 23	Leader: 0	Replicas: 0,2,1	Isr: 1,0,2

#### 10. Update deployment configuration.

Update the installation configuration with the desired number of Kafka nodes in the cluster. For clarity, only the Kafka cluster size is shown here.

```
global:
  kafka:
    clusterSize: 3
```

#### 11. Perform a helm upgrade using the updated configuration:

```
helm upgrade asm ocp-local/ibm-netcool-asm-prod --values=asm-config.yaml --tls
```

#### 12. Deprovision the additional storage (or persistent volumes).

See the related [“Scaling horizontally: Cassandra database” on page 325](#) topic for an example of storage deprovisioning.



**Warning:** Always make absolutely sure that you are cleaning the correct nodes.

## Scaling horizontally: Zookeeper synchronization service

Zookeeper synchronization service brokers can be scaled out and scaled in. As Zookeeper requires a majority, it is necessary to use an odd number of pods in the ensemble.

### About this task



**Warning:** Agile Service Manager Versions 1.1.5 and later do **not** support scaling for the Zookeeper component.

**Note:** To avoid an inconsistent lists of servers, you **must** perform scaling via `helm upgrade` rather than the Kubernetes `kubectl scale` command.

#### Scale out

To scale out Zookeeper, you provision extra storage, update the deployment configuration, and then perform a helm upgrade.

#### Scale in

To scale in Zookeeper, you update the deployment configuration, perform a helm upgrade, and then deprovision storage.

## Procedure

### Scale out

#### 1. Provision extra storage.

Before you scale out Zookeeper, you must provision extra storage (also described in more detail in the [“Scaling horizontally: Persistence” on page 324](#) topic). You must provision the additional storage

on worker nodes other than the current ones being used, so that the Zookeeper ensemble remains resilient to node failures, as in the following example:

```
$ source pak_extensions/common/kubhelper.sh
$ createPersistentVolume 172.16.190.218 asm netcool data-asm-zookeeper-3 15 /opt/ibm/netcool/asm/data/
zookeeper
$ createPersistentVolume 172.16.191.196 asm netcool data-asm-zookeeper-4 15 /opt/ibm/netcool/asm/data/
zookeeper
$ ssh root@172.16.190.218 mkdir -p /opt/ibm/netcool/asm/data/zookeeper
$ ssh root@172.16.191.196 mkdir -p /opt/ibm/netcool/asm/data/zookeeper
```

## 2. Update deployment configuration.

Update the installation configuration with the desired number of Zookeeper nodes in the cluster. For clarity, only the Zookeeper cluster size is shown here.

```
global:
  zookeeper:
    clusterSize: 5
```

## 3. Perform a helm upgrade using the updated configuration:

```
helm upgrade asm ocp-local/ibm-netcool-asm-prod --values=asm-config.yaml --tls
```

## 4. Once the upgrade completes, check that the additional Zookeeper pods are ready.

```
$ watch kubectl get pods -l release=asm,app=zookeeper --namespace=netcool
NAME          READY   STATUS    RESTARTS   AGE
asm-zookeeper-0 1/1     Running   1          17h
asm-zookeeper-1 1/1     Running   0          17h
asm-zookeeper-2 1/1     Running   1          17h
asm-zookeeper-3 1/1     Running   0          3m31s
asm-zookeeper-4 1/1     Running   0          3m31s
```

## 5. Also monitor the state of the Kafka pods.

```
$ watch kubectl get pod -lapp=kafka,release=asm
NAME          READY   STATUS    RESTARTS   AGE
asm-kafka-0   2/2     Running   0          30m
asm-kafka-1   2/2     Running   0          32m
asm-kafka-2   2/2     Running   0          33m
asm-kafka-3   2/2     Running   0          35m
asm-kafka-4   2/2     Running   0          36m
asm-kafka-5   2/2     Running   0          38m
```

## 6. Check that one of the Zookeeper nodes has been elected as leader:

```
for pod in `kubectl get pod -l app=zookeeper,release=asm | grep -v NAME | awk '{print $1}'`; do
  echo -n "$pod ";
  echo "stat" | kubectl exec -i $pod -- socat - tcp:localhost:2181 | grep Mode;
done
```

### Example output

```
asm-zookeeper-0 Mode: follower
asm-zookeeper-1 Mode: leader
asm-zookeeper-2 Mode: follower
asm-zookeeper-3 Mode: follower
asm-zookeeper-4 Mode: follower
```

## 7. To see the full Zookeeper node statistics, run the following command:

```
for pod in `kubectl get pod -l app=zookeeper,release=asm | grep -v NAME | awk '{print $1}'`; do
  echo "$pod stat:";
  echo "stat" | kubectl exec -i $pod -- socat - tcp:localhost:2181;
  echo;
done
```

## Scale in

Scaling in is almost the same operation as scaling out, but in reverse.

## 8. Update deployment configuration.

Update the installation config with the desired number of Zookeeper nodes in the ensemble. For clarity, only the Zookeeper ensemble size is shown here.

```
global:
  zookeeper:
    clusterSize: 3
```

9. Perform a helm upgrade using the updated configuration:

```
helm upgrade asm ocp-local/ibm-netcool-asm-prod --values=asm-config.yaml --tls
```

10. Deprovision the additional storage (or persistent volumes).

For an example, see the equivalent [storage deprovisioning step](#) in the related Cassandra database scaling topic.



**Warning:** Always make absolutely sure that you are cleaning the correct nodes.

#### Related information

[Zookeeper clustered \(multi-server\) setup](#)

[Zookeeper common problems](#)

## System health and logging

Agile Service Manager containers or pods have built-in health monitoring, which you can use to check if a service is still available, or to investigate your Agile Service Manager ecosystem. In addition, you can use configurable logging functionality to monitor system health and assist in troubleshooting.

### Performing a system health check from the UI

You access your Agile Service Manager deployment's system health information from the topology viewer. Your system's health is rendered as a topology, and you can use the topology viewer functionality to gain insights into your system's health.

#### Additional actions > View System Health

Open the **Additional actions** drop-down menu, and then use the **View System Health** option to access your Netcool Agile Service Manager deployment's system health information.

#### On-prem

Health for an on-prem system is reported by the Docker Observer.

The topology displayed is centered on the 'asm\_default' seed, and you can view the status for each docker container based on the most recent observer job run.

#### OCP

Health for an OCP system is reported by the Kubernetes Observer.

Separate topologies are displayed for each worker node.

#### Related concepts

[“Using Agile Service Manager” on page 245](#)

You use the Agile Service Manager UI to visualize your topology data. First you define a seed resource on which to base your view, then choose the levels of networked resources around the seed that you wish to display, before rendering the view. You can then further expand or analyze the displayed topology in real time, or compare it to previous versions within a historical time window.

#### Related tasks

[“Configuring Docker Observer jobs” on page 129](#)

Using the Docker Observer, you can discover Docker network resources, including Docker Swarm clusters, and then visualize (or model) this data as a topology view in the Agile Service Manager UI. You configure observer jobs from the **Observer Configuration UI**.

[“Configuring Kubernetes Observer jobs” on page 144](#)

Using this observer, you can configure jobs that discover the structure of your Kubernetes clusters, including pods, worker nodes and containers.

## Viewing the service logs (on-prem)

Logs for all Netcool Agile Service Manager services can be found in the `$ASM_HOME/logs/<service>` directories. You can set logging levels for the user-facing services, such as the observers and search, using scripts provided.

### About this task

You can change your current log levels to the following values (INFO is the default log level):

- OFF
- ERROR
- WARN
- INFO
- DEBUG
- TRACE
- ALL

### Default log housekeeping

When a log reaches its maximum size, it is compressed, date-stamped and versioned, and moved to a subdirectory, for example:

```
$ASM_HOME/logs/topology/tmp/topology-service-2017-03-23-0.log.gz
```

and then

```
$ASM_HOME/logs/topology/tmp/topology-service-2017-03-23-1.log.gz
```

### Tip:

#### Event integration on-prem

To avoid excessive disk usage with event integration enabled on an on-prem system, change the **LOG\_LEVEL** for the topology and status services to WARN.

### Network Discover Service logs:

#### Network Discovery logs

All log files start with 'nasm-net-disco'.

Discovery worker logs have one previous version stored and are rotated when the size of the file exceeds 10Mb.

Other Network Discovery log files are stored for five days and rotated when the size exceeds 10Mb, or when the total size of all log files exceeds 20Gb.

Table 66. Log names and directories for Netcool Agile Service Manager services		
Service	Directory	Log
Application Discovery service	<code>\$ASM_HOME/logs/app-disco</code>	pss (primary storage server) ds (discovery server) sss (secondary storage server)
Cassandra database	<code>\$ASM_HOME/logs/cassandra</code>	system.log
Elasticsearch	<code>\$ASM_HOME/logs/ elasticsearch</code>	elasticsearch.log

Table 66. Log names and directories for Netcool Agile Service Manager services (continued)

Service	Directory	Log
Gateway service	\$ASM_HOME/logs	gateway.log
Kafka message bus	\$ASM_HOME/logs/kafka	server.log
Network Discovery service	\$ASM_HOME/logs/net-disco	Log files start with 'nasm-net-disco'
Observer service	\$ASM_HOME/logs/observer	
Probe service	\$ASM_HOME/logs/probe	asm_probe.log
Search service	\$ASM_HOME/logs/search	search-service.log
Status service	\$ASM_HOME/logs/status	status-service.log
Topology service	\$ASM_HOME/logs/topology	topology-service.log
UI API service (see the <a href="#">UI API service logging example</a> )	\$ASM_HOME/logs/ui-api	ui-api.log
Zookeeper synchronization service	\$ASM_HOME/logs/zookeeper	zookeeper.log

**UI API service logging example:** The UI API logging script follows the same rules as all other logging scripts.

The Agile Service Manager UI uses the UI API to retrieve data from the other Agile Service Manager services. Data retrieval errors are recorded in the UI API log.

Use the following shell script to change your logging level:

```
/opt/ibm/netcool/asm/bin/ui_api_log_level.sh
```

At a log level of INFO or above, which is the default log level, the log only contains information about the service startup, as well as any high level warnings and errors. To diagnose a problem, a lower log level may be useful, such as DEBUG or TRACE, which logs all incoming requests as well as an indication of success; whereas at a setting of ALL, additional details are logged about how each REST request is handled.

## Viewing the logs (OCP)

You can see the logs for all Agile Service Manager OCP pods using the `kubectl logs` command. You can change the logging level to 'debug' to troubleshoot in more detail.

### About this task

All Netcool Agile Service Manager OCP containers are deployed in pods. You can identify logs either by their names or their labels.

#### Network Discovery service logs

In an OCP installation, the logs for the Network Discovery service are held in the container, and tailed by the nasm-net-disco-worker-logs sidecar container, which outputs to stdout.

#### Application Discovery service logs

For each component (primary storage server, discovery server and secondary storage server), a sidecar container is designated as log-collector.

To increase or decrease the log levels for a specific observer to allow you to investigate issues, you first use the `kubectl get deployments` command to display a list of all OCP deployments, then use the `kubectl edit deployment` command to change the logging levels to, for example, 'debug'.

## Procedure

### Viewing logs

1. To list all Agile Service Manager pods and the labels applied to the pods, run the `kubectl get pod --show-labels` command, as in the example below.
2. Display the logs using either the pod names or labels.

When running the `kubectl logs` command, you must also specify the namespace.

### View logs for a specific pod

```
kubectl logs asm-layout-88bd88bdb-htjfp --namespace=netcool
```

### View logs using the label

```
kubectl logs -l app=layout --namespace=netcool
```

### Example

When using the `kubectl get pod --show-labels` command to list all pods, the system will display the pod names, their status, number of restarts, age, and labels. The labels retrieved will also contain additional information, such as, for example, the release name, which is important if more than one Agile Service Manager deployment exists.

```
$ kubectl get pod --show-labels
NAME                                READY   STATUS    RESTARTS   AGE
LABELS
deploy-test-cassandra-0             1/1     Running   0           2d
app=cassandra,chart=cassandra,controller-revision-hash=deploy-test-cassandra-7d8f56b884,
heritage=Tiller,release=deploy-test,statefulset.kubernetes.io/pod-name=deploy-test-cassandra-0
deploy-test-ciscoaci-observer-8755fcc94-7cp7h 1/1     Running   0           2d
app=ciscoaci-observer,chart=ciscoaci-observer,heritage=Tiller,
pod-template-hash=431197750,release=deploy-test
deploy-test-contrail-observer-54f49cdc8c-bjdnq 1/1     Running   0           2d
app=contrail-observer,chart=contrail-observer,heritage=Tiller,
pod-template-hash=1090578747,release=deploy-test
deploy-test-dns-observer-7ffc598847-kdqlr     1/1     Running   0           2d
app=dns-observer,chart=dns-observer,heritage=Tiller,pod-template-hash=3997154403,
release=deploy-test
deploy-test-elasticsearch-0           1/1     Running   0           2d
app=elasticsearch,chart=elasticsearch,controller-revision-hash=deploy-test-elasticsearch-5bb4857dff,
heritage=Tiller,release=deploy-test,statefulset.kubernetes.io/pod-name=deploy-test-elasticsearch-0
deploy-test-file-observer-76c5869d8b-tm7xj     1/1     Running   0           2d
app=file-observer,chart=file-observer,heritage=Tiller,pod-template-hash=3271425846,
release=deploy-test
deploy-test-ibmcloud-observer-65497f5478-4nsr6 1/1     Running   0           2d
app=ibmcloud-observer,chart=ibmcloud-observer,heritage=Tiller,pod-template-hash=2105391034,
release=deploy-test
deploy-test-itnm-observer-85d7cc7878-2xdhf     1/1     Running   0           2d
app=itnm-observer,chart=itnm-observer,heritage=Tiller,pod-template-hash=4183773434,
release=deploy-test
deploy-test-kafka-0                  2/2     Running   0           2d
app=kafka,chart=kafka,controller-revision-hash=deploy-test-kafka-5c78c96dbc,
heritage=Tiller,release=deploy-test,statefulset.kubernetes.io/pod-name=deploy-test-kafka-0
deploy-test-kubernetes-observer-66cb697d7-ktgqt 1/1     Running   0           2d
app=kubernetes-observer,chart=kubernetes-observer,heritage=Tiller,
pod-template-hash=227625383,release=deploy-test
deploy-test-layout-66656685bd-wkp6b           1/1     Running   0           2d
app=layout,chart=layout,heritage=Tiller,pod-template-hash=2221224168,
release=deploy-test
deploy-test-merge-688468b7-q4ckq             1/1     Running   0           2d
app=merge,chart=merge,heritage=Tiller,pod-template-hash=24402463,
release=deploy-test
deploy-test-newrelic-observer-7c879f5545-72k4p 1/1     Running   0           2d
app=newrelic-observer,chart=newrelic-observer,heritage=Tiller,
pod-template-hash=3743591101,release=deploy-test
deploy-test-openstack-observer-5df8fffd56-kfdjc 1/1     Running   0           2d
app=openstack-observer,chart=openstack-observer,heritage=Tiller,
pod-template-hash=1894999812,release=deploy-test
deploy-test-rest-observer-789cd8699d-zb8vx     1/1     Running   0           2d
app=rest-observer,chart=rest-observer,heritage=Tiller,
pod-template-hash=3457842558,release=deploy-test
deploy-test-search-5dc4ccc99b-nmqz2           1/1     Running   0           2d
app=search,chart=search,heritage=Tiller,pod-template-hash=1870777556,
release=deploy-test
deploy-test-servicenow-observer-854bbff7dc-sl75x 1/1     Running   0           2d
app=servicenow-observer,chart=servicenow-observer,heritage=Tiller,
pod-template-hash=4106699387,release=deploy-test
deploy-test-taddm-observer-79dd5b556-bnzp2     1/1     Running   0           2d
app=taddm-observer,chart=taddm-observer,heritage=Tiller,
pod-template-hash=358816112,release=deploy-test
deploy-test-topology-75688cfc48-jjkfd         1/1     Running   0           2d
app=topology,chart=topology,heritage=Tiller,pod-template-hash=3124479704,
```

```

release=deploy-test
deploy-test-vmvcenter-observer-675bd88f5c-8kgnk      1/1      Running    0          2d
app=vmvcenter-observer,chart=vmvcenter-observer,heritage=Tiller,
pod-template-hash=2316844917,release=deploy-test
deploy-test-vmwarensx-observer-dc96946f4-7jcxj      1/1      Running    0          2d
app=vmwarensx-observer,chart=vmwarensx-observer,heritage=Tiller,
pod-template-hash=875250290,release=deploy-test
deploy-test-zookeeper-0                             1/1      Running    0          2d
app=zookeeper,chart=zookeeper,controller-revision-hash=deploy-test-zookeeper-d65f46875,
heritage=Tiller,release=deploy-test,statefulset.kubernetes.io/pod-name=deploy-test-zookeeper-0

```

### Setting log levels for observers to 'debug'

3. Run the following command to change the pod's debug level:

```
kubectl exec -it deploy-test-topology-75688cfc48-jjkfd -- /opt/ibm/netcool/asm/bin/log-level.py --level DEBUG
```

Example system output:

```

Configured logging level for com.ibm to DEBUG
Response code: 200

```

4. Verify that the observer log is in DEBUG mode.  
DEBUG will log all request/response operations to provide detailed diagnostic data.



---

## Chapter 8. Troubleshooting

Use the following topics to troubleshoot specific issues.

### Installation troubleshooting (on-prem)

---

See the following information to troubleshoot on-prem installation issues.

#### License issues

During a first installation of Netcool Agile Service manager, or when the license terms change, you will be prompted to accept the software license. If you do not complete this step, the software will not start and this error will occur:

```
ERROR: Couldn't find env file: /opt/ibm/netcool/asm/licenses/.accept_license
```

#### Workaround

To review and accept the license terms, use the following command:

```
/opt/ibm/netcool/asm/bin/license-review.sh
```

#### ASM\_HOME variable warnings and errors

When running any docker-compose commands, such as starting or stopping Agile Service Manager, the docker-compose service needs to load the `.env` and `docker-compose.yml` files located in `ASM_HOME`. If you do not run the docker-compose command from the `ASM_HOME` directory, warnings and errors like the following may occur:

```
WARNING: The ASM_HOME variable is not set. Defaulting to a blank string.
```

```
ERROR: .IOError: [Errno 2] No such file or directory: u'/etc/nasm-docker-observer.yml'
```

```
ERROR: Can't find a suitable configuration file in this directory or any parent. Are you in the right directory? Supported filenames: docker-compose.yml, docker-compose.yaml
```

#### Workaround

You must run docker-compose commands from the `ASM_HOME` directory.

To ensure that your present working directory is `ASM_HOME`, you can take the following steps:

##### Change your current directory

To change your current directory to the `ASM_HOME` directory, use the **cd** command, as in the following example:

```
$ cd /opt/ibm/netcool/asm
```

##### Check your current directory

To verify that your current directory is the `ASM_HOME` directory, use the **pwd** command, as in the following example:

```
$ pwd
/opt/ibm/netcool/asm
```

#### Apply Elasticsearch patch during upgrade

During an upgrade, once the statefulsets are upgraded, the Elasticsearch pods may crash.

**Workaround:** To prevent this, apply the following patch:

```
kubectl patch statefulset ${RELEASE_NAME}-elasticsearch -p '{"spec":{"template":{"spec":{"securityContext":{"runAsUser":1000}}}}}'
```

where `${RELEASE_NAME}` = the Agile Service Manager instance name being deployed.

## Probe and gateway troubleshooting

See the following information to troubleshoot probe and gateway issues.

### Probe and gateway continually restart

The probe and gateway may continually restart if the Netcool/OMNIbus ObjectServer is not accessible using the current configuration.

**Workaround:** Perform one or more of the following remedial steps.

1. Check that the ObjectServer host and port defined in `$ASM_HOME/integrations/omnibus/omni.dat` are as expected. If the appropriate `omni.dat` connection is not for 'AGG\_P', update the name in the property files:

#### Probe

`$ASM_HOME/integrations/omnibus/kafka/probe/probe.props`

#### Gateway

`$ASM_HOME/integrations/omnibus/kafka/gateway/G_ASM.props`

2. Check that the ObjectServer port can be reached using `nco_ping.sh`. For example, for an `omni.dat` file with a section for AGG\_P, use the following:

```
$ASM_HOME/bin/nco_ping.sh AGG_P
```

#### If the ping fails

If the Object Server port **cannot** be reached, check that the configured port can be accessed using a tool such as `telnet`, `nc` or `socat`. If inaccessible, check for firewalls.

If `omni.dat` contains a hostname rather than IP, check that name is resolvable using a tool such as `nslookup` or `getent hosts`. If unresolvable, change the DNS settings to make it resolvable, or change the `omni.dat` file.

#### If the ping succeeds

If the Object Server port **can** be reached, this does not yet validate any username and password settings.

3. Check the beginning of the startup.log files in `$ASM_HOME/logs/noi-gateway/` and `$ASM_HOME/logs/noi-probe/`. The startup logs contain the results of the ObjectServer ping and host lookup of all entries in the `omni.dat` file, for example:

#### Successful

```
Thu Mar 12 14:11:17 UTC 2020 Checking access to Object Server 'NCOMS'
NCO_PING: Server available.
Thu Mar 26 17:22:04 UTC 2020 Checking hostname resolution of 'resolvable.ibm.com'
1.2.3.4      resolvable.ibm.com
```

#### Unsuccessful

```
Thu Mar 12 14:11:15 UTC 2020 Checking access to Object Server 'AGG_P'
NCO_PING: Server unavailable.
Thu Mar 26 17:22:04 UTC 2020 Checking hostname resolution of 'unresolvable.ibm.com'
Thu Mar 26 17:22:04 UTC 2020 {anyMessageHere}
```

If the ping and host lookups work from the Agile Service Manager host, but fail inside the container, check how name resolution is configured. Also be aware that the startup log shows credentials being set up in order to access the secured Kafka broker, even if the ObjectServer connection is unsecured.

4. Check for Error messages in the runtime logs:

```
$ASM_HOME/logs/noi-probe/probe.log and $ASM_HOME/logs/noi-gateway/gateway.log
```

5. Check for erroneous security credentials. **Note** that having secure credentials configured when they are **not** required can cause problems similar to ones caused by omitting them when they **are** required.

#### **Username and password issues**

An incorrect password can result in the user being locked out from the ObjectServer.

Check the Object Server's alerts.login\_failures table.

