

Netcool/Impact
Version 6.1.0.1

User Interface Guide



Netcool/Impact
Version 6.1.0.1

User Interface Guide



Note

Before using this information and the product it supports, read the information in "Notices".

Edition notice

This edition applies to version 6.1.0.1 of IBM Tivoli Netcool/Impact and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright IBM Corporation 2006, 2011.

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

Contents

About this publication vii

Intended audience	vii
Publications	vii
Netcool/Impact library	vii
Accessing terminology online	vii
Accessing publications online	viii
Ordering publications	viii
Accessibility	viii
Tivoli technical training	viii
Support for problem solving	ix
Obtaining fixes	ix
Receiving weekly support updates	ix
Contacting IBM Software Support	x
Conventions used in this publication	xii
Typeface conventions	xii
Operating system-dependent variables and paths	xii

Chapter 1. Getting started 1

Globalization	1
Using SJIS or EUC Japanese character encoding	1
Tivoli Integrated Portal overview	1
Browser requirements	2
IPv6 support	2
Before you log in	2
Logging in	2
Navigating to Netcool/Impact components	4
Entry fields accessibility	5
Selecting a cluster and a project	6
Creating policies, data, and services for a project	6

Chapter 2. Working with roles 7

Assigning roles to Netcool/Impact users	8
Creating roles	8
Editing roles	9
Deleting custom roles	10
Managing roles for users	11
Managing roles for groups	11
Role properties	12

Chapter 3. Working with projects 15

Projects overview	15
TBSM-specific projects	15
Project components	16
Important differences between projects, and the global repository	16
Global repository	17
Viewing the global repository data	17
Adding an item to the global repository	17
Deleting an item from the global repository	17
Project panel controls	18
Creating a project	18
Project editor configuration window	19
Viewing project members	19
Editing a project	19
Deleting a project	20

Automated project deployment feature	20
Running the DeployProject policy	20
DeployProject policy runtime parameters window	21
Version control file locking	21
Unlocking all locked items	22

Chapter 4. Data model 23

Data model components	23
Setting up a data model	23
Accessing the data model tab	24
Data model menu controls	24
Data sources overview	25
Data source categories	25
List of data sources	26
Creating data sources	28
Editing data sources	28
Deleting data sources	28
Testing data source connections	28
Data types overview	28
Data type categories	29
Predefined data types overview	29
List of predefined data types	29
Viewing data types	30
Editing data types	30
Deleting data types	30
Data items overview	31
Links overview	31
Link categories	31

Chapter 5. Working with data sources 33

Data sources	33
SQL database DSA failover	33
SQL database DSA failover modes	33
SNMP data sources	33
Socket DSA data source	34
SQL database data sources	34
DB2 data source configuration	34
Creating flat file data sources	36
Flat file data source configuration	36
Informix data source configuration	37
MS-SQL Server data source configuration	39
MySQL data source configuration	40
ObjectServer data source configuration	42
ODBC data source configuration	43
Oracle data source configuration	45
Connecting to an Oracle data source using LDAP	47
Connecting to Oracle RAC cluster	47
PostgreSQL data source configuration	48
Sybase data source configuration	49
GenericSQL data sources	51
GenericSQL data source configuration	51
HSQLDB data source configuration	52
LDAP data sources	54
Creating LDAP data sources	54

LDAP data source configuration window	54
Mediator data sources	55
CORBAMediator DSA data source configuration window	55
DirectMediator DSA data source configuration window	56
Creating SNMPDirectMediator data sources	56
SNMPDirectMediator data source configuration window	56
JMS data source	57
JMS data source configuration properties	58

Chapter 6. Data types 61

Viewing data type performance statistics	61
Data type performance statistics	61
Data type caching	62
Data type caching types	62
Creating internal data types	63
Internal data type configuration window	63
External data types	64
Deleting a field	65
List of predefined data types	65
Predefined data types overview	65
Time range groups and schedules	65
ITNM DSA data type	69
SQL data types	70
Configuring SQL data types	70
SQL data type configuration window - Table Description tab	71
SQL data type configuration window - adding and editing fields in the table	73
SQL data type configuration window - Cache settings tab	75
Creating flat file data types	76
LDAP data types	76
Configuring LDAP data types	77
LDAP data type configuration window - LDAP Info tab.	77
Mediator DSA data types.	78
Viewing Mediator DSA data types.	78
SNMP data types	79
SNMP data types - configuration overview	79
Packed OID data types	80
Table data types	81
LinkType data types	83
Configuring LinkType data items	83
Document data types	84
Adding new Doc data items.	84
FailedEvent data types	84
Viewing FailedEvent data items	84
Hibernation data types	85
Working with composite data types	85
Creating composite data types	85
Creating linked fields	85
Configuring a linked field on a composite data type	86

Chapter 7. Data items 87

Viewing data items	87
Adding new data items	87

Filtering the view	88
Editing data items	88
Deleting data items.	88

Chapter 8. Links. 89

Dynamic links	89
Static links.	89
Dynamic links overview	90
Creating dynamic links	90
Editing dynamic links	92
Deleting dynamic links	92
Working with static links	93
Creating static links	93
Browsing dynamic links	93

Chapter 9. Working with policies 95

Policies overview	95
Accessing policies	95
Viewing policies	95
Policies panel controls	96
Writing policies	96
Policy wizards	96
XML policies	97
Writing custom policies	98
Writing policies using wizards	98
Writing policies using JavaScript	98
Editing policies	98
Deleting policies.	98
Recovering automatically saved policies	99
Policy editor	99
Policy editor toolbar controls	99
Policy syntax checking	101
Policy syntax highlighter	101
Optimizing policies	101
Running policies with parameters in the editor	102
Graphic view of a policy	102
Browsing data types	103
Setting policy runtime parameters in the editor	103
Adding functions to policy	104
List and overview of functions	104
Personalizing the policy editor.	110
Changing default font used in the policy editor	110
Using version control interface.	111
Uploading policies.	111
Predefined policies	112

Chapter 10. Working with services 115

Services overview	115
Accessing services.	115
Services panel controls	116
List of services	117
Personalizing services	119
Configuring services	119
Creating services	119
Starting services	120
Stopping services	120
Deleting services	120
Viewing services logs.	120
Services log viewer	120
Service log viewer results	121

Creating new tabs	122
Event mapping	122
Creating event filters	122
Event filter configuration window	123
Event mapping table	123
Editing filters	123
Reordering filters	124
Deleting filters	124
Filter analysis	124
Command execution manager service	124
Command line manager service	125
Command line manager service configuration window	125
Database event listener service	125
Database event listener service configuration window	125
E-mail sender service	126
E-mail sender service configuration window	126
Event processor service	126
Event processor service configuration window	127
Hibernating policy activator service	128
Hibernating policy activator configuration	128
Hibernating policy activator configuration window	128
Jabber service	129
Adding resources to the Jabber ID	129
Configuring Jabber service	130
Jabber service configuration window - general settings	130
Jabber transport accounts	131
Policy logger service	132
Policy logger configuration	133
Policy logger service configuration window	133
Policy log files	134
ITNM event listener service	135
Configuring ITNM event listener service	135
ITNM event listener service configuration window	135
Self monitoring service	136
Self monitoring service configuration window	137
Database event reader service	137
Configuring the database event reader service	138
Database event reader configuration window - general settings	138
Database event reader configuration window - event mapping	139
Email reader service	140
Email reader service configuration window	140
Event listener service	140
Event listener service configuration window	141
JMS message listener	141
JMS message listener service configuration properties	141
Jabber reader service	142
Jabber reader service configuration window	143
OMNIBus event listener service	143
Setting up the OMNIBus event listener service	144
OMNIBus event listener service configuration window	144
OMNIBus event reader service	145
OMNIBus event reader configuration	145

OMNIBus event reader service General Settings tab	145
OMNIBus event reader service Event Mapping tab	146
OMNIBus Event Reader event locking examples	148
Policy activator service	149
Policy activator configuration	149
Policy activator service configuration window	150
Web Services Notification Listener service	150
Web Services Notification Listener service configuration window	151

Chapter 11. Operator views 153

Viewing operator views	153
Operator views overview	153
Operator view types	154
Operator view components	155
Operator views panel controls	155
Layout options	155
Action panel policies	156
Information groups	156
Creating a basic operator view	156
Editing operator views	157
Deleting operator views	158
Displaying operator views in TIP	158

Chapter 12. Event Isolation and Correlation 161

Overview	161
Event Isolation and Correlation policies	161
Event Isolation and Correlation operator views	162
Configuring Event Isolation and Correlation data sources	162
Configuring Event Isolation and Correlation data types	162
Creating, editing, and deleting event rules	163
Creating an event rule	163
Configuring WebGUI to add a new launch point	165
Launching the Event Isolation and Correlation analysis page	165
Viewing the Event Analysis	166

Chapter 13. Reporting tools 167

Accessing reports	167
Viewing Reports	167
Reports toolbar	168
Action Efficiency report	169
Action Error report	169
Impact Profile report	170
Configuring Impact Profile report	170
Impact Profile Report data	170
Impact Profile Report rules editor	171
Impact ROI Efficiency report	172
Impact ROI Efficiency report business processes	174
Creating a sample Impact ROI Efficiency report	174
Node Efficiency report	176
Operator Efficiency report	176
Policy Efficiency report	177
Policy Error report	177

Chapter 14. Maintenance Window Management. 179

About MWM maintenance windows 179
Logging on to Maintenance Window Management 180
Creating a one time maintenance window 180
Creating a recurring maintenance window 180
Viewing maintenance windows 181

Chapter 15. Configuration documenter 183

Configuration documenter overview 183
Opening the configuration documenter. 183
Viewing the cluster status 184
Viewing the server status 184
Viewing data sources. 185
Viewing data types 185
Viewing policies 185
Viewing services 186

Appendix A. Accessibility 187

Appendix B. Notices 189

Trademarks 191

Glossary 193

A 193
B 193
C 193
D 193
E 194
F 195
G 195
H 195
I. 195
J. 196
K 196
L 196
M 197
N 197
O 197
P 197
S 197
U 199
V 199
W 199
X 199

Index 201

About this publication

The *Netcool/Impact User Interface Guide* contains information about the user interface in Netcool/Impact.

Intended audience

This publication is for users who use the Netcool/Impact user interface.

Publications

This section lists publications in the Netcool/Impact library and related documents. The section also describes how to access Tivoli® publications online and how to order Tivoli publications.

Netcool/Impact library

- *Quick Start Guide*, CF39PML
Provides concise information about installing and running Netcool/Impact for the first time.
- *Administration Guide*, SC23882904
Provides information about installing, running and monitoring the product.
- *User Interface Guide*, SC23883004
Provides instructions for using the Graphical User Interface (GUI).
- *Policy Reference Guide*, SC23883104
Contains complete description and reference information for the Impact Policy Language (IPL).
- *DSA Reference Guide*, SC23883204
Provides information about data source adaptors (DSAs).
- *Operator View Guide*, SC23885104
Provides information about creating operator views.
- *Solutions Guide*, SC23883404
Provides end-to-end information about using features of Netcool/Impact.
- *Integrations Guide*, SC27283402
Contains instructions for integrating Netcool/Impact with other IBM® software and other vendor software.
- *Troubleshooting Guide*, GC27283302
Provides information about troubleshooting the installation, customization, starting, and maintaining Netcool/Impact.

Accessing terminology online

The IBM Terminology Web site consolidates the terminology from IBM product libraries in one convenient location. You can access the Terminology Web site at the following Web address:

<http://www.ibm.com/software/globalization/terminology>

Accessing publications online

Publications are available from the following locations:

- The *Quick Start* DVD contains the publications that are in the product library. The format of the publications is PDF, HTML, or both. Refer to the readme file on the DVD for instructions on how to access the documentation.
- Tivoli Information Center web site at <http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.netcoolimpact.doc6.1/welcome.html>. IBM posts publications for all Tivoli products, as they become available and whenever they are updated to the Tivoli Information Center Web site.

Note: If you print PDF documents on paper other than letter-sized paper, set the option in the **File** → **Print** window that allows Adobe Reader to print letter-sized pages on your local paper.

- Tivoli Documentation Central at <http://www.ibm.com/developerworks/wikis/display/tivolidoccentral/Impact>. You can also access publications of the previous and current versions of Netcool/Impact from Tivoli Documentation Central.
- The Netcool/Impact wiki contains additional short documents and additional information and is available at <https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/Tivoli%20Netcool%20Impact>.

Ordering publications

You can order many Tivoli publications online at <http://www.elink.ibm.link.ibm.com/publications/servlet/pbi.wss>.

You can also order by telephone by calling one of these numbers:

- In the United States: 800-879-2755
- In Canada: 800-426-4968

In other countries, contact your software account representative to order Tivoli publications. To locate the telephone number of your local representative, perform the following steps:

1. Go to <http://www.elink.ibm.link.ibm.com/publications/servlet/pbi.wss>.
2. Select your country from the list and click **Go**.
3. Click **About this site** in the main panel to see an information page that includes the telephone number of your local representative.

Accessibility

Accessibility features help users with a physical disability, such as restricted mobility or limited vision, to use software products successfully. With this product, you can use assistive technologies to hear and navigate the interface. You can also use the keyboard instead of the mouse to operate all features of the graphical user interface.

For additional information, see Appendix A, “Accessibility,” on page 187.

Tivoli technical training

For Tivoli technical training information, refer to the following IBM Tivoli Education Web site at <http://www.ibm.com/software/tivoli/education>.

Support for problem solving

If you have a problem with your IBM software, you want to resolve it quickly. This section describes the following options for obtaining support for IBM software products:

- “Obtaining fixes”
- “Receiving weekly support updates”
- “Contacting IBM Software Support” on page x

Obtaining fixes

A product fix might be available to resolve your problem. To determine which fixes are available for your Tivoli software product, follow these steps:

1. Go to the IBM Software Support Web site at <http://www.ibm.com/software/support>.
2. Navigate to the **Downloads** page.
3. Follow the instructions to locate the fix you want to download.
4. If there is no **Download** heading for your product, supply a search term, error code, or APAR number in the search field.

For more information about the types of fixes that are available, see the *IBM Software Support Handbook* at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html>.

Receiving weekly support updates

To receive weekly e-mail notifications about fixes and other software support news, follow these steps:

1. Go to the IBM Software Support Web site at <http://www.ibm.com/software/support>.
2. Click the **My IBM** in the toolbar. Click **My technical support**.
3. If you have already registered for **My technical support**, sign in and skip to the next step. If you have not registered, click **register now**. Complete the registration form using your e-mail address as your IBM ID and click **Submit**.
4. The **Edit profile** tab is displayed.
5. In the first list under **Products**, select **Software**. In the second list, select a product category (for example, **Systems and Asset Management**). In the third list, select a product sub-category (for example, **Application Performance & Availability** or **Systems Performance**). A list of applicable products is displayed.
6. Select the products for which you want to receive updates.
7. Click **Add products**.
8. After selecting all products that are of interest to you, click **Subscribe to email** on the **Edit profile** tab.
9. In the **Documents** list, select **Software**.
10. Select **Please send these documents by weekly email**.
11. Update your e-mail address as needed.
12. Select the types of documents you want to receive.
13. Click **Update**.

If you experience problems with the **My technical support** feature, you can obtain help in one of the following ways:

Online

Send an e-mail message to erchelp@u.ibm.com, describing your problem.

By phone

Call 1-800-IBM-4You (1-800-426-4409).

World Wide Registration Help desk

For world wide support information check the details in the following link:

<https://www.ibm.com/account/profile/us?page=reghelpdesk>

Contacting IBM Software Support

Before contacting IBM Software Support, your company must have an active IBM software maintenance contract, and you must be authorized to submit problems to IBM. The type of software maintenance contract that you need depends on the type of product you have:

- For IBM distributed software products (including, but not limited to, Tivoli, Lotus®, and Rational® products, and DB2® and WebSphere® products that run on Windows or UNIX operating systems), enroll in Passport Advantage® in one of the following ways:

Online

Go to the Passport Advantage Web site at http://www-306.ibm.com/software/howtobuy/passportadvantage/pao_customers.htm.

By phone

For the phone number to call in your country, go to the IBM Worldwide IBM Registration Helpdesk Web site at <https://www.ibm.com/account/profile/us?page=reghelpdesk>.

- For customers with Subscription and Support (S & S) contracts, go to the Software Service Request Web site at <https://techsupport.services.ibm.com/ssr/login>.
- For customers with IBMLink, CATIA, Linux, OS/390®, iSeries®, pSeries®, zSeries®, and other support agreements, go to the IBM Support Line Web site at <http://www.ibm.com/services/us/index.wss/so/its/a1000030/dt006>.
- For IBM eServer™ software products (including, but not limited to, DB2 and WebSphere products that run in zSeries, pSeries, and iSeries environments), you can purchase a software maintenance agreement by working directly with an IBM sales representative or an IBM Business Partner. For more information about support for eServer software products, go to the IBM Technical Support Advantage Web site at <http://www.ibm.com/servers/eserver/techsupport.html>.

If you are not sure what type of software maintenance contract you need, call 1-800-IBMSERV (1-800-426-7378) in the United States. From other countries, go to the contacts page of the *IBM Software Support Handbook* on the Web at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html> and click the name of your geographic region for phone numbers of people who provide support for your location.

To contact IBM Software support, follow these steps:

1. "Determining the business impact" on page xi
2. "Describing problems and gathering information" on page xi
3. "Submitting problems" on page xi

Determining the business impact

When you report a problem to IBM, you are asked to supply a severity level. Use the following criteria to understand and assess the business impact of the problem that you are reporting:

Severity 1

The problem has a *critical* business impact. You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution.

Severity 2

The problem has a *significant* business impact. The program is usable, but it is severely limited.

Severity 3

The problem has *some* business impact. The program is usable, but less significant features (not critical to operations) are unavailable.

Severity 4

The problem has *minimal* business impact. The problem causes little impact on operations, or a reasonable circumvention to the problem was implemented.

Describing problems and gathering information

When describing a problem to IBM, be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently. To save time, know the answers to these questions:

- Which software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can you re-create the problem? If so, what steps were performed to re-create the problem?
- Did you make any changes to the system? For example, did you make changes to the hardware, operating system, networking software, and so on.
- Are you currently using a workaround for the problem? If so, be prepared to explain the workaround when you report the problem.

Submitting problems

You can submit your problem to IBM Software Support in one of two ways:

Online

Click **Submit and track problems** on the IBM Software Support site at <http://www.ibm.com/software/support/probsub.html>. Type your information into the appropriate problem submission form.

By phone

For the phone number to call in your country, go to the contacts page of the *IBM Software Support Handbook* at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html> and click the name of your geographic region.

If the problem you submit is for a software defect or for missing or inaccurate documentation, IBM Software Support creates an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. Whenever possible, IBM Software Support provides a workaround that you can implement until the APAR is resolved and a fix is delivered. IBM publishes resolved APARs on the Software Support Web site daily, so that other users who experience the same problem can benefit from the same resolution.

Conventions used in this publication

This publication uses several conventions for special terms and actions, operating system-dependent commands and paths, and margin graphics.

Typeface conventions

This publication uses the following typeface conventions:

Bold

- Lowercase commands and mixed case commands that are otherwise difficult to distinguish from surrounding text
- Interface controls (check boxes, push buttons, radio buttons, spin buttons, fields, folders, icons, list boxes, items inside list boxes, multicolumn lists, containers, menu choices, menu names, tabs, property sheets), labels (such as **Tip:**, and **Operating system considerations:**)
- Keywords and parameters in text

Italic

- Citations (examples: titles of publications, diskettes, and CDs)
- Words defined in text (example: a nonswitched line is called a *point-to-point line*)
- Emphasis of words and letters (words as words example: "Use the word *that* to introduce a restrictive clause."; letters as letters example: "The LUN address must start with the letter *L*.")
- New terms in text (except in a definition list): a *view* is a frame in a workspace that contains data.
- Variables and values you must provide: ... where *myname* represents....

Monospace

- Examples and code examples
- File names, programming keywords, and other elements that are difficult to distinguish from surrounding text
- Message text and prompts addressed to the user
- Text that the user must type
- Values for arguments or command options

Operating system-dependent variables and paths

This publication uses the UNIX convention for specifying environment variables and for directory notation.

When using the Windows command line, replace *\$variable* with *%variable%* for environment variables and replace each forward slash (/) with a backslash (\) in directory paths. The names of environment variables are not always the same in the Windows and UNIX environments. For example, *%TEMP%* in Windows environments is equivalent to *\$TMPDIR* in UNIX environments.

Note: If you are using the bash shell on a Windows system, you can use the UNIX conventions.

Chapter 1. Getting started

The graphical user interface (GUI) gives you immediate access to all projects, policies, reports, data types, operator views, services, and defined clusters.

Globalization

Netcool/Impact does not support Unicode names for databases, tables, schemas, and columns in foreign language data sources.

Using SJIS or EUC Japanese character encoding

You can input, display, and process Japanese characters in a policy by changing the encode option to Unicode in your browser. Use this following procedure to change the encode option on your browser.

Procedure

1. Open your browser.
2. Select **View > Encoding** or **View > Character Encoding**, depending on which browser you are using.
3. Select **Unicode (UTF-8)**.

Tivoli Integrated Portal overview

Web-based products built on the Tivoli Integrated Portal framework share a common user interface where you can launch applications and share information.

Tivoli Integrated Portal helps the interaction and secure passing of data between Tivoli products through a common portal. You can launch from one application to another and within the same dashboard view research different aspects of your managed enterprise.

Tivoli Integrated Portal is installed automatically with the first Tivoli product using the Tivoli Integrated Portal framework. Subsequent products may install updated versions of Tivoli Integrated Portal.

Tivoli Integrated Portal provides the following features:

- A Web based user interface for individual products and for integrating multiple products.
- A single, task-based navigation panel for multiple products. Users select actions based around the task that they want to complete, not by the product that supports that task.
- Single sign-on (SSO), consolidated user management, and a single point of access for different Tivoli applications.
- Aggregated views that span server instances, such as the Tivoli Netcool/OMNIbus ObjectServer and Tivoli Enterprise Portal Server.
- Inter-view messaging between products to support contextual linkage between applications.
- The ability to create customized pages and administer access to content by user, role, or group.

Browser requirements

Netcool/Impact 6.1 supports the following browsers:

- Microsoft Internet Explorer 7.0, 8.0 and 9.0 on Windows
- Mozilla Firefox 3.6 and Mozilla Firefox Extended Support Release 10.0 on the following operating systems:
 - Aix 6.1 Power system
 - Red Hat Enterprise Linux (RHEL) 5.0 Advanced Platform System z[®]
 - Red Hat Enterprise Linux (RHEL) 5.0 Advanced Platform x86-32 and x86-64
 - SUSE Linux Enterprise Desktop (SLED) 10.0 x86-32 and x86-64
 - SUSE Linux Enterprise Server (SLES) 10.0 System z
 - SUSE Linux Enterprise Server (SLES) 10.0 x86-32 and 86-64
 - Solaris 10 SPARC
 - Windows

For details of additional systems requirements needed for Netcool/Impact, see the following link <https://www.ibm.com/developerworks/wikis/display/tivolibsm/Overview+and+Planning>.

IPv6 support

If you operate a network that supports IPv6 you can access Netcool/Impact using its IPv6 address enclosed within square brackets.

For example:

`http://[2002:92a:111:440:9:42:30:213]:9080/nci`

If you are using a fully qualified domain name, there is no difference between IPv6 and IPv4. For example:

`http://cvtso105v6.tivlab.raleigh.ibm.com:9080/nci`

Before you log in

You need the following information to log in to Netcool/Impact in the Tivoli Integrated Portal.

- The IP address of the Tivoli Integrated Portal instance on which Netcool/Impact is installed.
- Your user name and password.
- One of the supported browsers.

Two or more users cannot log on using the same web browser on the same workstation at the same time.

Logging in

Log in to the portal whenever you want to start a work session.

Before you begin

The Tivoli Integrated Portal Server must be running before you can connect to it from your browser.

About this task

Complete these steps to log in:

Procedure

1. In a Web browser, enter the URL of the Tivoli Integrated Portal Server:
`http://host.domain:16310/ibm/console` or `https://host.domain:16311/ibm/console` if it is configured for secure access.
 - *host.domain* is the fully qualified host name or IP address of the Tivoli Integrated Portal Server (such as *MyServer.MySubdomain.MyDomain.com* or *9.51.111.121*, or *localhost* if you are running the Tivoli Integrated Portal Server locally).
 - 16310 is the default nonsecure port number for the portal and 16311 is the default secure port number. If your environment was configured with a port number other than the default, enter that number instead. If you are not sure of the port number, read the application server profile to get the correct number.
 - *ibm/console* is the default path to the Tivoli Integrated Portal Server, however this path is configurable and might differ from the default in your environment.
2. In the login page, enter your user ID and password and click **Log in**. This is the user ID and password that are stored with the Tivoli Integrated Portal Server.

Attention: After authentication, the web container used by the Tivoli Integrated Portal Server redirects to the last URL requested. This is usually `https://<host>:<port>/ibm/console`, but if you manually change the page URL, after being initially directed to the login page, or if you make a separate request to the server in a discrete browser window before logging in, you may be redirected unexpectedly.

Note: If you have more than one instance of the Tivoli Integrated Portal Server installed on your computer, you should not run more than one instance in a browser session, that is, do not log in to different instances on separate browser tabs.

Results

After your user credentials have been verified, the Welcome page is displayed. If you entered the *localhost* or port number incorrectly, the URL will not resolve. View the application server profile to check the settings for *localhost*, port, and user ID.

What to do next

Select any of the items in the navigation tree to begin working with the console.

While you are logged into the Tivoli Integrated Portal Server, avoid clicking the browser **Back** button because you will be logged out automatically. Click **Forward** and you will see that you are logged out and must resubmit your credentials to log in again.

Note: If you want to use single sign-on (SSO) then you must use the fully qualified domain name of the Tivoli Integrated Portal host.

Navigating to Netcool/Impact components

How to navigate to Netcool/Impact components when you log in to the Tivoli Integrated Portal.

When you log in for the first time to the Tivoli Integrated Portal, an introductory page displays. If you have multiple applications installed, you see a **Welcome** page with links to each product you have installed. Click a link to view an **About** page for a product. The About Tivoli Netcool/Impact page has three tabs:

- **Netcool/Impact:** shows information about the product and the version.
- **Support:** shows Impact, toolkits, browser, and server information that you can refer to in troubleshooting situations.
- **Learning Resources:** shows links to additional resources and information relating to the product.

To view this page at any stage, click the **Welcome** node in the navigation pane.

All the Netcool/Impact components are found in the navigation pane on the left side of the UI. Depending on the user permissions you are assigned, you have access to some or all of the following Netcool/Impact components.

Select **System Configuration > Event Automation** to locate the following components:

- **Data Model**
- **Policies**
- **Services**
- **Operator Views**
- **Event Isolation and Consolidation**

To locate **Reports** select **Reporting > Event Automation**. The following reports are available:

- **Action Efficiency Report**
- **Action Error Report**
- **Policy Efficiency Report**
- **Policy Error Report**
- **Node Efficiency Report**
- **Operator Efficiency Report**
- **Impact ROI Efficiency Report**
- **Impact Profile Report**

You can filter the Tivoli Integrated Portal menu to display the Netcool/Impact entries only. Above the Tivoli Integrated Portal menu, select **View > Tivoli Netcool/Impact**.

To locate the **Maintenance Window Management** component, select **Trouble Shooting and Support > Event Automation**

The following components are the most frequently used components:

Data Model

Set up a data model for your solution: data sources, data types, data items, and links.

For more information about the data model, see Chapter 4, “Data model,” on page 23.

Policies

Create policies to manipulate events and data from your data sources.

For more information about the policies, see Chapter 9, “Working with policies,” on page 95.

Services

Work with services: monitor event sources, send and receive e-mail notifications, and trigger policies.

For more information about the services, see Chapter 10, “Working with services,” on page 115.

Operator Views

View events and data in real time and run policies that are based on that data.

For more information about the operator views, see Chapter 11, “Operator views,” on page 153.

Reports

View information about your network and network operators, and assess the efficiency of your configuration.

For more information about the reports, see Chapter 13, “Reporting tools,” on page 167.

Event Isolation and Correlation

You can set up Event Isolation and Correlation to isolate the event that has caused the problem. You can also view the events dependent on the isolated event. For more information about configuring data sources and creating event rules, see Chapter 12, “Event Isolation and Correlation,” on page 161.

For more information about setting up and configuring Event Isolation and Correlation, see the *Netcool/Impact Solutions Guide*.

Maintenance Window Management

Maintenance Window Management (MWM) is an add-on for managing Netcool® OMNIbus maintenance windows.

For more information about using Maintenance Window Management, see Chapter 14, “Maintenance Window Management,” on page 179

For more information about setting up Maintenance Window Management, see the *Netcool/Impact Solutions Guide*.

Entry fields accessibility

The graphical user interface entry fields are accessible using the keyboard.

When you click a certain key combination on the keyboard, the system automatically takes you to the entry field you want. To access a field from the keyboard, press and hold ALT + SHIFT and the letter in parentheses that displays after the field name.

For example, in a configuration window that contains a field called **Username**, after the field name you see a letter in parentheses (**U**). To access the **Username** field using the keyboard, press and hold ALT + SHIFT + U.

Selecting a cluster and a project

Use the cluster and project lists to select the clusters and projects you want to use when working within various tabs.

1. When you log in to the Tivoli Integrated Portal GUI and select an item from the subnode, the **Cluster** and **Project** lists display at the top of each tab.
2. From the **Cluster** list, select the required cluster to change clusters. From the **Project** list, select the required project. When you change a cluster or project on one tab, the change applies only to that tab. Use the icons next to the projects menu to create, edit, and delete projects.

Note: You can use the Global project to view all the items in the selected tab.

3. If there is only one cluster it displays automatically and the **Cluster** list is disabled. You can click the arrow at the bottom of the toolbar to show or hide the cluster and project lists to save space. If no cluster is running, you cannot select any cluster or project from the lists.
4. You can switch between clusters and projects within each tab. Save any work in progress before switching clusters and projects to implement any changes. There is one exception, the **Reports** tab has only a cluster list because reports do not use projects. For example, if you want to create policies in different projects you can use the projects menu to switch between projects in the policy editor. For more information about working with projects, see Chapter 3, “Working with projects,” on page 15.

Creating policies, data, and services for a project

Start working with Netcool/Impact by creating a project for your data, policies, and services.

For information about creating the policies, data, and services for a new project, see Chapter 3, “Working with projects,” on page 15.

Chapter 2. Working with roles

Portal users are granted access to resources based on the role to which they have been assigned. In the navigation pane, click **Settings > Roles** to add and remove roles and to assign access to portlets, pages, and views.

To manage users and groups and assign them to roles, click **Users and Groups**.

After the portal is installed, there are some roles already defined to the server.

Attention: The “suppressmonitor” role is used to hide the tasks associated with the application server, including the tasks in the Security, Troubleshooting, and Users and Groups folders.

Access levels

The access level that a role has to a resource determines the actions that users within that role can perform on the resource.

Table 1. Access rights to portal resources based on access level

Resource	Access Level		
	“User”	“Privileged User”	“Editor”
Portlet	View and interact with the portlet and access portlet help	View and interact with the portlet, edit personal settings, and access portlet help	View and interact with the portlet, edit personal settings, edit global settings, and access portlet help
Page	Launch the node from the navigation		Launch the node from the navigation and edit the content and layout
Folder	Note: Folders are always available in the navigation if the user has access to at least one of its pages.		
External URL	Launch the node from the navigation		
View	Select the view		

For a given resource, if a role does not have one of these access level settings, then the role has no access to the resource.

Only users with “adminsecuritymanager” and “Administrator” role can create, delete or change the properties of a role. If you assign access for any other role to the Roles portlet, users in that role will only be able to view roles and change access to views and pages.

Note: The access control settings are not observed when using the administrative portlets under the **Settings** node. Users with access to these pages and portlets will be able to create, edit, and delete all custom pages, portlets, and views. For example, if a user has no access to “Page Two”, but has access to Pages, that user can edit all of the properties of “Page Two” and change access control settings. Keep this in mind when granting access to the **Settings** portlets for a role.

If a user is assigned to multiple roles, the user acquires the highest access level between these roles for a resource. For example, if a user belongs to the manager role with “Privileged User” access to a portlet and also belongs to the communications role with no access to the portlet, then the user has “Privileged User” access to the portlet.

Tasks

You can grant access for multiple roles while creating or editing a resource, such as a page or a portlet. You can also grant access to multiple pages or views while creating or editing a role.

Assigning roles to Netcool/Impact users

Assign security-based roles to Netcool/Impact users in accordance with their access requirements.

There are three roles already set up for Netcool/Impact users in the Tivoli Integrated Portal. You can use these roles to control users access to navigation pages in the GUI and also to permit users to unlock their own files. Users who have the **impactAdminUser** role can unlock files locked by other users. The user **tipadmin** can also unlock all files owned by other users. Files that are open for editing cannot be unlocked. These files must be saved and closed.

- **impactAdminUser**: Assign users this role to give access to the full Netcool/Impact menu navigation. Users with this role can unlock their own policies, services, and data model files, and any files locked by other users.
- **impactFullAccessUser**: Assign users this role to give access to the full Netcool/Impact menu navigation. However, users with this role can only unlock their own files.
- **impactOpViewUser**: Assign users this role to access the Operator View and Event Isolation and Configuration nodes in the navigation pane.

Follow the instructions for “Managing roles for users” on page 11 and “Managing roles for groups” on page 11 to assign these roles to users and groups.

Creating roles

Portal users are granted access to resources based on the role to which they have been assigned. All roles that are created in the portal have a resource type of Custom. This procedure describes creating a role for testing purposes. After completing these steps, you can remove or edit this role for production use.

Procedure

1. Click **Users and Groups > Roles** in the navigation. A list of all roles in the portal is displayed.
2. Click **New**. The properties panel for the new role is displayed.
3. Enter a descriptive name for the role.
4. Optional: Expand the **Users and Groups** section. Use this section to associate a role with one or more users and groups. The method to add users and groups is similar, so this topic describes adding users only. To associate a user with a role, follow these steps:
 - a. In the **Users** panel, click **Add**. A new page is displayed that allows you to search for and select users to be added to the role.

- b. Provide search filters in the relevant fields, select the maximum number of results that you want returned and click **Search** to return a list of users that match your criteria.

Tip: If you leave the search filter fields blank, the system returns all users (up to a limit of 1000).

- c. From the returned results, select the users that you want to associate with the role and click **Add**. The previous page is displayed listing the selected users in the **Users** panel.
5. Expand the **Access to Views** section. Use this section to grant access to one or more custom views for users who are assigned to the new role. If you have already created a custom view, follow these steps.
 - a. Click **Add**. A list of available views is displayed.
 - b. Select one or more views and click **OK**.
 - c. To make sure the role has access to all of the pages within the view, click **Grant to All**.
6. Expand the **Access to Pages** section. A list of pages that the role can access is displayed. However, this list is empty if you did not add a view and grant access to all of the pages within the view.
7. Optional: Click **Add** to grant access to additional pages.
8. For each page that is listed, verify that the **Access Level** is set correctly.
9. Click **Save** to save your changes and return to Roles.

Results

The new role is created with access to the views, users and groups, and pages that you indicated. To grant access to the portlets on those pages you must edit the portlets.

Editing roles

Portal users are granted access to resources based on the role to which they have been assigned. If you have sufficient authorization in the portal, you can change the name of custom roles. For all roles, you can change access to views and pages and set the access level to pages.

About this task

Procedure

1. In the navigation pane, click **Users and Groups > Roles**. A list of all roles in the portal is displayed.
2. Click the name of the role that you want to edit. The properties panel for the role is displayed. If this is a custom role, the only field you can edit is **Role Name**. For all other resource types, you cannot edit any of the role properties.
3. Optional: Expand the **Users and Groups** section. Use this section to associate a role with one or more users and groups. The method to add users and groups is similar, so this topic describes adding users only. To associate a user with a role, follow these steps:
 - a. In the **Users** panel, click **Add**. A new page is displayed that allows you to search for and select users to be added to the role.

- b. Provide search filters in the relevant fields, select the maximum number of results that you want returned and click **Search** to return a list of users that match your criteria.

Tip: If you leave the search filter fields blank, the system returns all users (up to a limit of 1000).

- c. From the returned results, select the users that you want to associate with the role and click **Add**. The previous page is displayed listing the selected users in the **Users** panel.
4. Expand the **Access to Views** section. Use this section to grant access to one or more custom views for users who are assigned to the new role. If you have already created a custom view, follow these steps.
 - a. Click **Add**. A list of available views is displayed.
 - b. Select one or more views and click **OK**.
 - c. To make sure the role has access to all of the pages within the view, click **Grant to All**.
5. Expand the **Access to Pages** section. A list of pages that the role can access is displayed. However, this list is empty if you did not add a view and grant access to all of the pages within the view.
6. Optional: Click **Add** to grant access to additional pages.
7. For each page that is listed, verify that the **Access Level** is set correctly.
8. Click **OK**.

Results

Your changes are saved and you are returned to the Roles page.

What to do next

For any pages that you added for the role, you should ensure that the role also has access to the portlets on the page..

Deleting custom roles

You can delete only roles with the resource type of Custom. These are roles created using the portal.

About this task

Attention: Before deleting a role, consider whether any users are actively using the role and any impacts this might have on services. If necessary, notify users in advance of any plans for changes that could affect their work.

Follow these steps to delete a custom role.

Procedure

1. Click **Users and Groups > Roles** in the navigation pane. The Roles page is displayed with the list of roles in the portal.
2. Select the custom role that you want to delete. You can select more than one custom role.
3. Click **Delete**. A message is displayed at the top prompting you to confirm the deletion.
4. Click **OK**.

Results

The custom role is removed from the list.

Managing roles for users

Administrators can search for users and manage their roles in the User Roles page.

About this task

To search for users and manage their roles:

Procedure

1. In the navigation pane, click **Users and Groups > User Roles**. The User Roles page is displayed.
2. In the search fields provided, you can enter search criteria by given name, surname, user ID, and e-mail address. If you do not have exact details for a particular item, all of the search fields support using an asterisk (*) as a wildcard character. For example, to return all user records with a given name that starts with "Mich", enter mich* in the **First name** field.

Tip: You can leave the search fields blank to return all user records.

3. From the **Number of results to display** list, select the number of records that you want returned and click **Search**.

Restriction: Returned records are displayed one page only. If more records are available than the setting you chose from the list, only a partial list is returned. To display all records you need to search again after selecting a larger number from the **Number of results to display** list.

A list of records that match your search criteria are listed in the grid.

4. Select a user from the **User ID** column. A list of available roles for the selected user is displayed on a new page. Those roles that are currently associated with the selected user are checked.
5. Modify the roles associated with the user as required, that is, check the roles that you want associated with the user and clear those that you do not.
6. Click **Save** to commit your changes, or **Reset** to reset the form to its initial state. Once you click **Save**, the User Roles page is displayed. The entry for the user in the **Roles** column is updated to reflect your changes.

What to do next

You can select another user from the search results and update their role settings, enter new search criteria to manage other user records, or close the User Roles page.

Managing roles for groups

Administrators can search for groups and manage their roles in the Group Roles page.

About this task

To search for user groups and manage their roles:

Procedure

1. In the navigation pane, click **Users and Groups > Group Roles**. The Group Roles page is displayed.
2. In the search fields provided, you can enter search criteria by group ID and description. If you do not have exact details for a particular item, both search fields support using an asterisk (*) as a wildcard character. For example, to return all group records with a group ID that starts with "tes", enter tes* in the **Group ID** field.

Tip: You can leave the search fields blank to return all records.

3. From the **Number of results to display** list, select the number of records that you want returned and click **Search**.

Restriction: Returned records are displayed one page only. If more records are available than the setting you chose from the list, only a partial list is returned. To display all records you need to search again after selecting a larger number from the **Number of results to display** list.

A list of records that match your search criteria are listed in the grid.

4. Select a group from the **Group Name** column. A list of available roles for the selected group is displayed on a new page. Those roles that are currently associated with the selected group are checked.
5. Modify the roles associated with the group as required, that is, check the roles that you want associated with the group and clear those that you do not.
6. Click **Save** to commit your changes, or **Reset** to reset the form to its initial state. Once you click **Save**, the Group Roles page is displayed. The entry for the group in the **Roles** column is updated to reflect your changes.

What to do next

You can select another group from the search results and update its role settings, enter new search criteria to manage other group records, or close the Group Roles page.

Role properties

This panel is used to edit the general properties of a role. The properties panel is displayed when you click one of the roles to edit it. This panel is also displayed when you create a new role.

To access this panel, click **Users and Groups > Roles** in the navigation pane. Then click the name of one of the roles that are listed, or click **New** to create a new role.

Role name

Enter a descriptive name for the role. This name should be informative enough to indicate the actions and resources that are available to the users in this role. The role name must be unique within the portal. The characters in a role name are restricted to letters, digits, blank spaces, and underscores. The role name cannot start with a digit.

Type of role

Displays either Custom, System, or Core. This field cannot be changed.

Users and Groups

This section shows two table displaying the users and groups that have been assigned to this role. The following options are unique for these tables.

Users This table lists the users with access to the role. The following column headings are unique to this table:

User ID

Displays the unique system ID associated with the user.

First Name

Displays the first name associated with the user ID.

Last Name

Displays the last name associated with the user ID.

Groups

This table lists the user groups with access to the role. The following column headings are unique to this table:

Group Name

Displays the name associated with this group.

Unique Name

Displays the unique system ID associated with the group.

Access to Views

This section provides a table displaying views that have been assigned to this role. The following options are unique for the views table.

View Name

Indicates the name of the view. You can sort the list of names by clicking the column heading.

Access Level

Indicates the level of access for role members in relation to the view.

Access to Child Resources

Grants access to all pages that are part of this view. After clicking this option, the **Access to Pages** table is updated with the new pages. If necessary, you can individually select pages from that table to remove access.

Access to Pages

This section provides a table displaying pages that have been assigned to this role. The following options are unique for the pages table.

Name Displays the name of the page or node as it appears in the navigation.

Unique Name

Displays the identifier that uniquely identified the page in the portal.

The following fields and controls are available in all sections.

Select all icon

Selects all items displayed in the table for deletion. If you are displaying only a filtered set of items, only those items are selected. You can deselect specific items before actually deleting.

Deselect all icon

Deselects all items displayed in the table.

Add Adds an item to the table.

Remove

Removes all selected items from the table. There is no warning prompt when you click **Remove**.

Filter Type in this field to quickly find an item in the table. This field is useful when there are a large number of items to look through.

Select Selects or deselects a single item in the table.

Access Level

Indicates the actions that users in the role can perform on the view or page.

Chapter 3. Working with projects

You use projects to organize and manage related data sources and types, operator views, policies, services, and wizards.

Using the GUI you can:

- Switch between clusters and projects
- Create projects
- Edit projects
- View project members
- Add and delete project members
- Delete projects
- Copy the data sources, data types, operator views, policies, and services in a project between two running server clusters on a network

Projects overview

A project is a view of a subset of the elements stored in the global repository.

You can use projects to manage your policies and their associated elements. They help you to remember which data types and services relate to each policy and how the policies relate to each other. Projects also help to determine whether a policy, or its associated data types or services, is still in use or must be deleted from the project.

Also, you can find policies and their associated data and services easily when they are organized by project. You can add any previously created policies, data types, operator views, and services to as many projects as you like. You can also remove these items when they are no longer needed in any project.

If you have not as yet created any projects, **Default** and **Global** projects and projects predefined by Netcool/Impact are the only projects listed in the **Projects** menu.

The **Global** project lists all items in the global repository. Any item that you create, for example a data type, is not stored in the project that is currently selected, it is automatically added to the **Global** project.

The **Default** project is an example, it works just like any project, you can add items to it edit, or delete it.

When you delete a project, the items that were assigned as project members remain in the global project and as members of any other projects they were assigned to.

Important: You cannot edit or delete the **Global** project.

TBSM-specific projects

The Tivoli Business Service Manager version of Netcool/Impact has additional Tivoli Business Service Manager specific projects.

- **TBSM_BASE** project contains all the predefined Netcool/Impact data sources, data types, services, and policies, that are specific to Tivoli Business Service Manager. Modifications to any of the items in this project must be done with caution.
- **TBSM** project contains all Netcool/Impact data sources, data types, services, and policies that are created using the Tivoli Business Service Manager configuration interface.
- **TBSM_SAMPLES** project contains sample policies used to interact with the Services Component Registry.
- **ForImpactMigration** project contains the data sources and data types necessary for a remote Impact Server to send events using the **PassToTBSM** function to Tivoli Business Service Manager. To send events to Tivoli Business Service Manager from a remote Impact Server, you must export the **ForImpactMigration** project from the Tivoli Business Service Manager server and import it into the Impact Server. For more information about **PassToTBSM**, see the *Netcool/Impact Solutions Guide*.

Project components

When you create a project, you can add any existing policies, data, documents, and services to it as project members.

A project can consist of the following components:

- Policies
- Data sources that are set up for project data types
- Data types that are associated with the policies
- Operator views that are related to the policies
- Services

Important differences between projects, and the global repository

Make sure that you are aware of the following differences when editing and deleting items in projects, and the global repository.

- When you select any project, and create an item on the selected tab, the item is automatically added to the global repository.
- If you select the **Global** project and create an item in the specified tab the item is added to the global repository only.
- Editing a policy, data, operator view, or service from the tab menu changes it in every project it is attached to and in the global repository.
- When you delete an item (a policy, data source, data type, operator view, or service) from the global repository, it is deleted (all versions of it) from the server and every project it is a member of.
- Deleting a policy, data model, service, or operator view from the tab menu deletes it everywhere. The item is deleted from the server, the global repository and from every project it was assigned to. You must be careful to delete only items that you want to delete globally.
- If you select the **Global** project, click an item in the selected tab and delete it. The item is removed from the server, global repository and from every project it was assigned to.
- The only safe way to delete an item from a project, without removing it permanently from the database, is to remove it in the project editor window.

Global repository

The global repository is the storage area for all the policies, data, operator views, and services for the cluster that you are connected to.

When you create an item on the **Data Model**, **Policies**, **Services** or **Operator View** tabs, the items are automatically added to the global repository.

You add new policies and their associated data and services to the global repository, just as you would to a project, but they are independent of any projects. You can attach added items to projects as project members at any time.

You must only edit and delete items that you want to change or delete globally. Deleting an item from the tab menu deletes it from the global repository and every project it is attached to.

A version control interface is provided so that you can use it to save data as revisions in a version control archive. You can also use the **Global** project to unlock all the items that you checked out.

Viewing the global repository data

All the available data sources, data types, policies, and services are stored in the global repository.

To view the items in the global repository select a specific tab, and select the **Global** option in the **Projects** menu. For example, open the **Operator View** tab and from the projects menu, select the **Global** project. You see all the operator views that are stored in the global repository.

Adding an item to the global repository

Each time you create an item, for example a data type, it is automatically added to the global repository.

The item that you created is listed automatically in the **Global** project in the specific tab.

Deleting an item from the global repository

Use this procedure to delete an item from the global repository.





Procedure

1. From the treeview, open the appropriate tab and find the item you want to delete.
2. From the **Projects** list, select **Global**.
3. Click the **Delete** icon on the tab menu bar.
4. Click **OK** in the confirmation box. The item is deleted from the global repository, and all projects it was assigned to.

Project panel controls

Use the project menu controls to create, edit, and delete projects.

Table 2. Project menu options

Project menu icons	Description
Project	Click to open a list of all defined projects for the selected cluster.
	Click this icon to create a project.
	Click this icon to edit the selected project. You can add or remove project members that are stored in the global repository to or from the project.
	Click this icon to delete the project. The project members remain in the global repository and in any other projects that they have been added to.
	The Unlock All icon is available only when you select the Global project. Click this icon to unlock all files you have locked. You can unlock only your own files. Refresh any open page tabs to update the display. See “Version control file locking” on page 21 and “Unlocking all locked items” on page 22 for more information.

Creating a project

Use this procedure to create and populate a project.

Procedure

1. In the navigation tree, expand **System Configuration > Event Automation**, click one of the links to open one of the tabs.
2. Select the required cluster and click the new project icon on the toolbar.
3. Use the project editor window to configure your new project.
 - a. In the **General Settings** section, type a unique name for the project in the **Project Name** field.

A default name is automatically given to the project but you can create a unique name for your project. You cannot edit a project name after the project is saved.
 - b. In the **Member Selection** section, you can add data sources, data types, policies, operator views, and services to your project.

From the **List By** list, select a group whose elements you want to add to your project.
 - c. In the **Global Repository** select an item, and click **Add** to include any of the members from the **Global Repository** pane to your project by moving them to the **Project Members** pane. To remove selected members from the project and return them to the **Global Repository**, click **Remove**. If you have not yet created any data sources, data types, policies, or services on your server, only predefined items are listed in the **Global Repository** pane.

For more information about the available options, see “Project editor configuration window” on page 19.
4. If you do not want to add any items to the project at this time, you can save the project.

Project editor configuration window

Use this information when you create a project, or configure an existing one.

1. If you are creating a project, assign a unique name for the project. A default name is automatically assigned to the project in the **Project Name** field. You cannot edit a project name after the project is saved.

Remember: If you use UTF-8 characters in the project name, make sure that the locale on the Impact Server where the project is saved is also set to the UTF-8 character encoding.

2. **List By** Contains a list of items including data sources, data types, policies, operator views, and services. When you select an item, for example, **Data Sources**, all the data sources that you have previously created, plus the predefined data sources are listed in the **Global Repository** pane. Similarly, when you select **Data types** all the data types that you have previously created, plus the predefined data types are listed in the **Global Repository**. Select the items you want to add to the project. To select more than one item at a time, press and hold the **Shift** key and then highlight the items you want.
3. To include existing data sources or data types in the current project, select them in Project members pane and add them to the project. If you are a Tivoli Business Service Manager user, select the **TBSM** project to add predefined data sources and data types to the project.
4. Click **Add** to include any of the members from the **Global Repository** pane to your project by moving them to the **Project Members** pane. Click **Remove** to remove selected members from the project and return them to the **Global Repository**.
5. Click **OK**.

Viewing project members

You can view the members of a single project by selecting the project name in the project menu for each open tab in the work area.

If you have not as yet created any projects, the **Default** and **Global** projects and predefined projects are the only projects listed in the **Projects** menu.

When you select a project, only the data sources, data types, policies, and services, that belong to the project are listed under each specific tab. For example, the **Policies** tab lists all the policies that are members of the project you selected.

Editing a project

Use this procedure to change the details of a project or to add or remove its members.

Procedure

1. From the **Project** list, select the project you want to edit.
2. Click the **Edit** icon to open the project editor window.

Edit the project, see “Project editor configuration window” if required. You cannot change the name of an existing project in the edit window. The only way to rename a project once it has been saved is to delete the project and create a new project with the new name.