**Note** that nco\_ping does **not** validate the username or password.

#### **TLS issues**

Check the Netcool/OMNIbus documentation to determine whether the probe's

**SSLServerCommonName** and the gateway's **Gate.Reader.CommonNames** properties are required.

### **Probe runs, but gateway continually restarts (IDUC reader issue)**

If the probe runs but the gateway continually restarts, this may be due to the gateway being unable to connect the IDUC reader thread.

Both probe and gateway connect to the host and port defined in the omni.dat file, but the gateway requires an additional reader connection to receive IDUC updates from the ObjectServer.

**Workaround:** Perform one or more of the following remedial steps, in order.

1. Check the gateway log for Error level messages about Iduc: \$ASM\_HOME/logs/noi-gateway/gateway.log
2. Check the Object Server properties. There are two key ObjectServer Iduc properties:

#### **Iduc.ListeningPort**

**If it is not configured**, the port will vary upon ObjectServer restart, which can cause problems if a firewall is present.

**If it is configured**, check that it can be reached from the Agile Service Manager host using a tool such as telnet, nc or socat.

If inaccessible, check for firewalls.

#### **Iduc.ListeningHostname**

**If configured**, this is the name that the gateway reader connection will try to connect to. If the gateway cannot resolve this name, the connection will fail.

One option to resolve this is to add a dummy entry to \$ASM\_HOME/integrations/omnibus/omni.dat and restart the probe or gateway.

```
[DUMMY]
{
  Primary: {theNameFromObjectServerPropsFile} 4100
}
```

The startup log will show whether it can be resolved, as described [here](#). If it cannot be resolved, change the value in the object server properties file or make the name resolvable, for example in /etc/hosts on the Agile Service Manager host.

### **Probe runs, but gateway continually restarts (Kafka topic issue)**

If the probe runs but the gateway continually restarts, this may be due to the Kafka topic that is required by the gateway not having been created yet.

**Workaround:** The topic **will** be created at startup by the status service (for Agile Service Manager Version 1.1.8 or later), or by the Event Observer (for Agile Service Manager Version 1.1.7 or earlier). The gateway service will continue to restart until the topic becomes available.

#### **Related information**

## Startup and UI troubleshooting

---

See the following information to troubleshoot issues occurring when launching the application.

### Browser certificate failure (Internet Explorer and Edge only) (on-prem and OCP)

When opening an Agile Service Manager UI page in DASH using Internet Explorer or Edge, a browser certificate failure can occur. If this happens, install the Agile Service Manager root certificate manually.

1. Copy the root certificate (`asm-ca.crt`) from the following location on the Agile Service Manager server, then save it locally:  
`/opt/ibm/netcool/asm/security/`

#### Tip: To obtain the root certificate using Firefox:

- a. Open the Firefox browser and log into DASH.
  - b. While logged into DASH, open a new browser tab and navigate to `https://asm-ui-host`, where 'asm-ui-host' is the host name of the server running the nasm-ui container. The **Agile Service Manager 'About'** page will be displayed.
  - c. Click the padlock icon in the browser address bar.
  - d. Click **Show Connection Details > More Information > View Certificate** to display the certificate dialog.
  - e. Click **Details**, then select the **ASM CA** item.
  - f. Click **Export**, and then save the root CA certificate as a local file.
2. On your local Windows machine, locate the Agile Service Manager root certificate using File Explorer, open the context (right-click) menu, then click **Install Certificate**. The **Import** wizard will be displayed.
  3. For the Store Location, select **This Machine**, then click **Next**.
  4. Select **Place all certificates in the following store**, then browse to choose **Trusted Root Certification Authorities**.
  5. Click **Next** and **Finish**.

The next time you open an Agile Service Manager UI page in DASH using Internet Explorer or Edge, the page should load successfully.

### Cassandra database startup issue (on-prem)

During startup, the topology service attempts to connect to the Cassandra datastore before it has fully started, thereby causing an error message such as the following:

```
ERROR [14:11:07.330] [main] c.i.i.t.g.ConnectionManager - Unexpected Throwable caught creating TitanGraphjava.lang.IllegalArgumentException: Could not instantiate implementation: com.thinkaurelius.titan.diskstorage.cassandra.astyanax.AstyanaxStoreManager
```

Permissions of data and logs directories can result in Agile Service Manager not coming up cleanly with some services continually restarting, e.g. if you delete the directories and then restart without the right environment.

#### Workaround

None required.

The topology service will try to connect to the Cassandra datastore again, and will succeed once Cassandra is up and running.

## Search troubleshooting

---

See the following information to troubleshoot Search service issues.

### Slow Search response time (following a previous query)

If your search query processes a large number of matches (several million), then it can take longer to return the results. Typically, such a search time could exceed 30 seconds instead of returning the results in under five seconds (more typical). This is expected behavior. However, if you then attempt a second query with a new search string that should return results more quickly, the Results page may continue to be slow, as it may still be dealing with remnants of the previous search query.

#### Workaround

Close the **Results** page from the slow query, and open a new search box.

### Re-indexing Search

If data in Elasticsearch is out of sync with data in the Cassandra database, resynchronize it by calling the rebroadcast API of the topology service. This triggers the rebroadcast of all known resources on Kafka, and the Search service will then index those resources in Elasticsearch.

#### Workaround

Call the rebroadcast API of the Topology service, specifying a tenantId:

```
https://master_fqdn/1.0/topology/swagger#!/Crawlers/rebroadcastTopology
```

### Elasticsearch versioning errors

While using the Search service, versioning errors are reported.

Agile Service Manager uses Elasticsearch version 6.5.3, and both search components have been updated accordingly (nasm-elasticsearch, nasm-search).

#### Workaround

Ensure that you have deployed the latest versions of nasm-elasticsearch and nasm-search.

### Elasticsearch running out of disk space

If Elasticsearch runs out of disk space (when the disk is 95% full or more), it places index blocks into read-only mode, which can result in the following error:

```
ClusterBlockException[blocked by: [FORBIDDEN/12/index read-only / allow delete (api)];]
```

This issue manifests itself as the inability to search for new or potentially updated resources that are present in the topology service because they have not been indexed by the search service. It is recommended to proactively monitor disk space and take preventative action as necessary should disk space become low.

#### Workaround

Make more space available on the disk, and then manually release the locked index by calling the 'unlock' API on the Search service.

```
curl -X POST --header 'Content-Type: application/json' --header 'Accept: application/json' --user <asm_user>:<asm_password> 'https://<asm_host>/1.0/search/index/unlock'
```

Restart Elasticsearch.

**Notice:** If this problem occurs when the Search service is restarted as part of an upgrade, you must access the container running ElasticSearch and then unlock the current Search service index.

1. Discover which index is current using the following command

```
curl -X GET "localhost:9200/_alias/searchservice?pretty" -H 'Content-Type: application/json'
```

This command finds the index associated with the alias.

2. Use the following command (using the discovered index name) to enable the Search service to start and complete the index schema migration or upgrade process:

```
curl -X -k -XPUT "localhost:9200/<index_name>/_settings" -H 'Content-Type: application/json' -d '{
  "index.blocks.read_only_allow_delete": "false"
}'
```

3. If required, execute the original workaround to unlock any other indexes that have been locked.

## Search returns data for up to 12 hours longer than it exists in the topology service

**Restriction:** Agile Service Manager interprets the timeToLive settings in the Topology service and Search service differently. While the topology service can remove deleted data at one minute intervals, the Search service removes data every 12 hours only. Therefore if both services have the same timeToLive settings, for example 48 hours, then the Search service will continue to 'find' that data for up to an additional 12 hours after it has been removed from the topology service.

### Workaround

If you want to ensure that Search can never return data that the topology service has already deleted, ensure that the number of timeToLive days in the Search service is one fewer than the number of days set in the topology service.

**Remember:** As a result of this workaround the Search service will potentially be unable to return some data that exists in the topology service for a period of up to 12 hours.

## Observer troubleshooting

See the following information to troubleshoot a variety of observer issues.

### OpenStack Observer certificate chaining error

A Certificate Chaining Error can occur when launching an OpenStack Observer job, as in the following example:

```
/opt/ibm/netcool/asm/logs/openstack-observer/openstack-observer.log has following
INFO [2019-11-01 14:48:50,609] [cfd95b7e-3bc7-4006-a4a8-a73a79c71255:OpenStack - Ericsson - ceevepc]
c.i.i.t.o.t.ObservationVertex - Backing up observation vertex Ericsson - VEPC
INFO [2019-11-01 14:48:50,617] [cfd95b7e-3bc7-4006-a4a8-a73a79c71255:OpenStack - Ericsson - ceevepc]
c.i.i.t.o.t.ObservationVertex - Existing backup observation vertex CTvJ5KIFQgaGNexrlJBsjA for Ericsson - VEPC.bak
INFO [2019-11-01 14:48:50,636] [cfd95b7e-3bc7-4006-a4a8-a73a79c71255:OpenStack - Ericsson - ceevepc/KeystoneV3IdentityTask]
c.i.i.t.o.o.j.r.v.t.AbstractTask - cfd95b7e-3bc7-4006-a4a8-a73a79c71255:OpenStack - Ericsson - ceevepc/KeystoneV3IdentityTask
- Starting...
INFO [2019-11-01 14:48:50,661] [cfd95b7e-3bc7-4006-a4a8-a73a79c71255:OpenStack - Ericsson - ceevepc]
c.i.i.t.o.o.j.r.OpenStackV3FullTopologyGetter - cfd95b7e-3bc7-4006-a4a8-a73a79c71255:OpenStack - Ericsson - ceevepc - cancel -
Cancelling Tasks, Shutting Down Executor...
ERROR [2019-11-01 14:48:50,663] [cfd95b7e-3bc7-4006-a4a8-a73a79c71255:OpenStack - Ericsson - ceevepc]
c.i.i.t.o.o.j.r.OpenStackV3FullTopologyGetter - cfd95b7e-3bc7-4006-a4a8-a73a79c71255:OpenStack - Ericsson - ceevepc -
OpenStack task error occurred, rethrowing...
java.util.concurrent.ExecutionException: com.ibm.itsm.topology.observer.openstack.job.OpenStackTaskProcessingException: An
error occurred while processing KeystoneV3IdentityTask:- javax.net.ssl.SSLHandshakeException: com.ibm.jsse2.util.h: PKIX path
building failed: java.security.cert.CertPathBuilderException: PKIXCertPathBuilderImpl could not build a valid CertPath.;
internal cause is:
  java.security.cert.CertPathValidatorException: The certificate issued by CN=IBMSubCA01, DC=IBM, DC=com, DC=Raleigh is
not trusted; internal cause is:
  java.security.cert.CertPathValidatorException: Certificate chaining error
at java.util.concurrent.FutureTask.report(FutureTask.java:133)
at java.util.concurrent.FutureTask.get(FutureTask.java:203)
at
com.ibm.itsm.topology.observer.openstack.job.rest.OpenStackV3FullTopologyGetter.waitForFutures(OpenStackV3FullTopologyGetter.jav
a:155)
at
com.ibm.itsm.topology.observer.openstack.job.rest.OpenStackV3FullTopologyGetter.go(OpenStackV3FullTopologyGetter.java:107)
at com.ibm.itsm.topology.observer.openstack.job.rest.FullRESTLoadJob.observe(FullRESTLoadJob.java:85)
at com.ibm.itsm.topology.observer.app.ObservationJob.call(ObservationJob.java:179)
at com.ibm.itsm.topology.observer.app.ObservationJob.call(ObservationJob.java:63)
at
com.ibm.itsm.topology.service.utils.InstrumentedVisibleExecutorService.wrapCallable(InstrumentedVisibleExecutorService.java:385)
at com.ibm.itsm.topology.service.utils.InstrumentedVisibleExecutorService.access
$400(InstrumentedVisibleExecutorService.java:65)
at com.ibm.itsm.topology.service.utils.InstrumentedVisibleExecutorService
$InstrumentedCallable.call(InstrumentedVisibleExecutorService.java:345)
at java.util.concurrent.FutureTask.run(FutureTask.java:277)
at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1160)
```

```

        at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:635)
        at java.lang.Thread.run(Thread.java:812)
Caused by: com.ibm.itsm.topology.observer.openstack.job.OpenStackTaskProcessingException: An error occurred while processing
KeystoneV3IdentityTask: - javax.net.ssl.SSLHandshakeException: com.ibm.jsse2.util.h: PKIX path building failed:
java.security.cert.CertPathBuilderException: PKIXCertPathBuilderImpl could not build a valid CertPath.; internal cause is:
java.security.cert.CertPathValidatorException: The certificate issued by CN=IBMSubCA01, DC=IBM, DC=com, DC=Raleigh is
not trusted; internal cause is:
java.security.cert.CertPathValidatorException: Certificate chaining error
        at
com.ibm.itsm.topology.observer.openstack.job.rest.v3.task.KeystoneV3IdentityTask.process(KeystoneV3IdentityTask.java:43)
        at com.ibm.itsm.topology.observer.openstack.job.rest.v2.task.AbstractTask.call(AbstractTask.java:45)
        at com.ibm.itsm.topology.observer.openstack.job.rest.v2.task.AbstractTask.call(AbstractTask.java:22)
        at java.util.concurrent.FutureTask.run(FutureTask.java:277)
        at java.util.concurrent.Executors$RunnableAdapter.call(Executors.java:522)
        at java.util.concurrent.FutureTask.run(FutureTask.java:277)
        at
com.ibm.itsm.topology.service.utils.InstrumentedVisibleExecutorService.wrapRunnable(InstrumentedVisibleExecutorService.java:406)
        at com.ibm.itsm.topology.service.utils.InstrumentedVisibleExecutorService.access
$200(InstrumentedVisibleExecutorService.java:65)
        at com.ibm.itsm.topology.service.utils.InstrumentedVisibleExecutorService
$InstrumentedRunnable.run(InstrumentedVisibleExecutorService.java:317)
... 3 common frames omitted

```

The problem can occur if not all OpenStack certificates have been loaded into Agile Service Manager, or the certificate has not been added to the trusted CA list on the Agile Service Manager server.

### Workaround

**To load all OpenStack certificates into Agile Service Manager**, obtain a copy of the root certificate(s) from the OpenStack host, and import them into the keystore.

**Note:** Ensure you obtain **all** certificates, if the host has more than one naming alias. Obtain the certificates directly from the OpenStack administrator or the Server (that is, **do not** generate them using the openssl command).

**To add the certificate to the trusted CA list on the Agile Service Manager server**, copy the ca.pem file as root certificate.

**Note:** See the following link for more information: <https://access.redhat.com/solutions/3220561>

## File Observer heap size issue

If a large number of events are being processed, the default Java Virtual Machine (JVM) memory settings may prove insufficient and processing errors may occur. These errors can generate WARNING logs, and processing of data may be suspended.

### Workaround

Increase the maximum Java heap size (Xmx) value to 6G.

1. Edit the ASM\_HOME/etc/nasm-file-observer.yml file and change the Xmx value in the following default argument to 6G:

```
JVM_ARGS: ${FILE_OBSERVER_JVM_ARGS:--Xms1G -Xmx2G}
```

2. Restart the service.

## Jenkins Observer troubleshooting

### Artifactory integration: script approval

The first time you use integration with Artifactory, your build may fail as a result of the Artifactory API code being called not yet being whitelisted (approved). In such a case the build log will suggest that you approve the API code.

**Workaround:** You can approve the scripts in Jenkins, on the **Manage Jenkins > In-process script approval** screen. Once approved, a No pending approvals message will be displayed.

### Culprits: getting the expected username

Depending on your build configuration, you may get a 'noreply' as the user in the culprits information reported by Jenkins.

To get the actual user ID as expected, you can modify your build configuration to make it use the actual author.

**Workaround:** Go to the Jenkins **Pipeline** tab inside the build configuration, select **Add** from the **Additional Behaviours** drop-down, then click **Use commit author in changelog**.



For more Jenkins-specific information on this issue and workaround, see the following location:  
<https://issues.jenkins-ci.org/browse/JENKINS-38698>

#### Git resources URLs

Topology tools expect that artifact properties contain a valid URL using HTTP rather than SSH.

If your current Jenkins pipeline performs the checkout operation using a SSH location (such as `git@github.domain:org/repo.git`), then the right-click links will not work.

**Workaround:** Modify your Jenkins pipeline to use HTTP checkout.

## Other troubleshooting

---

See the following information to troubleshoot a variety of service issues, such proxy server buffering warning.

### Conflicting Docker network

A conflicting Docker network can cause issues with the internal Agile Service Manager communication.

#### Diagnose:

1. Check current networks. The following command lists the current networks is use:

```
docker network ls
```

2. Confirm `asm_default` subnet:

```
docker network inspect asm_default | grep Subnet
```

3. Check that any other networks do not use the same subnet:

```
docker network inspect <NETWORK_NAME> | grep Subnet
```

#### Solution:



**Warning:** Before you remove any clashing networks or interfaces, ensure that they are not required.

- If the clashing network is not required, remove it:

```
docker network rm <NETWORK_NAME>
```

- If problems still occur, check `ifconfig` for clashing IP addresses. Remove unwanted network interfaces.

```
ifconfig <network interface> down
```

- If removing the docker network or clashing network interface isn't an option, you can configure the Agile Service Manager network to use a different subnet. Edit the `<ASM_HOME>/docker-compose.yml` file. Under the network section at the top of the file, modify the subnet value, then save the changes. Restart Agile Service Manager to use the new subnet.

### Netcool/OMNIBus integration event dataset error

If scope-based grouping is not enabled in Web GUI, an event dataset error is displayed when using the `Example_IBM_CloudAnalytics` view in the Event Viewer.

**Workaround:** To resolve this issue, enable scope-based grouping. For more information, see the [Installing scope-based event grouping](#) topic in the Tivoli Netcool/OMNIBus Knowledge Center.

### Scripts return HTTP error status 502 (bad gateway)

While running scripts in `$ASM_HOME/bin` an HTTP error status 502 (bad gateway) is returned.



One possible cause is the Nginx proxy running, but the service is not running, or has not yet started. You can verify that the service is running using the `$ASM_HOME/bin/asm_status.sh` script, and check the log file under `$ASM_HOME/logs/` for continuous startup errors.

Another possible cause in some host configurations can be a `localhost` resolution issue.

**Workaround:** Configure the scripts to use an explicit hostname by setting the **SERVICE\_HOST** environment variable before running them, for example:

```
export SERVICE_HOST='hostname'
```

## Scripts return HTTP error status 000 (no response)

While running scripts in `$ASM_HOME/bin` an HTTP error status 000 (no response) is returned.

A possible cause is the Nginx proxy not running, or Agile Service Manager processes not running. You can verify that the service is running using the `$ASM_HOME/bin/asm_status.sh` script.

**Workaround:** Run the Agile Service Manager start command (`$ASM_HOME/asm_start.sh`). If the error recurs, check the logs in `$ASM_HOME/logs/nginx/` for more information.

## Certificate failure: 'ssl3 alert handshake failure'

Due to a certificate not being returned, certain services may encounter connectivity problems that are subsequently recorded in the logs as 'ssl3 alert handshake failure' error.

You may encounter a system warning, such as the following AWS Observer example: Warning The AWS Observer cannot be reached. The observer might be stopped or might have been uninstalled. In this state it is not possible to view or submit jobs.

### Workaround

Restart your containers.

If the error persists, restart Agile Service Manager.

## Proxy service buffering warning

If large information payloads are sent to the Nginx proxy server service, the error log may record the following warning: `[warn]...a client request body is buffered to a temporary file...`

Such warnings indicate that Nginx is temporarily storing the payload in storage as opposed to using memory. While this does not affect the performance of Agile Service Manager much, these messages could flood the log file, making other debugging tasks more difficult.

### Workaround

To increase the limit at which Nginx uses memory rather than storage, open the `$ASM_HOME/etc/nginx/conf.d/general.conf` configuration file with a suitable text editor, and increase the value of the **client\_body\_buffer\_size** parameter as required.

Restart the proxy service using the following command:

```
$ASM_HOME/bin/docker-compose restart proxy
```

## OCP troubleshooting

---

See the following information to troubleshoot RedHat OpenShift Container Platform issues.

### Network discovery errors: 'cannot resolve bootstrap urls'

## 'Sherpa service nginx gateway timeout'

When these errors occur, restart dns pods using the following commands:

1. `kubectl get pods --namespace=openshift-dns`
2. `kubectl delete --all pods --namespace=openshift-dns`
3. `kubectl get pods --namespace=openshift-dns`

Ensure all pods are up and running.

## Services not binding to storage (after upgrade or uninstall)

Some services fail to bind to the provisioned storage, typically resulting in pods stuck in 'pending' state.

After removing a previous installation of Agile Service Manager and some of its PersistentVolumeClaim (PVC) objects, any associated PersistentVolume (PV) objects are placed in a 'Released' state. They are now unavailable for bonding, even if new PVCs that are part of a new Agile Service Manager installation have the same name and namespace. This is an important security feature to safeguard the previous PV data.

**Investigating the problem:** The following example lists the 'elasticsearch' pods and their status, and the result shows the 'pending' status, indicating the problem.

```
$ kubectl get pod -l app=elasticsearch
```

NAME	READY	STATUS	RESTARTS	AGE
asm-elasticsearch-0	0/1	ContainerCreating	0	4s
asm-elasticsearch-1	0/1	<b>Pending</b>	0	3s
asm-elasticsearch-2	0/1	<b>Pending</b>	0	3s

This example examines the state of the PersistentVolumeClaims and the (truncated) result indicates that the status is 'pending'.

```
$ kubectl get pvc -l app=elasticsearch
```

NAME	STATUS	VOLUME
data-asm-elasticsearch-0	Bound	asm-data-elasticsearch-0
data-asm-elasticsearch-1	<b>Pending</b>	
data-asm-elasticsearch-2	<b>Pending</b>	

This example examines the PersistentVolumes and the (truncated) result indicates that the status is 'released'.

```
$ kubectl get pv -l app=elasticsearch
```

NAME	CAPACITY	ACCESS MODES	RECLAIM POLICY	STATUS
asm-data-elasticsearch-0	75Gi	RWO	Retain	Bound
asm-data-elasticsearch-1	75Gi	RWO	Retain	<b>Released</b>
asm-data-elasticsearch-2	75Gi	RWO	Retain	<b>Released</b>

**Solution:** As admin user, remove the PV.Spec.ClaimRef.UID field from the PV objects to make the PV available again. The following (truncated) example shows a PV that is bound to a specific PVC:

```
apiVersion: v1
kind: PersistentVolume
spec:
  claimRef:
    apiVersion: v1
    kind: PersistentVolumeClaim
    name: data-asm-elasticsearch-1
    namespace: default
```

```
resourceVersion: "81033"
uid: 3dc73022-bb1d-11e8-997a-00000a330243
```

To solve the problem, you edit the PV object and remove the uid field, after which the PV status changes to 'Available', as shown in the following example:

```
$ kubectl get pv -l app=elasticsearch
```

NAME	CAPACITY	ACCESS MODES	RECLAIM POLICY	STATUS
asm-data-elasticsearch-0	75Gi	RWO	Retain	Bound
asm-data-elasticsearch-1	75Gi	RWO	Retain	<b>Available</b>
asm-data-elasticsearch-2	75Gi	RWO	Retain	<b>Available</b>

## User interface timeout errors

To prevent or mitigate UI timeout errors, you can increase the timeout values for the following parameters, which are defined in configmap:

- topologyServiceTimeout
- searchServiceTimeout
- layoutServiceTimeout

To change the timeout values of these (in seconds) edit the configmap using the following command:

```
kubectl edit configmap {{ .Release.Name }}-asm-ui-config
```

When done, restart the NOI webgui pod.

## System health check job fails (after rollback)

After rolling back Agile Service Manager from Version 1.1.10 to 1.1.9, the system health check job fails. While other pods are running, the noi-topology-system-health-scheduled job remains in an error state.

### Workaround

1. Obtain the Kubernetes Observer pod name:

```
oc get pods | grep kubernetes
```

Example output showing the pod name:

```
noi-full-topology-kubernetes-observer-5f9b8c8749-zfkf8 1/1 Running 0 30m
```

where 'noi-full' is the NOI release name, and 'noi-full-topology' is the Agile Service Manager release name.

2. Using the pod name, check the logs for the Kubernetes Observer pod:

```
oc logs noi-full-topology-kubernetes-observer-5f9b8c8749-zfkf8
```

Example log file error output containing the tenant ID and job ID:

```
Skipping scheduled execution of job 'kubernetes-observer:system-health-job' as current state is FINISHING for tenant cfd95b7e-3bc7-4006-a4a8-a73a79c71255
INFO [2020-12-03 15:27:49,518] [pool-12-thread-2] c.i.i.t.o.t.TopologyServiceTrackVertices - Associated 6CXp7zPwQVGZr7YdXN8cMg with kubernetes-observer:system-health-job
```

3. Verify the state of the Kubernetes Observer job using Swagger (see the [“Swagger reference”](#) on page 410 for more information):

```
https://<release-name>.<namespace>.apps.<hostname>/1.0/topology/swagger#!
/Management_Artifacts/getMgmtArtifact
```

Example output shows that hasState is stuck in a state of FINISHING (but **should** be FINISHED):

```
{
  "keyIndexName": "kubernetes-observer:system-health-job",
  "_id": "6CXp7zPwQVGZr7YdXN8cMg",
  "executionTime": 5,
  "hasStateDescription": "Job finished",
  "observedAt": "2020-12-03T15:27:49.478Z",
  "matchTokens": [
    "kubernetes-observer"
  ],
  "hasState": "FINISHING",
  "observedTime": 1607009269478,
  "tags": [
    "OBSERVATION_JOB",
    "kubernetes-observer"
  ],
  "schedule": {
    "interval": 5,
    "units": "minutes",
    "nextRunTime": 1607009568759
  },
  "path": "/jobs/local",
  "vertexType": "mgmtArtifact",
  "createTime": 1606952966418,
  "_created": "2020-12-02T23:49:26.418Z",
  "name": "system-health-job",
  "entityTypes": [
    "ASM_OBSERVER_JOB"
  ],
  "observerName": "kubernetes-observer",
  "parameters": {
    "connect_read_timeout_ms": 5000,
    "hide_terminated_pods": false,
    "namespace": "noi-on-ocp",
    "data_center": "icp"
  }
}
```

#### 4. Use Swagger to update the job:

```
https://<release-name>.<namespace>.apps.<hostname>/1.0/topology/swagger#!/
/Management_Artifacts/putMgmtArtifact
```

Using the tenant ID and job ID values obtained in a previous step, update the job state:

```
{
  "keyIndexName": "kubernetes-observer:system-health-job",
  "_id": "6CXp7zPwQVGZr7YdXN8cMg",
  "hasState": "FINISHED",
  "observerName": "kubernetes-observer"
}
```

Response code indicating success:

```
200
```

#### 5. Using the name obtained earlier, delete the Kubernetes Observer pod to recycle (recreate) it:

```
oc delete pod noi-full-topology-kubernetes-observer-5f9b8c8749-zfkf8
```

where 'noi-full' is the NOI release name, and 'noi-full-topology' is the Agile Service Manager release name.

#### Result

The Kubernetes Observer pod is recreated and loads successfully.

You can verify that the Kubernetes Observer job has run by checking that it is in a Finished state (described in step 3).

---

## Chapter 9. Reference

Use the following reference information to enhance your understanding of Netcool Agile Service Manager interfaces and functionality.

### Topology service reference

---

Use this introduction to the Netcool Agile Service Manager services to understand the most important topology service concepts and functions.

You can access the Swagger documentation for the topology service at the following location: `https://<your host>/1.0/topology/swagger`

**Remember:** IBM Netcool Agile Service Manager is cloud-born, and built on secure, robust and proven technologies. It is designed to be flexible and can be extended as needed using plug-in components and micro-services to cater for highly specific environments.

#### Refresher:

It is important that you are familiar with the following important terms introduced in the [Glossary](#), as they will be used and expanded on in this reference section:

#### resource

A resource is a node in an interconnected topology, sometimes also referred to as a vertex, or simply a node. It can be anything in a user-specific topology that has been designated as such, for example a hardware or virtual device, a location, a user, or an application.

#### edge

An edge is a relationship between resources, also simply referred to as the 'link' between resources. Edges have a *label*, which allocates them to a family of edges with specific behavior and governs how they are displayed in the UI, and an *edgeType*, which defines the relationship in real terms.

#### tenant

A tenant is represented by a globally unique identifier, its tenant ID.

The default tenant ID is: cfd95b7e-3bc7-4006-a4a8-a73a79c71255

#### provider

A provider is usually a single data source within the scope of a tenant.

**Note:** A provider's **uniqueId** property for a resource is unique only within the scope of a provider.

#### status

Status is a property of one or more resources, and a single resource can have different types of status.

Each status can be in one of three states: 'open', 'clear' or 'closed'.

The severity can be one of the following: 'clear', 'information', 'indeterminate', 'warning', 'minor', 'major' or 'critical'.

The status of a resource can be derived from events, in the case of the resource having been retrieved via the status service, or it can be supplied when resources are posted to the topology service.

**Additional:** For more information on how topologies are displayed in the UI, you can take a look at the topology viewer screen reference topic: [“Topology viewer reference” on page 393](#)

## Properties

The Topology Service has two categories of properties, generic and user. Generic properties have fixed data types, while user-defined properties do not.

### Generic properties

Generic properties are few in number and constrained to a fixed data type. They can also be subdivided into those which are read-write and those which are read-only.

#### uniqueId

The uniqueId is the string used to match resources from the same provider. It could be, for example, a UUID via which the provider can look up its own local data store for information about that device.

If you send the same resource with the same Id and the same provider more than once, it will be treated as the same resource. However, the uniqueId is only unique within the context of its provider.

#### matchTokens

These tokens are used to store strings which are significant with respect to that resource, and could match it to events.

#### name

The name string is **required** by the UI to display a resource.

This string does not have to be unique, and it should be short and memorable.

#### tags

Tags can be used to filter resources and store strings, which can later retrieve groups of related resources.

#### entityTypes

These are defined as a set, though with usually only a single member, of the type(s) this resource represents.

**Tip:** A set of predefined entityTypes with associated icons exist. See the [“Entity types” on page 360](#) topic for more information.

Table 67. Generic properties					
Name	Type	Cardinality	Alias	Read-only	Indexed
age	integer	single		no	y
aliasIds	Id	set	_aliasIds	yes	n
beginTime	long	single	_startedAt	yes	y
changeTime	long	single	_modifiedAt	yes	n
createTime	long	single	_createdAt	yes	y
deleteTime	long	single		yes	n
description	string	single		no	y
edgeTenantId	Id	single	_edgeTenantId	yes	n
edgeType	string	single	_edgeType	yes	n
endTime	long	single	_deletedAt	yes	y
entityTypes	string	set		no	y
eventId	string	single		yes	n
eventManager	string	single		yes	n
expireTime	long	single	_expiredAt	yes	n

Table 67. Generic properties (continued)

Name	Type	Cardinality	Alias	Read-only	Indexed
geolocation	GeoLocation	single		no	n
hasState	string	single		yes	n
icon	string	single		no	n
id	long	single		yes	n
keyIndexName	string	single		yes	n
label	string	single		yes	n
matchTokens	string	set		no	y
name	string	single		no	y
observedTime	long	single	_observedAt	yes	n
operation	string	single		yes	n
partOfExternal	Boolean	single		yes	n
prevBeginTime	long	single		yes	y
providerId	Id	single		yes	n
providerName	string	single		yes	n
reconciliation Tokens	string	set		yes	n
referenceId	Id	single		yes	y
referenceNo	long	single		yes	n
serialized HashMap	HashMap	single		yes	n
severity	string	single		no	n
speed	long	single		no	y
statusType	string	single		yes	n
tags	string	set		no	y
tenantIds	Id	set	_tenantIds	yes	y
uniqueId	string	single		no	y
uuid	Id	single	_id	yes	y
version	string	single		no	y
vertexType	string	single		yes	y

## User properties

User-defined properties are free-form, and are not constrained by any given data type. You can add new user properties as needed.

You can define any custom properties, such as, for example **ipAddress**.

**Note:** All user-defined properties such as `ipAddress` are not in the generic set, and are stored as a serialized 'blob' data type instead. The implication of this storage convention is that these properties cannot be filtered, as they are incompatible with the **\_filter** query parameter used in the REST API.

**Tip:** The Swagger documentation listing all properties can be found at the following default location:  
<https://<your host>:8080/1.0/topology/swagger#!/Schema/getDefaultProperties>

## Edge labels

The topology service defines a family of labels for the edges it supports.

**Note:** Most interactions with edges in the Topology Service are with edge *types* rather than edge *labels*. Edge types can be conceptualized as instances of the edge label classes, and are documented separately [here](#).

### aggregation

A relationship where the target vertex can exist independently of the source, but is needed for the source to function.

A 'parent-child' relationship where the parent (source) is aggregating the children (target).

The type of aggregation is determined by the value of the edge type.

Use this edge label to represent a resource in the UI that is composed of various elements, for example a book contains words.

- The direction is always from parent to child.
- Children can have multiple parents.

See [Table 68 on page 357](#) for information on the edge types associated with this edge label.

### association

A 'weak' relationship where both source and target vertex can exist independently and neither source nor target are required for the other to function.

The specific type of association is determined by the edge type.

Use this edge label to represent a general relationship between vertices in the UI, for example a person has a house.

The label itself has no direction, but a direction could be implied by the edge **type** used.

See [Table 69 on page 357](#) for information on the edge types associated with this edge label.

### composition

A 'part-whole' relationship where the child (source) is an integral part of the parent (target).

Use when you need to a short-cut from a node to its root, bypassing fine-grained relationships, for example a component is **partOf** a device.

- The direction is always child to parent.
- Children cannot have multiple parents.
- Composition relationships are hidden from the element-to-element view. They are used to construct the [element-to-host](#) view.

### dataFlow

A dataFlow label represents a data flow between a pair of vertices.

The specific type of data flow is qualified by properties on the edge type.

Use this label when you need to represent any form of data flow in the UI, for example a person emailing another person.

- The label itself has no direction, but a direction can be implied from the edge type used.

See [Table 70 on page 358](#) for information on the edge types associated with this edge label.

### dependency

A 'strong' relationship where the source depends on the target and cannot operate independently of it, and is not integral to it.

The specific type of dependency is determined by the edgeType.

Use this label when you need to represent the dependency of one resource on another in the UI, for example an application **dependsOn** a database.



The direction is always from the dependent resource to the independent resource.

See [Table 71 on page 359](#) for information on the edge types associated with this edge label.

## Edge types

All edges created in the Topology Service should have an edge **type** defined. The following section lists the edge types that are associated with each of the public-facing edge labels. If none of the default edge types suffice, you can create custom edge types.

### Edge types for public edge labels

**Remember:** Edge **types** can be thought of as being instances of the edge **label** classes, in this case the public-facing edge labels. Most interactions with edges in the Topology Service are with edge *types* rather than edge *labels*.

You can access the Swagger documentation for 'edge types' at the following default link: <https://localhost/1.0/topology/swagger#!/Schema/getEdgeTypes>

Table 68. Edge types for the **Aggregation** edge labels

Edge type	Description	Example
contains	The source resource is considered to contain the target resource	Slot contains a card
federates	The source resource is a federation of the target resources	Database federates nodes
members	The source resource has the target resources as its members	Subnet members are IP addresses

Table 69. Edge types for the **Association** edge labels

Edge type	Description	Example
aliasOf	Depicts that one resource is an alias of another; potentially from different providers	FQDN1 is an alias of FQDN2
assignedTo	When one resource has been assigned to another	The alarm is assignedTo an operator
attachedTo	When one resource is attached to another	The plug is attachedTo to the cable
classifies	When one resource is used to classify another	The government classifies the document
configures	When one resource configures or provides configuration for another resource	The file configures the application
deployedTo	Describes when one resource has been deployed to another	The application was deployedTo the server
exposes	When one resource exposes another	The application exposes the interface
has	Generalized relationship when one resource possesses, owns or has another resource	Host has component

Table 69. Edge types for the **Association** edge labels (continued)

Edge type	Description	Example
implements	When one resource implements another.	The class implements the interface
locatedAt	When one resource is physically located in/at another resource's location	Host is locatedAt data centre
manages	When one resource manages another	The boss manages the employee
monitors	When one resource monitors another	The application monitors the host
movedTo	When one resource has moved to a new and different resource	The service has movedTo the host
origin	Indicates the origin of a Vertex	Device's origin is a vendor
owns	Indicates ownership of one resource by another resource	The user owns the server
rates	Can be used when one resource rates another	The manager rates the employee
resolvesTo	When one resource resolves to another	The hostname resolvesTo an address
realizes	When one resource realizes another	The hypervisor realizes the virtual machine
segregates	When one resource segregates another	The firewall segregates the network
uses	When one resource takes, employs or deploys another resource as a means of achieving something	Application uses this disk

Table 70. Edge types for the **Data flow** edge labels

Edge type	Description	Example
accessedVia	One resource is accessed via another, typically remote.	Server is accessedVia a FQDN
bindsTo	A network layering relationship such that the source 'runs on top' of the target.	Logical interface bindsTo a physical port
communicatesWith	A relationship whereby one resource communicates with another	The sensor communicatesWith an application
connectedTo	A relationship whereby one resource is connected to another	The laptop is connectedTo the switch
downlinkTo	One resource is down-linked to another	The controller has a downlinkTo the sensor
reachableVia	A resource is reachable via another resource	The network is reachableVia the gateway

Table 70. Edge types for the <b>Data flow</b> edge labels (continued)		
Edge type	Description	Example
receives	A resource receives data from another	The Mail server receives an email
routes	A relationship whereby one resource routes data for another	The device routes the data
routesVia	A relationship whereby data from one resource routes via another	The traffic routesVia the device
loadBalances	One resource which load balances for others	The load balancer loadBalances to servers
resolved	When one resource resolved something for another	DNS server resolved IP address
resolves	Represents that one resource uses another to resolve it	FQDN resolves to the address
sends	A resource sends some data to another	The application sends an SMS message
traverses	Describes when one resource traverses another	The message traverses the network
uplinkTo	One resource is up-linked to another	The sensor has an uplinkTo the controller

Table 71. Edge types for the <b>Dependency</b> edge labels		
Edge type	Description	Example
dependsOn	A generic dependency between resources	One application dependsOn another
runsOn	A resource runs on (and therefore depends on) another resource	The service runsOn the host

## Custom edge types

If none of the default edge types suitably represent a certain relationship, you can define a custom edge type via the

<https://localhost:8080/1.0/topology/types/edge> **POST API**

**Important:** A custom edge type needs to be created **before** the observation job passes any edges with the type to the Agile Service Manager topology service.

### edgeType

Required

The edgeType name, which has to be unique and **cannot** match the edgeType name of a default edgeType, unless the edgeLabel also matches the corresponding default edge type's edgeLabel parameter.

**Restriction:** A scenario where both edgeType and edgeLabel match the fields of a default edge type is equivalent to manually creating a default edge types, which is not necessary as default edge types are created implicitly by the topology when needed.

### edgeLabel

Required

The edgeLabel of the custom edgeType, which has to be one of the following:

- dataFlow
- dependency
- association
- aggregation

#### description

Optional (but recommended)

This should be a meaningful description of the type of relationship this edge type represents.

#### Example:

```
curl -k -X POST --header 'Content-Type: application/json' --header 'Accept: application/json' --header 'X-TenantID: cfd95b7e-3bc7-4006-a4a8-a73a79c71255' -d '{
  "edgeType": "connectedTo",
  "edgeLabel": "dataFlow",
  "description": "Default relationship between two devices that exchange data"
}' 'https://localhost:8080/1.0/topology/types/edge'
```

## Entity types

The Topology Service allows you to group together resources of the same type using the `entityTypes` property, and identify them in the UI by their icon. Usually a resource has only a single entity type, however, as the property is a *Set*, it is possible for a resource to have more than one entity type. A number of pre-defined entity types are supplied, which are listed in this topic. In addition, you can create additional entity types as required.

### Pre-defined entity types

You can access the Swagger documentation for 'entity types' currently in use at the following default link: <http://localhost/1.0/topology/swagger#!/Schema/getTypes>

An entity type is used to map the resource vertex to an icon, and it also allows for more efficient searching via the **\_type** query parameter, which can be found in the Swagger documentation at the following default location: <http://localhost/1.0/topology/swagger#!/Resources/getResourceList>

**Remember:** An administrator user can create additional custom resource or entity types as required. Any observer may add new types.

**Note:** The topology service API only returns the resource or entity types in use.

The following table lists these entity types with links to their icons, if defined.

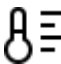



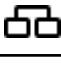

Table 72. Predefined entity types and icons, where defined	
Entity type	Icon
aircond	
alarm	
application	
backplane	
bridge	
building	

Table 72. Predefined entity types and icons, where defined (continued)

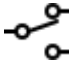




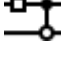






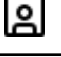

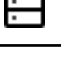
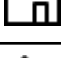
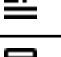
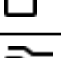
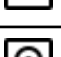
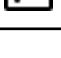
Entity type	Icon
bypass	
card	
cell	
channel	
chassis	
circuit	
cluster	
command	
component	
computer	
container	
cpu	
credential	
daemonset	
database	
datacenter	
deployment	
device	
directory	
disk	

Table 72. Predefined entity types and icons, where defined (continued)
















Entity type	Icon
domain	
emailaddress	
environment	
event	
fan	
file	
firewall	
flavor	
forum	
fqdn	WWW
gauge	
group	[...]
host	
hsrp	
hub	
image	
input	[←]
ipaddress	
ipam	[ $\frac{1}{0}$ ]

Table 72. Predefined entity types and icons, where defined (continued)





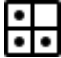
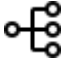



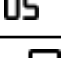
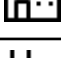
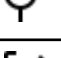
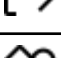
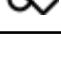


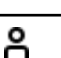


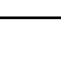
Entity type	Icon
loadbalancer	
location	
memory	
message	
namespace	
network	
networkaddress	
networkgear	
networkinterface	
operatingsystem	
organization	
outlet	
output	
patch	
path	
pdu	
peripherals	
person	
pod	
printer	

Table 72. Predefined entity types and icons, where defined (continued)










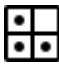







Entity type	Icon
process	
product	
project	
psu	
resource	
router	
routetable	
routetarget	
rsm	
rule	$f_x$
sector	
securitygroup	
sensor	
server	
service	
serviceaccesspoint	
slot	
snmpsystem	SNMP
stack	



Table 72. Predefined entity types and icons, where defined (continued)







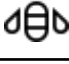
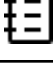
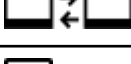


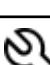






Entity type	Icon
statefulset	
status	
storage	
subnet	
subscription	
switch	
swarm	
task	
tcpudpport	
template	
tenant	
tier	
tool	
ups	
variable	$f_x$
version	$\begin{matrix} * \\ a \ b \end{matrix}$
vlan	VLAN
vm	
volume	
vpn	VPN
vrf	

Table 72. Predefined entity types and icons, where defined (continued)

Entity type	Icon
workload	

## Schema REST API

The Topology Service has an API call which will allow you to see all the instantiated entity types in a given topology. The default Swagger location is: <http://localhost/1.0/topology/swagger/#!/Schema/getTypes>

**Important:** To retrieve the type name, you must add the `_field=name` query parameter.

## REST API

Interactions with the Topology Service take place through a REST API. API calls and Swagger documentation links are listed in this section. Use the Swagger documentation to access more detailed information.

### REST API calls and default Swagger links

The REST API calls are grouped together into different sections. If installed on *localhost*, then the Swagger documentation can be accessed at the following link: <https://localhost/1.0/topology/swagger/>

#### Composites

The API calls in the composites section of the Merge Service allow you to view which resources have been merged into composites and to create and change them.

<https://localhost/1.0/merge/swagger/#!/Composites>

#### Groups

The API calls in the **groups** section allow you to create a vertex which represents a group and then to associate other resource vertices with that group.

<https://localhost/1.0/topology/swagger/#!/Groups>

#### Management artifacts

The Management artifacts calls provide the means to associate non-resource vertices, such as tests, with resource vertices.

[https://localhost/1.0/topology/swagger/#!/Management\\_Artifacts](https://localhost/1.0/topology/swagger/#!/Management_Artifacts)

#### Metadata

Metadata provides the means to store data in a metaData vertex, which can then be associated with a resource vertex using the *metadataFor* edge type.

<https://localhost/1.0/topology/swagger/#!/Metadata>

#### Resources

The most common calls.

These API calls are used to create, update and delete resource vertices in the topology. It also includes API calls to create, update and delete edges between resources.