Important: You cannot edit or delete the **Global** project.

3. Click **OK** to apply the changes. The changes you make to a project are unique to this project.

Deleting a project

Use this procedure to delete a project without removing the project members from other projects or from the global repository.

Procedure

1. From the **Project** menu, select the project you want to delete.
2. In the projects toolbar, click the **Delete Project** icon in the toolbar.

When you delete a project it is removed from the server. However, the project members that were assigned to it are not removed from other projects or from the global repository.

Important: You cannot edit or delete the **Global** project.

3. Click **OK** to confirm the deletion.

Automated project deployment feature

You can copy the data sources, data types, policies, and services in a project between two running server clusters on a network using the automated project deployment feature.

You can use this feature when moving projects from test environments into real-world production scenarios.

Important: Automated project deployment requires both server clusters to use the same Name Server.

When you copy data sources and types, policies, and services between clusters, you have the option of specifying a version control checkpoint label for the revisions of the items that you copy to the target server cluster. Two checkpoint labels are used for this process. The first is the label that you specify, which is applied to the copied versions of the project components. The second is the specified label with the string `_AFTER_DEPLOYMENT` appended. This label is applied to subsequent changes to the project components made using the GUI or CLI.

Revision checkpointing is supported only if you are using SVN and CVS version control system for Netcool/Impact. If you are using ClearCase or RCS, you can manually apply checkpoint labels to the revisions using the version control script or the tools provided by those applications.

To automatically deploy project, use one of the following options:

- Run the built-in **DeployProject** policy using the GUI. The **DeployProject** policy is a built-in policy that copies all the data sources, data types, policies, and services in a project between two running server clusters.
- Create and run a custom deployment policy that uses the **Deploy** function. See the *Policy Reference Guide* for more information.

Running the DeployProject policy

Use this procedure to run the **DeployProject** policy and copy all the data sources, data types, policies, and services in a project between two running server clusters.

Procedure

1. Select the server cluster from which you want to copy data from the main toolbar.
2. Select **DeployProject** from the list of policies.
The Policy Editor opens and displays the contents of the **DeployProject** policy.
3. Click **Trigger Policy** to open the Policy Runtime Parameters window.
For reference on the configuration options, see “DeployProject policy runtime parameters window.”
4. Click **OK** to save the configuration and close the window.
After you run the **DeployProject** policy, you can check the contents of the policy log for the results of the project deployment.

DeployProject policy runtime parameters window

Use this information to configure the DeployProject policy parameters.

Table 3. DeployProject policy runtime parameters window

Window element	Description
TargetCluster	Enter the name of the destination server cluster.
Username	Enter a valid user name.
Password	Enter a valid password.
Project	Enter the name of the project to copy.
Checkpoint ID	If you are using Subversion as the version control system you can type a checkpoint label. This label is applied to all project components when checked into the version control system for the target cluster. If you are not using Subversion or you do not want to use a checkpoint label, accept the default value for this field, which is NULL.

Version control file locking

Netcool/Impact is installed with a version control interface that you can use to save data as revisions in a version control archive.

When you create a policy, data source, data type, or service, a corresponding element is created in the version control system. When you open one of these items for viewing or editing, the item is automatically locked. When an item is locked, other users can view the item, but cannot edit the item until the lock has been released. When you save and close an item, the lock is automatically released and is available for editing by other users.

If required, for example, after the system goes down, you can use the global project to unlock the locked files. You can unlock only the items that you have checked out. If you have an item open for editing you cannot unlock it. Save and close the item.

Important: You cannot unlock an item if the lock belongs to another user. If you open a file locked by another user, the file will open in read-only mode. Only the lock owner or administrators with the **impactAdminUser** role can unlock the item in exceptional circumstances.

For details about unlocking your own locked files, see “Unlocking all locked items” on page 22.

Unlocking all locked items

Use this procedure to unlock all items that you have checked out.

Procedure

1. From the **Projects** menu, select the **Global** project.
2. Click **Unlock All** to unlock all the items that you have checked out. You can unlock only your own items. If you want to unlock an item that is owned by another user, contact an administrator assigned the **tipadmin** or **impactAdmin** user role.
3. A confirmation message shows when the files are unlocked. Click the **Refresh** icon at the top of each open page tab to update the display in the open page tabs.

Chapter 4. Data model

A data model is a model of the business data and metadata used in an Netcool/Impact solution.

DSA (Data Source Adapter) data models are sets of data sources, data types, and data items that represent information managed by the internal data repository or an external source of data. For each category of DSA, the data model represents different structures and units of data that are stored or managed by the underlying source. For example, for SQL database DSAs, data sources represent databases; data types represent database tables; and data items represents rows in a database table.

Data model components

A data model is made up of components that represent real world sources of data and the actual data inside them.

Data sources

Data sources are elements of the data model that represent real world sources of data in your environment.

Data types

Data types are elements of the data model that represent sets of data stored in a data source.

Data items

Data items are elements of the data model that represent actual units of data stored in a data source.

Links Links are elements of the data model that define relationships between data types and data items.

Event sources

Event sources are special types of data sources. Each event source represents an application that stores and manages events.

Setting up a data model

To set up a data model, you must first determine what data you need to use in your solution and where that data is stored. Then, you create a data source for each real world source of data and create a data type for each structural element that contains the data you need.

Procedure

1. Create data sources

Identify the data you want to use and where it is stored. Then, you create one data source for each real world source of data. For example, if the data is stored in one MySQL database and one LDAP server, you must create one MySQL and one LDAP data source.

2. Create data types

After you have set up the data sources, you create the required data types. You must create one data type for each database table (or other data element,

depending on the data source) that contains data you want to use. For example, if the data is stored in two tables in an Oracle database, you must create one data type for each table.

3. Optional: Create data items

For most data types, the best practice is to create data items using the native tools supplied by the data source. For example, if your data source is an Oracle database, you can add any required data to the database using the native Oracle tools. If the data source is the internal data repository, you must create data items using the GUI.

4. Optional: Create links

After you create data types, you can define linking relationships between them using dynamic links. You can also define linking relationships between internal data items using static links. That makes it easier to traverse the data programmatically from within a policy. Use of links is optional.

5. Create event sources

Most process events are retrieved from a Netcool/OMNIBus ObjectServer. The ObjectServer is represented in the data model as an event source.

Accessing the data model tab

Use this procedure to access the data model tab.

Procedure

1. In the navigation tree, expand **System Configuration > Event Automation** click **Data Model** to open the **Data Model** tab.
2. From the **Cluster** list, select the cluster you want to use.
3. From the **Project** list, select the project you want to use. The data sources that are available to the project are displayed in the **Data Model** tab.

Data model menu controls

This topic gives an overview of the controls that are used in the data model menu.

Table 4. Data model menu controls








Icon	Description
	Click this icon to create a data source. Select one of the available data source types from the list. After you create a data source, you can right-click the data source and click New Data Type to create an associated data type.
	Select a data source and click this icon to create a data type for the selected data source. After you create a data type, it is listed under its data source. Alternatively, you can right-click a data source and select New Data Type to create a data type for this data source.
	Select an element in the list and click this icon to edit it. Alternatively, right-click an item in the list and select Edit in the menu.
	Click to view the selected data type in the editor panel. Select the View Data Items option to view the data items for the data type, or the View Performance Report option to review a performance report for the data type. Alternatively, you can view the data items or the performance report for a data type by right-clicking the data type.

Table 4. Data model menu controls (continued)

Icon	Description
	Click this icon to test the connection to the data source. Alternatively, right-click an item in the list and select Test Connection in the menu. Important: If you see an error message telling you that the data source cannot establish a connection to a database because a JDBC driver was not found, it means that a required JDBC driver is missing in the shared library directory. To fix this, place a licensed JDBC driver in the shared library directory and restart the server. For more information see, the “SQL database DSAs” chapter in the <i>DSA Reference Guide</i> .
	Click the Delete icon to delete a data source or type from the server. Alternatively, you can right-click a data source or type and select Delete . This action deletes an item permanently from the database. To safely remove a data type from only one project and not from the database, use the project editor.
	This icon is visible when a data source or data type item is locked, or the item is being used by another user. Hover the mouse over the locked item to see which user is working on the item. You can unlock your own items but not items locked by other users. If you have an item open for editing you cannot unlock it. Save and close the item. To unlock an item you have locked, right click on the item name and select Unlock . The tipadmin user and users who are assigned the impactAdminUser role are the only users who can unlock items that are locked by another user in exceptional circumstances.

Data sources overview

Data sources provide an abstract layer between Netcool/Impact and real world source of data.

Internally, data sources provide connection and other information that Netcool/Impact uses to access the data. When you create a data model, you must create one data source for every real world source of data you want to access in a policy.

The internal data repository of Netcool/Impact can also be used as a data source.

Data source categories

Netcool/Impact supports four categories of data sources.

SQL database data sources

An SQL database data source represents a relational database or another source of data that can be accessed using an SQL database DSA.

LDAP data sources

The LDAP data source represent LDAP directory servers.

Mediator data sources

Mediator data sources represent third-party applications that are integrated with Netcool/Impact through the DSA Mediator.

JMS data sources

A JMS data source abstracts the information that is required to connect to a JMS Implementation.

List of data sources

You can create you own data sources from the following list:

Table 5. User-defined data sources

















Data source	Type	Description
 CorbaMediator	Mediator	The Mediator data source represents third-party applications that are integrated with Netcool/Impact through the DSA Mediator.
 DB2	SQL database	You use the DB2 DSA to access information in an IBM DB2 database.
 DirectMediator	Mediator	The Mediator data source represents third-party applications that are integrated with Netcool/Impact through the DSA Mediator.
 Flat File	SQL database	You use the Flat File DSA to read information in a character-delimited text file. The flat file DSA is read only which means that you cannot add new data items in GUI. To create a flat file data source you need a text file that is already populated with data.
 Generic SQL	SQL database	You use the Generic SQL DSA to access information in any database application through a JDBC driver.
 Informix	SQL database	You use the Informix [®] DSA to access information in an IBM Informix database.
 JMS	Messaging API	A JMS data source abstracts the information that is required to connect to a JMS Implementation.
 LDAP	Directory Server	The LDAP data source represent LDAP directory servers. The LDAP DSA supports only non-authenticating data sources.
 MS SQL Server	SQL database	You use the MS-SQL Server DSA to access information in a Microsoft SQL Server database.
 MySQL	MySQL database	You use the MySQL DSA to access information in a MySQL database.
 ObjectServer	SQL database	The ObjectServer data source represents the instance of the Netcool/OMNIBus ObjectServer that you monitor using the OMNIBus event listener service, or OMNIBus event reader service.
 ODBC	SQL database	You use the ODBC DSA to access information in an ODBC data source.
 Oracle	SQL database	You use the Oracle DSA to access information in an Oracle database.










Table 5. User-defined data sources (continued)

Data source	Type	Description
PostgreSQL 	SQL database	You use the PostgreSQL DSA to access information in an PostgreSQL database.
SNMPDirectMediator	Mediator	The Mediator data source represents third-party applications that are integrated with Netcool/Impact through the DSA Mediator.
Sybase 	SQL database	You use the Sybase DSA to access information in a Sybase database.
HSQldb 	SQL database	You use the HSQL DSA to access information in a HSQL database.

List of predefined data sources

Initially, the data sources listed in the global project are predefined data sources.

Table 6. predefined data sources

Data source	Description
defaultobjectserver 	The default ObjectServer data source. The defaultobjectserver data source is configured during the installation, when you create an instance of the Impact Server.
ITNM 	The ITNM data source is used with ITNM and the ITNM DSA.
ReportsHSQldb 	ReportsHSQldb represents the database where the reporting data is stored.
SocketMediatorDataSource 	The SocketMediator data source is used with the Socket DSA.
XmlDsaMediatorDataSource 	The XmlDsaMediator data source is used with the XML DSA.
URL 	The URL data source contains the predefined data type document. You cannot edit the URL data source but you can add additional data types.
Schedule 	The Schedule data source contains the predefined data type schedule. You cannot edit the schedule data source but you can add additional data types.
Statistics 	The Statistics data source contains the hibernation data type. You cannot edit the statistics data source or add additional data types.
Internal 	The Internal data source contains the following predefined data types, TimeRangeGroup, LinkType, and FailedEvent.

Creating data sources

Use this procedure to create a user-defined data source.

Procedure

1. In the navigation tree, expand **System Configuration** > **Event Automation** click **Data Model** to open the **Data Model** tab.
2. From the **Cluster** and **Project** lists, select the cluster and project you want to use.
3. In the **Data Model** tab, click the **New Data Source** icon in the toolbar. Select a template for the data source that you want to create. The tab for the data source opens.
4. Complete the required information, and click **Save** to create the data source.

Editing data sources

Use this procedure to configure an existing data source.

Procedure

1. In the **Data Model** tab, double-click the name of the data source that you want to edit. Alternatively, right click the data source and click **Edit**.
2. Make the changes and click **Save** to apply them.

Deleting data sources

Before deleting a data source, you must first delete any data types listed under the data source.

If you do not delete them, you get an error message when you try to delete the data source. When you delete a data source from within a project, it is also deleted from any other projects that use it and from the global repository. To remove a data source from one project, use the editor window for that project. For more information about removing data sources from a project, see "Project editor configuration window" on page 19.

In the **Data Model** tab, select the data source you want to delete, and click the delete icon on the toolbar. Alternatively, right click the data source and select **Delete**.

Testing data source connections

If you have defined a backup data source, both the primary and backup data source connections are tested.

If the test succeeds for the primary connection, you get a message indicating that it was successful. If the test fails for the primary source, the backup source is then tested. If the backup succeeds, you get a message that the connection was successful. It is only when both the primary and backup tests fail that you receive a message that the connection could not be made.

Data types overview

Data types describe the content and structure of the data in the data source table and summarize this information so that it can be accessed during the execution of a policy.

Data types provide an abstract layer between Netcool/Impact and the associated set of data in a data source. Data types are used to locate the data you want to use in a policy. For each table or other data structure in your data source that contains information you want to use in a policy, you must create one data type. To use a data source in policies, you must create data types for it.

Attention: Some system data types are not displayed in the GUI. You can manage these data types by using the Command Line Interface (CLI).

The structure of the data that is stored in a data source depends on the category of the data source where the data is stored. For example, if the data source is an SQL database, each data type corresponds to a database table. If the data source is an LDAP server, each data type corresponds to a type of node in the LDAP hierarchy.

A data type definition contains the following information:

- The name of the underlying table or other structural element in the data source
- A list of fields that represent columns in the underlying table or another structural element (for example, a type of attribute in an LDAP node)
- Settings that define how Netcool/Impact caches data in the data type

Data type categories

Netcool/Impact supports four categories of data types.

SQL database data types

SQL database data types represent data stored in a database table.

LDAP data types

LDAP data types represent data stored at a certain base context level of an LDAP hierarchy.

Mediator data types

Mediator data types represent data that is managed by third-party applications such as a network inventory manager or a messaging service.

Internal data types

You use internal stored data types to model data that does not exist, or cannot be easily created, in external databases.

Predefined data types overview

Predefined data types are special data types that are stored in the global repository.

You can edit some predefined data types by adding new fields, but you cannot edit or delete existing fields. You can view, edit, create, and delete data items of some predefined data types by using the GUI. You cannot delete predefined data types except for the **FailedEvent** predefined data type.

List of predefined data types

An overview of the predefined data types available in the global project.

Table 7. Predefined data types








Data type	Type	Description
 Schedule	Editable	Schedules define a list of data items associated with specific time ranges, or time range groups, that exist.

Table 7. Predefined data types (continued)

Data type	Type	Description
 Document	Editable	Custom URL Document data types are derived from the predefined Doc data type.
 FailedEvent	Editable	The FailedEvent data type, together with the ReprocessedFailedEvents policy, provides you with a way to deal with failed events that are passed from the ObjectServer.
 ITNM	Editable	This data type is used with ITNM and the ITNM DSA.
 TimeRangeGroup	Non-editable	A time range group data type consists of any number of time ranges.
 LinkType	Non-editable	The LinkType data type provides a way of defining named and hierarchical dynamic links.
 Hibernate	Non-editable	When you call the Hibernate function in a policy, the policy is stored as a Hibernate data item for a certain number of seconds.

Viewing data types

You view data types in the data navigator panel.

Before you have created any data types, you see only the data source type selection list. Each time you create a data type, you first create the data source you want it to connect to. After you configure a data type, it is listed in the data connections panel under the associated data source.

Editing data types

Use this procedure to edit an existing data type.

Procedure

1. In the navigation tree, expand **System Configuration > Event Automation > Data Model**, to open the **Data Model** tab.
2. Expand the data source that contains the data type you want to edit, select the data type, double-click the name of the data type that you want to edit. Alternatively, right-click the data source and click **Edit**.
3. Make the required changes in the **Data type** tab.
4. Click **Save** to apply the changes.

Deleting data types

Use the following procedure to delete a data type.

Procedure

1. From the list of data sources and types, locate the data type you want to delete.
2. Select the data type, right-click and select **Delete**, or click the **Delete** icon on the toolbar.

Attention: When you delete a data type from within project or the global repository, it is also deleted from any other projects that use it. To remove a data type from one project, open the editor window for that project.

Data items overview

Data items are elements of the data model that represent actual units of data stored in a data source.

You create internal data items individually in the data items viewer. External data items are created automatically when a policy references the data type to which they belong, by a lookup in the external database.

Attention: The LDAP data type, which uses the LDAP DSA, is a read-only data type. Therefore you cannot edit or delete LDAP data items from within the GUI.

Links overview

Links are an element of the data model that defines relationships between data items and between data types.

They can save time during the development of policies because they allow you to define a data relationship once and then reuse it several times when you need to find data related to other data in a policy. Links are an optional part of a data model. Dynamic links and static links are supported.

Link categories

Netcool/Impact provides two categories of links.

Static links

Static links define a relationship between data items in internal data types.

Dynamic links

Dynamic links define a relationship between data types.

Chapter 5. Working with data sources

Using the GUI you can view, create, edit, and delete data sources.

Data sources

Data sources are elements of the data model that represent real world sources of data in your environment.

These sources of data include third-party SQL databases, LDAP directory servers, or other applications such as messaging systems and network inventory applications.

Data sources contain the information that you need to connect to the external data. You create a data source for each physical source of data that you want to use in your Impact solution. When you create an SQL database, LDAP, or Mediator data type, you associate it with the data source that you created. All associated data types are listed under the data source in the Data Sources and Types task pane. When you create a data type, you simply select the data source it must use.

SQL database DSA failover

Failover is the process by which an SQL database DSA automatically connects to a secondary database server (or other data source) when the primary server becomes unavailable.

This feature ensures that Netcool/Impact can continue operations despite problems accessing one or the other server instance. You can configure failover separately for each data source that connects to a database using an SQL Database DSA.

SQL database DSA failover modes

Standard failover, failback, and disabled failover are supported failover modes for SQL database DSAs.

Standard failover

Standard failover is a configuration in which an SQL database DSA switches to a secondary database server when the primary server becomes unavailable and then continues using the secondary until Netcool/Impact is restarted.

Failback

Failback is a configuration in which an SQL database DSA switches to a secondary database server when the primary server becomes unavailable and then tries to reconnect to the primary at intervals to determine whether it has returned to availability.

Disabled failover

If failover is disabled for an SQL database DSA the DSA reports an error to Netcool/Impact when the database server is unavailable and does not attempt to connect to a secondary server.

SNMP data sources

SNMP data sources represent an agent in the environment.

The data source configuration specifies the host name and port where the agent is running, and the version of SNMP that it supports. For SNMP v3, the configuration also optionally specifies authentication properties.

You can either create one data source for each SNMP agent that you want to access using the DSA, or you can create a single data source and use it to access all agents. You can create and configure data sources using the GUI. After you create a data source, you can create one or more data types that represent the OIDs of variables managed by the corresponding agent.

Socket DSA data source

The socket DSA data source is named `SocketMediatorDataSource`.

This data source is inserted into Netcool/Impact automatically at installation. Do not change the data source configuration properties at any time.

SQL database data sources

An SQL database data source represents a relational database or another source of data that can be accessed using an SQL database DSA.

Most commonly used commercial relational databases are supported, such as Oracle, Sybase, and Microsoft SQL Server. In addition, freely available databases like MySQL, and PostgreSQL are also supported. The Netcool/OMNIBus ObjectServer is also supported as a SQL data source.

The configuration properties for the data source specify connection information for the underlying source of data. Some examples of SQL database data sources are:

- A DB2 database
- A MySQL database
- An application that provides a generic ODBC interface
- A character-delimited text file

You create SQL database data sources using the GUI. You must create one such data source for each database that you want to access. When you create an SQL database data source, you need to specify such properties as the host name and port where the database server is running, and the name of the database. For the flat file DSA and other SQL database DSAs that do not connect to a database server, you must specify additional configuration properties.

Note that SQL database data sources are associated with databases rather than database servers. For example, an Oracle database server can host one or a dozen individual databases. Each SQL database data source can be associated with one and only one database.

DB2 data source configuration

Use this information to create a DB2 data source.

Table 8. DB2 Data Source configuration

Window element	Description
General Settings	

Table 8. DB2 Data Source configuration (continued)

Window element	Description
Data Source Name	Type a unique name to identify the data source. Only letters, numbers, and the underscore character must be used in the data source name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data source is saved is set to the UTF-8 character encoding.
Username	Type a user name with which you can access the database.
Password	Type a password that allows you access to the database. As you type, the characters are replaced with asterisks (*).
Maximum SQL Connection	<p>When you configure the maximum number of threads for the event processor service, for optimal performance the number of processing threads should be greater than, or equal to the size of the connection pool of the SQL data sources used in the policies being triggered. For information about viewing existing thread and connection pool information, see the information in the <i>Administration Guide</i> in the section <i>Command-Line tools, Event Processor commands</i>. See the Select PoolConfig from Service where Name='EventProcessor';</p> <p>Important: In a clustered environment, the event processor configuration is not replicated between servers. You must run the Select PoolConfig from Service where Name='EventProcessor'; command on the primary and the secondary servers.</p> <p>Limiting the number of concurrent connections manages performance. Type the maximum number of connections allowed to the database at any one time. That number has to be greater than, or equal to, the number of threads running in the Event Processor. See “Event processor service configuration window” on page 127.</p> <p>The default value is 5.</p>
Database Failure Policy	<p>Select the failover option. Available options are Fail over, Fail back, and Disable Backup.</p> <p>For more information about failover options, see “SQL database DSA failover modes” on page 33.</p>
Primary Source	
Host Name	Type the host name. Default value is localhost.
Port	Type or select a port number. The default number is 50000.
Database	Type the name of the database to connect to.
Test Connection	Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.
Backup Source (optional)	

Table 8. DB2 Data Source configuration (continued)

Window element	Description
Host Name	Type the host name. The default value is localhost.
Port	Type or select a port number. The default value is 50000.
Database	Type the name of the database to connect to.
Test Connection	Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.

Creating flat file data sources

Flat file DSA is read only which means that you cannot add new data items in GUI.

Procedure

- To create a flat file data source you need a text file that is already populated with data.
For example, create a /home/impact/myflatfile.txt file with the following content:
Name, Age
Ted, 11
Bob, 22
- In the **Data Model** tab, click the **New Data Source** icon and click **Flat File**. The **New Flat File** tab opens.
- Enter the required information
 - Enter a unique name for your data source name, for example MyFlatFileDataSource.
 - In the **Directory** field, provide the path to your flat file, for example /home/impact.
 - In the **Delimiters** field, specify the delimiters that you used in your flat file, for example ','. Use only single quotation marks '. The header row of the flat file supports the use of the following characters ;:/+|,\t\n\r\f and <space>. The remaining rows of the flat file support the use of the following characters ;:/+|,\t and <space>.
- Click **Save** to finish creating a new flat file data source.

What to do next

Use the data source that you just created to create a flat file data type. For more information about creating flat file data types, see "Creating flat file data types" on page 76.

Flat file data source configuration

Use this information to create a flat file data source.

Window element	Description
General Settings	

Window element	Description
Data Source Name	Type a unique name to identify the data source. Only letters, numbers, and the underscore character must be used in the data source name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data source is saved is set to the UTF-8 character encoding.
Source	
Directory	The path to the directory that contains the flat file.
Delimiters	Characters that separate the information tokens in the flat file. The characters must be enclosed in single quotation marks, for example: ' ; ; - + / '. The header row of the flat file supports the use of the following characters, ; : / + , \ t \ n \ r \ f and <space>. All other rows support the use of the following characters, ; : / + - , \ t and <space>.

Informix data source configuration

Use this information to create an Informix data source.

Table 9. Informix Data Source window

Window element	Description
General Settings	
Data Source Name	Type a unique name to identify the data source. Only letters, numbers, and the underscore character must be used in the data source name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data source is saved is set to the UTF-8 character encoding.
Username	Type a user name with which you can access the database.
Password	Type a password that allows you access to the database. As you type, the characters are replaced with asterisks (*).

Table 9. Informix Data Source window (continued)

Window element	Description
Maximum SQL Connection	<p>When you configure the maximum number of threads for the event processor service, for optimal performance the number of processing threads should be greater than, or equal to the size of the connection pool of the SQL data sources used in the policies being triggered. For information about viewing existing thread and connection pool information, see the information in the <i>Administration Guide</i> in the section <i>Command-Line tools, Event Processor commands</i>. See the Select PoolConfig from Service where Name='EventProcessor';</p> <p>Important: In a clustered environment, the event processor configuration is not replicated between servers. You must run the Select PoolConfig from Service where Name='EventProcessor' command on the primary and the secondary servers.</p> <p>Limiting the number of concurrent connections manages performance. Type or select the maximum number of connections allowed to the database at one time. That number has to be greater than or equal to the number of threads running in the Event Processor. See “Event processor service configuration window” on page 127.</p> <p>The default value is 5.</p>
Database Failure Policy	<p>Select the failover option. Available options are Fail over, Fail back, and Disable Backup.</p> <p>For more information about failover options, see “SQL database DSA failover modes” on page 33.</p>
Primary Source	
Host Name	Type the host name. Default value is localhost.
Port	Select a port number. The default number is 1526.
Server	Type the name of the server where the database resides.
Database	Type the name of the database to connect to.
Test Connection	Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.
Backup Source (optional)	
Host Name	Type the host name. The default value is localhost.
Port	Select a port number. The default value is 1526.
Server	Type the name of the server where the database resides.
Database	Type the name of the database to connect to.

Table 9. Informix Data Source window (continued)

Window element	Description
Test Connection	Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.

MS-SQL Server data source configuration

Use this information to create an MS_SQL Server data source.

Table 10. MS-SQL Server Data Source window

Window element	Description
General Settings	
Data Source Name	Type a unique name to identify the data source. Only letters, numbers, and the underscore character must be used in the data source name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data source is saved is set to the UTF-8 character encoding.
User name	Type a user name with which you can access the database.
Password	Type a password that allows you access to the database. As you type, the characters are replaced with asterisks (*).
Maximum SQL Connection	<p>When you configure the maximum number of threads for the event processor service, for optimal performance the number of processing threads should be greater than, or equal to the size of the connection pool of the SQL data sources used in the policies being triggered. For information about viewing existing thread and connection pool information, see the information in the <i>Administration Guide</i> in the section <i>Command-Line tools, Event Processor commands</i>. See the Select PoolConfig from Service where Name='EventProcessor';</p> <p>Important: In a clustered environment, the event processor configuration is not replicated between servers. You must run the Select PoolConfig from Service where Name='EventProcessor'; command on the primary and the secondary servers.</p> <p>Limiting the number of concurrent connections manages performance. Type or select the maximum number of connections allowed to the database at one time. That number has to be greater than or equal to the number of threads running in the Event Processor. See “Event processor service configuration window” on page 127.</p> <p>The default value is 5.</p>

Table 10. MS-SQL Server Data Source window (continued)

Window element	Description
Database Failure Policy	Select the failover option. Available options are Fail over , Fail back , and Disable Backup . For more information about failover options, see “SQL database DSA failover modes” on page 33.
Primary Source	
Host Name	Type the host name. The default value is localhost.
Port	Select a port number. The default number is 1433.
Database	Type the name of the database to connect to.
Test Connection	Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later. Important: If you see an error message telling you that the data source cannot establish a connection to a database because a JDBC driver was not found, it means that a required JDBC driver is missing in the shared library directory. To fix this, place a licensed JDBC driver in the shared library directory and restart the server. For more information see, the “SQL database DSAs” chapter in the <i>DSA Reference Guide</i> .
Backup Source (optional)	
Host Name	Type the host name. The default value is localhost.
Port	Select a port number. The default value is 1433.
Database	Type the name of the database to connect to.
Test Connection	Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.

MySQL data source configuration

Use this information to create a new MySQL data source.

Table 11. MySQL Data Source window

Window element	Description
General Settings	
Data Source Name	Type a unique name to identify the data source. Only letters, numbers, and the underscore character must be used in the data source name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data source is saved is set to the UTF-8 character encoding.
Username	Type a valid user name with which you can access the database.
Password	Type a valid password that allows you access to the database. As you type, the characters are replaced with asterisks (*).

Table 11. MySQL Data Source window (continued)

Window element	Description
Maximum SQL Connection	<p>When you configure the maximum number of threads for the event processor service, for optimal performance the number of processing threads should be greater than, or equal to the size of the connection pool of the SQL data sources used in the policies being triggered. For information about viewing existing thread and connection pool information, see the information in the <i>Administration Guide</i> in the section <i>Command-Line tools, Event Processor commands</i>. See the Select PoolConfig from Service where Name='EventProcessor';</p> <p>Important: In a clustered environment, the event processor configuration is not replicated between servers. You must run the Select PoolConfig from Service where Name='EventProcessor'; command on the primary and the secondary servers.</p> <p>Limiting the number of concurrent connections manages performance. Type or select the maximum number of connections allowed to the database at one time. For best performance this number should be greater than or equal to the maximum number of event processor threads. See “Event processor service configuration window” on page 127.</p> <p>The default value is 5.</p>
Database Failure Policy	<p>Select the failover option. Available options are Fail over, Fail back, and Disable Backup.</p> <p>For more information about failover options, see “SQL database DSA failover modes” on page 33.</p>
Primary Source	
Host Name	Type the host name or IP address of the system where the data source is located. The default value is localhost.
Port	Select the port number used by the data source. The default number is 3306.
Database	Type the name of the database to connect to.
Test Connection	<p>Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.</p> <p>Important: If you see an error message telling you that the data source cannot establish a connection to a database because a JDBC driver was not found, it means that a required JDBC driver is missing in the shared library directory. To fix this, place a licensed JDBC driver in the shared library directory and restart the server. For more information see, the “SQL database DSAs” chapter in the <i>DSA Reference Guide</i>.</p>
Backup Source (optional)	
Host Name	Type the host name or IP address of the system where the backup data source is located. Optional. The default value is localhost.

Table 11. MySQL Data Source window (continued)

Window element	Description
Port	Select a port number used by the backup data source. Optional. The default value is 3306.
Database	Type the name of the database to connect to.
Test Connection	Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.

ObjectServer data source configuration

Use this information to create a new ObjectServer data source.

Table 12. New ObjectServer Data Source window

Window element	Description
General Settings	
Data Source Name	Type a unique name to identify the data source. Only letters, numbers, and the underscore character must be used in the data source name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data source is saved is set to the UTF-8 character encoding.
User name	Type a user name with which you can access the database.
Password	Type a password that allows you access to the database.
Maximum SQL Connection	<p>When you configure the maximum number of threads for the event processor service, for optimal performance the number of processing threads should be greater than, or equal to the size of the connection pool of the SQL data sources used in the policies being triggered. For information about viewing existing thread and connection pool information, see the information in the <i>Administration Guide</i> in the section <i>Command-Line tools, Event Processor commands</i>. See the Select PoolConfig from Service where Name='EventProcessor';</p> <p>Important: In a clustered environment, the event processor configuration is not replicated between servers. You must run the Select PoolConfig from Service where Name='EventProcessor'; command on the primary and the secondary servers.</p> <p>Limiting the number of concurrent connections manages performance. Type or select the maximum number of connections allowed to the database at one time. That number has to be greater than or equal to the number of threads running in the Event Processor. See “Event processor service configuration window” on page 127.</p> <p>The default value is 5.</p>

Table 12. New ObjectServer Data Source window (continued)

Window element	Description
Database Failure Policy	Select the failover option. Available options are Fail over , Fail back , and Disable Backup . For more information about failover options, see “SQL database DSA failover modes” on page 33.
Primary Source	
Host Name	Type the host name. The default value is localhost.
Port	Select a port number. The default number is 4100.
SSL Mode: Enable	Select if this data source connects to the ObjectServer through SSL.
Test Connection	Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.
Backup Source (optional)	
Host Name	Type the host name. The default value is localhost.
Port	Select a port number. The default number is 4100.
SSL Mode: Enable	Select if this data source connects to the ObjectServer through SSL.
Test Connection	Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.

ODBC data source configuration

Use this information to create a new ODBC data source.

Table 13. ODBC Data Source window

Window element	Description
General Settings	
Data Source Name	Type a unique name to identify the data source. Only letters, numbers, and the underscore character must be used in the data source name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data source is saved is set to the UTF-8 character encoding.
User name	Type a user name that allows you access to the database.
Password	Type a password that allows you access to the database. As you type, the characters are replaced with asterisks (*).

Table 13. ODBC Data Source window (continued)

Window element	Description
Maximum SQL Connection	<p>When you configure the maximum number of threads for the event processor service, for optimal performance the number of processing threads should be greater than, or equal to the size of the connection pool of the SQL data sources used in the policies being triggered. For information about viewing existing thread and connection pool information, see the information in the <i>Administration Guide</i> in the section <i>Command-Line tools, Event Processor commands</i>. See the Select PoolConfig from Service where Name='EventProcessor';</p> <p>Important: In a clustered environment, the event processor configuration is not replicated between servers. You must run the Select PoolConfig from Service where Name='EventProcessor'; command on the primary and the secondary servers.</p> <p>Limiting the number of concurrent connections manages performance. Type or select the maximum number of connections allowed to the database at one time. That number has to be greater than or equal to the number of threads running in the Event Processor. See “Event processor service configuration window” on page 127.</p> <p>The default value is 5.</p>
Database Failure Policy	<p>Select the failover option. Available options are Fail over, Fail back, and Disable Backup.</p> <p>For more information about failover options, see “SQL database DSA failover modes” on page 33.</p>
Primary Source	
ODBC Name	Type the ODBC name.
Test Connection	Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.
Backup Source	
ODBC Name	When you select the Database Failure Policy as either Fail over or, Fail back , you must specify a Backup Source . If you select the Database Failure Policy as Disable Backup , the Backup Source field is not required.
Test Connection	Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.

Oracle data source configuration

Use this information to create a new Oracle data source.

Table 14. Oracle Data Source window

Window element	Description
General Settings	
Data Source Name	Type a unique name to identify the data source. Only letters, numbers, and the underscore character must be used in the data source name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data source is saved is set to the UTF-8 character encoding.
User name	Type a user name that allows you access to the database.
Password	Type a password that allows you access to the database. As you type, the characters are replaced with asterisks (*).
Maximum SQL Connection	<p>When you configure the maximum number of threads for the event processor service, for optimal performance the number of processing threads should be greater than, or equal to the size of the connection pool of the SQL data sources used in the policies being triggered. For information about viewing existing thread and connection pool information, see the information in the <i>Administration Guide</i> in the section <i>Command-Line tools, Event Processor commands</i>. See the Select PoolConfig from Service where Name='EventProcessor';</p> <p>Important: In a clustered environment, the event processor configuration is not replicated between servers. You must run the Select PoolConfig from Service where Name='EventProcessor'; command on the primary and the secondary servers.</p> <p>Limiting the number of concurrent connections manages performance. Type the maximum number of connections allowed to the database at one time. That number has to be greater than or equal to the number of threads running in the Event Processor. See “Event processor service configuration window” on page 127.</p> <p>The default value is 5.</p>
Connection Options	<p>Type of connection to an Oracle data source. Select one of the following options:</p> <ul style="list-style-type: none">• General Settings - default settings.• Customized URL - integration with an Oracle RAC cluster is supported. For more information, see “Connecting to Oracle RAC cluster” on page 47.• LDAP Data Source - connect to data source bound in a Naming Service using LDAP. For more information, see “Connecting to an Oracle data source using LDAP” on page 47.

Table 14. Oracle Data Source window (continued)

Window element	Description
Context Factory	<p>The class name to initialize the context. It is dependent on which Naming Service is used. For example, <code>com.sun.jndi ldap.LdapCtxFactory</code>.</p> <p>This option is displayed only if you choose LDAP Data Source in the Connection Options.</p>
Provider URL	<p>The URL used to connect to the Naming service. For example, <code>ldap://localhost:389/dc=abc</code>.</p> <p>This option is displayed only if you choose LDAP Data Source in the Connection Options.</p>
Binding Name	<p>The name to which the Oracle Data Source object is bound. For information about the binding name, refer the docs of the Naming Service provider. For example, <code>cn=myDataSource</code>.</p> <p>This option is displayed only if you choose LDAP Data Source in the Connection Options.</p>
Database Failure Policy	<p>Select the failover option. Available options are Fail over, Fail back, and Disable Backup.</p> <p>For more information about failover options, see “SQL database DSA failover modes” on page 33.</p>
Primary Source	
Host Name	Type a primary host name. The default value is <code>localhost</code> .
Port	Select a primary port number. The default value is set to a common port number: 1521.
SID	Type a primary Oracle service identifier. The default value is <code>ORCL</code> . For more information, see your Oracle documentation.
Test Connection	<p>Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.</p> <p>Important: If you see an error message telling you that the data source cannot establish a connection to a database because a JDBC driver was not found, it means that a required JDBC driver is missing in the shared library directory. To fix this, place a licensed JDBC driver in the shared library directory and restart the server. For more information see, the “SQL database DSAs” chapter in the <i>DSA Reference Guide</i>.</p>
Backup Source	
Host Name	Type a backup host name. The default value is <code>localhost</code> . Backup host name is optional.
Port	Select a secondary port number. The default value is set to a common port number: 1521. Backup port number is optional.
SID	Type a backup SID. The default value is <code>ORCL</code> . For more information, see your Oracle documentation. Backup SID is optional.

Table 14. Oracle Data Source window (continued)

Window element	Description
Test Connection	<p>Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.</p> <p>This button is disabled when the backup source information is left blank.</p>

Connecting to an Oracle data source using LDAP

Use this information to connect to an Oracle data source bound in a Naming Service using LDAP.

Select the **LDAP DataSource** option in the **Connection Option** list. The user name and password that you use to connect to an Oracle data source bound in a Naming Service using LDAP are the credentials that are required to access the Naming Service (not the database login credentials).

The Oracle data source should have the necessary information required to access the database (like user name, password, SID, host, port) already configured in it. This information is used behind the scenes to connect to the datasource. If the connection to the database is successful Connection OK message is displayed.

For more information about Oracle data sources, refer to the *Oracle JDBC Developer's Guide and Reference*.

For example, you use OpenLDAP as the Naming Service, and an Oracle data source is already bound to a logical name (cn=myDataSource in the **Binding Name** field). When you click **TestConnection**, the first connection is made to the naming service using LDAP and when the connection is established, Netcool/Impact looks for an Oracle data source for the logical name cn=myDataSource.

Connecting to Oracle RAC cluster

Netcool/Impact supports integration with an Oracle RAC cluster.

If you choose to connect to an Oracle RAC cluster, in the **URL** field, enter the URL, preceded by jdbc:oracle:thin:@. For example:

```
jdbc:oracle:thin:@
(DESCRIPTION = (ADDRESS = (PROTOCOL = TCP)(HOST = host1)(PORT = port1))
(ADDRESS = (PROTOCOL = TCP)(HOST = host2)(PORT = port2))
(LOAD_BALANCE = yes)
(CONNECT_DATA =
(SERVER = DEDICATED)
(SERVICE_NAME = service-name)
(FAILOVER_MODE =(TYPE = SELECT)(METHOD = BASIC)(RETRIES = 180)(DELAY = 5))
)
)
```

PostgreSQL data source configuration

Use this information to create a new PostgreSQL data source.

Table 15. PostgreSQL Data Source window

Window element	Description
General Settings	
Data Source Name	Type a unique name to identify the data source. Only letters, numbers, and the underscore character must be used in the data source name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data source is saved is set to the UTF-8 character encoding.
User name	Type a user name that allows you access to the database.
Password	Type a password that allows you access to the database. As you type, the characters are replaced with asterisks (*).
Maximum SQL Connection	<p>When you configure the maximum number of threads for the event processor service, for optimal performance the number of processing threads should be greater than, or equal to the size of the connection pool of the SQL data sources used in the policies being triggered. For information about viewing existing thread and connection pool information, see the information in the <i>Administration Guide</i> in the section <i>Command-Line tools, Event Processor commands</i>. See the Select PoolConfig from Service where Name='EventProcessor';</p> <p>Important: In a clustered environment, the event processor configuration is not replicated between servers. You must run the Select PoolConfig from Service where Name='EventProcessor'; command on the primary and the secondary servers.</p> <p>Limiting the number of concurrent connections manages performance. Type the maximum number of connections allowed to the database at one time. That number has to be greater than or equal to the number of threads running in the Event Processor. See “Event processor service configuration window” on page 127.</p> <p>The default value is 5.</p>
Database Failure Policy	<p>Select the failover option. Available options are Fail over, Fail back, and Disable Backup.</p> <p>For more information about failover options, see “SQL database DSA failover modes” on page 33.</p>
Primary Source	
Host Name	Type the host name. The default value is localhost.
Port	Select a port number. The default number is 5432.
Database	Type the name of the database to connect to.

Table 15. PostgreSQL Data Source window (continued)

Window element	Description
Test Connection	Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later. Important: If you see an error message telling you that the data source cannot establish a connection to a database because a JDBC driver was not found, it means that a required JDBC driver is missing in the shared library directory. To fix this, place a licensed JDBC driver in the shared library directory and restart the server. For more information see, the “SQL database DSAs” chapter in the <i>DSA Reference Guide</i> .
Backup Source (optional)	
Host Name	Type the host name. The default value is localhost.
Port	Select a port number. The default value is 5432.
Database	Type the name of the database to connect to.
Test Connection	Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.

Sybase data source configuration

Use this information to create a new Sybase data source.

Table 16. Sybase Data Source window

Window element	Description
General Settings	
Data Source Name	Type a unique name to identify the data source. Only letters, numbers, and the underscore character must be used in the data source name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data source is saved is set to the UTF-8 character encoding.
User name	Type a user name that allows you access to the database.
Password	Type a unique password. As you type, the characters are replaced with asterisks (*).

Table 16. Sybase Data Source window (continued)

Window element	Description
Maximum SQL Connection	<p>When you configure the maximum number of threads for the event processor service, for optimal performance the number of processing threads should be greater than, or equal to the size of the connection pool of the SQL data sources used in the policies being triggered. For information about viewing existing thread and connection pool information, see the information in the <i>Administration Guide</i> in the section <i>Command-Line tools, Event Processor commands</i>. See the Select PoolConfig from Service where Name='EventProcessor';</p> <p>Important: In a clustered environment, the event processor configuration is not replicated between servers. You must run the Select PoolConfig from Service where Name='EventProcessor'; command on the primary and the secondary servers.</p> <p>Limiting the number of concurrent connections manages performance. Type the maximum number of connections allowed to the database at one time. That number has to be greater than or equal to the number of threads running in the Event Processor. See “Event processor service configuration window” on page 127.</p> <p>The default value is 5.</p>
Database Failure Policy	<p>Select the failover option. Available options are Fail over, Fail back, and Disable Backup.</p> <p>For more information about failover options, see “SQL database DSA failover modes” on page 33.</p>
Primary Source	
Host Name	Type the host name. The default value is localhost.
Port	Select a port number. The default number is 5000.
Test Connection	<p>Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.</p> <p>Important: If you see an error message telling you that the data source cannot establish a connection to a database because a JDBC driver was not found, it means that a required JDBC driver is missing in the shared library directory. To fix this, place a licensed JDBC driver in the shared library directory and restart the server. For more information see, the “SQL database DSAs” chapter in the <i>DSA Reference Guide</i>.</p>
Backup Source (optional)	
Host Name	Type the host name. The default value is localhost.
Port	Select a port number. The default value is 5000.
Test Connection	<p>Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.</p>

GenericSQL data sources

Before creating a GenericSQL data source, you need to install the JDBC driver for your database, and specify its directories in the Impact Server properties file.

For details, see the *Administration Guide*

GenericSQL data source configuration

Use this information to configure a GenericSQL data source.

Table 17. GenericSQL Data Source window

Window element	Description
General Settings	
Data Source Name	Type a unique name to identify the data source. Only letters, numbers, and the underscore character must be used in the data source name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data source is saved is set to the UTF-8 character encoding.
JDBC Driver Name	Type the name of the JDBC driver for the database.
Username	Type a user name with which you can access the database.
Password	Type a password that allows you access to the database. As you type, the characters are replaced with asterisks (*).
Maximum SQL Connection	<p>When you configure the maximum number of threads for the event processor service, for optimal performance the number of processing threads should be greater than, or equal to the size of the connection pool of the SQL data sources used in the policies being triggered. For information about viewing existing thread and connection pool information, see the information in the <i>Administration Guide</i> in the section <i>Command-Line tools, Event Processor commands</i>. See the Select PoolConfig from Service where Name='EventProcessor';</p> <p>Important: In a clustered environment, the event processor configuration is not replicated between servers. You must run the Select PoolConfig from Service where Name='EventProcessor'; command on the primary and the secondary servers.</p> <p>Limiting the number of concurrent connections manages performance. Type the maximum number of connections allowed to the database at one time. That number has to be greater than or equal to the number of threads running in the Event Processor. See “Event processor service configuration window” on page 127.</p> <p>The default value is 5.</p>
Database Failure Policy	<p>Select the failover option. Available options are Fail over, Fail back, and Disable Backup.</p> <p>For more information about failover options, see “SQL database DSA failover modes” on page 33.</p>
Primary Source	

Table 17. GenericSQL Data Source window (continued)

Window element	Description
Host Name	Type the host name. The default value is localhost.
Port	Select a port number. The default number is 5432.
URL	
Test Connection	<p>Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.</p> <p>Important: If you see an error message telling you that the data source cannot establish a connection to a database because a JDBC driver was not found, it means that a required JDBC driver is missing in the shared library directory. To fix this, place a licensed JDBC driver in the shared library directory and restart the server. For more information see, the “SQL database DSAs” chapter in the <i>DSA Reference Guide</i>.</p>
Backup Source (optional)	
Host Name	Type the host name. The default value is localhost.
Port	Select a port number. The default value is 5432.
URL	
Test Connection	<p>Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.</p>

HSQldb data source configuration

Use this information to create a new HSQldb data source.

Table 18. HSQldb Data Source window

Window element	Description
General Settings	
Data Source Name	Type a unique name to identify the data source. Only letters, numbers, and the underscore character must be used in the data source name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data source is saved is set to the UTF-8 character encoding.
Username	Type a user name with which you can access the database.
Password	Type a password that allows you access to the database. As you type, the characters are replaced with asterisks (*).

Table 18. HSQLDB Data Source window (continued)

Window element	Description
Maximum SQL Connection	<p>When you configure the maximum number of threads for the event processor service, for optimal performance the number of processing threads should be greater than, or equal to the size of the connection pool of the SQL data sources used in the policies being triggered. For information about viewing existing thread and connection pool information, see the information in the <i>Administration Guide</i> in the section <i>Command-Line tools, Event Processor commands</i>. See the Select PoolConfig from Service where Name='EventProcessor';</p> <p>Important: In a clustered environment, the event processor configuration is not replicated between servers. You must run the Select PoolConfig from Service where Name='EventProcessor'; command on the primary and the secondary servers.</p> <p>Limiting the number of concurrent connections manages performance. Type the maximum number of connections allowed to the database at one time. That number has to be greater than or equal to the number of threads running in the Event Processor. See “Event processor service configuration window” on page 127.</p> <p>The default value is 5.</p>
Database Failure Policy	<p>Select the failover option. Available options are Fail over, Fail back, and Disable Backup.</p> <p>For more information about failover options, see “SQL database DSA failover modes” on page 33.</p>
Primary Source	
Host Name	Type the host name. The default value is localhost.
Port	Select a port number. The default number is 9001.
Database	Type the name of the database to connect to.
Test Connection	Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.
Backup Source (optional)	
Host Name	Type the host name. The default value is localhost.
Port	Select a port number. The default value is 9001.
Database	Type the name of the database to connect to. The default value is Impact.
Test Connection	Click to test the connection to the host to ensure that you entered the correct information. Success or failure is reported in a message box. If the host is not available at the time you create the data source, you can test it later.

LDAP data sources

The LDAP data source represent LDAP directory servers.

Netcool/Impact supports the OpenLDAP, and Microsoft Active Directory servers.

You create LDAP data sources in the GUI Server. You must create one such data source for each LDAP server that you want to access. The configuration properties for the data source specify connection information for the LDAP server, and any required security or authentication information.

Creating LDAP data sources

The LDAP DSA supports only non-authenticating data sources.

You can make them authenticating, however, using the Netcool/Impact properties file. For information about authenticating LDAP data sources, see the *DSA Reference Guide*.

Do not specify authentication parameters for the LDAP data source unless the underlying LDAP server is configured to require them. If you specify authentication parameters and they are not required by the LDAP server, Netcool/Impact fails to connect to the data source.

LDAP data source configuration window

Use this information to configure an LDAP data source.

Table 19. LDAP Data Source window

Window element	Description
General Settings	
Data Source Name	Type a unique name to identify the data source. Only letters, numbers, and the underscore character must be used in the data source name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data source is saved is set to the UTF-8 character encoding.
Source	
LDAP Server	Type the server name where the LDAP database resides. The default is localhost.
Port	Select a port number. The default value is set to a common port number: 389.
Security Protocol	Optional. Type the security protocol to use when connecting to the LDAP server. Supported security protocols are ssl and sasl. If you do not specify a security protocol, none is used.
Service Provider	Optional. Type the service provider to use when connecting to the LDAP server. To use the default Java LDAP provider, do not specify any value for this property. If you do not want to use the default Java LDAP provider, enter the fully qualified package and class name of the initial context factory class for the LDAP provider you want to use.
Authentication	

Table 19. LDAP Data Source window (continued)

Window element	Description
Authentication Mechanism	Optional. Type the authentication type to use when connecting to the LDAP server. Basic authentication types are none, anonymous and simple. Other types of authentication as described in the LDAP v2 and v3 specifications are also supported. If the LDAP server does not have authentication enabled, do not specify a value for this property. For more information about authentication types, see the documentation provided by the LDAP server.
User name	For simple authentication, enter the fully qualified LDAP user name. For authentications none and anonymous, leave this field blank.
Password	For simple authentication, enter a valid LDAP password. For authentication types of none and anonymous, leave this field blank. Restriction: Do not specify authentication parameters for the LDAP data source unless the underlying LDAP server is configured to require them. If you specify authentication parameters and they are not required by the LDAP server, Netcool/Impact will fail to connect to the data source.