<https://localhost/1.0/topology/swagger/#!/Resources>

**Note:** The topology service has a history model which allows it to retain information on the historical resource properties and edges for 30 days. The Resources methods will take an `_at` query parameter, which will cause them to return what the topology previously looked like at a specific point in time. This allows the UI to visualize the topology as it was in the past.

#### Rules

The API calls in the rules section of the Merge Service allow you to view which merge rules have been defined and to create or update merge rules.

<https://localhost/1.0/merge/swagger/#!/Rules>

## Schema

The schema API calls can be used to query the Topology Service for information about the types of entities which exist within the topology.

<https://localhost/1.0/topology/swagger/#/Schema>

## Service info

The Service info API calls include a health check call, and a call to return the current Topology Service version.

[https://localhost/1.0/topology/swagger/#/Service\\_Info](https://localhost/1.0/topology/swagger/#/Service_Info)

## Status

The status API provides methods to associate and manipulate the status that is associated with a given resource.

<https://localhost/1.0/topology/swagger/#/Types>

## Tenants

The Tenants API provides a mechanism by which resource, metadata and management artifacts can be made globally readable by multiple tenants.

<https://localhost/1.0/topology/swagger/#/Tenants>

## Types

These return information on the entity types which have been instantiated.

**Tip:** This includes the **\_include\_count** query parameter, which you can use to return a time-stamped count of both the number of types, and number of instances of each type.

<https://localhost/1.0/topology/swagger/#/Types>

# Status (and state)

Resources have specific statuses, and in turn each status has a state of open, clear or closed.

## Status

A single status can affect multiple resources, and a single resource can have multiple different statuses. For example, a single *link down* event can generate the status for both the interface resource and the host resource; or a single host could have *CPU* or *disk usage* status in addition to any *link down* status.

Resource status can be viewed in the Topology Viewer, and can be set or retrieved via the topology service REST API. An 'open' event is one with a severity other than 'clear'.

You can access the Swagger documentation for 'status' at the following default link: <https://localhost/1.0/topology/swagger/#/Status>

**Important:** When modeling resources, you must consider [Status assignment from events](#).

**Tip:** The Topology Service stores the event **Severity**, and nodes in the UI are colored based on severity, which is always one of the following:

- clear
- indeterminate
- information
- warning
- minor
- major
- critical

Take a look at the severity icons in the topology viewer reference topic: [Severity icons table](#)

## Status assignment

The status of a single resource can be supplied, alongside other resource properties, when creating or updating a resource. Alternatively, an event can generate the status of one or more resources.

**Remember:** A status **always** has one of the following three states:

### Open

Always has a severity other than 'clear'

An active issue that may require your attention

### Clear

Working as expected

Also has a severity of 'clear'

### Closed

No longer active or relevant, a deleted event

Also has a severity of 'clear'

## Status assignment from events

The status service receives events and tries to find matching resources in the topology service. A resource with no match tokens defined will not have events matched to it, but if found, the status of those resources is set from the event data. The assigned status depends on the following event data:

### matchTokens

This property must be used to list any data that can identify (match) the resource.

Each field may be globally unique, or may be unique within the scope of a composition. In other words, a resource modeled via a composition relationship, such as `partOf`, can be distinguished from other children within the composition using these fields.

For example, either a globally unique and fully qualified domain name or IP address, or a locally unique interface name (that is, local within the scope of the host), can be used to identify the resource.

### partOf composition relationship

The status service uses composition relationships to match fields that are unique only within the scope of a parent.

For example, an IP address can be used to find a main node, and an interface name can be used to identify an interface within that main node.

## Timestamps

Both vertex and edge graph elements can have multiple timestamps, documented here.

### Timestamp formats

The following formats apply across all Agile Service Manager APIs when specifying timestamps:

#### Millisecond Epoch

Example:

```
$ date "+%s%3N"
1595261016885
```

#### RFC-3339 / ISO-8601

Format:

```
YYYY-MM-DDThh:mm:ss.sssZ
```

Example:

```
$ date --universal "+%FT%T.%3NZ"  
2020-07-20T16:03:00.382Z
```

## beginTime

The beginTime timestamp records the beginning of a period of time for which the resource was valid, with endTime marking the end of that period.

**Tip:** A given resource may have multiple begin times in its historic record, and there may be gaps in that record if the resource was offline for periods.

- All resources and historic resources that are representations of the same thing have a distinct beginTime
- Resource beginTime together with endTime is used in historic graph traversals, that is, when the **\_at** parameter is supplied. The period during which a resource is valid is defined as:

```
atTime >= beginTime && atTime < endTime
```

- A vertex or edge which has the beginTime equal to the endTime can be used to store audit information, such as the provider which deleted a given resource. However, because it takes up zero time it does not form part of the history and is ignored by the above equation.

## prevBeginTime

If history exists for a given resource then this property will be set to the beginTime of the most recent historical resource.

## changeTime

The changeTime timestamp records when the properties of the element last changed.

Its value may be less than the observedTime, which is updated on a POST or PUT even if no property values have changed.

## createTime

The createTime timestamp records when the element was first created in the topology service.

- Historical resources do not store createTime, as it is shared with the anchor.
- This is needed when looking for something older than 30 days, that is, when there is no beginTime this old because the historical resources have timed out.

## endTime

The endTime timestamp records when the element was deleted.

- All resources and historic resources that are representations of the same thing have a distinct endTime.
- Resource endTime is used in historic graph traversals, that is, when the **\_at** parameter is supplied.
- For current resources, endTime is LONG\_MAX. This is sometimes hidden via the REST API.

## observedTime

The observedTime timestamp records when the element was last observed, that is, when data was last input to the topology service for the element.

## Netcool Agile Service Manager cookbook

The Netcool Agile Service Manager cookbook is a collection of 'recipes' and best-practice guidelines compiled by Netcool Agile Service Manager SMEs, developers and testers. The purpose of this section is

to provide you with practical information, such as implementation examples and code samples, that you can use to get started.

**Restriction:** Recipes provided here must be amended and adjusted to suit your own specific Netcool Agile Service Manager implementation.

## Virtual machine recipe

One of the most important goals of Netcool Agile Service Manager is to support the assurance and provisioning of modern IT, network and storage environments. These environments all make extensive use of increasingly nested virtualization technologies that need to be modeled. The following recipe introduces such an IT Virtualization scenario, and describes an OpenStack response that provides a solution.

### IT Virtualization

The Netcool Agile Service Manager model of a nested virtualization scenario can extend the traditional types of models provided by other solutions.

This model can represent a multi-domain view of the world that links IT, network, storage, applications and services. In addition, it can incorporate concepts such as OpenStack's Heat Orchestration and Slack collaboration relative to traditional IT resources.

Some of the benefits of this approach are:

#### To provide additional context

Increasingly, teams are more multi-disciplined and no longer operate in informational or functional silos. For example, network teams may include IT Virtualization specialists.

Such teams need access to additional context when needed in order to answer some of their business-critical questions, such as:

- What storage volume is a VM attached to?
- Which orchestration step realized a network port?
- Who collaborated with whom for a particular incident?
- Which applications and services are supported by a network subnet?
- Which VM instances were shutdown as part of a scale-in activity 1 hour ago?
- What is the impact of removing a given Hypervisor from the environment?
- Which fixed IP addresses have a floating IP address been bound to in the last week?


#### To provide a data-rich base

Value-added services can be bolted onto a base system, provided the information exists, and the system has an architecture that allows for rapid extension.

For example, when building analytics on the topology data, the availability of information such as seasonality can provide additional insights.

The following diagram depicts the nested layers of virtualization, including networking, between these layers and technologies such as Docker and LXC or LXD.

**Note:** The services exposed can be applications or appear to be traditionally physical services such as network routers, switches, firewalls and load balancers (a key goal of NFV).



## OpenStack

OpenStack is a free and open-source platform for cloud computing, typically deployed as an IaaS (Infrastructure-as-a-Service) capability. The software platform consists of interrelated components that control diverse, multi-vendor hardware pools of processing, storage, and network resources throughout and between data centers.

OpenStack provides a number of projects, and related services and APIs, that are summarized here, as they speak directly to the need to have a multi-domain view of the environment. For more information, see the OpenStack project navigator and documentation at the following location: <https://www.openstack.org/software/project-navigator/>

OpenStack **core services** include the following:

### Nova

*Compute* manages the lifecycle of compute instances in an OpenStack environment.

### Neutron

*Networking* enables network connectivity as a service for other OpenStack services.

### Swift

*Object Storage* stores and retrieves arbitrarily unstructured data via a REST API.

### Cinder

*Block Storage* provides persistent storage to running instances.

### Keystone

*Identity* provides authentication and authorization services to OpenStack services and a service catalog.

### Glance

*Image Service* stores and retrieves virtual machine disk images for use by Nova.

OpenStack **optional services** include the following:

### Horizon

dashboarding

### Ceilometer

telemetry

### Heat

orchestration

## Sahara

Elastic Map Reduce

## Designate

DNS

## Barbican

Key Management

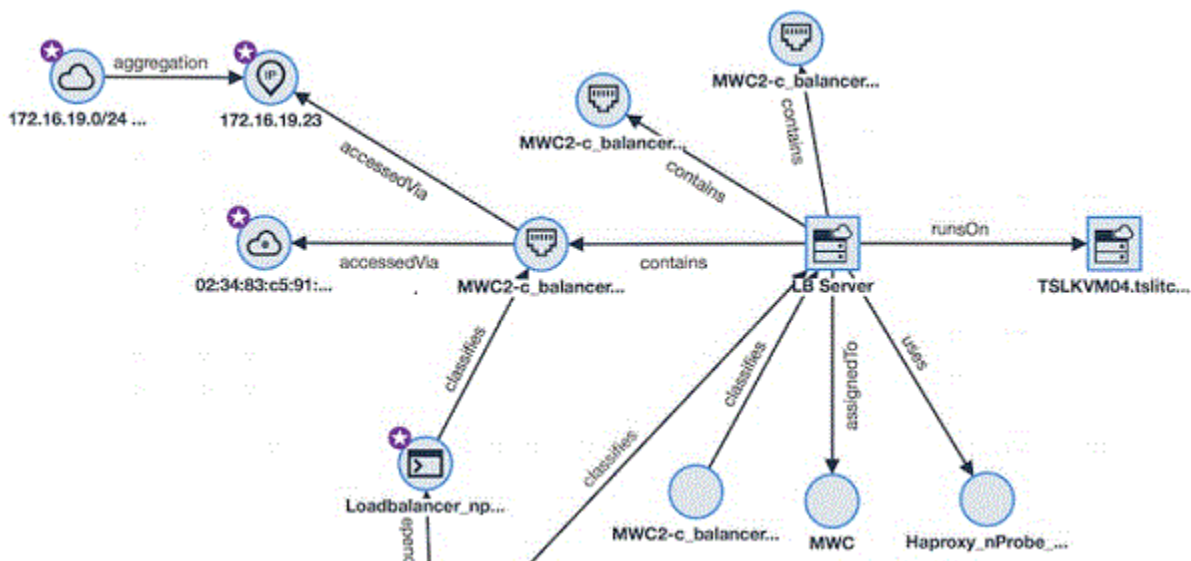
Netcool Agile Service Manager provides an Observer that makes extensive use of OpenStack's core and Heat APIs to build an end-to-end (or multi-domain) model.

## IT Virtualization OpenStack scenario

The following example of an OpenStack environment accessed through Netcool Agile Service Manager provides insights into any environment consisting of, for example, a combination of physical services and storage combined with virtualization from VMware (or similar).

### IT virtualization

The following figure depicts a Hypervisor running a VM (LB Server), that is assigned to a tenant, has classification data and a number of network interfaces. Each network interface has a MAC address, one or more IP addresses (v4 and/or v6) and each IP address is considered to be part of an IP subnet. It is also related to orchestration data from Heat that helps identify how the instance was created.



### Property use guidance

Although the topology service is flexible, you should follow the following guidelines when setting property values to ensure elements are appropriately represented:

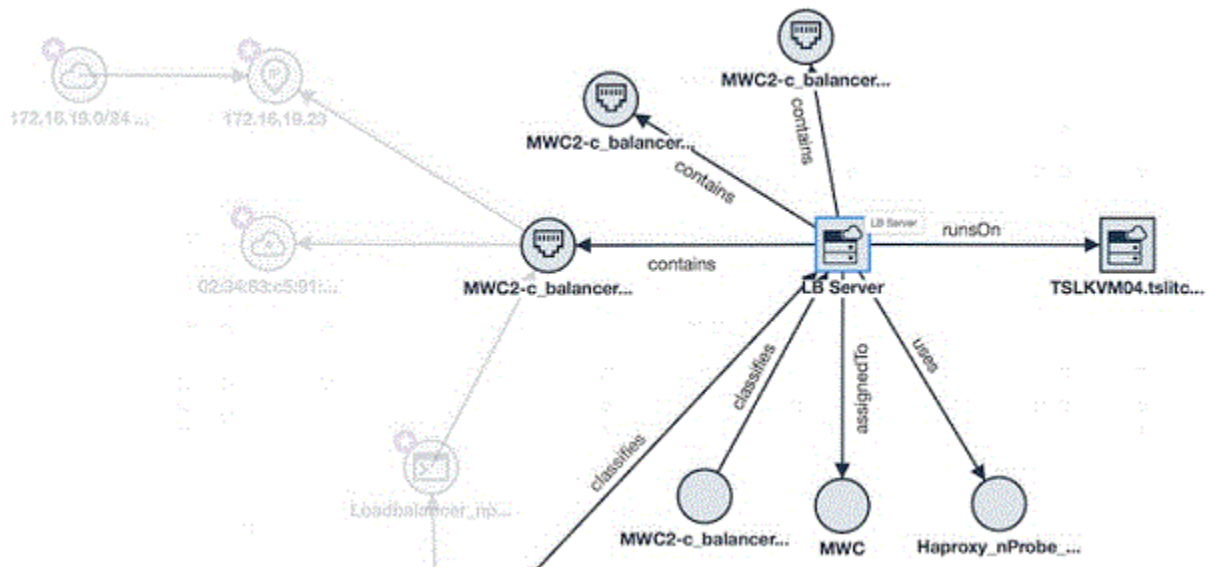
- Set entityType to one of the pre-defined entity types with an associated icon where possible.
- Set name and description to be user-friendly values wherever possible.
- Make use of the default generic properties to represent generally applicable characteristics of an element and to provide a basic degree of labeling and filtering. For a list of generic properties, see the following topic: [“Properties” on page 354](#)

### Model patterns - Part one

Stepping through each of the sections of the example of a multi-domain topology helps to identify reusable patterns.



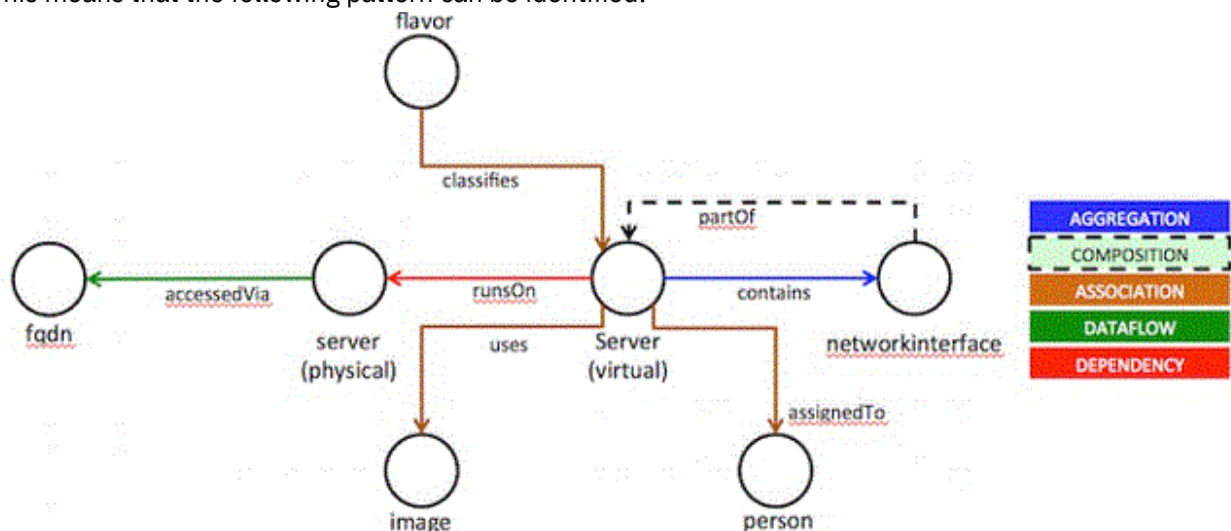
In the following figure, the Hypervisor 'TSLKVM04' is running a VM 'LB Server'.



The LB Server in this image is:

- Assigned to the 'MWC' tenant
- Contains three network interfaces
- Uses the Haproxy\_nProbe... image (that is, the OpenStack image)
- Is classified as an MWC2-c\_balancer (which is the OpenStack flavor)

This means that the following pattern can be identified:



### Usage tips

The associations shown for flavor, image and person are optional.

Network interfaces and other components of the device must be associated with it via a partOf relationship. The exception is if an IP or MAC address is known independently of any device, for example flow data would expose those but the device would be unknown.

The FQDN (hostname, short or full DNS name) is associated directly with the server in this case as nothing else is known about the Hypervisor.

**Note:** This is **not** shown in the topology fragment.

The partOf composition is not shown by the topology GUI, but must be created where the relationship between a component and a device is known. This is in addition to relationships such as contains, which the GUI will show.

### Hypervisor example JSON

```
{
  "_executionTime": 3,
  "createTime": 1501741671066,
  "name": "TSLKVM04",
  "uniqueId": "CYooventry_DC1:MWC/ComputeHost/TSLKVM04",
  "observedTime": 1501776262368,
  "_startedAt": "2017-08-03T06:27:51.066Z",
  "entityTypes": [
    "server"
  ],
  "beginTime": 1501741671066,
  "_id": "KNN6TCGhKyM4MCI6jooGWg",
  "_observedAt": "2017-08-03T16:04:22.368Z",
  "_modifiedAt": "2017-08-03T06:27:51.066Z",
  "_createdAt": "2017-08-03T06:27:51.066Z",
  "changeTime": 1501741671066,
  "matchTokens": [
    "Coventry_DC1:MWC/ComputeHost/TSLKVM04",
    "TSLKVM04"
  ]
}
```

**Note:** Many of the properties starting with \_ are internal (such as timestamps). Also:

- The uniqueId is a composite of a number of fields: data center, tenant, classname and name of the Hypervisor because the ID is a highly ambiguous integer.
- The name is the name of the instance from OpenStack.
- The entityType is set to 'server' to ensure correct classification and icon use.

### Virtual machine example JSON

```
{
  "instanceName": "LB server",
  "tenantId": "2f79c691570c4a598be386325ea01da8",
  "launchedAt": 1501741751194,
  "_executionTime": 4,
  "userId": "48c7cd25ad0842be8e9b84390de0e587",
  "imageName": "None Available",
  "availabilityZone": "nova",
  "createTime": 1501741733817,
  "flavorName": "m1.nano",
  "name": "LB server",
  "uniqueId": "1e35c68a-86b0-445f-9741-e581121a0577",
  "serverStatus": "active",
  "observedTime": 1501741752717,
  "_startedAt": "2017-08-03T06:29:12.717Z",
  "entityTypes": [
    "server",
    "vm"
  ],
  "beginTime": 1501741752717,
  "flavorId": "42",
  "vmState": "active",
  "_id": "3nDmTkKfrvNhZinSDCYHDw",
  "_observedAt": "2017-08-03T06:29:12.717Z",
  "_createdAt": 1501741733000,
  "_modifiedAt": "2017-08-03T06:29:12.717Z",
  "_created": "2017-08-03T06:28:53.817Z",
  "changeTime": 1501741752717,
  "matchTokens": [
    "1e35c68a-86b0-445f-9741-e581121a0577",
    "LB server"
  ]
}
```

**Note:** Many of the properties starting with \_ are internal (such as timestamps). Also:

- The uniqueId in this case is the UUID of the instance from OpenStack.
- The name is the name of the instance from OpenStack.

- The entityType is set to 'server' and 'vm' to ensure correct classification and icon use.
- The isVirtual boolean is set to true.

### Network interface example JSON

```
{
  "executionTime": 3,
  "isAdminStateUp": true,
  "createTime": 1501741717674,
  "name": "aab47c85-3110-401a-8bad-960b7c4bcd7b",
  "uniqueId": "aab47c85-3110-401a-8bad-960b7c4bcd7b",
  "observedTime": 1501741718855,
  "_startedAt": "2017-08-03T06:28:38.855Z",
  "entityTypes": [
    "networkinterface"
  ],
  "beginTime": 1501741718855,
  "isPortSecurityEnabled": true,
  "_id": "kkKpDH6yLn7cmNVX0ddImg",
  "_observedAt": "2017-08-03T06:28:38.855Z",
  "_modifiedAt": "2017-08-03T06:28:38.855Z",
  "_createdAt": "2017-08-03T06:28:37.674Z",
  "changeTime": 1501741718855,
  "matchTokens": [
    "aab47c85-3110-401a-8bad-960b7c4bcd7b"
  ]
}
```

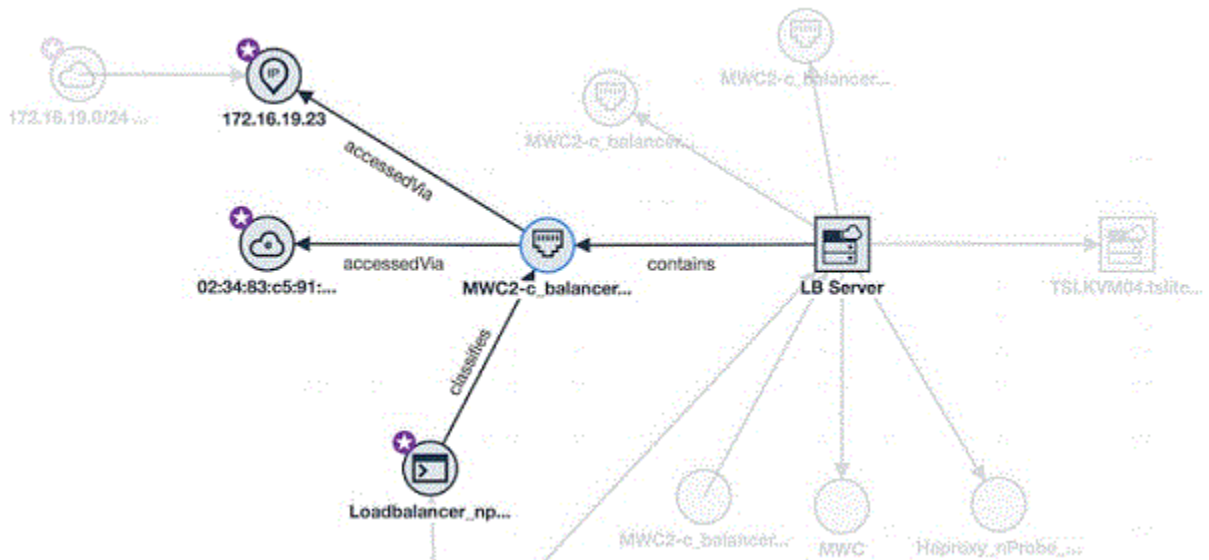
**Note:** As with the LB Server data, many of the properties are internal. Also:

- The uniqueId in this case is the UUID of the instance from OpenStack.
- The name is the name of the instance from OpenStack.
- The entityType is set to 'networkinterface' to ensure correct classification and icon use.
- The type of the interface (such as `ifType` from ITNM) is unknown.