Mediator data sources

Mediator data sources represent third-party applications that are integrated with Netcool/Impact through the DSA Mediator.

These data sources include a wide variety of network inventory, network provisioning, and messaging system software. In addition, providers of XML and SNMP data can also be used as mediator data sources.

Typically Mediator DSA data sources and their data types are installed when you install a Mediator DSA. The data sources are available for viewing and, if necessary, for creating or editing.

Attention: For a complete list of supported data source, see your IBM account manager.

CORBAMediator DSA data source configuration window

Use this information to configure a **CORBAMediator** DSA data source.

Table 20. CORBAMediator DSA data source window

Window element	Description
General Settings	
Data Source Name	Type a unique name to identify the data source. Only letters, numbers, and the underscore character must be used in the data source name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data source is saved is set to the UTF-8 character encoding.
Source Complete either the Name Service fields or the IOR File Location field.	

Table 20. CORBAMediator DSA data source window (continued)

Window element	Description
Name Service Host	Add the Name Service Host
Name Service Port	Add the Name Service Port
Name Service Context	Add the Name Service Context
Name Service Object Name	Add the Name Service Object Name
IOR File Location	Add the IOR File Location

DirectMediator DSA data source configuration window

Use this information to configure a DirectMediator Data Source.

Table 21. DirectMediator DSA data source window

Window element	Description
General Settings	
Data Source Name	Type a unique name to identify the data source. Only letters, numbers, and the underscore character must be used in the data source name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data source is saved is set to the UTF-8 character encoding.
Source	
Mediator Class Name	Add the Mediator Class Name

Creating SNMPDirectMediator data sources

When you have an SNMP DSA installed, you need to create any required SNMP data sources.

You can either create one data source for each SNMP agent that you want to access using the DSA, or you can create a single data source and use it to access all agents.

If you plan to use the standard data-handling functions `AddDataItem` and `GetByFilter` to access SNMP data, you must create a separate data source for each agent.

Important: To create a data source with SNMP v3 authentication, specify the properties described in Table 22 and then enter the information for the agent to authenticate the DSA as an SNMP user. The authentication parameters can be overridden by calls to the SNMP functions in the Impact Policy Language.

SNMPDirectMediator data source configuration window

Use this information to configure a SNMPDirectMediator data source.

Table 22. SNMPDirectMediator Data Source Configuration window

Window element	Description
General Settings	

Table 22. *SNMPDirectMediator Data Source Configuration window (continued)*

Window element	Description
Data Source Name	Type a unique name to identify the data source. Only letters, numbers, and the underscore character must be used in the data source name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data source is saved is set to the UTF-8 character encoding.
Data Source Settings	
Mediator Class Name	The following class name appears in this field: <code>com.micromuse.dsa.snmpdsa.Snmp</code>
SNMP Agent Settings	
Host Name	If you are creating this data source for use with the standard data-handling functions <code>AddDataItem</code> and <code>GetByFilter</code> , enter the host name or IP address. If you are creating this data source for use with the new SNMP functions, accept the default value.
Read Community	Type the name of the SNMP read community. The default is public.
Write Community	Type the name of the SNMP write community. The default is public
Timeout	Type or select a timeout value in seconds. When the DSA connects to an agent associated with this data source, it waits for the specified timeout period before returning an error to Netcool/Impact.
Port	If you are creating this data source for use with the standard data-handling functions <code>AddDataItem</code> and <code>GetByFilter</code> , select or enter the port number. If you are creating this data source for use with the new SNMP functions, accept the default value.
Version	Select the correct version, 1, 2 or 3. If you select SNMP version 3, the SNMP V3 section of the window activates.
SNMP V3	
User	The name of an SNMP v3 authentication user.
Authentication Protocol	Select a protocol. The default is MD5
Authentication Password	Password for the authentication user.
Privacy Protocol	Select a protocol.
Privacy Password	Type a privacy password.
Context ID	Type a context ID.
Context Name	Type a context name.

JMS data source

A JMS data source abstracts the information that is required to connect to a JMS Implementation.

This data source is used by the `JMSMessageListener` service, and the `SendJMSMessage`, and `ReceiveJMSMessage` functions.

JMS data source configuration properties

Use this information to configure a JMS data source.

Table 23. JMS data source window

Window element	Description
General Settings	
Data Source Name	Type a unique name to identify the data source. Only letters, numbers, and the underscore character must be used in the data source name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data source is saved is set to the UTF-8 character encoding.
Source	
JNDI Factory Initial	<p>Type the name of the JNDI initial context factory. The JNDI initial context factory is a Java object that is managed by the JNDI provider in your environment. The JNDI provider is the component that manages the connections and destinations for JMS.</p> <p>OpenJMS, BEA WebLogic, and Sun Java Application Server distribute a JNDI provider as part of their JMS implementations. The required value for this field varies by JMS implementation. For OpenJMS, the value of this property is org.exolab.jms.jndi.InitialContextFactory. For other JMS implementations, see the related product documentation.</p>
JNDI Provider URL	Type the JNDI provider URL. The JNDI provider URL is the network location where the JNDI provider is located. The required value for this field varies by JMS implementation. For OpenJMS, the default value of this property is tcp://hostname:3035 , where host name is the name of the system where OpenJMS is running. The network protocol (TCP or, RMI,) must be specified in the URL string. For other JMS implementations, see the related product documentation.
JNDI URL Packages	Type the Java package prefix for the JNDI context factory class. For OpenJMS, BEA WebLogic, and Sun Java Application Server, you are not required to type a value in this field.
JMS Connection Factory Name	Type the name of the JMS connection factory object. The JMS connection factory object is a Java object that is responsible for creating new connections to the messaging system. The connection factory is a managed object that is administered by the JMS provider. For example, if the provider is BEA WebLogic, the connection factory object is defined, instantiated, and controlled by that application. For the name of the connection factory object for your JMS implementation, see the related product documentation.
JMS Destination Name	Type the name of a JMS topic or queue. This is the name of the remote topic or queue where the JMS message listener listens for new messages.

Table 23. JMS data source window (continued)

Window element	Description
JMS Connection User Name	Type a JMS user name. If the JMS provider requires a user name to listen to remote destinations for messages, type the user name in this field. JMS user accounts are controlled by the JMS provider.
JMS Connection Password	In the JMS Connection Password field type a JMS password. If the JMS provider requires a password to listen to remote destinations for messages, type the password in this field.
Test Connection	Test the connection to the JMS Implementation. If the test is successful you see a message: "JMS: Connection OK".

Chapter 6. Data types

Data types are elements of the data model that represent sets of data stored in a data source.

The structure of data types depends on the category of data source where it is stored. For example, if the data source is an SQL database, each data type corresponds to a database table. If the data source is an LDAP server, each data type corresponds to a type of node in the LDAP hierarchy.

Viewing data type performance statistics

You can use the Performance Statistics report to determine whether the caching enabled for the data type is working efficiently.

Procedure

1. Make sure that the data for this report are collected.
You must set the performance measurements settings in the data types **Caching** tab. See “SQL data type configuration window - Cache settings tab” on page 75.
2. Select a project.
3. In the **Data Model** tab, locate the data type for which you want performance statistics.
4. Click **View Performance Statistics** next to the data type.
For more information about the statistics reported in the window, see “Data type performance statistics.”
5. Close the window.
6. If you determine that caching needs to be reworked, see “SQL data type configuration window - Cache settings tab” on page 75.

Data type performance statistics

The definitions of data type performance statistics.

Table 24. Performance statistics

Statistic type	Description
Performance Averages	
Number of Queries	Average number of queries calculated over the time interval (seconds).
Number of Inserts	Average number of inserts calculated over the time interval (seconds).
Number of Updates	Average number of updates calculated over the time interval (seconds).
Number of Rows	Average number of rows retrieved (either from the cache or from the database) by the number of queries over the query interval.
Time to Execute Each Query	Average time it took to run each query calculated over the query interval.

Table 24. Performance statistics (continued)

Statistic type	Description
Time to Read Results of Each Query	Average time it took to read the results of each query over the query interval.
Averages are calculated over time interval	The time interval.
Cache Status	
Number of Queries (% of total)	Actual number of queries and the percentage of queries retrieved from the query cache per query interval.
Number of Data Items (% of total)	Actual number of data items and the percentage of data items loaded from the data cache per query interval.
Number of Data Items in Use	The number of data items loaded from the data cache referred by queries in the query cache.
Time Spent Clearing the Cache	The time it took to clear the cache.
Percentages are calculated over query interval	Query interval.

Data type caching

You can use data type caching to reduce the total number of queries that are made against a data source for performance or other reasons.

Caching helps you to decrease the load on the external databases used by Netcool/Impact. Data caching also increases system performance by allowing you to temporarily store data items that have been retrieved from a data source.

Important: Caching works best for static data sources and for data sources where the data does not change often.

Caching works when data is retrieved during the processing of a policy. When you view data items in the GUI, cached data is retrieved rather than data directly from the data source.

You can specify caching for external data types to control the number of data items temporarily stored while policies are processing data. Many data items in the cache uses significant memory but can save bandwidth and time if the same data is referenced frequently.

Important: Data type caching works with SQL database and LDAP data types. Internal data types do not require data type caching.

You configure caching on a per data type basis within the GUI. If you do not specify caching for the data type, each data item is reloaded from the external data source every time it is accessed.

Data type caching types

You can control the following aspects of data type caching.

Data caching

Use data caching to temporarily store individual data items retrieved from a data source.

When a policy uses the `GetByKey` function, data caching defines the number of records that can be held in the cache. You can configure both the maximum number of data items to cache and the expiration time for data items in the cache.

Query caching

You can use query caching to temporarily store sets of data items that are retrieved during individual queries to a data source.

When a policy uses the `GetByFilter` function, query caching defines the number of completed queries allowed in the cache (not the number of data items).

Important: You have to set data caching for query caching to work.

Count caching

Count caching is used to temporarily store the count values obtained in a policy. Count caching uses the `GetByFilter` function with the `CountOnly` parameter set to `True`.

This type of caching is for compatibility with earlier versions only, do not use it unless it is necessary.

Creating internal data types

Overview of the tabs in the internal data type editor.

Table 25. Internal data type editor tabs

Tab	Description
Custom Fields	In this tab, you can add any number of fields to form a database table.
Dynamic Links	<p>In this tab you can create links to other data types, both external and internal, to establish connections between information.</p> <p>Links between individual data items can represent any relationship between the items that policies need to be able to look up. For example, a node linked to an operator allows a policy to look up the operator responsible for the node.</p>

Internal data type configuration window

Use this information to configure an internal data type.

Table 26. New Internal Data Type Editor Custom Fields tab

Editor element	Description
General settings	
Data Type Name	<p>Type a unique name to identify the data type. Only letters, numbers, and the underscore character must be used in the data type name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data type is saved is set to the UTF-8 character encoding.</p> <p>If you receive an error message when saving a data type, check the Global tab for a complete list of data type names for the server. If you find the name you have tried to save, you need to change it.</p>

Table 26. New Internal Data Type Editor Custom Fields tab (continued)

Editor element	Description
State: Persistent	<p>Leave the box checked as Persistent (permanent) to permanently store the data items created for this data type. When the server is restarted, the data is loaded. If the box is cleared, the data is held in memory, but only while the server is running. When the server restarts, the data is lost because it was not backed up in a file. This feature is useful if you need data only on a temporary basis and then want to discard it.</p> <p>Persistent data types are always written to file. Therefore making internal data types temporary is faster.</p>
New Field	Click to add a new field to the table.
Additional fields (New Field window)	
ID	Type a unique ID for the field.
Field Name	<p>Type the actual field name. This can be the same as the ID. You can reference both the ID field and the Field Name field in policies.</p> <p>If you do not enter a Display Name (see below), Netcool/Impact uses the ID field name by default.</p>
Format	Select a format for the field from the Format list:
Display Name Field:	You can use this field to select a field from the menu to label data items according to the field value. Choose a field that contains a unique value that can be used to identify the data item for example, ID . To view the values on the data item you need to go to View Data Items for the data type and select the Links icon. Click the data item to display the details.
Description	Type some text that describes the field.

External data types

External data types use an external data source to access information in vendor acquired databases, such as SQL, or LDAP databases, as well as DSAs.

By definition, an external data type is the lookup method used to find data from the external data source. An external data type contains all the fields (data items) in its data source that meet the lookup criteria. When the database is accessed, the fields from the database schema are assigned to the data type. You can also add additional fields to the type, for example, if a field was added to the data source after you created the data type. You can delete fields that you do not need to have as part of your data type.

Creating data types from external data sources is similar to creating internal data types, except that the external data type editor has a **Table Description** tab instead of a **Custom Fields** tab and an additional data caching part to regulate the flow of data between Netcool/Impact and the external data source. The fields in the **Table Description** tab are not custom fields that you create. These fields identify the required data from the external data source.

All data types must belong to a data source. Before you create an external data type, create a data source to associate with the data type.

Deleting a field

You can use the Delete function to limit which fields are updated, inserted, and selected from the data source.








Remember: When you delete a field from the data type, it is not deleted from the data source.

Using a subset of the database fields can speed performance of the data type.

List of predefined data types

An overview of the predefined data types available in the global project.

Table 27. Predefined data types

Data type	Type	Description
 Schedule	Editable	Schedules define a list of data items associated with specific time ranges, or time range groups, that exist.
 Document	Editable	Custom URL Document data types are derived from the predefined Doc data type.
 FailedEvent	Editable	The FailedEvent data type, together with the ReprocessedFailedEvents policy, provides you with a way to deal with failed events that are passed from the ObjectServer.
 ITNM	Editable	This data type is used with ITNM and the ITNM DSA.
 TimeRangeGroup	Non-editable	A time range group data type consists of any number of time ranges.
 LinkType	Non-editable	The LinkType data type provides a way of defining named and hierarchical dynamic links.
 Hibernation	Non-editable	When you call the Hibernate function in a policy, the policy is stored as a Hibernation data item for a certain number of seconds.

Predefined data types overview

Predefined data types are special data types that are stored in the global repository.

You can edit some predefined data types by adding new fields, but you cannot edit or delete existing fields. You can view, edit, create, and delete data items of some predefined data types by using the GUI. You cannot delete predefined data types except for the **FailedEvent** predefined data type.

Time range groups and schedules

The Schedule and Time Range Group data types have special data items that are similar to internal data items, but they are used specifically for defining scheduling information.

Policies typically use schedules and time range groups to look up the availability of another item, for example, whether an administrator is on call at the time the policy is run.

Schedules contain time ranges associated with data items. You can group time ranges so that you can easily reuse them so that you do not have to enter the information each time.

Time range group data types

A time range group data type consists of any number of time ranges.

There are three types of time ranges, as described below:

Table 28. Time range specifications

Time range type	Description
Positive	The time range is active when the current time is within the time range, unless it is overlapped by a Negative or an Override.
Negative	The time range is inactive for the specified range. This time range is useful, for example, to exclude a lunch hour from a Positive time range.
Override	The time range is always active within the range, regardless of any negative ranges.

You can specify any combination of the time ranges as described below:

Table 29. Time Range Combinations

Time range	Description
Daily	A time range between a starting time and an ending time for every day of the week, for example, 9 a.m. to 5 p.m.
Weekly	A range between a starting time on a specified day and ending on a specified day every week, for example Monday 9 a.m. to Friday 5 p.m.
Absolute	<p>A range of time between two specific dates, for example, March 3, 2004 to March 4, 2004.</p> <p>One way this time range is useful is for server maintenance. If a server is due to be down for maintenance on a specific day and you do not want it to show up as an alarm, you could define an Absolute range and use it in an Event Suppression policy.</p>

Configuring time range groups:

Use this procedure to create a new time range group.

Procedure

1. In the **Data Model** tab select the **Global** project, from the **Project** menu.
2. In the list of data sources, and data types, click the plus sign next to the **Internal** data source to view its data types.
3. Select the **TimeRangeGroup** data type.
4. Click **View Data Items** to open the TimeRangeGroup window.
5. Click **New**.
6. In the **Group Name** field, type a unique name to describe the group.

7. From the **Available** list, select the type of range you want to define **Daily**, **Weekly**, or **Absolute**, and click **New**.
See “Adding daily time ranges,” “Adding weekly time ranges,” and “Adding absolute time ranges” on page 68.
8. From the **Groups** list, select the groups to add to this time range.
Click the plus icon to add the selected group as a new time range. When you finish adding the new time range, it displays in the time ranges viewer.
9. Click **OK**.

Adding daily time ranges:

In the Time Range Group Viewer, after you select **Daily**, click **New** to open the Daily Time Range window.

Enter the information in the window, using this table as a guide:

Table 30. Daily Time Range window

Window element	Description
Start Time: hour/min	Using the 24-hour clock, enter the start time.
EndTime: hour/min	Using the 24-hour clock, enter the end time.
Time Zone	Select the appropriate time zone from the list.
Positive	See Table 29 on page 66.
Negative	
Override	

Adding weekly time ranges:

Use this procedure to add weekly time ranges.

Procedure

In the Time Range Group window, after you select **Weekly**, click **New** to open the Weekly Time Range window.

Enter the information in the window, using this table as a guide:

Table 31. Weekly Time Range window

Window element	Description
Start	Select the day of the week to indicate the beginning day of the time range.
hour/min	Type or select the time of day to start the time range.
End	Select the day of the week to indicate the end of the time range.
hour/min	Type or select the time of day to end the time range.
Time Zone	Select the appropriate time zone from the list.
Positive	See Table 29 on page 66.
Negative	
Override	

Adding absolute time ranges:

Use this procedure to add absolute time ranges.

Procedure

1. In the Time Range Group window, after you select **Absolute**, click **New** to open the following Absolute Time Range window
Enter the information in the window, using this table as your guide:
2. **Start Date:**Click the calendar icon to select the start date and time.
3. **End Date:** Click the calendar icon to select the end date and time.
4. **Time Zone:**Select the appropriate time zone from the list.
5. Select the Effect from the list
Positive
Negative
Override
6. Click **OK**.

Schedules overview

Schedules define a list of data items associated with specific time ranges, or time range groups, that exist.

You can use links between Schedule data items and other data items to schedule any items, for example, the hours when a departmental node is business critical or to identify who is currently on call when an alert occurs.

Adding schedules:

You can define multiple schedules.

Procedure

1. To add a schedule open the **Schedule** data source in the **Data Model** tab.
2. Select the **Schedule** data type.
3. Click **Edit**.
4. In the Schedule editor, type a unique name for the schedule in the **Data Type Name** field.
5. See “Configuring schedules”

Configuring schedules:

Use this procedure to create a new schedule.

Procedure

1. Open the **Global** tab and locate the Schedules data type.
2. To create a new schedule, click **Create New Schedule**.
3. Enter the following information in the window:

Table 32. Schedule Editor window

Window element	Description
General Settings	
Schedule Name	Type a unique name for the schedule.

Table 32. Schedule Editor window (continued)

Window element	Description
Description	Type a description for the schedule.
Members and Time Ranges	
Schedule Members	Members are displayed in this window after selecting the data type and adding members using the instructions below.
Edit Members by Type	Select a data type.
Edit	See Step 5.

- After selecting a data type, click **Edit** to open the Select Schedule Members window.
- Enter information in the **Select Schedule Members** window using this table as your guide:

Table 33. Select Schedule Members window

Window element	Description
Filter	Type a filter in the text field to limit the number of displayed member candidates.
Filter	Click to apply the filter to the member candidates.
Member Candidates	Highlight a candidate from the list.
Add	Click to add the candidate to the Members list.
Members	Highlight a candidate from the list.
Remove	Click to remove the candidate from the Member list.

- After you have selected the data type and added members to the schedule, the Time Ranges for Schedule Member pane opens on the right side of the window.
- Highlight a member in the Schedule Members pane.
The name you selected appears at the top of the Time Ranges for Schedule Member pane.
- Enter time ranges for the candidate using the instructions in “Configuring time range groups” on page 66.
Note that the green light next to the On Call Status for the current member indicates that the administrator is on call. The 'traffic light' would be red if the administrator were not on call.
- Repeat for each schedule member.
- Click **OK**.
The new schedule data item displays as a row in the table.
For information about editing and deleting data items, see Chapter 7, “Data items,” on page 87.

ITNM DSA data type

The ITNM data type is the only one that works with the ITNM DSA.

You cannot rename an ITNM data type.

When the DSA queries the ITNM database, the records returned are data items of the ITNM data type. Each field in the records is turned into an attribute of the corresponding data item.

For example, a record can contain fields such as:

- ObjectId
- EntityName
- Address
- Description
- ExtraInfo

To access the values, you can directly access the attributes just like any other data items using the following command:

```
log("Description is " + DataItem.Description);
```

This command prints out the Description field string that was on the ITNM record returned by the query.

SQL data types

SQL data types define real-time dynamic access to data in tables in a specified SQL database.

When the database is accessed, the fields from the database schema are assigned to the data type. Some of the SQL data sources automatically discover the fields in the table. Others do not support automatic table discovery; for these data sources, you must enter the table name to see the names of the fields.

The editor contains three tabs.

Table 34. External data type editor tabs

Tab	Description
Table Description	Name the data type, change the data source, if necessary, and add any number of fields from the data source to form a database table.
Dynamic Links	<p>In this tab you can create links to other data types, both external and internal, to establish connections between information.</p> <p>Links between individual data items can represent any relationship between the items that policies must be able to look up. For example, a node linked to an operator allows a policy to look up the operator responsible for the node.</p> <p>For more information about dynamic links tab, see Chapter 8, “Links,” on page 89.</p>
Cache Settings	<p>In this tab, you can set up caching parameters to regulate the flow of data between Netcool/Impact and the external data source.</p> <p>Use the guidelines in “SQL data type configuration window - Cache settings tab” on page 75, plus the parameters for the performance report for the data type to configure data and query caching.</p>

Important: SQL data types in Netcool/Impact require all columns in a database table to have the Select permission enabled to allow discovery and to enable the save option when creating data types.

Configuring SQL data types

Use this procedure to configure an SQL data type.

Procedure

- Provide a unique name for the data type.
- Specify the name of the underlying data source for the data type.
- Specify the name of the database and the table where the underlying data is stored.
- Auto-populate the fields in the data type.
- Select a display name for the data type.
- Specify key fields for the data type.
- Optional: Specify a data item filter.
- Optional: Specify which field in the data type to use to order data items.
- Optional: Specify the direction to use when ordering data items.

What to do next

After you have saved the data type, you can close the Data Type Editor or you can configure caching and dynamic links for the data type.

SQL data type configuration window - Table Description tab

Use this information to configure the SQL data type.

Table 35. New External Data Type editor - Table Description tab

Editor element	Description
General Settings	
Data Type Name	Type a unique name to identify the data type. Only letters, numbers, and the underscore character must be used in the data type name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data type is saved is set to the UTF-8 character encoding. Data type names must be unique globally, not just within a project. If you receive an error message when saving a data type, check the Global project tab for a complete list of data type names for the server. If you find the name you have tried to save, you need to change it.
Data Source: Name	This field is automatically populated, based on the data source you selected in the Data Sources and Types task pane. However, if you have other SQL data sources configured to use with Netcool/Impact, you can change the name to any of the SQL data sources in the list, if necessary. If you enter a new name, a message window prompts you to confirm your change. Click OK to confirm the change. If you change your mind about selecting a different data source, click Cancel .
State: Enabled	Leave the State check box checked to activate the data type so that it is available for use in policies.
Table Description	

Table 35. New External Data Type editor - Table Description tab (continued)

Editor element	Description
Base Table	<p>Specify the underlying database and table where the data in the data type is stored.</p> <p>The names of all the databases and tables are automatically retrieved from the data source so that you can choose them from a list.</p> <p>Type the name of the database and the table in the Base Table lists. The first list contains the databases in the data source. The second list contains the tables in the selected database, for example, alerts, and status.</p>
Refresh	<p>Click Refresh to populate the table.</p> <p>The table columns are displayed as fields in a table. To make database access as efficient as possible, delete any fields that are not used in policies.</p>
Add Deleted Fields	<p>If you have deleted fields from the data type that still exist in the SQL database, these fields do not show in the user interface. To restore the fields to the data type, mark the Add Deleted Fields check box and click Refresh.</p>
New Field	<p>Use this option if you need to add a new field to the table from the data source database. For example, in the case where the field was added to the database after you created the data type.</p> <p>Make sure that the field name you add has the same name as the field name in the data source.</p> <p>Important: Any new fields added to this table are not automatically added to the data source table. You cannot add fields to the database table in this way.</p> <p>For more information, see “SQL data type configuration window - adding and editing fields in the table” on page 73.</p>
Key field	<p>Key fields are used when you retrieve data from the data type in a policy using the <code>GetByKey</code> function. They are also used when you define a <code>GetByKey</code> dynamic link.</p> <p>Important: You must define at least one key field for the data type, even if you do not plan to use the <code>GetByKey</code> functionality in your policy. If you do not, Netcool/Impact will not function properly.</p> <p>Generally, the key fields you define correspond to key fields in the underlying database table.</p> <p>To specify a key field, click the check box in the appropriate row in the Key Field column. You can add multiple key fields.</p>
Display Name Field	<p>You can use this field to select a field from the menu to label data items according to the field value. Choose a field that contains a unique value that can be used to identify the data item for example, ID. To view the values on the data item you need to go to View Data Items for the data type and select the Links icon. Click the data item to display the details.</p>

Table 35. New External Data Type editor - Table Description tab (continued)

Editor element	Description
Automatically Remove Deleted Fields	Mark the Automatically Remove Deleted Fields check box to remove any fields from the data type that have already been removed from the SQL database. This happens automatically when a policy that uses this data type is run.
Data Filtering and Ordering	
Filter	Type a restriction clause to limit the types of data items seen for the data type. For example, to limit the rows in a field called <i>City</i> to <i>New York</i> , you would enter: City = "New York" For example, to limit the rows to the <i>New York</i> or <i>Athens</i> , you would enter: City = "New York" OR City = "Athens" You can use any sql Where clause syntax.
Order By	Enter the names of one or more fields to use when sorting data items retrieved from the data source.