### Model patterns - Part two

**Remember:** We are stepping through each of the sections of the example of a multi-domain topology to identify reusable patterns.

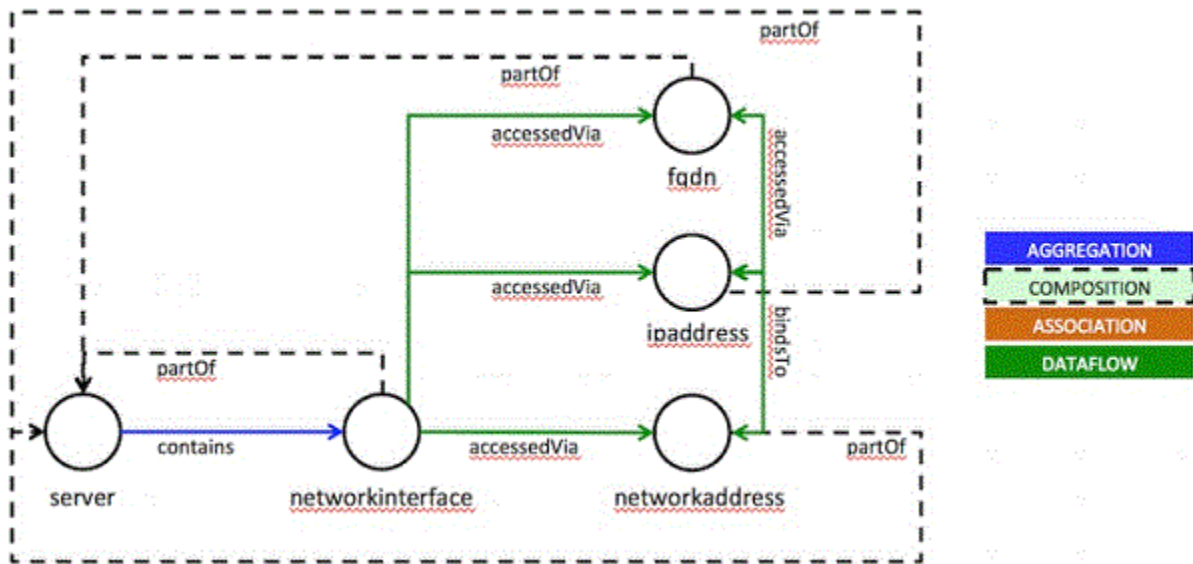
In the following figure, the LB Server VM contains a network interface.



The network interface in this image is:

- Accessed via an IP address and a MAC address
- Classified by a Heat orchestration resource

This means that the following pattern can be identified:



### Usage Tips

The Heat orchestration element and relationship is optional. Such things should be added if known to provide additional context.

The network interface must be contained by and **partOf** the device. Contains is more fine-grained and may reference intermediate cards (for example) within the containment hierarchy of the device, such as

```
json device--contains-->card--->contains--->port
```

If an IP address, FQDN or MAC address is known to be related to a specific device, then they must be associated with the device via a **partOf** relationship *in addition* to the **accessedVia** relationship.

If an IP address, FQDN or MAC address are known independently of a device, then no **partOf** relationship from them is necessary.

If an IP address is known to resolve to an FQDN, then relationships between them should be created to depict that one resolves to another and one accesses another (**accessedVia** shown in the example).

### IP address example JSON

```
{
  "_executionTime": 3,
  "createTime": 1501741717656,
  "name": "172.24.4.5",
  "uniqueId": "172.24.4.5",
  "ipNumber": 2887255045,
  "addressSpace": "Coventry_DC1:MWC",
  "observedTime": 1501741731565,
  "_startedAt": "2017-08-03T06:28:51.565Z",
  "entityTypes": [
    "ipaddress"
  ],
  "beginTime": 1501741731565,
  "version": "IPv4",
  "_id": "jPLc72DU-UvPeTQE_7YdPQ",
  "_observedAt": "2017-08-03T06:28:51.565Z",
  "_modifiedAt": "2017-08-03T06:28:51.565Z",
  "_createdAt": "2017-08-03T06:28:37.656Z",
  "changeTime": 1501741731565,
  "matchTokens": [
    "172.24.4.5",
    "IPv4:2887255045"
  ]
}
```

### Note:



Elements accessed via an IP should be related to it accordingly, e.g. a network interface and/or service or process.

If an FQDN is known to resolve to an IP address (and vice-versa), then they should be related.

If a MAC address is known to bind to an IP address and vice-versa, they should be related.

If an IP address, MAC address or FQDN are known to relate to a device, they should be considered partOf it; otherwise they are independent.

### IP subnet example JSON

```
{
  "uniqueId": "c009ff59-13b9-48dc-8863-cd0c75070d99",
  "name": "172.24.4.0/24 (public-subnet)",
  "entityTypes": [
    "subnet"
  ],
  "matchTokens": [
    "172.24.4.0/24 (public-subnet)",
    "172.24.4.0/24",
    "c009ff59-13b9-48dc-8863-cd0c75070d99"
  ],
  "_id": "u0sjaHcumtg5A4DR11fyAQ",
  "_references": [
    {
      "_id": "9duzx1-3o52g-ys5-47si0",
      "_edgeType": "contains",
      "_label": "aggregation",
      "_fromId": "u0sjaHcumtg5A4DR11fyAQ",
      "_toId": "1SJQs8JmYzDQ-wU0fvJbg",
      "_fromUniqueId": "c009ff59-13b9-48dc-8863-cd0c75070d99",
      "_toUniqueId": "172.24.4.12",
      "createTime": 1501838680418,
      "_observedAt": "2017-08-04T09:24:40.418Z",
      "createdAt": "2017-08-04T09:24:40.418Z",
      "beginTime": 1501838680418,
      "startedAt": "2017-08-04T09:24:40.418Z",
      "observedTime": 1501838680418
    },
    {
      "_executionTime": 11,
      "_modifiedAt": "2017-08-04T09:24:40.356Z",
      "isDhcpEnabled": false,
      "dnsNames": "None Available",
      "_observedAt": "2017-08-04T09:24:40.356Z",
      "gatewayIp": "172.24.4.1",
      "_startedAt": "2017-08-04T09:24:40.356Z",
      "observedTime": 1501838680356,
      "changeTime": 1501838680356,
      "ipVersion": "V4",
      "createTime": 1501838680356,
      "_createdAt": "2017-08-04T09:24:40.356Z",
      "cidr": "172.24.4.0/24",
      "networkId": "3e2b5d07-653a-4fc8-8224-45801d9d113f",
      "beginTime": 1501838680356,
      "allocationPools": "172.24.4.2-to-172.24.4.254"
    }
  ]
}
```

#### Note:

- The uniqueId in this case is the UUID of the subnet from OpenStack.
- The name in this case is set to CIDR notation plus the ID of the subnet.
- The entityType is set to subnet to ensure appropriate classification and icon usage.
- Some example relationships are shown: For example, an IP address that is part of the subnet is visible, and the subnet's use of an allocation pool is also depicted.

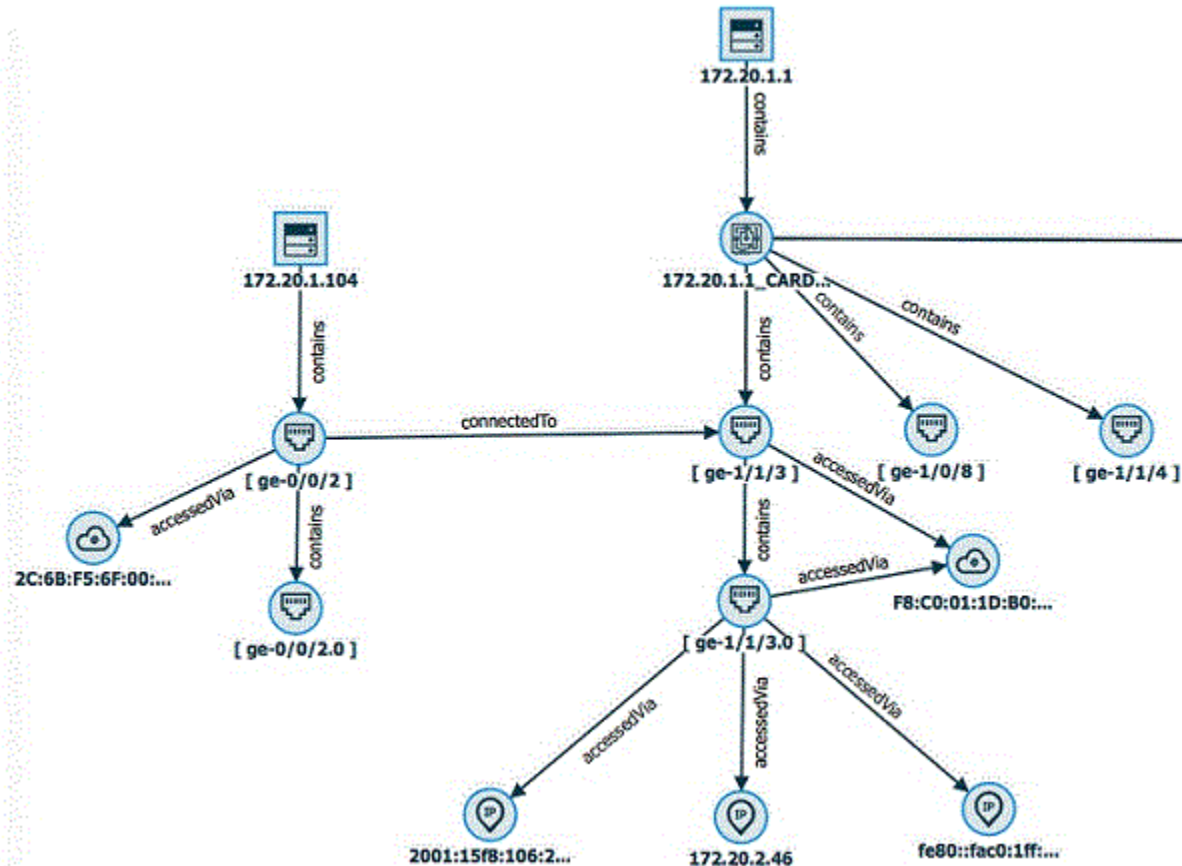
## Physical device recipe

The following example of an ITNM environment accessed through Netcool Agile Service Manager provides insights into any environment that makes use of physical network devices.

### Network physical devices

The following figure depicts two physical devices, 172.20.1.104 and 172.20.1.1, which are connected to each other.


Server 172.20.1.1 has three Gigabit Ethernet ports on the same card, one of which has a sub-interface with an `ifName` of 'ge-1/1/3.0'. That sub-interface shares the same MAC address as its physical parent and has two IPv6 addresses and one IPv4 address associated with it.




### Outline pattern

The topology service has an edge type of `partOf`, which is a member of the composition edge label family. See the [“Edge labels” on page 356](#) topic for more information.

The following image illustrates how this relationship is used to tie the card, interfaces and network addresses to the hosting device.



The `partOf` relationship is not shown as an explicit edge in the topology UI. However, it is used to determine which resources should be hidden in the host-to-host view, which in the context of this scenario would show just the hosts 172.20.1.104 and 172.20.1.1, as well as the `connectedTo` edge between them, as depicted in the following part of the image.



### Host example JSON

The following JSON extract is an example of the properties that you might choose to include when creating a host vertex. The properties which start with an underscore character are aliases for some of the read-only generic properties, and there are also a few read-write generic properties, such as `name` and `uniqueId`; however, the majority of the properties in this example are free from 'User' properties.

```
{
  "uniqueId": "NCOMS:172.20.1.1",
  "name": "172.20.1.1",
  "entityTypes": [
    "host"
  ],
  "_createdAt": "2017-03-03T16:02:40.845Z",
  "_observedAt": "2017-03-03T16:04:18.636Z",
  "_id": "y7EX0KrHud21CWySCyMsBg",
  "_href": "/1.0/topology/resources/y7EX0KrHud21CWySCyMsBg",
  "_nodeType": "resource",
  "_executionTime": 4,
  "_modifiedAt": "2017-03-03T16:02:40.845Z",
  "matchTokens": [
    "sbk-pe1-jrmx80.southbank.eu.test.lab"
  ],
  "sysObjectId": "1.3.6.1.4.1.2636.1.1.1.2.90",
  "entityChangeTime": "2017-03-03T14:13:17.000Z",
  "className": "JuniperMSeries",
  "services": "datalink(2) network(3)",
  "_startedAt": "2017-03-03T16:02:40.845Z",
  "manual": 0,
  "cdmAdminState": 0,
  "cdmType": 2,
  "sysDescription": "Juniper Networks, Inc. mx5-t internet router,"
}
```



```

kernel JUNOS 15.1F4.15, Build date: 2015-12-23 20:50:37 UTC
Copyright (c) 1996-2015 Juniper Networks, Inc.",
"sysName": "sbk-pe1-jrmx80.southbank.eu.test.lab",
"sysLocation": "The Mad Hatter Hotel 3-7 Stamford St
London SE1 9NY UK, -0.10499474,51.50711477",
"interfaceCount": 73,
"entityCreateTime": "2017-03-03T14:13:17.000Z",
"isIpForwarding": "forwarding",
"accessIPAddress": "172.20.1.1",
"entityDiscoveryTime": "2017-03-03T14:11:09.000Z",
"sysContact": "williamking@uk.ibm.com",
"tenantIds": [
  "Moa1dcmKHfx3dlyJnGm6JQ"
],
"accessProtocol": "IPv4"
}

```

**Note:** Some of the generic properties in this example are the following:

- The `uniqueId` in this case is a string, which uniquely identifies this resource to ITNM.
- The name is the name of the host and will be used in the UI.
- The `entityType` is set to 'host'.
- The `matchTokens` contains the FQDN name of the host.

### Network interface example JSON

```

{
  "uniqueId": "NCOMS:172.20.1.1[ ge-1/1/3.0 ]",
  "name": "[ ge-1/1/3.0 ]",
  "entityTypes": [
    "networkinterface"
  ],
  "_createdAt": "2017-03-03T16:03:03.889Z",
  "_observedAt": "2017-03-03T16:03:03.939Z",
  "_id": "v8aFVG6JoigYxTrlFSDkLw",
  "_href": "/1.0/topology/resources/v8aFVG6JoigYxTrlFSDkLw",
  "_nodeType": "resource",
  "_executionTime": 5,
  "operationalStatus": "started",
  "ifIndex": 537,
  "ifAdminStatus": "up",
  "_modifiedAt": "2017-03-03T16:03:03.940Z",
  "ifType": 53,
  "matchTokens": [
    "ge-1/1/3.0",
    "ifEntry.537"
  ],
  "ifName": "ge-1/1/3.0",
  "ifTypeString": "propVirtual",
  "connectorPresent": "false",
  "_startedAt": "2017-03-03T16:03:03.939Z",
  "speed": 10000000000,
  "mtu": 1500,
  "accessIpAddress": "172.20.2.46",
  "operationalDuplex": "FullDuplex",
  "promiscuous": false,
  "physicalAddress": "F8:C0:01:1D:B0:13",
  "ifDescription": "ge-1/1/3.0",
  "ifOperStatus": "up",
  "tenantIds": [
    "Moa1dcmKHfx3dlyJnGm6JQ"
  ],
  "accessProtocol": "IPv4"
}

```

**Note:** Some of the generic properties in this example are the following:

- The `uniqueId` in this case is a string, which uniquely identifies this resource to ITNM.
- The name is the name of the host and will be used in the UI.
- The `entityType` is set to 'networkinterface'.
- The `matchTokens` contains both the `ifName` and a string denoting the `ifIndex` in the `ifEntry` table.

## XML Gateway reference [deprecated from V. 1.1.6.1]

---

Agile Service Manager resource status can be generated from Netcool/OMNIbus events. The gateway must be configured to post XML events to the Event Observer. **For the latest probe and gateway documentation, see [“Configuring the probe and gateway services” on page 18](#)**

### Prerequisites

**Remember:** For up-to-date information on the version of the XML Gateway required, see [“Software requirements” on page 10](#).

#### Location

The default \$NCHOME install location is /opt/IBM/tivoli/netcool and the \$OMNIHOME install location is \$NCHOME/omnibus.

**Tip:** Export OMNIHOME as an environment variable, as it is repeatedly used in the scripts.

#### Standard gateway configuration

You must create a \$NCHOME/etc/omni.dat entry for the gateway (which in these examples is assumed to be **G\_ASM**):

```
[G_ASM]
{
  Primary: nasm-test1 4300
}
```

Run \$NCHOME/bin/nco\_igen

Generate a key file with nco\_keygen.

### Minimum XML gateway configuration requirements

For the XML gateway to post XML events to the Event Observer, you must edit the following files as a minimum:

#### XML Gateway properties file

If this file does not exist, you must create it in the \$OMNIHOME/etc directory.

For example, the XML Gateway properties file for a gateway called G\_ASM would be \$OMNIHOME/etc/G\_ASM.props

You define a number of properties, such as the name of the Netcool/OMNIbus Object Server, in the XML Gateway properties file.

You also reference the transformers XML file and the transport properties file here.

#### XML Gateway transport properties file

The file name of the XML Gateway transport properties file must match the one referenced in the XML Gateway properties file.

Here you define as a minimum the URL of the Event Observer to which XML events are posted, the batch header and footer, the maximum number of events in a single batch, the maximum waiting period before sending the events, and access to the HTTPS (TLS) truststore.

Default location and name: \$OMNIHOME/java/conf/asm\_httpTransport.properties

#### XML Gateway transformer XML file

The file name of the XML Gateway transformer XML file must match the one referenced in the XML Gateway properties file.

Here you define as a minimum the URL of the Event Observer to which XML events are posted.

Default location and name: \$OMNIHOME/java/conf/asm\_Transformers.xml

## Additional information

For more information on configuring the XML gateway, see the following section in the Netcool/OMNIbus Knowledge Center: [https://www.ibm.com/support/knowledgecenter/en/SSSHTQ/omnibus/gateways/xmlintegration/wip/concept/xmlgw\\_intro.html](https://www.ibm.com/support/knowledgecenter/en/SSSHTQ/omnibus/gateways/xmlintegration/wip/concept/xmlgw_intro.html)

For additional gateway configuration information, see the following IBM developerWorks discussion: <https://developer.ibm.com/answers/questions/256154/how-is-the-xml-message-bus-probe-and-gateway-confi.html>

### Important:

The gateway must be run with this environment variable set:

```
export JAVA_TOOL_OPTIONS=-Dhttps.protocols=TLSv1.2
```

## Gateway properties file

You create and/or edit the XML Gateway properties file: `$OMNIHOME/etc/<your_gateway>.props` and then define at least the following properties:

- The name of the Netcool/OMNIbus ObjectServer
- The name of the transport properties file
- The name of the transformer XML file

The following sample code is for a `$OMNIHOME/etc/G_ASM.props` gateway properties file, retrieving data from the AGG\_V ObjectServer via the G\_ASM gateway.

```
# Standard properties
Gate.Reader.Server : 'AGG_V'

# Properties defining XML messages over HTTP
Gate.MapFile: '$OMNIHOME/gates/xml/asm_xml.map'
Gate.StartupCmdFile: '$OMNIHOME/gates/xml/xml.startup.cmd'
Gate.Reader.TblReplicateDefFile: '$OMNIHOME/gates/xml/asm_xml.reader.tblrep.def'
Gate.XMLGateway.TransformerFile: '$OMNIHOME/java/conf/asm_transformers.xml'
Gate.XMLGateway.TransportFile: '$OMNIHOME/java/conf/asm_httpTransport.properties'
Gate.XMLGateway.TransportType: 'HTTP'

# The event observer requires the timestamp in this format, including the timezone
Gate.XMLGateway.DateFormat : 'yyyy-MM-dd\'T\'HH:mm:ssZ'

# To flush events to the gateway from the object server at 5s intervals, use this
Gate.Reader.IducFlushRate : 5

#####
# Security credentials required for the proxy service
# For full FIPS compliance, alter the java.security file as per the Omnibus documentation
#####
# This algorithm must be AES_FIPS
ConfigCryptoAlg: 'AES_FIPS'
# Secure key file generated using nco_keygen
ConfigKeyFile: '/opt/IBM/netcool/core/omnibus/etc/crypto.key'
```

**Important:** Do not use the `$OMNIHOME` variable in `ConfigKeyFile`.

### Example mapping (minimum fields required)

**Note:** The name of the gateway map file must match the one specified by the `Gate.MapFile` property in the gateway properties file.

```
CREATE MAPPING StatusMap
(
  'Agent'          = '@Agent',
  'AlertGroup'     = '@AlertGroup',
  'Class'          = '@Class',
  'Customer'       = '@Customer',
  'EventId'        = '@EventId',
  'Identifier'     = '@Identifier',
  'LastOccurrence' = '@LastOccurrence',
  'LocalPriObj'    = '@LocalPriObj',
```

```

'LocalRootObj'    =    '@LocalRootObj',
'Manager'         =    '@Manager',
'Node'           =    '@Node',
'NodeAlias'       =    '@NodeAlias',
'ServerName'      =    '@ServerName',
'ServerSerial'    =    '@ServerSerial',
'Severity'        =    '@Severity',
'StateChange'     =    '@StateChange',
'Summary'         =    '@Summary',
'Type'           =    '@Type'
);

```

## Gateway transport properties file

**Note:** The name of the gateway transport properties file must match the one specified by the `Gate.XMLGateway.TransportFile` property in the gateway properties file.