SQL data type configuration window - adding and editing fields in the table

Use this information to add or edit a field to the table for a SQL data type.

In the Table tab, in the **New Field** area, click **New** to add a field to the data type, or select the edit icon next to an existing field that you want to edit.

Table 36. External Data Type Editor - New Field window

Window element	Description
ID	By default, the ID is the same as the column name in the database. You can change it to any other unique name. For example, if the underlying column names in the data source are difficult to use, the ID field to provide an easier alias for the field.
Field Name	Type a name that can be used in policies. It represents the name in the SQL column. Type the name so that it is identical to how it appears in the data source; otherwise, Netcool/Impact reports an error when trying to access the data type.

Table 36. External Data Type Editor - New Field window (continued)

Window element	Description
Format	<p>For SQL database data types, Netcool/Impact auto-discovers the columns in the underlying table and automatically detects the data format for each field when you set up the data type. For other data types, you must manually specify the format for each field that you create. For more information about formats, see the Working with Data Types chapter in the Solutions Guide.</p> <p>Select a format from the following list:</p> <ul style="list-style-type: none"> • STRING • LONG_STRING • INTEGER • PASSWORD_STRING • LONG • FLOAT • DOUBLE • DATE • BOOLEAN • CLOB
Display Name	<p>You can use this field to select a field from the menu to label data items according to the field value. Choose a field that contains a unique value that can be used to identify the data item for example, ID. To view the values on the data item you need to go to View Data Items for the data type and select the Links icon. Click the data item to display the details.</p> <p>If you do not enter a display name, Netcool/Impact uses the ID field name by default.</p>
Description	Type some text that describes the field. This description is only visible when you edit the data type using the GUI.
Default Value	Type a default expression for the field. It can be any value of the specified format see the format row, or it can be a database-specific identifier such as an Oracle pseudonym; for example, <code>sequence.NEXTVAL</code> .

Table 36. External Data Type Editor - New Field window (continued)

Window element	Description
Insert Statements: Exclude this field	<p>When you select the Exclude this Field check box Netcool/Impact does not set the value for the field when inserting and updating a new data item into the database. This field is used for insert and update statements only, not for select statements.</p> <p>Sybase data types:</p> <p>You must select this option when you map a field to an Identity field or a field with a default value in a Sybase database. Otherwise, Netcool/Impact overwrites the field on insert with the specified value or with a space character if no value is specified.</p> <p>ObjectServer data types:</p> <p>The Tally field automatically selects the Exclude this Field check box to be excluded from inserts and updates for the objectserver data type since this field is automatically set by Netcool OMNIBus to control deduplication of events.</p> <p>The Serial field automatically selects the Exclude this Field check box to be excluded from inserts and updates when an ObjectServer data type points to alerts.status.</p>
Type Checking: Strict	<p>Click to enable strict type checking. When you enable strict type checking on the field, Netcool/Impact checks the format of the value of the field on insert or update to ensure that it is of the same format as the corresponding field in the data source. If it is not the same, Netcool/Impact does not perform the insert or update and a message to that effect is displayed in the server log. If you do not enable strict type checking, all type checking and format conversions are done at the data source level.</p>

SQL data type configuration window - Cache settings tab

Use this information to configure caching for a SQL data type.

Table 37. External Data Type Cache Settings tab - caching types

Cache type	Description
Enable Data Caching	This check box toggles data caching on and off.
Maximum number of data items	Set the total number of data items to be stored in the cache during the execution of the policy.
Invalidate Cached Data Items After	Set to invalidate the cached items after the time periods selected.
Enable Query Caching	This check box toggles query caching on and off.
Maximum number of queries	Set the maximum number of database queries to be stored in the cache.
Invalidate Cached Queries After	Set to invalidate the cached items after the time periods selected.
Enable Count Caching	Do not set. Available for compatibility with earlier versions only.
Performance Measurements Intervals	Use this option to set the reporting parameters for measuring how fast queries against a data type are executed.

Table 37. External Data Type Cache Settings tab - caching types (continued)

Cache type	Description
Polling Interval	Select a polling interval for measuring performance statistics for the data type.
Query Interval	Select the query interval for the performance check.

Creating flat file data types

Use this procedure to create a flat file data type.

Procedure

1. Before you can create a flat file data type you must create a flat file data source.
For more information about creating flat file data sources, see “Creating flat file data sources” on page 36.
2. Click the **Projects** tab.
3. Click the **Data Sources And Types** list to expand it.
4. Click **Create a new data type** next to the flat file data source that you created earlier, for example MyFlatFileDataSource.
5. In the new data type window, provide the required information.
 - a. In the **Data Type Name:** field type a unique name for your data type name.
For example, MyFlatFileDataType.
Your data source, MyFlatFileDataSource, should already have been preselected in the **Data Source Name:** list. If not, select it from the list.
 - b. In the **Base Label:** field, enter the name of your flat file that you created for your flat file data source, for example myflatfile.txt.
 - c. Click **Refresh** to load field names from your text file.
 - d. Select the check boxes in the **Key Field** column.
 - e. Save your flat file data type.

Results

If you open the data items viewer, you can see the entries from your flat file.

LDAP data types

An LDAP data type represents a set of entities in an LDAP directory tree.

The LDAP DSA determines which entities are part of this set in real time by dynamically searching the LDAP tree for those that match a specified LDAP filter within a certain scope. The DSA performs this search in relation to a location in the tree known as the base context.

The LDAP Data Type editor contains three tabs.

Table 38. LDAP Data Type editor tabs

Tab	Description
LDAP Info	In this tab, you configure the attributes of the data type. For more information about these attributes, see “LDAP data type configuration window - LDAP Info tab” on page 77.

Table 38. LDAP Data Type editor tabs (continued)

Tab	Description
Dynamic Links	<p>In this tab you can create links to other data types, both external and internal, to establish connections between information. Links between individual data items can represent any relationship between the items that policies need to be able to look up. For example, a node linked to an operator allows a policy to look up the operator responsible for the node.</p> <p>For more information about creating links to other data types, see Chapter 8, “Links,” on page 89.</p>
Cache Settings	<p>In this tab, you can set up caching parameters to regulate the flow of data between Netcool/Impact and the external data source.</p> <p>For more information about, cache settings see “SQL data type configuration window - Cache settings tab” on page 75.</p>

Important: You must create one LDAP data type for each set of entities that you want to access. The LDAP data type is a read-only data type which means that you cannot edit or delete LDAP data items from within the GUI.

Configuring LDAP data types

Use this procedure to configure an LDAP data type.

Procedure

- Provide a unique name for the data type.
- Specify the name of the underlying data source for the data type.
- Specify the base context level in the LDAP hierarchy where the elements you want to access are located.
- Specify a display name field.
- Optional: Specify a restriction filter.

LDAP data type configuration window - LDAP Info tab

Use this information to configure LDAP information for a LDAP data type.

Table 39. LDAP Data Type editor - LDAP Info Tab

Editor element	Description
General Settings	
Data Type Name	Type a unique name to identify the data type. Only letters, numbers, and the underscore character must be used in the data type name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data type is saved is set to the UTF-8 character encoding.
State: Enabled	Leave checked to enable the data type so that it can be used in policies.
LDAP Info	

Table 39. LDAP Data Type editor - LDAP Info Tab (continued)

Editor element	Description
Data Source Name	Type the name of the underlying data source. This field is automatically populated, based on your data source selection in the Data Types task pane of the Navigation panel. However, if you have more than one LDAP data source configured for use with Netcool/Impact, you can select any LDAP data source in the list, if necessary. If you enter a new name, a message window asks you to confirm your change.
Search scope	Select the search scope: <ul style="list-style-type: none"> • OBJECT_SCOPE • ONLEVEL_SCOPE • SUBTREE_SCOPE
Base Context	Type the base context that you want to be used when searching for LDAP entities. For example:ou=people,o=companyname.com.
Key Search Field	Type the name of a key field, for example, dn.
Display Name Field	You can use this field to select a field from the menu to label data items according to the field value. Choose a field that contains a unique value that can be used to identify the data item for example, ID. To view the values on the data item you need to go to View Data Items for the data type and select the Links icon. Click the data item to display the details.
Restriction Filter:	Optionally, type a restriction filter. The restriction filter is an LDAP search filter as defined in Internet RFC 2254. This filter consists of one or more Boolean expressions, with logical operators prefixed to the expression list. For more information, see the <i>LDAP Filter</i> information in the <i>Policy Reference Guide</i> .
Attribute Configuration	
New Field	For each field that you want to add to the data type, click New .

Mediator DSA data types

Mediator DSA data types are typically created using scripts or other tools provided by the corresponding DSA.

Usually the data types, and their associated data sources are installed when you install the Mediator DSA (Corba or Direct), so you do not have to create them. The installed data types are available for viewing and, if necessary, for editing.

For more information about the Mediator data types used with a particular DSA, see the DSA documentation.

Viewing Mediator DSA data types

Use this information to configure the Mediator DSA Data Type.

The *DSA Data Type* editor contains three tabs, as described in the following table:

Table 40. DSA Data Type editor tabs

Tab	Description
DSA Mediator	This tab contains the attributes of the data type. See your DSA documentation for more information.
Dynamic Links	<p>In this tab you can create links to other data types, both external and internal, to establish connections between information.</p> <p>Links between individual data items can represent any relationship between the items that policies need to be able to look up. For example, a node linked to an operator allows a policy to look up the operator responsible for the node.</p> <p>For more information about dynamic links tab, see Chapter 8, "Links," on page 89.</p>
Cache Settings	In this tab, you can set up caching parameters to regulate the flow of data between Netcool/Impact and the external data source.

SNMP data types

If you are using an SNMP DSA, once you have created an SNMP data source, you can use the GUI to create SNMP data types.

See "Creating SNMPDirectMediator data sources" on page 56.

If you plan to use the standard data-handling functions `AddDataItem` and `GetByFilter` to access SNMP data, create a separate data type for each set of variables (packed OID data types) or each set of tables (table data types) that you want to access. If you plan to use the SNMP functions provided with the DSA, you can create a single data type for each data source and use it to access all the variables and tables associated with the agent. For more detailed information about SNMP data types, see the *DSA Reference Guide*.

SNMP data types - configuration overview

An overview of the SNMP data type configuration window.

The *SNMPDirectMediator Data Type* editor contains three tabs.

Table 41. DSA Data Type editor tabs

Tab	Description
DSA Mediator	This tab contains the attributes of the data type. See your DSA documentation for more information.
Dynamic Links	<p>In this tab you can create links to other data types, both external and internal, to establish connections between information.</p> <p>Links between individual data items can represent any relationship between the items that policies need to be able to look up. For example, a node linked to an operator allows a policy to look up the operator responsible for the node.</p> <p>For more information about dynamic links tab, see Chapter 8, "Links," on page 89.</p>
Cache Settings	In this tab, you can set up caching parameters to regulate the flow of data between Netcool/Impact and the external data source.

Packed OID data types

Packed OID data types reference the OIDs of one or more variables managed by a single agent.

You use this category of data type when you want to access single variables or sets of related variables. When you create a packed OID data type, you specify the name of the associated data source, the OID for each variable and options that determine the behavior of the DSA when connecting to the agent.

Packed OID SNMPDirectMediator data type - configuration window

Use this information to configure the Packed OID SNMPDirectMediator data type.

Table 42. *SNMPDirectMediator Data Type editor - DSA Mediator tab*

Editor element	Description
General Settings	
Data Type Name	Type a unique name to identify the data type. Only letters, numbers, and the underscore character must be used in the data type name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data type is saved is set to the UTF-8 character encoding.
Data Source Name	<p>This field is automatically populated, based on your data source selection in the Data Types task pane of the Navigation panel. However, if you have other SQL data sources configured to use with Netcool/Impact, you can change it to any of the SQL data sources in the list, if necessary.</p> <p>If you enter a new name, a message window asks you to confirm your change.</p> <p>Click OK to confirm the change. If you change your mind about selecting a different data source, click Cancel.</p>
Display Icon: Browse	<p>Click to display an icon selection window. The one you select associates it with the data type and all its related data items throughout the GUI.</p> <ul style="list-style-type: none">Click over an icon to select it and close the window.Click Close to close the window without choosing an icon.
SNMP Settings	
Oid Configuration	Select Packed OID data types from the Oid Configuration list.
New Attribute	<p>If you are creating this data type for use with the standard data-handling functions <code>AddDataItem</code> and <code>GetByFilter</code>, create a new attribute on the data type for each variable you want to access. To create an attribute, click New Attribute and specify an attribute name and the OID for the variable.</p> <p>If you are creating this data source for use with the new SNMP functions, you do not need to explicitly create attributes for each variable. In this scenario, you pass the variable OIDs when you make each function call in the Impact policy.</p>

Table 42. *SNMPDirectMediator Data Type editor - DSA Mediator tab (continued)*

Editor element	Description
Get Bulk: Enabled	<p>If you want the DSA to retrieve table data from the agent using the SNMP GETBULK command instead of an SNMP GET, select Get Bulk. The GETBULK command retrieves table data using a continuous GETNEXT command. This option is suitable for retrieving data from very large tables.</p> <p>When you select Get Bulk, you can control the number of variables in the table for which the GETNEXT operation is performed using the specified Non-Repeaters and Max Repetitions values.</p>
Max Repetitions	Max Repetitions specifies the number of repetitions for each of the remaining variables in the operation.
Nonrepeaters	The Non-Repeaters value specifies the first number of non-repeating variables.

Table data types

Table data types reference the OIDs of one or more tables managed by a single agent.

You use this category of data type when you want to access SNMP tables. When you create a table data type, you specify the name of the associated data source, the OID for each table and options that determine the behavior of the DSA when connecting to the agent.

Creating table data types

Use this procedure to create a table data type.

Procedure

1. In the data types tab, select an SNMP data source from the list.
2. Click the **New Data Type** button to open the New Data Type editor.
3. Type a name for the data type in the **Data Type Name** field.

Important:

The data type name must match the table name that will be queried, for example, `ifTable`, or `ipRouteTable`.

4. Select an SNMP data source from the **Data Source Name** field. By default, the data source you chose in step 2 above is selected.
5. Select Table from the **OID Configuration** list.
6. If you are creating this data type for use with the standard data-handling functions `AddDataItem` and `GetByFilter`, you must create a new attribute on the data type for each table you want to access. To create an attribute, click the **New Attribute** button and specify an attribute name and the OID for the table.

Important:

The attributes are the column names in each table. For example, in the following `ifTable`, the attributes will be `ifIndex`, `ifDescr` and other column names:

Column Names	OID
ifIndex	.1.3.6.1.2.1.2.2.1.1
ifDescr	.1.3.6.1.2.1.2.2.1.2
...	...

If you are creating this data source for use with the new SNMP functions, you do not need to explicitly create attributes for each table. In this scenario, you pass the table OIDs when you make each function call in the Netcool/Impact policy.

7. If you want the DSA to retrieve table data from the agent using the SNMP GETBULK command instead of an SNMP GET, select **Get Bulk**.

The GETBULK command retrieves table data using a continuous GETNEXT command. This option is suitable for retrieving data from very large tables.

8. If you have selected **Get Bulk**, you can control the number of variables in the table for which the GETNEXT operation is performed using the specified **Non-Repeaters** and **Max Repetitions** values.

The **Non-Repeaters** value specifies the first number of non-repeating variables and **Max Repetitions** specifies the number of repetitions for each of the remaining variables in the operation.

9. Click **Save**.

Table data types configuration window

Use this information to configure a table data type.

Table 43. *SNMPDirectMediator Data Type editor - DSA Mediator tab*

Editor element	Description
General Settings	
Data Type Name	Type a unique name to identify the data type. Only letters, numbers, and the underscore character must be used in the data type name. If you use UTF-8 characters, make sure that the locale on the Impact Server where the data type is saved is set to the UTF-8 character encoding.
Data Source Name	<p>This field is automatically populated, based on your data source selection in the Data Types task pane of the Navigation panel. However, if you have other SQL data sources configured to use with Netcool/Impact, you can change it to any of the SQL data sources in the list, if necessary.</p> <p>If you enter a new name, a message window asks you to confirm your change.</p> <p>Click OK to confirm the change. If you change your mind about selecting a different data source, click Cancel.</p>
Display Icon: Browse	<p>Click to display an icon selection window. The one you select associates it with the data type and all its related data items throughout the GUI.</p> <ul style="list-style-type: none"> • Click over an icon to select it and close the window. • Click Close to close the window without choosing an icon.
SNMP Settings	
Oid Configuration	Select Table from the list.

Table 43. *SNMPDirectMediator Data Type editor - DSA Mediator tab (continued)*

Editor element	Description
New Attribute	<p>If you are creating this data type for use with the standard data-handling functions <code>AddDataItem</code> and <code>GetByFilter</code>, you must create a new attribute on the data type for each variable you want to access. To create an attribute, click New Attribute and specify an attribute name and the OID for the variable.</p> <p>If you are creating this data source to use with the new SNMP functions, you do not need to explicitly create attributes for each table. In this scenario, you pass the variable OIDs when you make each function call in the Impact policy.</p>
Get Bulk: Enabled	<p>If you want the DSA to retrieve table data from the agent using the SNMP <code>GETBULK</code> command instead of an SNMP <code>GET</code>, select Get Bulk. The <code>GETBULK</code> command retrieves table data using a continuous <code>GETNEXT</code> command. This option is suitable for retrieving data from very large tables.</p> <p>When you select Get Bulk, you can control the number of variables in the table for which the <code>GETNEXT</code> operation is performed using the specified Non-Repeaters and Max Repetitions values.</p>
Max Repetitions	Max Repetitions specifies the number of repetitions for each of the remaining variables in the operation.
Nonrepeaters	The Non-Repeaters value specifies the first number of non-repeating variables.

LinkType data types

The LinkType data type provides a way of defining named and hierarchical dynamic links.

To reference links directly from a policy, you can specify the link type directly instead of the target data type name.

You can create hierarchies between data types, for example, using the source as a parent to multiple target children (for example, one customer to multiple servers).

Linktype data items are useful when you want to create several dynamic links between the same target and source data type for use in several policies. For example, in one policy you might want to filter the severity level for events for the target data type. In another policy, you might want to filter the server names for the target data type. You would create a LinkType data item for each scenario and select the appropriate one when creating the link.

For more information, see “Dynamic links” on page 89.

Configuring LinkType data items

Use the following procedure to create a new LinkType data item:

Procedure

1. In the data navigator locate the Linktype data type.
2. Click **Create New LinkType Data Item** to create a new LinkType data item.
3. Select the source and target data types for the new link type.

The new data item appears in the **Available LinkType Data Items** table.

When you create dynamic links, the LinkType data type is available for selection. See Chapter 8, “Links,” on page 89 for more information.

Document data types

Custom URL Document data types are derived from the predefined Doc data type.

You can add additional fields to the predefined Doc type and you can add data items. You cannot modify or delete the built-in fields in a custom URL Doc data type.

Adding new Doc data items

Use the following procedure to add a new Doc data item:

Procedure

1. To create a new Doc data item, click **Create a New Doc Data Item**.

The Create Doc Data Item window opens.

2. Type a Document name.
3. Type a description for the document.
4. Type the IP address of the document.
5. Click **OK**.

The new Doc data item is displayed in the table.

FailedEvent data types

The FailedEvent data type, together with the ReprocessedFailedEvents policy, provides you with a way to deal with failed events that are passed from the ObjectServer.

Both the FailedEvent data type and the ReprocessedFailedEvents policy are predefined and are stored in the global repository.

Note: The best practice to deal with failed events is to run PolicyActivatorService at regular intervals.

Viewing FailedEvent data items

Each FailedEvent data item row includes four fields.

- Key
- EventContainerString
- Policy Name
- EventReader name

You can use this information to re-create the EventContainer and send it back to the original policy that caused the error.

Hibernation data types

When you call the Hibernate function in a policy, the policy is stored as a Hibernation data item for a certain number of seconds.

You typically do not need to create or modify Hibernation data items using the GUI. However, you can delete stored hibernations if an error condition occurs and the hibernations are not woken up by the policy activator or another policy. See the *Solutions Guide* for more information about handling hibernations.

Working with composite data types

Composite data types are data types that have one or more fields dynamically linked to fields in another data type.

Composite data types are useful for creating a single data type that references information in more than one data source or that references more than one table or other structure in a single data source. You can use composite data types to retrieve and update data in data sources. You cannot use composite data types to insert new data or delete data.

Complete the following steps to create a composite data type:

- Create a composite internal or external data type.
- Edit the data type and create static or dynamic link from base data type to the target data type.
- Create a linked field for the base data type.

Creating composite data types

To create a composite data type, you create a base data type. The base data type can be an internal or external data type.

Before you begin

See the following links about creating data types to determine the type of composite data type you want to create:

- For information about creating internal data types, see “Creating internal data types” on page 63.
- For information about creating external data types, see “External data types” on page 64.

Procedure

1. In the navigation tree, expand **System Configuration > Event Automation** click **Data Model** to open the **Data Model** tab.
2. Select the data source from the data sources list.
3. Click the **New Data Type** icon. A new **Data Type Editor** tab opens.
4. Create your chosen data type.

Creating linked fields

To create a linked field in a composite data type, you create a link from the base data type to the target data type. Then, add a field to the data type by using a linking expression as the name of the field. The linking expression specifies which field in the target data type you want the linked field to reference.

The linking expression syntax is as follows:

```
links.type.item.field
```

- **type** is the name of the target data type
- **field** is the name of the field you want to reference
- **item** identifies the OrgNode in the array returned by the linking expression as first, last, or array[n]

An array of OrgNodes is a zero-based array, where **array[0]** is the first item. For example, the following linking expression references the value of the Name field in the first Customer OrgNode returned when a link is evaluated:

```
links.Customer.first.Name
```

The following linking expression references the value of the Location field in the second Node OrgNode returned when a link is evaluated:

```
links.Node.array[1].Location
```

Configuring a linked field on a composite data type

Complete the following steps to create a dynamic or static link and a linked field from the base data type to the target data type.

Before you begin

See the following sections about creating links to determine which type of link you want to create for your composite data type.

- For information about creating dynamic links, see “Creating dynamic links” on page 90.
- For information about creating static links, see “Creating static links” on page 93.

Procedure

1. In the navigation tree, expand **System Configuration > Event Automation** click **Data Model** to open the **Data Model** tab.
2. Expand the data source that contains the data type you want to edit, select the data type, double-click the name of the data type. Alternatively, right-click the data source and click **Edit**.
3. Create a dynamic or static link, from the base data type to the target data type.
4. In the **New Field** area of the **Table description** tab, click **New** to open the Field properties window to create a field for the base data type: Complete the following steps to create the linked field:
 - a. In the **ID field**, give the field a unique name
 - b. In the **Field Name** field, add a linking expression as the field name.
 - c. From the **Format** list, select the type of data to be held in this field.
 - d. In the **Display name** field, add the display name.
 - e. In the **Description** field, add the description.

Note: If using a link by key and the data type is internal, the field referenced as the key must match the key field in a row in the target data type. Otherwise, **NULL** is returned.

- f. Click **OK**. The field you created shows in the list of fields in the **Table Description** tab.
5. Click **Save** to add the changes to the data type.

Chapter 7. Data items

Data items are elements of the data model that represent actual units of data stored in a data source.

The structure of this unit of data depends on the category of the associated data source. For example, if the data source is an SQL database data type, each data item corresponds to a row in a database table. If the data source is an LDAP server, each data item corresponds to a node in the LDAP hierarchy.

Viewing data items

Use this procedure to view the data items for a data type.

Procedure

1. Locate the data type in the data connections list.
2. Select a data type and click the **View Data Items** next to the data type. If you have multiple data items open and you select **View Data Items** on a data type you opened already, the tab switches to the existing open data item tab.

When viewing data items, Netcool/Impact has a built-in threshold mechanism to control how much data gets loaded. The default threshold limit is 10000. If the underlying table which the data type points has more than 10000 rows which match the data type filter, Netcool/Impact shows a warning message indicating that the number of rows for the data type exceeds the threshold limit.

Note: The threshold limit is set in `NCHOME/impact/etc/server.props` using the property, `impact.dataitems.threshold`. To view data exceeding the threshold limit, the `impact.dataitems.threshold` property would need to be modified and the server restarted. The higher the value is set, the more memory is consumed. The heap settings for both the **ImpactProfile** and the **TIPProfile** would have to be increased from the default values. For more information about setting the minimum and maximum heap size limit, see the chapter on *Self Monitoring* in the *Netcool/Impact Administration Guide*.

Adding new data items

Use this procedure to add a new data item.

Procedure

1. In the **Data Model** tab, select the appropriate data type and click **View Data Items**.
2. To add a new data item to the table, click the **New Data Item** in the toolbar. The window that opens depends on the data type configuration.
3. Enter the information in the window.
4. Click **OK** to save the data item and close the window.

The new data item is listed in the table.

Filtering the view

You can limit the number of data items shown by entering a search string in the Filter field.

The filter syntax depends on the type of data type. For information about entering filter syntaxes, see the *Administration Guide*.

Editing data items

Use this procedure to edit a data item.

Procedure

1. To edit a data item, select the data item and click **Edit**.
The edit window you see depends on the data type configuration.
2. Change the information as necessary.
3. Click **OK** to save the changes and close the window.

Deleting data items

Use this procedure to delete a data item.

Procedure

1. In the **Data Model** tab, select the data items that you want to delete. Check marks are placed in the check boxes next to the selected data items and the data items are highlighted.
If you want to delete all the data items in the table, click the All link. Checkmarks are placed in every check box in the **Select:** column and the data items are highlighted. You can clear individual data items if you decide you do not want to delete all of them.
2. Click the **Delete** icon to delete the selected data items.

Chapter 8. Links

Links are elements of the data model that define relationships between data types and data items.

Static links define relationships between data items, and dynamic links define relationships between data types. Links are an optional component of the Netcool/Impact data model.

Dynamic links

Dynamic links define a relationship between data types.

This relationship is specified when you create the link and is evaluated in real time when a call to the `GetByLinks` function is encountered in a policy. Dynamic links are supported for internal, SQL database and LDAP data types.

The relationships between data types are resolved dynamically at run time when you traverse the link in a policy or when you browse links between data items. They are dynamically created and maintained from the data in the database.

The links concept is similar to the JOIN function in an SQL database. For example, there might be a 'Table 1' containing customer information (name, phone number, address, and so on) with a unique Customer ID key. There may also be a 'Table 2' containing a list of servers. In this table, the Customer ID of the customer that owns the server is included. When these data items are kept in different databases, Netcool/Impact permits the creation of a link between Table 1 and Table 2 through the **Customer ID** field, so that you can see all the servers owned by a particular customer.

You can use dynamic links only at the database level. (When relationships do not exist at the database level, you need to create static links.) You can create dynamic links for all types of data types (internal, external, and predefined). See Chapter 6, "Data types," on page 61 for information about the kinds of data type.

Dynamic links are unidirectional links, configured from the source to the target data type.

Static links

Static links define a relationship between data items in internal data types.

Static links are supported for internal data types only. Static links are not supported for other categories of data types, such as SQL database and LDAP types, because the persistence of data items that are stored externally cannot be ensured.

A static link is manually created between two data items when relationships do not exist at the database level.

With static links, the relationship between data items is static and never changes after they have been created. You can traverse static links in a policy or in the user interface when you browse the linked data items. Static links are bi-directional.

Dynamic links overview

Dynamic links use a specified method to link data items of the source data type to the data items of a target data type.

The linking methods are described below:

Table 44. Linking Methods

Link By:	Description
Key	This method evaluates an expression from one data type and matches this to the Key field of the target data type.
Filter	This method uses a filter expression to describe the link between any fields in the source type to any fields of the target data type.
Policy	This method runs a specified policy to look up data items in the target and link all the retrieved data items to data items of the source type.

Creating dynamic links

Use this procedure to create a dynamic link.

Procedure

1. To open the Data Type editor, click a data type name.
2. In the Data Type editor, select the **Dynamic Links** tab.
3. You can create the following types of dynamic links:
 - Link By Filter. For more information about creating links by filter, see “Adding new links by filter.”
 - Link By Key. For more information about creating links by key, see “Adding new links by key” on page 91.
 - Link By Policy. For more information about creating links by policy, see “Adding new links by policy” on page 92.
4. Click **OK** and click **Save** on the main to tab to implement the changes.

Adding new links by filter

Use the following procedure to add a new link by filter.

Procedure

1. Click **New Link by Filter**.
2. Enter the information in the New Link By Filter window, using this table as your guide:

Table 45. New Link by Filter Window

Window element	Description
Target Data Type	Select the target data type from the list.
Exposed Link Type	Select a link to follow from the list. The target data type name (in other words the exposed link) and the link type data items that match this source and target. See “LinkType data types” on page 83.

Table 45. New Link by Filter Window (continued)

Window element	Description
Filter	A filter is an expression that specifies which fields in the source and target types have to match in order for a link to exist. It can be either a simple expression (source name = target name) or a complex expression defined by a Boolean operator that indicates the order of the operation: (Custname = '%customer%') AND (device_num = %DeviceNumber%)

The link appears in the **New Link By Filter** table in the **Dynamic Links** tab.

3. Click **OK** and click **Save** on the main to tab to implement the changes.

Adding new links by key

When you define a Link by Key dynamic link, you specify a field in the source and target data types that contains a matching value.

Procedure

1. Click **New Link by Key**.
2. Enter the information in the window, using this table as your guide:

Table 46. New Link by Key window

Window element	Description
Target Data Type	Select the target data type from the list. For example, User.
Exposed Link Name	Select a link to follow from the list. For example, User. The target data type name (in other words the exposed link) and the link type data items that match this source and target.
Foreign Key Expression	Type the foreign key expression, for example: LastName + ", " + FirstName For more information about foreign key expression, see "Foreign key expressions."

The new link appears as a row in the **New Link By Key** table in the **Dynamic Links** tab.

3. Click **OK** and click **Save** on the main to tab to implement the changes.

Foreign key expressions

You can build the expression from one or more fields.

Type a field name or combination of field names in the source type that match the **Key** field in the target type. For example, if you want the key into the source type to be a field called 'NodeName', you enter NodeName. You can enter more than one field by entering the characters '+' '+' to join them.

For example, if the source type has a **FirstName** field and a **LastName** field and the target **Key** field is **Name**, you can create the link by entering the following expression:

FirstName + ' ' + LastName

The above expression is applied to the following field value pairs, for example, if in the source the fields are:

FirstName = 'John'
LastName = 'Doe'

The resulting value for the target **Key** field (Name in this case) is:

Name = 'John Doe'

this matches to:

'John' + ' ' + 'Doe' = 'John Doe'

Adding new links by policy

Use this procedure to add a new link by policy.

Procedure

1. Click **New Link by Policy**.

The New Link By Policy window opens.

2. Enter the information in the window, using this table as your guide:

Table 47. Internal Data Type - New Link by Policy window

Window element	Description
Target Data Type	Select the target data type from the list. For example, LinkPolicy.
Exposed Link Type	Select a link to follow from the list. For example, LinkPolicy. The target data type name (in other words the exposed link) and the link type data items that match this source and target.
Policy to execute to find links	Select a policy from the list of available policies. For example, GetPolicy.

The new link appears as a row in the table in the **Dynamic Links** tab.

3. Click **OK** and click **Save** on the main to tab to implement the changes.

Editing dynamic links

Use this procedure to edit a dynamic link.

Procedure

1. To edit a link, click the **Edit** in the row of the link you want to edit.
2. Make any necessary changes.
See “Dynamic links overview” on page 90 sections for more details.
3. Click **OK** and click **Save** on the main to tab to implement the changes.

Deleting dynamic links

Use this procedure to delete a dynamic link.

Procedure

1. In the **Select:** column, select the links that you want to delete. Check marks are placed in the check boxes next to the selected links and the links are highlighted.
If you want to delete all the links in the table, click the **All** link. Check marks are placed in every check box in the **Select:** column and the data links are highlighted. You can clear the check boxes for the individual data links if you decide you do not want to delete them.
2. Click the **Delete** link to delete the selected links.

Working with static links

You can view static links and create a static link between the data items of internal data types.

Creating static links

Use this procedure to create a static link for an internal data type.

Procedure

1. Open the **Global** tab.
2. In the data navigator, locate the source internal data type, and click **View Data Items**.
The Data Item editor opens in the Main Work panel.
3. Click the **Links** icon in the **Links** column next to one of the data item rows.
The Link Editor window opens.
4. Select **Target Type of Linked Items** from the selection list.
Only Internal and Predefined data types show in the list.
5. To add a link, highlight the data items you want that are listed in the **Unlinked Data Items** list and click **Add**.
The items move to the **Linked Data Items and LinkTypes** list.
6. To remove a link, highlight the data items that you want to remove from the **Linked Data Items** list and click **Remove**.
The data items are returned to the **Unlinked Data Items** list.
7. Click **OK** to save and close the editor.

Browsing dynamic links

You can browse dynamic links by using the Dynamic Links Tree Viewer.

Procedure

1. To browse links for internal, external and pre-defined data types, locate the data type you want to browse.
2. Click **View Data Items** next to the data type to open the Data Items Editor in the Main Work panel.
3. Locate the data item you want to see links for and click **View Linked Data Items Browser** to open the browser.
4. Click the plus sign next to the top-level links to view child links.
5. Click a data item to see more detailed information about the event.
6. Click the x to close the browser.

Chapter 9. Working with policies

You use the policy editor to create, manipulate, save, delete and edit policies.

You can create new policies from scratch, or use a policy wizard. Policy wizards present a series of windows that help you through the policy creation process.

Policies overview

Policies consist of a series of function calls that manipulate events and data from your supported data sources.

A policy, for example, can contain a set of instructions to automate alert management tasks, defining the conditions for sending an e-mail to an administrator, or sending instructions to the ObjectServer to clear an event.

You use the policy editor to create, manipulate, save, delete and edit policies. You can create new policies from scratch, or use a policy wizard. Policy wizards present a series of windows that help you through the policy creation process.

Accessing policies

Use this procedure to access the policies.

Procedure

1. In the navigation tree, expand **System Configuration > Event Automation**, click **Policies** to open the **Policies** tab.
2. From the **Cluster** and **Project** lists, select the cluster and project you want to use.

Viewing policies


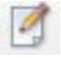




You view your policies in the **Policies** tab.

Before you create any policies, the **Policies** tab is empty. To display a list of policies for a project select a project from which you want to view the policies. If you want to display a list of all your policies, not just those associated with a particular project, you can access the entire list in the global repository. You can also create a new policy in the global repository if you do not want to add it to a project at the current time. It can be added to a project later. For more information about the global repository, see “Global repository” on page 17.

Policies panel controls

An overview of the policies task pane icons and indicators.

Table 48. Policy task pane controls

icon	Description
	Click the New Policy icon to create an IPL policy. To create a policy using JavaScript select the JavaScript Policy option. To create a policy using one of the policy wizards, select Use Wizard . Remember: If you use UTF-8 characters in the policy name, make sure that the locale on the Impact Server where the policy is saved is set to the UTF-8 character encoding.
	Select a policy and use this icon to edit it. Alternatively, you can edit a policy by right clicking its name and selecting Edit in the menu.
	Select a policy and use this icon to delete it from the database. Alternatively, you can delete a policy by right clicking its name and selecting Delete in the menu.
	Click the icon to open a window where you can recover an auto-saved policy. When the Enable Autosave option is selected, a temporary copy of the policy that you are working on is saved periodically. This feature saves your work in instances of a session timeout, browser crash, or other accident. Automatically saved policies are not shown in the policies navigation panel and are not replicated among clusters/import. You must first recover and save the drafted policy before you run it. For more information about recovering auto-saved policies, see “Recovering automatically saved policies” on page 99.
	Upload a Policy File. Click the icon to open the Upload a Policy window. You can upload policy and policy parameters files that you wrote in an external editor or files that you created previously.
	This icon is visible when a policy is locked, or the item is being used by another user. Hover the mouse over the locked item to see which user is working on the item. You can unlock your own items but not items locked by other users. If you have an item open for editing you cannot unlock it. Save and close the item. To unlock an item you have locked, right click on the item name and select Unlock . The tipadmin user and users who are assigned the impactAdminUser role are the only users who can unlock items that are locked by another user in exceptional circumstances.

Writing policies

You write policies in the policy editor using one of the following methods:

- You can write them from scratch. For more information, see “Writing custom policies” on page 98.
- You can use a policy wizard. For more information, see “Writing policies using wizards” on page 98.
- You can use JavaScript. For more information, see “Writing policies using JavaScript” on page 98

Policy wizards

You use policy wizards to create simple policies without having to manually create data types and add functions.

The wizards consist of a series of windows that guide you through the policy creation process. At the end of the process, you can run the policy immediately

without any further modification. However, if you want to modify the policy at any time, you can do so using the Policy editor.

Note: The OMNIbus event reader service must be running before you can use all wizards, except for the Web Services and XML DSA wizards.

You can use the following policy wizards:

Event Enrichment

Event enrichment is the process by which Netcool/Impact monitors an event source for new events, looks up information related to them in an external data source and then adds the information to them.

Event Notification

Event notification is the process by which Netcool/Impact monitors an event source for new events and then notifies an administrator or users when a certain event or combination of events occurs. Event Notification policies notify you that an event has occurred. Before you can use the Event Notification policy wizard, configure the e-mail sender service.

Event Relocation

Event Relocation policies allow you to send an event from one central ObjectServer to another ObjectServer.

Event Suppression

Event Suppression policies set a flag in an event in response to a database query. This flag can then be used in a filter to prevent the event from appearing in the Event List.

XinY X events in Y time is the process in which Netcool/Impact monitors an event source for groups of events that occur together and takes the appropriate action based on the event information. X Events in Y policies suppress events until a certain number of identifiable events occur within a specified time period.

You can configure two main parameters in the wizard.

- The number of incidents (N) in an event that will trigger a violation.
- The length of the rolling time window in which these (N) incidents must occur in order to trigger the violation.

The XinY policy wizard tracks how many times a single event with a single identifier is inserted or updated during the time window. If this number reaches (N) then it sends an event to Netcool OMNIbus indicating that the threshold for incidents has been exceeded for the particular event. The XinY policy wizard tracks incidents separately for each of the events that match the filter that triggers the policy.

XML XML policies are used to read and to extract data from any well-formed XML document.

Web Services

Web Services DSA policies are used to exchange data with external systems, devices, and applications using Web Services interfaces.

XML policies

XML policies are used to read and to extract data from any well-formed XML document.

The XML DSA can read XML data from files, from strings, and from HTTP servers via the network (XML over HTTP). The HTTP methods are **GET** and **POST**. **GET** is selected by default. In the XML wizard you can specify the target XML source and the schema file, to create the corresponding data source and data types for users. The wizard also updates the necessary property files and creates a sample policy to help you start working with XML DSA. When choosing the XML String option in the XML DSA wizard, ensure that the xml string you copy and paste does not contain references to stylesheet-related tags.

Writing custom policies

Use this procedure to develop a policy from scratch.

Procedure

In the **Policies** tab, select **New Policy > IPL Policy**.

Writing policies using wizards

Use this procedure to develop a policy using a wizard.

Procedure

1. In the **Policies** tab, select **New Policy > Use Wizard**.
2. Click on the wizard name to open the first wizard window.
3. Follow the on-screen instructions and click **Next**.
4. At the final window, click **Finish** to create the policy.

Writing policies using JavaScript

Use this procedure to develop a policy using JavaScript.

Procedure

In the **Policies** tab, select **New Policy > JavaScript Policy**.

Editing policies

Use this procedure to edit an existing policy.

Procedure

1. In the **Policies** tab, select a policy name in the list.
2. Right click on the policy and select **Edit** or click the **Edit** icon in the toolbar.

Deleting policies

Use this procedure to delete a policy.

Procedure

- Select the policy in the policies pane and click the **Delete Policy** icon in the toolbar.
- You can also delete a policy by right clicking its name in the policies pane and selecting **Delete** in the menu.

Recovering automatically saved policies

When the autosave option is selected you can recover and save an automatically saved policy.

Procedure

1. In the **Policies** tab, click the **Auto-Save version** icon in the toolbar.
2. Choose one auto-saved policy from the **Drafted Policy** list.
3. Click **Open** to view the drafted policy in the editor.
4. Click **Save** to save the drafted policy.

The auto saved copy is removed when you click **Save**.

Policy editor

The GUI provides a policy editor that you can use to create and edit policies.

The policy editor offers a text editor with syntax highlighting, a function browser, a syntax checker, a tree viewer, and other utilities to make it easy to manage policies. You can also write policies in an editor of your choice and then upload them into Netcool/Impact. After they are uploaded, you can edit them and check the syntax using the policy editor.

Note: If you create and edit a policy using an external editor of your choice, you must check its syntax using the `nci_policy` script before you run it. For more information about the `nci_policy` script, see the *Administration Guide*.

Policy editor toolbar controls

An overview of the policy editor toolbar controls.

Table 49. Policy Editor toolbar options






Icon	Description
	The Save icon saves the current policy. Use the Save with comments option to save your policy with comments. To save a policy with a different file name click Save as... Remember: If you use UTF-8 characters in the policy name, check that the locale on the Impact Server where the policy is saved is set to the UTF-8 character encoding.
	Restore your work to its state before your last action, for example, add text, move or, delete. Undo works for one-level only.
	Restore your work to its state before you selected the Undo action. Redo works for one-level only.
	Cut highlighted text. In some instances, due to browser limitations, the Cut icon cannot be activated. Use the keyboard short cut <code>Ctrl + x</code> instead.
	Copy highlighted text. In some instances due to browser limitations, the Copy icon cannot be activated. Use the keyboard short cut <code>Ctrl + c</code> instead.

Table 49. Policy Editor toolbar options (continued)














Icon	Description
	<p>Use this icon to paste cut, or copied text to a new location. In some instances due to browser limitations, the Paste icon cannot be activated. Use the keyboard short cut Ctrl + v instead.</p> <p>To copy and paste rich text formatted content, for example from a web page or document file:</p> <ol style="list-style-type: none"> 1. Paste the content into a plain text editor first to remove the rich text formatting. 2. Copy the content from the plain text editor into the policy editor.
	<p>Use this icon to find and replace text in a policy. Search for a text string. Type the text that you want to find, choose if you want to run a case-sensitive search, and choose the direction of the search.</p> <p>Search for text and replace it with a text you specify. Type the text that you want to search for. Type the replacement text. Choose if you want to run a case-sensitive search, and choose the direction of the search.</p>
	<p>Click the Go To icon to show a Go To Line field in the policy editor. Type the number of the line you want the cursor to go to. Click Go.</p>
	<p>Insert a selected function, an action function, or a parser function, in your policy. Add additional parameters for the function if required.</p> <p>The toolbar selection lists provide you with a set of functions to use in your policy.</p>
	<p>Access a list of data types. The Data Type Browser icon simplifies policy development by showing available data types and details including field name and type information. You do not have to open the data type viewer to get the data type information.</p>
	<p>The Check Syntax icon checks the policy for syntax errors. If there are errors, the error message locates the error by the line number. If there are no errors, a message to that effect is shown.</p>
	<p>Check the sanity of the policy that you specified in the Policy Name: field.</p>
	<p>Click the Run Policy icon to start the policy. After removing all syntax errors, you can run the policy to ensure that it produces the result you wanted. To run your policy with additional parameters, use the Run with Parameters option. You can use this option after you configure runtime parameters for your policy.</p>
	<p>Use this icon to set the runtime parameters for the policy. For more information, see “Setting policy runtime parameters in the editor” on page 103.</p>
	<p>Click the View Version History icon to view the history of changes made to policies, and compare different versions of policies. For more information about version history interface, see “Using version control interface” on page 111.</p> <p>Important:</p> <p>The View Version History icon is disabled for new and drafted policies and it becomes active after the policy is committed to server.</p> <p>This option is supported only with the embedded SVN version control system.</p>

Table 49. Policy Editor toolbar options (continued)

Icon	Description
	Click this icon to view the policy logs in the log viewer. For more information about the policy log viewer, see “Services log viewer” on page 120.
	The Graphic View is not available for JavaScript policies. For information about viewing policies in the graphic view, see “Graphic view of a policy” on page 102.
	Click this icon to manually enable or disable the syntax highlighter. For information about automatically configuring the syntax highlighter, see “Policy syntax highlighter.”

Policy syntax checking

While you are creating your policy, you can check to ensure that the syntax is correct.

When you select the **Check Syntax** icon, a list of errors are shown at the bottom of the policy editor. If there are no errors in the policy, the following message is displayed:

Syntax checking successful. No error found.

If the checker finds errors, you will see a table listing all the errors that were found.

The **Type** column of the table contains an error indicator, either Warning or Error.

The **Line** column of the table contains the line number where the error occurred. To find the error, click the line number. The editor scrolls to that line in the script.

Policy syntax highlighter

When you create a policy, the syntax highlighter can be configured to automatically toggle itself off at startup if the policy exceeds a predefined character limit. When working with large policies in the policy editor, disabling the syntax highlighter can alleviate possible performance slowdowns.

Procedure

1. Open a policy, in the policy editor toolbar, click the toggle icon to manually enable or disable the syntax highlighter.
2. The syntax highlighter can be configured to automatically toggle itself off at startup when the policy exceeds a specified character limit.
 - a. Open the **Policies** tab. On the right side of the tab next to the **Refresh** icon, click the **Edit options** icon (the down arrow) and select **Personalize**.
 - b. In the **Character limit for syntax highlighting** field, type the character limit for the policies. When a policy reaches this character limit, the syntax highlighter is automatically turned off.
 - c. Click **Save**.
 - d. Reopen the **Policies** tab to implement the changes

Optimizing policies

After you create your policy, you can check to see whether there is a way to improve it.

Procedure

1. Click the **Optimize** icon.

The Optimization handles three functions:

- Hibernate
- GetByKey
- GetByFilter

For the Hibernate function, the optimization checks to make sure that you have a RemoveHibernate function with the same hibernation key and notifies you if you do not. For the GetByKey and GetByFilter functions, the optimization checks the data type and sees what fields are returned from a data type. It then checks the policy to see if all of the fields are being used. When all of the fields from the data type are not being used, you receive a message showing which fields are not being used. You can change the data type fields if required.

2. Click **Save** to implement any changes. When you change a policy and you want to click **Optimize** again you must save the policy first. The optimize feature works from the saved version and not the modified version.

Running policies with parameters in the editor

If you specified any runtime parameters for the policy you can run the policy with these parameters in the policy editor.

Procedure

1. Click the **Run with Parameters** icon to open the Policy Runtime Parameters window.

Note: The fields you see in the Policy Runtime Parameters window depend on the runtime parameters and values you specified for the policy. If you have not set a default value for a parameter you must provide it now, otherwise a NULL value will be passed.

For more information about setting runtime parameters, see “Setting policy runtime parameters in the editor” on page 103.

2. Click **Execute** to run the policy with parameters.

Graphic view of a policy

Use the graphic view feature to view the code lines that contain variable assignments, and log statements of a policy. The graphic view feature is not available for JavaScript policies.

1. In the policy editor menu, click the **Graphic View** icon to open the graphic view window.
2. Select the **Assignments** check box to show the assignments in the graphical view. The graphical view shows the code lines that contain variable assignments. For example, `Type = 'REPORT_Policy';`.
3. Select the **Logging** check box to show logging details in the graphical view. The graphical view now shows the code lines that contain `log()` statements. For example, `log("Start Add PolicyProcessMapping");`.

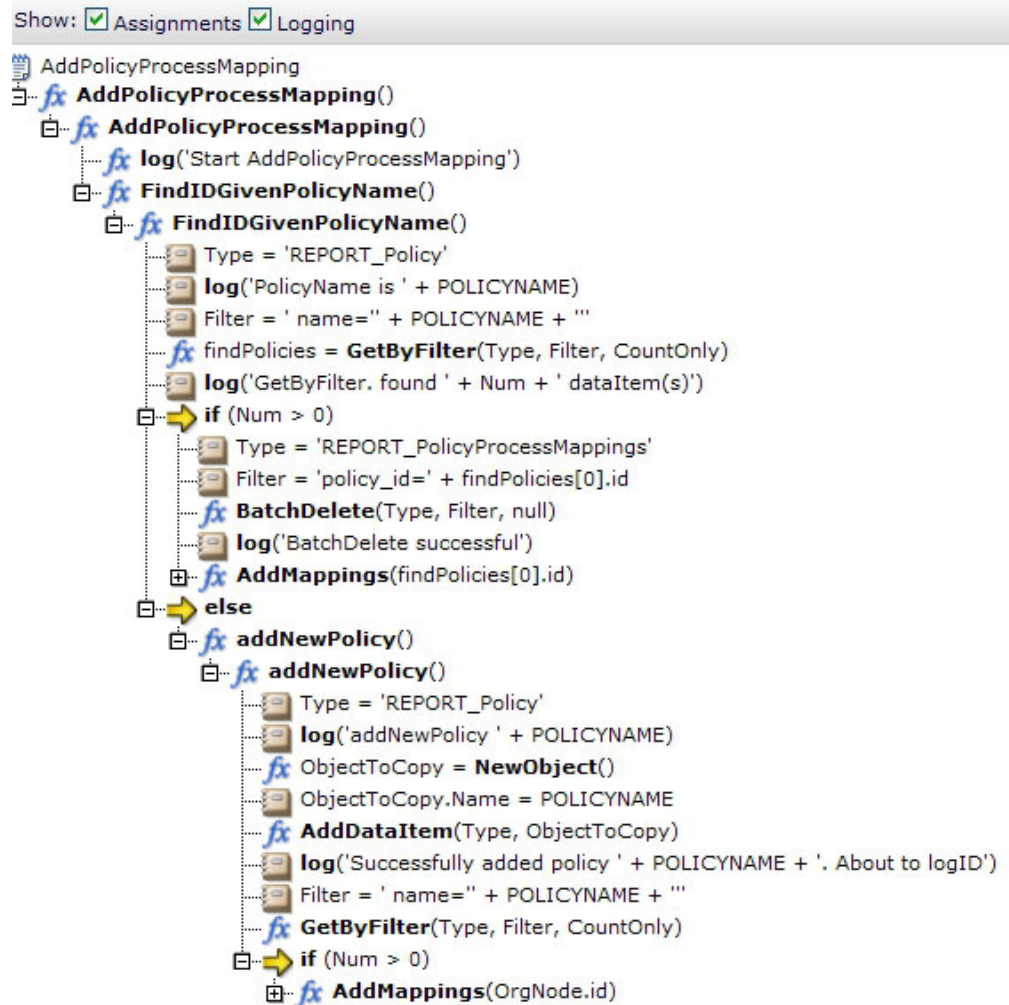


Figure 1. Display graphical view of policy window

Browsing data types

Use this procedure to view available data types and their details directly from the policy editor.