The gateway transport properties file (in these examples `$OMNIHOME/java/conf/asm_httpTransport.properties`) **must** specify at least the following properties:

- The user's authentication credentials.
- The batch header and footer, as well as the size and flush time, which specify the maximum number of events in a single XML batch file. For example:

```

batchHeader=<?xml version="1.0" encoding="UTF-8"?><tns:netcoolEventList xmlns:tns=
"http://item.tivoli.ibm.com/omnibus/netcool">
batchFooter=</tns:netcoolEventList>
bufferSize=10
bufferFlushTime=15

```

Without the **batchHeader** and **batchFooter** the gateway does not work.

- The maximum wait, in seconds, before sending the events
- The proxy service username and password. Encrypt the proxy service password (and optionally the username):

```
nco_aes_crypt -c AES_FIPS -k /opt/IBM/netcool/core/omnibus/etc/crypto.key <password>
```

- Add the username and password to the `asm_httpTransport.properties` file, for example:

```

username=asm
password=@44:9WxiH51VqMNHNYOLvoShaX001KwBLqXtGqtB/ZGCYPo=@

```

**Tip:** You only edit the java security file for FIPS compliance.

- For gateway access to the truststore, you need to complete the following steps:
  1. Create a truststore from the ASM CA certificate, and copy it to the Netcool/OMNIBus host (if different).

```

keytool -import \
        -alias asm-ca \
        -file $ASM_HOME/security/asm-ca.crt \
        -keystore gateway_truststore.jks \ <---- this file needs to go in the gw config
        -storetype JKS \
        -noprompt

```

While running the command, you are prompted for a password.

2. Add the truststore and password to the Gateway transport properties file. When completed, the gateway transport properties file should contain the following:

```

- trustStore=/fullPath/gateway_truststore.jks
- trustStorePassword={passwordGivenInPreviousStep}

```

```

$ grep ^trust /opt/IBM/netcool/core/omnibus/java/conf/asm_httpTransport.properties
trustStore=/opt/ibm/netcool/asm/security/gateway_truststore.jks
trustStorePassword=changeit

```

Optionally, you can also define:

- **httpTimeout**- The timeout is the amount of time in seconds that an HTTP client waits before aborting the connection.
- **retryLimit** - The retry limit is the number of times an HTTP client tries to connect.
- **retryWait** - The retry wait time is the amount of time (in seconds) an HTTP client waits before attempting to reconnect.

### Example

To use the following example, you modify the values for your host, username and encrypted password.

```
clientURL=http://<your host>/1.0/event-observer/netcool/list
batchHeader=<?xml version="1.0" encoding="UTF-8"?><tns:netcoolEventList
xmlns:tns="http://item.tivoli.ibm.com/omnibus/netcool">
batchFooter=</tns:netcoolEventList>
bufferSize=10
bufferFlushTime=15

username=<username>
password=<encryptedPassword>
```

## Gateway transformer XML file

The gateway transformer XML file (\$OMNIHOME/java/conf/asm\_Transformers.xml) must specify at least the URL of the Event Observer (endpoint), to which XML events are posted.

**Note:** The name of the gateway transformer XML file must match the one specified by the Gate.XMLGateway.TransformerFile property in the gateway properties file.

In the following example, the your host part of the URL specified in endpoint will be specific to your installation.

```
<?xml version="1.0" encoding="UTF-8"?>
<tns:transformers
  xmlns:tns="http://item.tivoli.ibm.com/omnibus/netcool/transformer"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">
  <tns:transformer name="netcoolEvents" type="northbound"
    endpoint="https://<your-asm-host>/1.0/event-observer/netcool/list"
    className="com.ibm.tivoli.netcool.integrations.transformer.XSLTTransformer">
    <tns:property name="xsltFilename" type="java.lang.String"
      value="{OMNIHOME}/java/conf/netcoolstripxmlheader.xsl"
      description="XSLT file for converting Netcool events to NC events"/>
  </tns:transformer>
</tns:transformers>
```

## State and status derived from Netcool/OMNIBus

The Event Observer derives the status of resources from individual fields of the event.

Table 73. General event state rules		
State	Meaning	Netcool/OMNIBus event mapping
closed	An active issue, may require attention	Active event with Severity > 0
open	Current state, working as expected	Cleared event with Severity = 0
clear	No longer relevant	Deleted event

Table 74. Use of Netcool/OMNIBus alerts.status event fields by Agile Service Manager	
alerts.status fields	Use by Agile Service Manager
Agent	Provider name for events generated from Agile Service Manager
AlertGroup	Type of Agile Service Manager event

*Table 74. Use of Netcool/OMNIBus alerts.status event fields by Agile Service Manager (continued)*

<b>alerts.status fields</b>	<b>Use by Agile Service Manager</b>
Class	45111 for Agile Service Manager events (should be mapped in alerts.conversions)
Customer	TenantId for events generated from Agile Service Manager
EventId	Status [type] for events generated from Agile Service Manager
Identifier	Determines the type of status, populating the status field
LastOccurrence	Used for the observedTime of open events
LocalPriObj	Resource lookup
LocalRootObj	Resource lookup
Manager	Observer name for events generated from Agile Service Manager
Node	Resource lookup
NodeAlias	Resource lookup
ServerName	Used to generate the unique eventId
ServerSerial	Used to generate the unique eventId
Severity	Severity 0 events represent a clear state
StateChange	Used for the observedTime of clear events
Summary	Used for the status description, shown in the UI
Type	Only Type 1 (Problem), Type 13 (Information) and Type 20 (ITMProblem) events are processed. All others are ignored.

*Table 75. Netcool/OMNIBus event data mapped onto Topology Service status*

<b>Topology Service status field</b>	<b>Netcool/OMNIBus source</b>
description	alerts.status Summary
eventId	<ServerName>/<ServerSerial>
eventManager	"netcool"
observedTime	closed - time event received by observer clear - alerts.status StateChange open - alerts.status ObservedTime
severity	alerts.status Severity
state	closed - deleted events clear - Severity 0 events open - none of the above
status	alerts.status Identifier

## Filtering events

**Note:** The name of the table replication definition file must match the one specified by the `Gate.Reader.TblReplicateDefFile` property in the gateway properties file.

To improve performance and prevent unnecessary events from being displayed in the topology viewer, you can filter events by type, and then refine these further by extracting only the fields of interest.

For example, you may want to include only the following event types:

- Problem (Type 1)
- Information (Type 13)
- ITMProblem (Type 20)

You may also want to remove all Netcool/OMNIBus self-monitoring events, that is, class 99999 events.

To include only problem (type 1), information (type 13), and ITMProblem (type 20) event types, and exclude Netcool/OMNIBus self-monitoring events (class 99999), use the following code:

```
asm_xml.reader.tblrep.def: |-
REPLICATE ALL FROM TABLE 'alerts.status'
USING MAP 'StatusMap'
FILTER WITH 'Type IN (1, 13, 20) AND Class != 99999';
```

## Probe for Message Bus reference [deprecated from V 1.1.6.1]

The Netcool/OMNIBus Message Bus probe must be configured to receive status from Agile Service Manager in JSON format via HTTP, and generate corresponding events in the Netcool/OMNIBus Event Viewer. These events are then fed back to the Agile Service Manager via the Netcool/OMNIBus XML Gateway, which updates the Agile Service Manager status via the Event Observer with the eventId. **For the latest probe and gateway documentation, see [“Configuring the probe and gateway services” on page 18](#)**

### Prerequisites

#### Location

The default OMNIHOME install location is `/opt/IBM/tivoli/netcool/omnibus`

### Probe for Message Bus configuration requirements

**Tip:** You can use the `topology_service_probe_list.sh` script (run without credentials) to list configured probes.

For the probe to receive status from Agile Service Manager, you must edit the following files as a minimum:

#### Probe properties file

Create and edit the probe property file.

In the following example, a non-default property file is used, which requires the `-propsfile` option when running the probe.

```
cd $OMNIHOME/probes/linuxx86/
cp message_bus.props asm_message_bus.props
```

Edit the `asm_message_bus.props` file as in the following example:

```
# Tell the probe to expect json over REST
MessagePayload : 'json'
TransformerFile : ''
TransportFile : '$OMNIHOME/java/conf/probe_httpTransport.properties'
TransportType : 'HTTP'

# Tell the probe how to parse the json payload, such that each member
of its variable-length
# _status array is processed as a separate message, with top-level
```

```

properties also included
MessageHeader           : 'json'
MessagePayload          : 'json._status'

# standard probe properties
MessageLog               : '$OMNIHOME/log/asm_probe.log'
RulesFile                : '$OMNIHOME/probes/linux2x86/
asm_message_bus.rules'

```

### Probe transport file

Create and edit the probe transport file.

The name of the probe transport file must match the name given in the probe properties, in this example 'probe\_httpTransport.properties'

Create a new file if necessary:

```

cd $OMNIHOME
cp java/conf/httpTransport.properties java/conf/probe_httpTransport.properties

```

This file needs to specify at least the URL of the probe, where it will accept JSON status as input; for example:

```
serverPort=http:18080
```

This port number is required when registering the probe URL.

### Probe rules file

You use the supplied [probe rules file](#) (\$ASM\_HOME/integrations/omnibus/asm\_message\_bus.rules).

The name of the probe rules file must match the name given in the probe properties, which in this example is 'asm\_message\_bus.rules'

The probe rules transform the input into events suitable for the Netcool/OMNIBus alerts.status table. The name of the file must be given as a probe property or command line option.

Create a new file if necessary, by copying and editing the supplied file:

```

cd $OMNIHOME/probes/linux2x86/
cp message_bus.rules asm_message_bus.rules

```

## Registering a probe URL to which status is exported

**Important OCP Note: Do not** configure an event sink when using the OCP version of Agile Service Manager, as this has been configured during installation.

To register a probe URL as an event sink, you run the following command:

```

cd $ASM_HOME/bin
topology_service_probe_register.sh -url http://{probeHost}:{probePort}

```

For example:

```
topology_service_probe_register.sh -url http://probe-host.ibm.com:18080
```

**Tip:** Docker comes with a default bridge network, docker0. This allows access to the docker host from within the docker container. By default, the docker host is available as 172.17.0.1, so an omnibus probe running on port 18080 of the Docker host could be configured for use via:

```
topology_service_probe_register.sh -url http://172.17.0.1:18080
```

The Agile Service Manager observer framework automatically emits status if an ASM\_EVENT\_SINK management artifact (stored in the topology service graph) has been configured, as depicted in the following example. The URL must include the probe's serverPort.

```

{
  "keyIndexName": "<your probe on your-system>",    <- any unique name

```



```

    "entityTypes": [
      "ASM_EVENT_SINK"                                     <- must be exactly this
    ],
    "tags": [
      "ASM_OBSERVER_CONFIG"                               <- must be exactly this
    ],
    "url": "http://<your.hostname>:18080"                 <- must match the probe's
  configured serverPort
}

```

## Additional information

For more information on configuring the probe, see the following section in the Netcool/OMNIBus Knowledge Center: [https://www.ibm.com/support/knowledgecenter/en/SSSHTQ/omnibus/probes/message\\_bus/wip/concept/messbuspr\\_intro.html](https://www.ibm.com/support/knowledgecenter/en/SSSHTQ/omnibus/probes/message_bus/wip/concept/messbuspr_intro.html)

For information on using the probe and the gateway as a single implementation, see the following section in the Netcool/OMNIBus Knowledge Center: [https://www.ibm.com/support/knowledgecenter/en/SSSHTQ/omnibus/probes/message\\_bus/wip/concept/messbuspr\\_integration\\_intro.html](https://www.ibm.com/support/knowledgecenter/en/SSSHTQ/omnibus/probes/message_bus/wip/concept/messbuspr_integration_intro.html)

## Example probe rules file

The following is an example of a rules file.

### probe.rules

```

#-----
# Licensed Materials - Property of IBM
# 5725-Q09
#
# (C) Copyright IBM Corporation 2017, 2018 All Rights Reserved.
# US Government Users Restricted Rights - Use, duplication
# or disclosure restricted by GSA ADP Schedule Contract
# with IBM Corp.
#-----

#-----
#
# Rules file intended to generate events from ASM, under Class 45111,
# from kafka topic itsm.status.json
#
#-----

if( match( @Manager, "ProbeWatch" ) )
{
  switch(@Summary)
  {
    case "Running ...":
      @Severity = 1
      @AlertGroup = "probestat"
      @Type = 2
    case "Going Down ...":
      @Severity = 5
      @AlertGroup = "probestat"
      @Type = 1
    case "Start resynchronization" | "Finish resynchronization":
      @Severity = 2
      @AlertGroup = "probestat"
      @Type = 13
    case "Connection to source lost":
      @Severity = 5
      @AlertGroup = "probestat"
      @Type = 1
    default:
      @Severity = 1
  }
  @AlertKey = @Agent
  @Summary = @Agent + " probe on " + @Node + ": " + @Summary
}
else
{
  #####
  # Input from ASM
}

```

```

#
# guaranteed json fields:
#
#     statusId          - the topology service status _id
#     resources.0._id   - the topology service resource _id
#
#     tenantId          - the topology service tenant _id
#     providerName      - the name of the resource provider
#
#     status            - the type of status affecting the resource
#     state             - the current resource state wrt this status
#
#     resources.0.uniqueId - provider's id for the resource
#
# optional json fields:
#
#     resources.0.name   - resource name, as shown in the UI (falls back to uniqueId)
#     observerName      - the name of the observer generating this status
#     description        - human readable description of the status (falls back to the
status type)
#     severity          - current severity of the status (defaults to 'indeterminate')
#     eventType         - type of event (defaults to 'ASM Status')
#     expiryTimeSeconds - optional expiryTime for the event
#
#####

#####
# @AlertGroup      # Purpose                                     # @Type      #
#####
# ASM Status      # Status about observed resources                         # Problem / Resolution #
# ASM Self Monitoring # Status about ASM itself                                   # Problem / Resolution #
# ASM Resource Creation # Identifies newly created ASM resources                 # Information          #
# ASM Resource Deletion # Identifies deleted ASM resources                 # Information          #
#####

@EventId      = $status
@Manager      = $observerName
@Customer     = $tenantId
@Agent        = $providerName
@NodeAlias    = $(resources.0.uniqueId)

@LocalNodeAlias = $(resources.0._id)
if ( exists($statusId) )
{
    @AsmStatusId    = $statusId
}

@Node            = $(resources.0.uniqueId)
if ( exists( $(resources.0.name) ) )
{
    # This is a user-friendly string identifying the resource
    @Node=$(resources.0.name)
}

@AlertGroup = "ASM Status"
if ( exists($eventType) )
{
    @AlertGroup=$eventType
}

@ExpireTime = 0
if ( exists($expiryTimeSeconds) )
{
    @ExpireTime=$expiryTimeSeconds
}

switch(@AlertGroup)
{
case "ASM Status" | "ASM Self Monitoring":
    switch($state)
    {
    case "open":
        @Type = 1
    case "clear":
        @Type = 2
    case "closed":
        @Type = 2
    default:
        @Type = 1
    }
case "ASM Resource Creation":
    @Type = 13
}

```

```

case "ASM Resource Deletion":
    @Type = 13
default:
    @Type = 13
}

@Severity=1
if ( exists($severity) )
{
    switch($severity)
    {
        case "clear":
            @Severity = 1
        case "indeterminate":
            @Severity = 1
        case "warning":
            @Severity = 2
        case "minor":
            @Severity = 3
        case "major":
            @Severity = 4
        case "critical":
            @Severity = 5
        default:
            @Severity = 1
    }
}

@Summary=$status
if ( exists($description) )
{
    @Summary=$description
}

if ( exists($observedTime) )
{
    # The Object Server uses seconds, whereas ASM uses milliseconds
    $seconds = regreplace($observedTime, "(.*)\d\d\d$", "\1")
    $milliseconds = regreplace($observedTime, ".*?(\d\d\d)$", "\1")
    @LastOccurrence = $seconds
    @LastOccurrenceUsec = int($milliseconds) * 1000
}

@AlertKey=$uniqueId + "->" + $status + "->" + @Agent + "->" + @Customer
@Identifier=@AlertKey + @Type

@Class = 45111

```

## Status service reference

This topic contains reference information for the Agile Service Manager status service.

### IBM Netcool Operations Insight integration

#### Remember:

**IBM Netcool Operations Insight Version 1.6.0.1:** Up to (and including) this version, Agile Service Manager required separate installations of the IBM Tivoli Netcool/OMNIBus XML Gateway for Message Bus and the IBM Tivoli Netcool/OMNIBus Message Bus probe. These worked together with the Event Observer to integrate Agile Service Manager with Netcool/OMNIBus.

**IBM Netcool Operations Insight Version 1.6.0.2:** From this version onwards, you deploy containerized and pre-configured versions of the gateway and probe while installing Agile Service Manager, and the previous versions of the probe and gateway are deprecated.

**IBM Netcool Operations Insight Version 1.6.1:** From this version onwards, the Event Observer has been deprecated and its functions are now performed by the new status service.

#### Matching Netcool/OMNIBus events to resources (on-prem and OCP)

Resources are matched to events based on their matchTokens property.

You can use a (comma-separated) list of Netcool/OMNIBus ObjectServer alerts.status field names to identify top-level resources.

If you **do not** supply a list, this defaults to "Node, NodeAlias".

If you **do** supply a list, you must list all alerts.status field names.

**For on-prem**, set the environment variable NETCOOL\_EVENT\_NODE\_FIELDS in the \$ASM\_HOME/etc/nasm-status.yml file.

**For OCP**, edit the status.nodeFields values.

**Tip:** You can define extra event properties to be added to the status displayed in the Topology Viewer using the extra\_status\_fields property, for example 'LocalNodeAlias'. You can then define topology viewer status tools that reference these. These fields must be passed through by the gateway (as configured in \$ASM\_HOME/integrations/omnibus/kafka/gateway/field\_filter.map).

## Additional information

For more information on the merge service and the creation of matchTokens, see the following topic: [“Defining rules” on page 309](#)

For more information on configuring the XML gateway, see the following section in the Netcool/OMNIBus Knowledge Center: [https://www.ibm.com/support/knowledgecenter/en/SSSHTQ/omnibus/gateways/xmlintegration/wip/concept/xmlgw\\_intro.html](https://www.ibm.com/support/knowledgecenter/en/SSSHTQ/omnibus/gateways/xmlintegration/wip/concept/xmlgw_intro.html)

For additional gateway configuration information, see the following IBM developerWorks discussion: <https://developer.ibm.com/answers/questions/256154/how-is-the-xml-message-bus-probe-and-gateway-confi.html>

## State and status derived from Netcool/OMNIBus

The status of resources is derived from individual fields of the event.

Table 76. General event state rules	
Agile Service Manager State	Netcool/OMNIBus event mapping
open	Active event with Severity > 0
closed	Closed event with Severity = 0
clear	Deleted event

Table 77. Use of Netcool/OMNIBus alerts.status event fields by Agile Service Manager	
alerts.status fields	Use by Agile Service Manager
Agent	Provider name for events generated from Agile Service Manager
AlertGroup	Type of Agile Service Manager event
AsmStatusId	Identifies the unique ID of the topology service status node. Used as a performance improvement to bypass resource lookup.
Class	45111 for Agile Service Manager events (should be mapped in alerts.conversions)
Customer	TenantId for events generated from Agile Service Manager
EventId	Status [type] for events generated from Agile Service Manager
Identifier	Determines the type of status, populating the status field
LastOccurrence	Used for the observedTime of open events
LastOccurrenceUSec	Used for sub-second event clearing by the updated generic_clear automation, available in \$ASM_HOME/integrations/omnibus/updated-generic-clear.sql

<i>Table 77. Use of Netcool/OMNIBus alerts.status event fields by Agile Service Manager (continued)</i>	
<b>alerts.status fields</b>	<b>Use by Agile Service Manager</b>
LocalPriObj	Resource lookup
LocalRootObj	Resource lookup
Manager	Observer name for events generated from Agile Service Manager
Node	Resource lookup
NodeAlias	Resource lookup
ServerName	Used to generate the unique eventId
ServerSerial	Used to generate the unique eventId
Severity	Severity 0 events represent a clear state
StateChange	Used for the observedTime of clear events
Summary	Used for the status description, shown in the UI
Type	Only Type 1 (Problem), Type 13 (Information) and Type 20 (ITMProblem) events are processed. All others are ignored.

<i>Table 78. Netcool/OMNIBus event data mapped onto Topology Service status</i>	
<b>Topology Service status field</b>	<b>Netcool/OMNIBus source</b>
description	alerts.status Summary
eventId	<ServerName>/<ServerSerial>
eventManager	"netcool"
observedTime	closed - time event received by observer clear - alerts.status StateChange open - alerts.status ObservedTime
severity	alerts.status Severity
state	closed - deleted events clear - Severity 0 events open - none of the above
status	alerts.status Identifier

## Topology viewer reference

This reference topic describes the Netcool Agile Service Manager UI screen elements and associated functionality.

### Search page

When you access Agile Service Manager, the **Search** page is displayed immediately.

### Search for a resource

The seed resource of the topology visualization.

You define the seed resource around which a topology view is rendered using the **Search for a resource** field. As you type in a search term related to the resource that you wish to find, such as name or server, a drop-down list is displayed with suggested search terms that exist in the topology service.

If the resource that you wish to find is unique and you are confident that it is the first result in the list of search results, then instead of selecting a result from the suggested search terms, you can choose to click the shortcut in the **Suggest** drop-down, which will render and display the topology for the closest matching resource.

If you select one of the suggested results, the **Search Results** page is displayed listing possible resource results.

The Results are listed under separate **Resources** and **Topologies** tabs.

#### **Note:**

##### **Deleted resources**

You can select the **Include deleted resources** toggle in order to include deleted resources in the Results page (unless an administrator has disabled that feature).

Deleted resources are marked as such, and the date and time that they were deleted is displayed.

If you click **View topology**, the resource is displayed in the Topology Viewer with the history timeline open, and set to 1ms before deletion. When closed, the topology is cleared.

When you search for resources with a previous point in time set in the history timeline, deleted resources are **not** returned and the toggle to include them is not available.

Deleted templates or defined topologies are **not** included in any search results.

##### **Defined topologies**

Defined topologies **must** be defined by an administrator user before they are listed.

If you are an administrator defining topology templates in the **Topology template builder**, search results are listed under separate **Resources** and **Templates** tabs.

To add a defined topology search result to the collection of topologies accessible in the Topology Dashboard, tag it as a favorite by selecting the **star** icon next to it.

For each result, the name, type and other properties stored in the Elasticsearch engine are displayed.

If a status other than clear exists for a search result, the maximum severity is displayed in the information returned, and a color-coded information bar above each result displays all non-clear statuses (in proportion).

You can expand a result in order to query the resource or defined topology further and display more detailed, time-stamped information, such as its state and any associated severity levels, or when the resource was previously updated or replaced (or deleted).

You can click the **View Topology** button next to a result to render the topology.

##### **Defined topology restriction:**

- When you load a predefined topology, it is displayed in a 'defined topology' version of the Topology Viewer, which has restricted functionality. You are unable to follow its neighbors, or change its hops, or make use of its advanced filters.
- You can recenter the defined topology from the context menu, which loads it in the Topology Viewer with all its standard functionality.

## **Navigation toolbar**

The navigation toolbar is displayed at the top of the Topology Viewer and provides access to the following functionality, or information.

### **Topology Search**

If you conduct a resource search from the navigation toolbar with a topology already loaded, the search functionality searches the loaded topology as well as the topology database.

As you type in a search term, a drop-down list is displayed that includes suggested search results from the displayed topology listed under the **In current view** heading.

If you hover over a search result in this section, the resource is highlighted in the topology window.

If you click on a search result, the topology view zooms in on that resource and closes the search.

### No. Hops

The number of relationship hops to visualize from the seed resource, with the default set at 'one'.

You define the number of relationship hops to be performed, which can be from zero to four, unless this setting has been customized. See the [“Defining advanced topology settings” on page 285](#) topic for more information on customizing the maximum hop count.

### Type of Hop

The type of graph traversal used.

The options are:

#### Element to Element hop type

This type performs the traversal using all element types in the graph.

#### Host to Host hop type

This type generates a view showing host to host connections.

#### Element to Host hop type

This type provides an aggregated hop view like the Host to Host type, but also includes the elements that are used to connect the hosts.

**Tip:** The URL captures the hopType as 'e2h'. When launching a view using a direct URL, you can use the hopType=e2h URL parameter.

### Filter toggle

Use this icon to display or hide the filter toolbar. You can filter resources that are displayed in the topology, or set filters before rendering a topology to prevent a large, resource-intensive topology from being loaded.

If a filter has been applied to a displayed topology, the text 'Filtering applied' is displayed in the status bar at the bottom of the topology.

### Render

This performs the topology visualization action, rendering the topology based on the settings in the navigation toolbar.

Once rendered, the topology will refresh on a 30 second interval by default. You can pause the auto-update refresh, or select a custom interval.

**Tip:** The UI can time out if a large amount of data is being received. See the [timeout troubleshooting](#) section in the following topic for information on how to address this issue, if a timeout message is displayed: [“Rendering \(visualizing\) a topology” on page 253](#)

### Sharing options

You can share a topology either by obtaining a direct URL linking to the topology view, or by exporting a view of the topology as an image.

#### Obtain Direct URL

Open the **Sharing options** drop-down menu, and then use the **Obtain Direct URL** option to display the **Direct Topology URL** dialog.

The displayed URL captures the current topology configuration, including layout type (layout orientation is not tracked).

Click **Copy** to obtain a direct-launch URL string, then click **Close** to return to the previous screen.

Use the direct-launch URL for quick access to a given topology view within DASH.

**Tip:** You can share this URL with all DASH users with the required permissions.

#### Export as PNG / SVG

You can share a snapshot of a topology in either PNG or SVG format, for example with someone who does not have DASH access.

Open the **Sharing options** drop-down menu, and then use either the **Export as PNG** or the **Export as SVG** option.

Specify a name and location, then click **Save** to create a snapshot of your topology view.

You can now share the image as required.

#### **Additional actions > View System Health**

Open the **Additional actions** drop-down menu, and then use the **View System Health** option to access your Netcool Agile Service Manager deployment's system health information.

#### **Additional actions > Edit User Preferences**

Open the **Additional actions** drop-down menu, and then use the **Edit User Preferences** option to access the **User Preferences** window. Click **Save**, then **Close** when done.

You can customize the following user preferences to suit your requirements:

#### **Updates**

##### **Default auto refresh rate (seconds)**

The rate at which the topology will be updated.

The default value is 30.

You must reopen the page before any changes to this user preference take effect.

##### **Maximum number of resources to load with auto refresh enabled**

When the resource limit set here is reached, auto-refresh is turned off.

The maximum value is 2000, and the default is set to 500.

**Tip:** If you find that the default value is too high and negatively impacts your topology viewer's performance, reduce this value.

#### **Auto render new resources**

Enable this option to display new resources at the next scheduled or ad-hoc refresh as soon as they are detected.

#### **Remove deleted topology resources**

Enable this option to remove deleted resources at the next scheduled or ad-hoc refresh.

#### **Layout**

Set **Default layout type** including the layout orientation for some of the layout types. You can also configure a default layout in User Preferences.

You can choose from a number of layout types, and also set the orientation for layouts 4, 6, 7 and 8.

**Tip:** A change to a layout type is tracked in the URL (layout orientation is not tracked). You can manually edit your URL to change the layout type display settings.

The following numbered layout types are available:

##### **Layout 1**

A layout that simply displays all resources in a topology without applying a specific layout structure.

##### **Layout 2**

A circular layout that is useful when you want to arrange a number of entities by type in a circular pattern.

##### **Layout 3**

A grouped layout is useful when you have many linked entities, as it helps you visualize the entities to which a number of other entities are linked. This layout helps to identify groups of interconnected entities and the relationships between them.

##### **Layout 4**

A hierarchical layout that is useful for topologies that contain hierarchical structures, as it shows how key vertices relate to others with peers in the topology being aligned.



### Layout 5

A force-directed (or 'peacock') layout is useful when you have many interlinked vertices, which group the other linked vertices.

### Layout 6

A planar rank layout is useful when you want to view how the topology relates to a given vertex in terms of its rank, and also how vertices are layered relative to one another.

### Layout 7

A rank layout is useful when you want to see how a selected vertex and the vertices immediately related to it rank relative to the remainder of the topology (up to the specified amount of hops). The root selection is automatic.

For example, vertices with high degrees of connectivity outrank lower degrees of connectivity. This layout ranks the topology automatically around the specified seed vertex.

### Layout 8

A root rank layout similar to layout 7, except that it treats the selected vertex as the root. This layout is useful when you want to treat a selected vertex as the root of the tree, with others being ranked below it.

Ranks the topology using the selected vertex as the root (root selection: Selection)

### Layout orientation

**For layouts 4, 6, 7 and 8**, you can set the following layout orientations:

- Top to bottom
- Bottom to top
- Left to right
- Right to left

### Misc

#### Information message auto hide timeout (seconds)

The number of seconds that information messages are shown for in the UI.

The default value is 3.

**Tip:** If you are using a screen reader, it may be helpful to increase this value to ensure that you do not miss the message.

#### Screen reader support for graphical topology

You can enable the display of additional Help text on screen elements, which can improve the usability of screen readers.

You must reopen the page before any changes to this user preference take effect.

#### Enhanced client side logging, for problem diagnosis

If enabled, additional debug output is generated, which you can use for defect isolation.

**Tip:** Use this for specific defect hunting tasks, and then disable it again. If left enabled, it can reduce the topology viewer's performance.

You must reopen the page before any changes to this user preference take effect.

## Visualization toolbar

The Topology Viewer visualization toolbar is displayed below the navigation toolbar, and provides you with access to functionality to manipulate the topology visualization.

### Select tool submenu

When you hover over the Select tool icon, a submenu is displayed from which you can choose the **Select**, **Pan** or **Zoom Select** tool.

#### Select tool

Use this icon to select individual resources using a mouse click, or to select groups of resources by creating a selection area (using click-and-drag).

**Pan tool**

Use this icon to pan across the topology using click-and-drag on a blank area of the visualization panel.

**Zoom Select tool**

Use this icon to zoom in on an area of the topology using click-and-drag.

**Zoom In**

Use this icon to zoom in on the displayed topology.

**Zoom Out**

Use this icon to zoom out of the displayed topology.

**Zoom Fit**

Use this icon to fit the entire topology in the current view panel.

**Overview Toggle**

Use this icon to create the overview mini map in the bottom right corner.

The mini map provides an overview of the entire topology while you zoom in or out of the main topology. The mini map displays a red rectangle to represent the current topology view.

**Layout**

Use this icon to recalculate, and then render the topology layout again.

You can choose from a number of layout types and orientations.

**Layout 1**

A layout that simply displays all resources in a topology without applying a specific layout structure.

**Layout 2**

A circular layout that is useful when you want to arrange a number of entities by type in a circular pattern.

**Layout 3**

A grouped layout is useful when you have many linked entities, as it helps you visualize the entities to which a number of other entities are linked. This layout helps to identify groups of interconnected entities and the relationships between them.

**Layout 4**

A hierarchical layout that is useful for topologies that contain hierarchical structures, as it shows how key vertices relate to others with peers in the topology being aligned.

**Layout 5**

A peacock layout is useful when you have many interlinked vertices, which group the other linked vertices.

**Layout 6**

A planar rank layout is useful when you want to view how the topology relates to a given vertex in terms of its rank, and also how vertices are layered relative to one another.

**Layout 7**

A rank layout is useful when you want to see how a selected vertex and the vertices immediately related to it rank relative to the remainder of the topology (up to the specified amount of hops). The root selection is automatic.

For example, vertices with high degrees of connectivity outrank lower degrees of connectivity. This layout ranks the topology automatically around the specified seed vertex.

**Layout 8**

A root rank layout similar to layout 7, except that it treats the selected vertex as the root. This layout is useful when you want to treat a selected vertex as the root of the tree, with others being ranked below it.

Ranks the topology using the selected vertex as the root (root selection: Selection)

**Layout orientation**

For layouts 4, 6, 7 and 8, you can set the following layout orientations:

- Top to bottom
- Bottom to top
- Left to right
- Right to left

### History toggle

Use this to open and close the Topology History toolbar. The topology is displayed in history mode by default.

### Configure Refresh Rate

When you hover over the **Refresh Rate** icon, a submenu is displayed from which you can configure the auto-update refresh rate.

You can pause the topology data refresh, or specify the following values: 10 seconds, thirty seconds (default), one minute, or five minutes.

### Resource display conventions

**Deleted:** A minus icon shows that a resource has been deleted since last rendered.

Displayed when a topology is updated, and in the history views.

**Added:** A purple plus (+) icon shows that a resource has been added since last rendered.

Displayed when a topology is updated, and in the history views.

**Added (neighbors):** A blue asterisk icon shows that a resource has been added using the 'get neighbors' function.

## Topology visualization panel

The main panel under the visualization toolbar displays the topology.

The displayed topology consists of resource nodes and the relationship links connecting the resources. You can interact with these nodes and links using the mouse functionality.

### Dragging a node

Click and drag a node to move it.

### Selecting a node

Selection of a node highlights the node, and emphasizes its first-order connections by fading all other resources.

### Context menu (right-click)

You open the context menu using the right-click function. The context menu provides access to the resource-specific actions you can perform.

For resource entities, you can perform the following:

#### Resource Details

When selected, displays a dialog that shows all the current stored properties for the specified resource in tabular and raw format.

When selected while viewing a topology history with Delta mode **On**, the properties of the resource at both the reference time and at the delta time are displayed.








#### Resource Status

If statuses related to a specific resource are available, the resource will be marked with an icon depicting the status severity level, and the Resource Status option will appear in the resource context menu.

When selected, Resource Status displays a dialog that shows the time-stamped statuses related to the specified resource in table format. The Severity and Time columns can be sorted, and the moment that Resource Status was selected is also time-stamped.

In addition, if any status tools have been defined, the status tool selector (three dots) is displayed next to the resource's statuses. Click the status tool selector to display a list of any status tools that have been defined, and then click the specific tool to run it. Status tools are only displayed for the states that were specified when the tools were defined.

The **severity** of a status ranges from 'clear' (white tick on a green square) to 'critical' (white cross on a red circle).

Table 79. Severity levels	
Icon	Severity
	clear
	indeterminate
	information
	warning
	minor
	major
	critical

### Comments

When selected, this displays any comments recorded against the resource.

By default, resource comments are displayed by date in ascending order. You can sort them in the following way:

- Oldest first
- Newest first
- User Id (A to Z)
- User Id (Z to A)

Users with the inasm\_operator role can view comments, but not add any. Users with inasm\_editor or inasm\_admin roles can also add new comments. See the [“Configuring DASH user roles” on page 32](#) topic for more information on assigning user roles.

To add a new comment, enter text into the New Comment field, and then click **Add Comment** to save.

### Get Neighbors

When selected, opens a menu that displays the resource types of all the neighboring resources. Each resource type lists the number of resources of that type, as well as the maximum severity associated with each type.

You can choose to get all neighbors of the selected resource by adding a **Filter tag**, or only the neighbors of a specific type by adding a **Filter neighbor tag**. This lets you expand the topology in controlled, incremental steps.

Selecting **Get Neighbors** overrides any existing filters.

You can **Undo** the last neighbor request made.

### Follow Relationship

When selected, opens a menu that displays all adjacent relationship types.

Each relationship type lists the number of relationships of that type, as well as the maximum severity associated with each type.

You can choose to follow all relationships, or only the neighbors of a specific type.

### Show last change in timeline

When selected, will display the history timeline depicting the most recent change made to the resource.

**Show first change in timeline**

When selected, will display the history timeline depicting the first change made to the resource.

**Find path**

When selected, opens the **Path tool** dialog, from where you can define the criteria for the system to find the shortest path between two resources in your topology. See the [“Using the Path tool” on page 267](#) topic for more information.

**Recenter View**

When selected, this updates the displayed topology with the specified resource as seed.

**Information bar**

A section at the bottom of the screen displays the current status of the rendered topology.

A timestamp on the left of the information bar indicates the time of the most recent refresh. If two time periods are being compared, both will be indicated.

Additional information on the right describes the number of resources rendered, their relationships, whether they were added or removed since the last refresh, and whether a filter has been applied.

**Filter toolbar**

Open and close the Filter toolbar using the **Filter** toggle in the Navigation toolbar (on the top). When you have filtered your topology, click **Close** to remove the toolbar from view.

The Filter toolbar is displayed as a panel on the right-hand side of the page, and consists of a **Simple** and an **Advanced** tab. If selected, each tab provides you with access to lists of resource types and relationship types. Any types that have been excluded from view in order to optimize performance when rendering topologies are listed here. The seed resource of a topology **can not** be excluded.

**Note:** You can override the exclusion using the context menu functionality.

- **If you are filtering a topology before rendering it:** All resource types and relationship types are displayed, except those that have been excluded. After rendering the topology, you can toggle the **Show all types** switch so that only types relevant to your topology are displayed.
- **If you are filtering a topology already displayed in the viewer:** Only types relevant to your topology are displayed, for example **host**, **ipaddress**, or **operatingsystem**. You can toggle the **Show all types** switch so that all types are listed (apart from the excluded types).

**Simple tab**

When you use the Simple tab to filter out resource or relationship types, all specified types are removed from view, including the seed resource.

It **only** removes the resources matching that type, leaving the resources below, or further out from that type, based on topology traversals.

By default, all types are **On**. Use the **Off** toggle to remove specific types from your view.

**Advanced tab**

The Advanced tab performs a server-side topology-based filter action.

It can exclude or include types:

- **Exclude** removes the resources matching that type, **as well as** all resources below that type. However, the seed resource is **not** removed from view, even if it is of a type selected for removal.
- **Include** displays the selected types in the topology.

**Tips**

**Reset or invert all filters:** Click **Reset** to switch all types back on, or click **Invert** to invert your selection of types filtered.

**Hover to highlight:** When a topology is displayed, hover over one of the filtering type options to highlight them in the topology.

## Topology History toolbar

Open and close the Topology History toolbar using the **History** button in the Topology Visualization toolbar (on the left). You can hide the Topology History toolbar by clicking **Close**, which also returns the topology to update mode.

### Update mode

The topology is displayed in update mode by default with Delta mode set to **Off**.

While viewing the timeline in update mode with Delta mode set to **On**, any changes to the topology history are displayed on the right hand side of the timeline, with the time pins moving apart at set intervals. By clicking **Render**, you reset the endpoint to 'now' and the pins form a single line again.

While viewing the timeline in update mode with Delta mode set to **Off**, only a single pin is displayed.

### Delta mode

You toggle between delta mode **On** and **Off** using the Delta switch above the topology.

When Delta mode is **On** with Update mode also **On**, differences in topology are displayed via purple plus or minus symbols next to the affected resource.

When Delta mode is **On** with History mode **On** (that is, Update mode set to **Off**), you can compare two time points to view differences in topology. Historical change indicators (blue dots) are displayed next to each affected resource.

**Note:** For efficiency reasons, historical change indicators are only displayed for topologies with fifty or fewer resources. You can reduce (but not increase) this default by changing the Historical Change Threshold as described in [“Defining advanced topology settings” on page 285](#).

### Lock time pin

Click the **Lock** icon on a time pin's head to lock a time point in place as a reference point, and then use the second time slider to view topology changes.

### Compare resource properties

Click **Resource Properties** on a resource's context menu to compare the resource's data at the two selected time points. You can view and compare the resource's property names and values in table format, or raw JSON format.

### History timeline

You open the Topology History toolbar using the **History** toggle in the Topology Visualization toolbar (on the left).

**Note:** Resources with a large amount of history may take longer to display. To avoid timeout issues, you can [adjust the default timeout values](#).

You use the time pins to control the topology shown. When you move the pins, the topology updates to show the topology representation at that time.

While in delta mode you can move both pins to show a comparison between the earliest pin and the latest. The timeline shows the historic changes for a single selected resource, which is indicated in the timeline title. You can lock one of the time pins in place to be a reference point.

When you first display the history timeline, coach marks (or tooltips) are displayed, which contain helpful information about the timeline functionality. You can scroll through these, or switch them off (or on again) as required.

To view the timeline for a different resource, you click on it, and the heading above the timeline changes to display the name of the selected resource. If you click on the heading, the topology centers (and zooms into) the selected resource.

The history timeline is displayed above a secondary time bar, which displays a larger time segment and indicates how much of it is depicted in the main timeline. You can use the jump buttons to move back and forth along the timeline, or jump to the current time.

You can use the time picker, which opens a calendar and clock, to move to a specific second in time.

To view changes made during a specific time period, use the two time sliders to set the time period. You can zoom in and out to increase or decrease the granularity using the + and - buttons on the right, or by double-clicking within a time frame. The most granular level you can display is an interval of one

second. The granularity is depicted with time indicators and parallel bars, which form 'buckets' that contain the recorded resource change event details.

The timeline displays changes to a resource's state, properties, and its relationships with other resources. These changes are displayed through color-coded bars and dash lines, and are elaborated on in a tooltip displayed when you hover over the change. You can exclude one or more of these from display.

#### **Resource state changes**

The timeline displays the number of state changes a resource has undergone.

#### **Resource property changes**

The timeline displays the number of times that resource properties were changed.

Each time that property changes were made is displayed as one property change event regardless of whether one or more properties were changed at the time.

#### **Resource relationship changes**

The number of relationships with neighboring resources are displayed, and whether these were changed.

The timeline displays when relationships with other resources were changed, and also whether these changes were the removal or addition of a relationship, or the modification of an existing relationship.

## **Update manager**

If auto-updates have been turned off, the Update Manager informs you if new resources have been detected. It allows you to continue working with your current topology until you are ready to integrate the new resources into the view.

The Update Manager is displayed in the bottom right of the screen.

The Update Manager provides you with the following options:

#### **Show details**

Displays additional resource information.

#### **Render**

Integrates the new resources into the topology.

Choosing this option will recalculate the topology layout based on your current display settings, and may therefore adjust the displayed topology significantly.

#### **Cogwheel icon**

When clicked, provides you with quick access to change your user preferences:

- **Enable auto-refresh:** Switches auto-refresh back on, and disables the Update Manager.
- **Remove deleted resources:** Removes the deleted resources from your topology view when the next topology update occurs.

#### **Hide**

Reduces the Update Manager to a small purple icon that does not obstruct your current topology view.

When you are ready to deal with the new resources, click on the icon to display the Update Manager again.

#### **Related tasks**

[“Defining advanced topology settings” on page 285](#)

As a system administrator, you can define advanced topology settings, such as the URLs for trusted sites, the required properties for tooltips or relationship type styles, or the maximum hop numbers a user can

choose in the Agile Service Manager topology viewer. You do this from the **Advanced topology settings** page accessed through DASH.

## Topology tools reference

---

This reference topic describes the Netcool Agile Service Manager Topology tools functionality.

### Topology Tools - Details

The **Topology Tools - Details** page is displayed when you select a right-click tool to edit it, or when you create a new tool. Here you define a tool's name and label as a minimum.

#### Name

Unique name used as an internal reference.

Required.

#### Menu label

The menu label is the text displayed in the context menu.

This can be the same name as used by other tools, which is why the unique name is required.

Required

#### Description

A description to help administrator users record the tool's purpose.

Not displayed in the context menu.

Optional.

#### Menu priority

The menu priority slider defines where in the context menu the tool is displayed.

For example, tools with a priority of two will be displayed higher in the menu than tools that have a priority of four.

Available values are one to ten.

Optional.

#### Navigation

You can move to the next page by using the page selector.

The minimum requirement to save the tool and open the **Topology Tools - Implementation** page is the name and label.

### Topology Tools - Implementation

The **Topology Tools - Implementation** page is displayed after you have completed the **Topology Tools - Details** page. Here you define the tool using valid JavaScript. To help you create tools, you have access to the following custom helper functions:

#### asmProperties

The tool implementation has access to the properties of the relevant **resource**, **relationship** or **status** via the `asmProperties` JavaScript object, which contains all the properties.

You can access the properties using standard JavaScript, but you must protect against a value not being present.

For example if you intend to use the property 'latitude', you must verify that it is present before using it. To do so, use the following check command:

```
asmProperties.hasOwnProperty('latitude')
```

If the property is present, the Boolean value `true` will be returned.

#### Status tools properties

When creating **status** tools, you use JavaScript that is similar to the script that you use when creating **resource** or **relationship** tools. However, the properties you use in your status tool



scripts, such as `asmProperties`, reference the properties for the **status** item; unlike the properties you use in your resource or relationship tool scripts, which reference the properties for the resources or relationships. For example, if you use `asmProperties.location` in a status tool script, there must be a corresponding 'location' property in the status record.

When creating status tools, the `asmProperties` object has a property that takes the form of an array called **resources**, which represents the resources in the topology with which this status is associated. Each item in the resources array is an object with properties that represent the properties of that resource. For example, if a status is associated with two resources, the **uniqueId** property of the first of those two resources could be referenced in the script by using `asmProperties.resources[0].uniqueId`

In addition, you can access the properties of a resource against which you are running a status tool by using the **asmSourceProperties** object when scripting the status tool.

### **asmSourceProperties**

You can access information about the source properties of any **relationships** or **status** the custom tool is acting on via the `asmSourceProperties` JavaScript object.

Example of using the source resource properties in a custom relationship stroke definition:

```
if (asmSourceProperties.myProp === 'high') {  
    return 'blue';  
} else {  
    return 'black';  
}
```

**Remember:** The arrows indicating a relationship point from the source to the target.

### **asmTargetProperties**

You can access information about the target properties of **relationships** the custom tool is acting on via the `asmTargetProperties` JavaScript object.

### **asmFunctions**

You can use a number of other helper functions, which are accessed from the `asmFunctions` object, which includes the following:

#### **showConfirmationPopup(title, message, onOk)**

Creates a popup confirmation allowing the tool to confirm an action.

Takes a title and message, which is displayed on the popup, and a function definition, which is run if the user clicks the OK button on the popup.

#### **showToasterMessage(status, message)**

Shows a popup toaster with the appropriate status coloring and message.

#### **showPopup(title, text)**

Shows a popup with a given title and text body (including markdown), which can be generated based on the properties of the resource or relationship.

**Tip:** The `asmFunctions.showPopup` helper function lets you use markdown to create more sophisticated HTML popups. For more information on markdown syntax, consult a reputable markdown reference site.

#### **showIframe(url)**

Displays a popup filling most of the page which wraps an iframe showing the page of the given URL.

Allows you to embed additional pages.

#### **sendPortletEvent(event)**

Allows you to send DASH portlet events from the Topology Viewer that can be used to manipulate other DASH portlets, such as the Event Viewer within IBM Tivoli Netcool/OMNIBus Web GUI.

**Note:** You can send events to other DASH portlets only if you are running Agile Service Manager within DASH (rather than in a direct-launch browser window), and if the receiving DASH portlets subscribe to the types of events being sent. See the [“sendPortletEvent examples” on page 277](#) topic for more information.

**getResourceStatus(<resource\_id>, <callback\_function>, [<time\_stamp>])**

Allows you to request status information from a tool definition for a given resource using its **\_id** parameter.

**resource\_id**

Required

Can be obtained from a resource via `asmProperties._id` and from a relationship using `asmSourceProperties._id` or `asmTargetProperties._id`

**callback\_function**

Required

Is called once the status data has been collected from the topology service, with a single argument containing an array of status objects

**time\_stamp**

Optional

Unix millisecond timestamp to get the status from a given point in history

The following example prints the status information of a source resource from a relationship context to the browser console log:

```
let printStatusCallback = function(statuses) {
    statuses.forEach(function(status) {
        console.log('status:', status.status,
                    'state:', status.state,
                    'severity:', status.severity,
                    'time:', new Date(status.time));
    })
}
asmFunctions.getResourceStatus(asmSourceProperties._id,
printStatusCallback);
```

**sendHttpRequest(url, options)**

Lets you send an HTTP or HTTPS request to a remote web server using the Agile Service Manager backend server rather than the browser, thereby avoiding any browser domain-blocking.

**url**

Required

The full URL of the remote site to be accessed.

For example:

```
https://data-svr-01.uk.com/inv?id=1892&offset=0
```

**Restriction:** You must add any websites referenced by the url parameter to a list of trusted sites as described in the [“Defining advanced topology settings” on page 285](#) topic (in this example `data-svr-01.uk.com`).

**options**

Optional

**method**

HTTP method used:

- GET
- POST
- PUT
- DELETE

The default is GET

**headers**

An object defining any special request headers needed

**body**

A string containing the body data for the request

POST and PUT requests only

**autoTrust**

A flag to indicate if the remote web server can be automatically trusted

This flag is **required** if the web site uses a self-signed SSL certificate with no CA, or a CA that is unknown to the Agile Service Manager server.

True or false, with a default of false

**onSuccess**

A callback function to run if the HTTP request is successful

This function will be passed the following three parameters:

- Response text
- HTTP status code
- Response headers

**onError**

A callback function to run if the HTTP request fails

This function will be passed the following three parameters:

- Response text
- HTTP status code
- Response headers

**Options parameter script sample:**

```
{
  method: 'GET',
  headers: {
    Content-Type: 'application/json',
    X-Locale: 'en'
  },
  body: '{ "itemName": "myData1" }',
  autoTrust: true,
  onSuccess: _onSuccessCallback,
  onError: _onErrorCallback
}
```

## Topology Tools - Conditions

The **Topology Tools - Conditions** page is displayed after you have completed the **Topology Tools - Implementation** page. Here you select the resource, relationship or status that will display the tool in their context menus.

**Applicable item type for tool definition**

From this drop-down, select the types to which the tool is applicable: **Resource**, **Relationship**, **Resource and Relationship**, or **Status**.

Depending on your selection, a number of check boxes are displayed, which you use to configure which resources, relationships or states are included.

**All types / All states**

Select this option if you want the tool to be displayed for all resource and relationship types, or all states (for Status).

The tool will also be displayed for any specific types not listed here.

**Resource types**

Select one or more resource types from the list displayed.

**Relationship types**

Select one or more relationship types from the list displayed.

**Status**

Select from the following possible states for which the tool will be available:

- **Open**
- **Clear**
- **Closed**

**Remember:** When creating status tools, the properties you use in your status tool scripts reference the properties for the status item, while the properties you use in your resource or relationship tools reference the properties for the resources or relationships.

## Custom icons reference

---

This reference topic describes the Netcool Agile Service Manager Custom Icons functionality.

### Custom Icons

The **Custom Icons** page is displayed when you select **Administration** from the DASH menu, and then click **Custom Icons** under the Agile Service Management heading.

The **Custom Icons** page displays the following buttons.

#### New

Opens the **Configure Custom Icon** page

#### 'Refresh' symbol

Reloads the icon information from the topology service

In addition, the **Custom Icons** page displays the following icon information in table format.

#### Name

Unique icon name

#### Icon

The icon itself

If you hover over a custom icon, it will be enlarged and displayed inside a circle to show what it will look like within a topology view.

#### Last Updated

Date and timestamp

#### Size (KB)

Size of the icon SVG in KB

#### 'Edit' symbol

Opens the **Configure Custom Icon** page

#### 'Bin' symbol

Deletes an icon.

If assigned to a resource type, a warning is displayed.

#### Category

Sorts icons by category

### Configure Custom Icon

The **Configure Custom Icon** page is displayed when you select an icon on the **Custom Icons** page to edit it, or when you create a new icon. Here you define an icon's name and SVG XML (both required) using the provided SVG XML editor.

#### Name

Each icon must have a name, which uniquely identifies the icon when assigning it to a type.

You cannot change the name of an existing icon. If you want an icon to have a different name, create a new icon, then delete the old one.

#### SVG XML

Use the XML editor to enter or edit the SVG text.

Each icon definition must be valid svg xml with a given viewBox, which is important to ensure scaling of the image. The SVG editor rejects any invalid XML entered.

The svg definition must include inline styling of the image, such as stroke color and fill color. If style classes are used, naming must be unique for each svg image to prevent class definitions from being overwritten.

The XML editor includes a **Preview** area where the results of your SVG edits are displayed.

### Category

Optionally, each icon can be assigned to a category. You can use categories to group icons of the same type or function together.

If you sort the full list of icons by Category, icons with the same category are displayed together.

**Example:** Use the following definition for the 'disk' icon as guidance:

```
<svg xmlns="http://www.w3.org/2000/svg" viewBox="0 0 64 64">
  <ellipse style="fill-opacity:0;stroke:currentColor;stroke-width:12.12270069;"
    id="path4139" cx="33.627117" cy="32.949142" rx="16.803904" ry="17.210684"/>
  <circle cx="33.827423" cy="33.055576" r="3.3037829"/>
</svg>
```

## Example sysctl.conf file

The following example of a `sysctl.conf` file shows settings that have been used in testing.

### /etc/sysctl.conf

**Tip:** Optimize Cassandra and Elasticsearch Kernel parameters by either disabling Swap, or setting the Kernel **vm.swappiness** parameter to 1.

To customize your `sysctl.conf` file, first back-up the original file, then edit it, before restarting the system. The default location is `/etc/sysctl.conf`

```
# sysctl settings are defined through files in
# /usr/lib/sysctl.d/, /run/sysctl.d/, and /etc/sysctl.d/.
#
# Vendors settings live in /usr/lib/sysctl.d/.
# To override a whole file, create a new file with the same in
# /etc/sysctl.d/ and put new settings there. To override
# only specific settings, add a file with a lexically later
# name in /etc/sysctl.d/ and put new settings there.

vm.swappiness = 1
vm.dirty_background_ratio = 3
vm.dirty_ratio = 80
vm.dirty_expire_centisecs = 500
vm.dirty_writeback_centisecs = 100

kernel.shmmax = 4398046511104
kernel.shmall = 1073741824

kernel.sem = 250 256000 100 16384

net.core.rmem_default = 262144
net.core.rmem_max = 4194304
net.core.wmem_default = 262144
net.core.wmem_max = 1048576

fs.aio-max-nr = 1048576

kernel.panic_on_oops = 1

fs.file-max = 6815744

net.ipv4.tcp_tw_recycle = 1
net.ipv4.tcp_tw_reuse = 1

net.ipv4.tcp_max_syn_backlog = 4096
net.ipv4.tcp_syncookies = 1

net.core.somaxconn = 1024
```

```
kernel.shmmni = 16384
net.ipv4.ip_local_port_range = 9000 65535
kernel.msgmnb = 65536
kernel.msgmax = 65536
kernel.shmmax = 540971653120
kernel.shmall = 4294967296
```

## Swagger reference

Specific links to Agile Service Manager Swagger documentation are included in many of the topics, as and when useful. This topic summarizes some of that information in a single location, for example by listing the default ports and Swagger URLs for each Agile Service Manager service.

### Swagger overview

Swagger is an open source software framework, which includes support for automated documentation and code generation. You can find more information on the Swagger website: <https://swagger.io/docs/>

#### Swagger

Agile Service Manager uses Swagger for automated documentation generation and utilizes a Swagger server for each micro-service.

You can access and explore the REST APIs of the topology service and observers using Swagger via the proxy service.

#### For example

- To access the **Topology service** via Swagger, use the following URL: `https://<your host>/1.0/topology/swagger`
- To access the **AWS Observer** service via Swagger, use the following URL: `https://<your host>/1.0/aws-observer/swagger`

#### Important:

#### On-prem

For the **on-prem** version of Agile Service Manager, you access the micro-services through the proxy service (nasm-nginx), which requires a proxy user and password to be configured for Nginx.

The default values for the user name and password are:

#### username

asm

#### password

asm

#### On OCP

When installed **on OCP** as part of NOI, the Agile Service Manager release name becomes `{noiReleaseName}-topology`.

The name of the Agile Service Manager secret is therefore `{noiReleaseName}-topology-asm-credentials`.

To obtain credentials on OCP, use the following example command, which prints out a secret in plain text (for the credentials, it will print the username and password in plain text):

```
kubectl get secret <noiReleaseName-topology>-asm-credentials -o go-template='{{range $k,$v := .data}}{{printf "%s: " $k}}{{if not $v}}{{v}}{{else}}{{v | base64decode}}{{end}}{{"\n"}}{{end}}'
```

**Example:**

```
kubectl get secret noi1-topology-asm-credentials -o go-template='{{range $k,$v := .data}}
{{printf "%s: " $k}}{{if not $v}}{{v}}{{else}}{{v | base64decode}}{{end}}{{"\n"}}{{end}}'
password: xxxxxxxxxxxxxxxx
username: xxx
```

Where {noi1} is the NOI release name, and {noi1-topology} is the Agile Service Manager release name.

**Default Swagger URLs**

The following tables show the default Swagger URLs for Agile Service Manager services and observers.

**OCP Note:** On OCP the value of *<your host>* is:

```
<release-name>-topology.<namespace>.apps.<my host>
```

Therefore the Swagger URL for the ALM Observer is:

```
https://<release-name>-topology.<namespace>.apps.<my host>/1.0/alm-observer/swagger
```

**Example:**

```
https://noi1-topology.master204.apps.noi161-ocp449.cp.fyre.acme.com/1.0/alm-observer/swagger
```

Table 80. Default Swagger URLs for Agile Service Manager services	
Service	Swagger URL
layout	https://<your host>/1.0/layout/swagger
merge	https://<your host>/1.0/merge/swagger
search	https://<your host>/1.0/search/swagger
topology	https://<your host>/1.0/topology/swagger
Observer Service	https://<your host>1.0/observer/swagger
Application Discovery service	<p><b>On prem:</b> https://&lt;your host&gt;-app-disco/rest/discover/swagger</p> <p><b>On OCP:</b> The value of &lt;your host&gt; for application discovery is:</p> <ul style="list-style-type: none"> <li>https://&lt;release-name&gt;-topology-&lt;namespace&gt;-app-disco.apps.&lt;my host&gt;</li> </ul> <p>The Swagger reference is:</p> <ul style="list-style-type: none"> <li>https://&lt;your host&gt;/rest/discover/swagger</li> </ul>
Network Discovery service	<p><b>On prem:</b> https://&lt;your host&gt;:9443/docs/swaggerui</p> <p><b>On OCP:</b> The value of &lt;your host&gt; for network discovery is:</p> <ul style="list-style-type: none"> <li>https://&lt;release-name&gt;-topology-net-disco.&lt;namespace&gt;.&lt;my host&gt;</li> </ul> <p>The Swagger reference is:</p> <ul style="list-style-type: none"> <li>https://&lt;your host&gt;/docs/swaggerui</li> </ul>
Status service	https://<your host>1.0/status/swagger

*Table 81. Default Swagger URLs for Agile Service Manager observers*

<b>Observer</b>	<b>Swagger URL</b>
alm-observer	<a href="https://&lt;your host&gt;/1.0/alm-observer/swagger">https://&lt;your host&gt;/1.0/alm-observer/swagger</a>
ansibleawx-observer	<a href="https://&lt;your host&gt;/1.0/ansibleawx-observer/swagger">https://&lt;your host&gt;/1.0/ansibleawx-observer/swagger</a>
aws-observer	<a href="https://&lt;your host&gt;/1.0/aws-observer/swagger">https://&lt;your host&gt;/1.0/aws-observer/swagger</a>
appdisco-observer	<a href="https://&lt;your host&gt;/1.0/appdisco-observer/swagger">https://&lt;your host&gt;/1.0/appdisco-observer/swagger</a>
appdynamics-observer	<a href="https://&lt;your host&gt;/1.0/appdynamics-observer/swagger">https://&lt;your host&gt;/1.0/appdynamics-observer/swagger</a>
azure-observer	<a href="https://&lt;your host&gt;/1.0/azure-observer/swagger">https://&lt;your host&gt;/1.0/azure-observer/swagger</a>
bigfixinventory-observer	<a href="https://&lt;your host&gt;/1.0/bigfixinventory-observer/swagger">https://&lt;your host&gt;/1.0/bigfixinventory-observer/swagger</a>
cienablueplanet-observer	<a href="https://&lt;your host&gt;/1.0/cienablueplanet-observer/swagger">https://&lt;your host&gt;/1.0/cienablueplanet-observer/swagger</a>
ciscoaci-observer	<a href="https://&lt;your host&gt;/1.0/ciscoaci-observer/swagger">https://&lt;your host&gt;/1.0/ciscoaci-observer/swagger</a>
contrail-observer	<a href="https://&lt;your host&gt;/1.0/contrail-observer/swagger">https://&lt;your host&gt;/1.0/contrail-observer/swagger</a>
dns-observer	<a href="https://&lt;your host&gt;/1.0/dns-observer/swagger">https://&lt;your host&gt;/1.0/dns-observer/swagger</a>
docker-observer	<a href="https://&lt;your host&gt;/1.0/docker-observer/swagger">https://&lt;your host&gt;/1.0/docker-observer/swagger</a>
dynatrace-observer	<a href="https://&lt;your host&gt;/1.0/dynatrace-observer/swagger">https://&lt;your host&gt;/1.0/dynatrace-observer/swagger</a>
file-observer	<a href="https://&lt;your host&gt;/1.0/file-observer/swagger">https://&lt;your host&gt;/1.0/file-observer/swagger</a>
googlecloud-observer	<a href="https://&lt;your host&gt;/1.0/googlecloud-observer/swagger">https://&lt;your host&gt;/1.0/googlecloud-observer/swagger</a>
ibmcloud-observer	<a href="https://&lt;your host&gt;/1.0/ibmcloud-observer/swagger">https://&lt;your host&gt;/1.0/ibmcloud-observer/swagger</a>
itnm-observer	<a href="https://&lt;your host&gt;/1.0/itnm-observer/swagger">https://&lt;your host&gt;/1.0/itnm-observer/swagger</a>
jenkins-observer	<a href="https://&lt;your host&gt;/1.0/jenkins-observer/swagger">https://&lt;your host&gt;/1.0/jenkins-observer/swagger</a>
juniopercso-observer	<a href="https://&lt;your host&gt;/1.0/juniopercso-observer/swagger">https://&lt;your host&gt;/1.0/juniopercso-observer/swagger</a>
kubernetes-observer	<a href="https://&lt;your host&gt;/1.0/kubernetes-observer/swagger">https://&lt;your host&gt;/1.0/kubernetes-observer/swagger</a>
netdisco-observer	<a href="https://&lt;your host&gt;/1.0/netdisco-observer/swagger">https://&lt;your host&gt;/1.0/netdisco-observer/swagger</a>
newrelic-observer	<a href="https://&lt;your host&gt;/1.0/newrelic-observer/swagger">https://&lt;your host&gt;/1.0/newrelic-observer/swagger</a>
openstack-observer	<a href="https://&lt;your host&gt;/1.0/openstack-observer/swagger">https://&lt;your host&gt;/1.0/openstack-observer/swagger</a>



Table 81. Default Swagger URLs for Agile Service Manager observers (continued)	
Observer	Swagger URL
rancher-observer	https://<your host>/1.0/rancher-observer/swagger
rest-observer	https://<your host>/1.0/rest-observer/swagger
servicenow-observer	https://<your host>/1.0/servicenow-observer/swagger
taddm-observer	https://<your host>/1.0/taddm-observer/swagger
vmvcenter-observer	https://<your host>/1.0/vmvcenter-observer/swagger
vmwarensx-observer	https://<your host>/1.0/vmwarensx-observer/swagger
zabbix-observer	https://<your host>/1.0/zabbix-observer/swagger

**Important:** Ensure that the body is structured correctly. When posting the body, information included in the body after the closing `}` that matches an opening `{` is ignored, and no error is recorded.

## Installation parameters

This topic lists the installation parameters you can override during a Helm installation.

### Configurable Helm installation parameters

You override the Helm installation parameters by adding them to the Helm install command as follows:

```
--set key=value[,key=value]
```

Table 82. Helm installation parameters		
Parameter	Description	Default
asm.almObserver.enabled	Option to install the Agile Lifecycle Manager observer	false
asm.awsObserver.enabled	Option to install the Amazon Web Services observer	false
asm.bigfixinventoryObserver.enabled	Option to install the BigFix Inventory observer	false
asm.cienablueplanetObserver.enabled	Option to install the Ciena Blue Planet observer	false
asm.ciscoaciObserver.enabled	Option to install the Cisco ACI observer	false
asm.contrailObserver.enabled	Option to install the Juniper Contrail observer	false
asm.dnsObserver.enabled	Option to install the DNS observer	false
asm.dockerObserver.enabled	Option to install the Docker observer	false

Table 82. Helm installation parameters (continued)

Parameter	Description	Default
asm.dynatraceObserver.enabled	Option to install the Dynatrace observer	false
asm.fileObserver.enabled	Option to install the File observer	false
asm.ibmcloudObserver.enabled	Option to install the IBM Cloud observer	false
asm.itnmObserver.enabled	Option to install the ITNM observer	false
asm.newrelicObserver.enabled	Option to install the New Relic observer	false
asm.openstackObserver.enabled	Option to install the OpenStack observer	false
asm.restObserver.enabled	Option to install the REST observer	false
asm.servicenowObserver.enabled	Option to install the ServiceNow observer	false
asm.taddmObserver.enabled	Option to install the TADDM observer	false
asm.vmvcenterObserver.enabled	Option to install the VMware vCenter observer	false
asm.vmwarensxObserver.enabled	Option to install the VMware NSX observer	false
asm.zabbixObserver.enabled	Option to install the Zabbix observer	false
license	Have you read and agree to the License agreement? set to 'accept'	not-accepted
noi.releaseName	The name of the Helm release of NOI to connect to.	noi
global.image.repository	Docker registry to pull ASM images from	
global.ingress.api.enabled	Option to enable the creation of ingress objects for the application endpoints	true
global.ingress.domain	Optional hostname to bind to the ingress rules, which must resolve to an proxy node. Multiple deployments of this chart will need to specify different values.	
global.ingress.tlsSecret	Optional TLS secret for the ingress hostname.	

Table 82. Helm installation parameters (continued)

Parameter	Description	Default
global.persistence.enabled	Option to disable the requests for PersistentVolumes, for test and demo only.	true
global.persistence.storageSize.cassandradata	Option to configure the requested amount of storage for Cassandra	50Gi
global.persistence.storageSize.kafkadata	Option to configure the requested amount of storage for Kafka	15Gi
global.persistence.storageSize.zookeeperdata	Option to configure the requested amount of storage for Zookeeper	5Gi
global.persistence.storageSize.elasticdata	Option to configure the requested amount of storage for Elasticsearch	75Gi
global.cassandraNodeReplicas	The number of instances to run for Cassandra	3
global.elasticsearch.replicaCount	The number of instances to run for Elasticsearch	3
global.environmentSize	'size0' requests fewer resources and is suitable for test and demo. Choose 'size1' for a production deployment.	size1
global.kafka.clusterSize	The number of instances to run for Kafka	3
global.zookeeper.clusterSize	The number of instances to run for Zookeeper	3



## Notices

---

This information applies to the PDF documentation set for IBM Netcool Agile Service Manager.

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 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 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](http://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](http://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.

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.