Procedure

1. Click the **Types Browser** icon.
2. Click a data type to see the details. The details are shown under the list.

Setting policy runtime parameters in the editor

Use this procedure to set the runtime parameters for your policy in the policy editor.

Procedure

1. In the policy editor toolbar, click the **Configure Runtime Parameters** icon to open the policy runtime parameter editor.
2. Click **New Runtime Parameter** to open the Create a New Policy Runtime Parameter window.

Enter the information in the new runtime parameter configuration window. Required fields are marked with an asterisk (*).

3. To edit an existing runtime parameter, select the check box next to the parameter and select **edit** in the corresponding cell of the **Edit** column.
4. Click **OK** to save the changes to the parameters and close the window.

Policy runtime parameter configuration window

Policy runtime parameters take a set of attributes.

Table 50. List of attributes that are used with a policy runtime parameters

Attribute	Description
Name	Type a name to describe the parameter.
Label	Type a label that will appear in the Policy Trigger window.
Format	Choose a format.
Default Value	Type a default value that will always display in the Policy Trigger window, to avoid entering it each time.
Description	Type some text to describe the parameter.

Adding functions to policy

Use this procedure to add a function to a policy.

Procedure

1. Click the **Insert function** icon and select one of the functions.
2. Enter the required parameters in the new function configuration window.

Note: When entering a string, check that all string literals are enclosed in quotation marks ("*string*"), to distinguish them from variable names, which do not take quotation marks.

List and overview of functions

A list of all functions with a short overview.

Table 51. List of functions

Name	Type	Description
Activate	Policy	The Activate function runs another policy.
ActivateHibernation	Policy	The ActivateHibernation function continues running a policy that was previously put to sleep using the Hibernate function. You must also run the RemoveHibernation function to remove the policy from the hibernation queue and to free up memory resources.
AddDataItem	Database, Internal	The AddDataItem function adds a data item to a data type.
BatchDelete	Database	The BatchDelete function deletes a set of data items from a data type.
BatchUpdate	Database	The BatchUpdate function updates field values in a set of data items in a data type.

Table 51. List of functions (continued)

Name	Type	Description
BeginTransaction	Database	The BeginTransaction is a local transactions function that is used in SQL operations.
CallDBFunction	Database	CallDBFunction calls an SQL database function.
CallStoredProcedure	Database	The CallStoredProcedure function calls a database stored procedure.
ClassOf	Context	The ClassOf function returns the data type of a variable.
CommandResponse	Systems	Use the CommandResponse function to run interactive and non-interactive programs on both local and remote systems.
CommitChanges	Database	Used only in connection with GetByFilter, and GetByKey functions to force updates in a database.
CommitTransaction	Database	The CommitTransaction function is a local transactions function that is used in SQL operations.
CurrentContext	Context	The CurrentContext function returns the current policy context.
Decrypt	String	The Decrypt function decrypts a string that has been previously encrypted using Encrypt or the nci_crypt tool.
DeleteDataItem	Database, Internal	The DeleteDataItem function deletes a single data item from a data type.
Deploy	Miscellaneous	The Deploy function copies data sources, data types, policies, and services between server clusters.
DirectSQL	Database	The DirectSQL function runs an SQL operation against the specified database and returns any resulting rows to the policy as data items.
DataItems	Keys	Provided for backward-compatibility only.
Distinct	Array	The Distinct function returns an array of distinct elements from another array.
Encrypt	String	The Encrypt function encrypts a string.
Eval	Context	The Eval function evaluates an expression using the given context.
EvalArray	Array, Context	The EvalArray function evaluates an expression using the given array.
Exit	Policy	You use the Exit function to stop a function anywhere in a policy or to exit a policy.
Extract	String	The Extract function extracts a word from a string.
FindFilters	Database	Provided for backward-compatibility only.

Table 51. List of functions (continued)

Name	Type	Description
Float	Numeric	The Float function converts an integer, string, or Boolean expression to a floating point number.
FormatDuration	Time	The FormatDuration function converts a duration in seconds into a formatted date/time string.
GetByFilter	Database, Internal, ITNM, LDAP, XML	The GetByFilter function retrieves data items from a data type using a filter as the query condition.
GetByKey	Database, Internal, LDAP	The GetByKey function retrieves data items from a data type using a key expression as the query condition.
GetByLinks	Database, Internal, XML	The GetByLinks function retrieves data items in target data types that are linked to one or more source data items.
GetByXPath	XML	The GetByXPath function provides a way to parse an XML string or get an XML string through an URL specified as parameter.
GetClusterName	Variables	You use the GetClusterName function inside a policy to identify which cluster is running the policy.
GetDate	Time	The GetDate function returns the date/time as the number of seconds expired since the start of the UNIX epoch.
GetFieldValue	Java	Use this function to get the value of static, or non static fields. For non static fields, use the variable <i>FieldName</i> for a Java class or <i>TargetObject</i> for a Java object. For a static Java class field, use the variable <i>ClassName</i> .
GetGlobalVar	Variables	This function retrieves the global value saved by previous SetGlobalVar calls.
GetHTTP	REST	You can use the GetHTTP function to retrieve any HTTP URL.
GetHibernatePolicies	Policy	The GetHibernatePolicies function retrieves data items from the Hibernation data type by performing a search of action key values.
GetScheduleMember	Time	The GetScheduleMember function retrieves schedule members associated with a particular time range group and time.
GetServerName	Variables	You use the GetServerName function inside a policy to identify which server is running the policy.
GetServerVar	Variables	You use this function to retrieve the global value saved by previous SetServerVar.

Table 51. List of functions (continued)

Name	Type	Description
Hibernate	Policy	The Hibernate function causes a policy to hibernate.
Int	Numeric	The Int function converts a float, string, or Boolean expression to an integer.
JRExecAction	Systems	The JRExecAction function executes an external command using the JRExec server.
JavaCall	Java	You use this function to call the method <code>MethodName</code> in the Java object <code>TargetObject</code> with parameters, or, to call the static method <code>MethodName</code> in the Java class <code>ClassName</code> with parameters.
Keys	Context	The Keys function returns an array of strings that contain the field names of the given data item.
Length	Array, String	The Length function returns the number of elements or fields in an array or the number of characters in a string
Load	JavaScript	You use this function to load a JavaScript library into your JavaScript policy.
LocalTime	Time	The LocalTime function returns the number of seconds since the beginning of the UNIX epoch as a formatted date/time string.
Log	Policy	The Log function prints a message to the policy log.
Merge	Context	The Merge function merges two contexts or event containers by adding the member variables of the source context or event container to the those of the target.
NewEvent	Context, Database	The NewEvent function creates a new event container.
NewJavaObject	Java	The NewJavaObject function is used to call the constructor for a Java class.
NewObject	Context	The NewObject function creates a new context.
ParseDate	Time	The ParseDate function converts a formatted date/time string to the time in seconds since the beginning of the UNIX epoch.
PassToTBSM	TBSM	Use the PasstoTBSM function to send event information from Netcool/Impact to TBSM.
RemoteTBSMShell	TBSM	A stand-alone implementation of Netcool/Impact can run <code>RADShell</code> commands from a policy in Tivoli Business Service Manager.

Table 51. List of functions (continued)

Name	Type	Description
RExtract	String	The RExtract function uses regular expressions to extract a substring from a string.
RExtractAll	String	The RExtractAll function uses regular expression matching to extract multiple substrings from a string.
Random	Numeric	The Random function returns a random integer between zero and the given upper bound.
ReceiveJMSMessage	JMS	The ReceiveJMSMessage function retrieves a message from the specified JMS destination.
RemoveHibernation	Policy	The RemoveHibernation function deletes a data item from the Hibernation data type and removes it from the hibernation queue.
Replace	String	The Replace function uses regular expressions to replace a substring of a given string.
ReturnEvent	Policy	The ReturnEvent function inserts, updates, or deletes an event from an event source.
RollbackTransaction	Database	The RollbackTransaction function rolls back any changes done by an SQL operation.
SendEmail	Notifications	The SendEmail function sends an email that uses the email sender service.
SendInstantMessage	Notifications	The SendInstantMessage function sends an instant message using the Jabber service.
SendJMSMessage	JMS	The SendJMSMessage function sends a message to the specified destination using the JMS DSA.
SetFieldValue	Java	Use the SetFieldValue function to set the field variable in the Java class to some value.
SetGlobalVar	Variables	The SetGlobalVar function creates in a policy a global variable which can be accessed from any local functions, library functions, and exception handlers in a policy.
SetServerVar	Variables	The SetServerVar function creates a server-wide global variable in a policy.
SnmpGetNextAction	SNMP, Sytems	The SnmpGetNextAction function retrieves the next SNMP variables in the variable tree from the specified agent.
SnmpGetAction	SNMP, Systems	The SnmpGetAction function retrieves a set of SNMP variables from the specified agent

Table 51. List of functions (continued)

Name	Type	Description
SnmpSetAction	SNMP	The SnmpSetAction function sets variable values on the specified SNMP agent.
SnmpTrapAction	SNMP	The SnmpTrapAction function sends a trap (for SNMP v1) or a notification (for SNMP v2) to an SNMP manager.
Split	String	The Split function returns an array of substrings from a string using the given delimiters.
String	String	The String function converts an integer, float, or boolean expression to a string.
Strip	String	The Strip function strips all instances of the given substring from a string.
Substring	String	The Substring function returns a substring from a given string using index positions.
Synchronised	Policy	Use the Synchronized function to write thread-safe policies for use with a multi-threaded event processor using IPL or JavaScript.
ToLower	String	The ToLower function converts a string to lower case characters.
TBSMShell	TBSM	This topic describes the TBSMShell action function which lets you put RADshell commands in a policy. With the TBSMShell function, you can change the TBSM configuration in a policy.
ToUpper	String	The ToUpper function converts a string to upper case characters.
Trim	String	The Trim function trims leading and trailing white space from a string.
URLDecode	String, REST	The URLDecode function returns a URL encoded string to its original representation.
URLEncode	String, REST	The URLEncode function converts a string to a URL encoded format.
UpdateEventQueue	Database	The UpdateEventQueue function updates or deletes events in the event reader event queue.
WSDMGetResourceProperty	Web Services	The WSDMGetResourceProperty function retrieves the value of a management property associated with a WSDM (Web Services Distributed Management) managed resource.
WSDMInvoke	Web Services	The WSDMInvoke function sends a web services message to a WSDM (Web Services Distributed Management) managed resource.

Table 51. List of functions (continued)

Name	Type	Description
WSDMUpdateResourceProperty	Web Services	The WSDMUpdateResourceProperty function updates the value or values of a management property associated with a WSDM (Web Services Distributed Management) managed resource.
WSInvoke	Web Services	Provided for backward-compatibility only.
WSInvokeDL	Web Services	The WSInvokeDL function is used to make Web services calls when a WSDL file is compiled with nci_compilewsdl, or when a Web services DSA policy wizard is configured.
WSNewArray	Web Services	The WSNewArray function creates a new array of complex data type objects or primitive values, as defined in the WSDL file for the Web service.
WSNewEnum	Web Services	The WSNewEnum function returns an enumeration value to a target Web service.
WSNewObject	Web Services	The WSNewObject function creates a new object of a complex data type as defined in the WSDL file for the Web service.
WSNewSubObject	Web Services	The WSNewSubObject function creates a new child object that is part of its parent object and has a field or attribute name of ChildName.
WSSetDefaultPKGName	Web Services	The WSSetDefaultPKGName function sets the default package used by WSNewObject and WSNewArray.

For more details about each of these functions, see the *Policy Reference Guide*.

Personalizing the policy editor

You can change some of the default behavior of the policy editor.

Procedure

1. Click the **Edit options** icon in the upper right corner of the policy editor panel and select **Personalize** from the menu.
2. Select the options that you want to personalize.
 - Select **Show Line Number** to view the line numbers for the policy editor.
 - Select **Enable Autosave** and the policy is automatically saved every 10 minutes while you are editing it.
3. Click **Save** to save the changes.

Changing default font used in the policy editor

Use this procedure to change the default font in the policy editor.

Procedure

1. Open the `$IMPACT_HOME/eWAS/profiles/ImpactProfile/installedApps/ImpactCell/guiserver.ear/netcool.war/styles/editor.css` file.
2. Update the values of the following entries with your own values:
 - `font-family`
 - `font-size`
 - `line-height`
3. (Firefox) Update the values of the following lines in the `$IMPACT_HOME/eWAS/profiles/ImpactProfile/installedApps/ImpactCell/guiserver.ear/netcool.war/scripts/editor.js` file with your own values:

```
this.editor.contentWindow.document.body.style.fontSize="16px";
this.editor.contentWindow.document.body.style.lineHeight="18px";
this.editor.contentWindow.document.body.style.fontFamily="Verdana,
Arial, Helvetica, sans-serif";
```
4. Refresh the browser to apply the changes.
It is also recommended to clear the browser cache.

Using version control interface

Use this procedure to view the version history of your policies.

Procedure

1. Open a policy in the policy editor.
2. Click the **View Version History** icon in the policy editor toolbar to open the version control interface.
You see the following columns:

Table 52. Version control interface columns

Column	Description
Version	Version number of the policy.
Author	Person updating the policy.
Date	Date the change was committed.
Comments	Any comments submitted with the version of the policy.

3. Click a version of the policy to view its contents.
 - To view the differences between versions of the policy click **View Differences**.
 - To revert to an older version of the policy select the version to which you want to revert and click **Revert**.

Uploading policies

You can upload policies and policy parameters files that you have written previously to the Impact Server.

Procedure

1. In the **Policies** tab, from the policy menu, click the **Upload a Policy File** icon. The Upload a Policy File window opens.
2. Select the check box for each type of file you want to upload, a policy file, or parameters file. You can upload both file types at the same time. The file

extension must end with .ipl for an IPL policy or .js for a JavaScript policy. Policy parameter file extensions must end with .params.

3. Type the path and file name, or click **Browse** to locate and select the policy or parameter file.
4. From the **Encoding** list menu, click the arrow to select the original encoding of the file you are uploading. The default option is **Unicode UTF-8**.
5. Click **Upload**.
6. The policy is added to the selected project in the **Policies** tab. The policies list refreshes automatically and shows the added policy in the policy list. The uploaded policy parameters file is stored in the Impact Server in \$IMPACT_HOME/policy.

Predefined policies

You can find predefined policies in the global repository, by selecting the **Global** in the project selection. No configuration is required for predefined policies.

Table 53. Predefined policies

Policy	Description
AddPolicyProcessMapping	This policy is used in reports. You do not need to change this policy.
DefaultExceptionHandler	<p>This policy is used to handle failed events if the policy failure is not handled locally using the Exception Handler. You can write your own policy if you need to. If you do not write your own, the provided policy is used by default.</p> <p>The DefaultExceptionHandler policy prints a log of the Events that failed to execute. To configure a customized error handling policy, refer to the Policy Logger information in the <i>Solutions Guide</i>.</p>
DeployProject	You can use this policy to copy the data sources, data types, policies, and services in a project between two running server clusters on a network. You can use this feature when moving projects from test environments into real-world production scenarios. For more information about automated project deployment, see “Automated project deployment feature” on page 20.
Export	<p>This policy is used by the nci_export script during the export of the Netcool/Impact configuration to another server.</p> <p>It is recommended that you do not change this policy.</p>

Table 53. Predefined policies (continued)

Policy	Description
FailedEventExceptionHandler	<p>When errors occur during the execution of a policy, the Policy Logger service executes the appropriate error handling policy, and temporarily stores the events as data items in a predefined data type called FailedEvent.</p> <p>FailedEvent is an internal data type and all data that is stored internally consumes memory. When you have resolved the reasons for the event failures you can reduce the amount of memory consumed using one of the following options:</p> <ul style="list-style-type: none"> • reprocess the failed events using the ReprocessFailedEvent. • delete the events from FailedEvent data type. <p>See “FailedEvent data types” on page 84 for more information.</p>
Import	<p>This policy is used during the import of the Netcool/Impact configuration from another server.</p> <p>It is recommended that you do not change this policy.</p>
ReprocessFailedEvent	<p>This policy is used to reprocess failed events. For more information about failed events, see “FailedEvent data types” on page 84.</p>

Chapter 10. Working with services

Services are runnable components of the Impact Server that you start and stop using both the GUI and the CLI.

Services overview

Services perform much of the functionality associated with the Impact Server, including monitoring event sources, sending and receiving e-mail, and triggering policies.

The most important service is the OMNIbus event reader, which you can use to monitor an ObjectServer for new, updated or deleted events. The event processor, which processes the events retrieved from the readers and listeners is also important to the function of Netcool/Impact.

Internal services control the application's standard processes, and coordinate the performed tasks, for example:

- Receiving events from the ObjectServer and other external databases
- Executing policies
- Responding to and prioritizing alerts
- Sending and receiving e-mail and instant messages
- Handling errors

Some internal services have defaults, that you can enable rather than configure your own services, or in addition to creating your own. For some of the basic internal services, it is only necessary to specify whether to write the service log to a file. For other services, you need to add information such as the port, host, and startup data.

User defined services are services that you can create for use with a specific policy.

Generally, you set up services once, when you first design your solution. After that, you do not need to actively manage the services unless you change the solution design.

To set up services, you must first determine what service functionality you need to use in your solution. Then, you create and configure the required services using the GUI. After you have set up the services, you can start and stop them, and manage the service logs.

Accessing services

How to access services in the Tivoli Integrated Portal


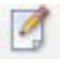


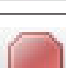




Procedure

1. In the navigation tree, expand **System Configuration** > **Event Automation** click **Services** to open the **Services** tab.
2. From the **Cluster** and **Projects** lists, select the cluster and project you want to use. A list of services related to the selected project is displayed.

Services panel controls

An overview of the services panel icons and indicators.

Table 54. Service Status panel icons and indicators

Element	Description
	Click the Create New Service icon to create a user-defined service using one of the available service templates.
	Click the Edit Service icon to edit an existing service using one of the available service templates. You can also double click on the service to open the service for editing.
	Click the View Service Log icon to access the log for the selected service. You can also view the log for a selected service by right clicking its name and selecting View Log .
	Select a stopped service and click the Start service icon to start it. Alternatively, you can start a service by right clicking its name and selecting Start .
	Select a running service and click the Stop Service icon to stop it. Alternatively, you can stop a service by right clicking its name and selecting Stop .
	Click the Delete Service icon to delete a user-defined service. Alternatively, you can delete a user-defined service by right clicking its name and selecting Delete . Important: You cannot delete a running service, you must stop it first.
	This indicator next to a service name indicates that the service is running.
	This indicator next to a service name indicates that the service is stopped.
	Source control locking for the service. This icon is visible when the service is locked or the item is being used by another user. Hover the mouse over the locked item to see which user is working on the item. You can unlock your own items but not items locked by other users. If you have an item open for editing you cannot unlock it. Save and close the item. To unlock an item you have locked, right click on the item name and select Unlock . The tipadmin user and users who are assigned the impactAdminUser role are the only users who can unlock items that are locked by another user in exceptional circumstances.

The following table describes the service icons used in Netcool/Impact.

Table 55. Service icons










Icon	Description
	Indicates a generic service icon used to indicate the following services: <ul style="list-style-type: none">• CommandExecutionManager• EventProcessor• PolicyLogger• SelfMonitoring• WSNotificationListener
	Indicates the CommandLineManager service.
	Indicates database-related services DatabaserEventReader , and ImpactDataBase .

Table 55. Service icons (continued)

Icon	Description
	Indicates a mail and message-related service including the following services: <ul style="list-style-type: none"> • EmailReader • EmailSender • JabberReader • JMSListener • JabberService
	Indicates policy activator services, including PolicyActivator , and MWMActivator .
	Indicates a hibernation service, HibernationPolicyActivator .
	Indicates Omnibus-related services, OMNIBusEventReader , and OMNIBusEventListener .
	Indicates an EventListener and EventReader service.
	Indicates an ITNMEventListener service.

List of services

A list of internal and user-defined Netcool/Impact services.

Table 56. Impact Services

Service	Type	Description
CommandExecutionManager	internal	The command execution manager is the service responsible for operating the command and response feature.
CommandLineManager	internal	Use the command-line manager service to access the Impact Server from the command line to configure services parameters as well as start and stop services.
DatabaseEventListener	internal	The database event listener service monitors an Oracle event source for new, updated, and deleted events.
DefaultEmailReader	internal	The email reader service reads incoming email, and runs policies based on the contents of the email.
DefaultJabberReader	internal	The jabber reader service is the instant message listener service.
DefaultPolicyActivator	internal	The policy activator service activates policies at startup or at the intervals you specify for each selected policy.
DatabaseEventReader	user defined	The database event reader is a service that polls supported, external SQL data sources at regular intervals to get business events in real time.

Table 56. Impact Services (continued)

Service	Type	Description
EmailReader	user defined	The email reader service reads incoming email, and runs policies based on the contents of the email.
EventListener	user defined	Event listeners monitor non-ObjectServer event source events.
EmailSender	internal	The e-mail sender is a service that sends e-mail through an external SMTP service.
EventProcessor	internal	The event processor manages the incoming event queue and is responsible for sending queued events to the policy engine for processing.
HibernatingPolicyActivator	internal	The hibernating policy activator service monitors hibernating policies and awakens them at specified intervals.
ITNMEventListener	internal	The ITNM event listener service listens for events sent from ITNM.
ImpactDatabase	internal	The Netcool Database Server runs as a service in the Impact Server
JMSMessageListener	internal, user defined	The JMS message listener service runs a policy in response to incoming messages sent by external JMS message providers.
JabberReader	user defined	The jabber reader service is the instant message listener service.
JabberService	internal	The jabber service acts as a Jabber client and is responsible for logging on to the external Instant Messaging services and sending instant messages.
MWMAActivator	internal	Maintenance Window Management service. An add-on for managing Netcool/OMNIbus maintenance windows.
OMNIbusEventListener	user defined	The OMNIbus event listener service is used to integrate with Netcool/OMNIbus and receive immediate notifications of fast track events.
OMNIbusEventReader	internal, user defined	OMNIbus event readers are services that monitor a Netcool/OMNIbus ObjectServer event source for new, updated, and deleted alerts and then runs policies when the alert information matches filter conditions that you define.
PolicyActivator	user defined	The policy activator service activates policies at startup or at the intervals you specify for each selected policy.
PolicyLogger	internal	The policy logger service is responsible for managing the policy log.
SelfMonitoring	internal	The self monitoring service is used to send messages about the internal state of Impact Server to an ObjectServer.

Table 56. Impact Services (continued)

Service	Type	Description
WSNotificationListener	user defined	Netcool/Impact provides the implementation for the consumer part of the WS-notification standard through the Web Services Notification Listener service.

Personalizing services

You can change the refresh period for the services tab.

Procedure

1. Click the **Edit options** icon in the upper right corner of the policy editor panel and select **Personalize** from the menu.
2. Select the options that you want to personalize.
 - Select the **Enable auto refresh** check box to automatically refresh the services.
 - Select the **Refresh interval** period. The services are automatically refreshed at time interval you select.
3. Click **Save**.

Configuring services

How to configure an existing service.

Procedure

1. In the **Services** tab, select the service you want to edit.
Right click and select **Edit** or click the **Edit** icon on the toolbar to open the service.
2. Enter the required information for the values and fields.
3. Click the **Save** icon to implement the changes or close the tab without saving to cancel any changes.

Creating services

How to create a user-defined service.

Procedure

1. In the navigation tree, expand **System Configuration > Event Automation** click **Services** to open the **Services** tab.
2. From the **Cluster** and **Projects** lists, select the cluster and project you want to use.
3. In the **Services** tab, click the **Create New Service** icon.
4. From the menu, select a template for the service that you want to create.
5. In the service configuration tab, provide the necessary information to create the service.
6. Click the **Save Service** icon
7. To edit a service, you can double click on the service, or right click on the service and select **Edit**.
8. Make the necessary changes to the service.

- Click **Save** to implement the changes.
- To cancel the changes, close the tab without clicking **Save**.

Important: You can create a user-defined service by using the defaults that are stored in the **Global** project.

Starting services

How to start a service that is stopped.

Procedure

- Select the service in the services pane and click **Start**.
- You can also start a service by right clicking its name in the services pane and selecting **Start** in the menu.

Stopping services

How to stop a service that is running.

Procedure

- Select the service in the services pane and click **Stop**.
- You can also stop a service by right clicking its name in the services pane and selecting **Stop** in the menu.

Deleting services

How to delete a user defined service.

Procedure

- Select the service in the services pane and click **Delete Service**.
- You can also delete a service by right clicking its name in the services pane and selecting **Delete**.

Important: Do not delete the default services. If you delete one, you cannot create new services of the type you delete. Deleting a user-defined service from the services panel, deletes it permanently from the database. If you want to remove it from a project, but retain it in the database, use the project editor.

Viewing services logs

Use this procedure to display the service log for a service.

Procedure

- Select a service in the services tab and click **View Service Log**.
- You can also view a service log by right clicking the service name in the services pane and selecting **View Log** in the menu.




Services log viewer

You can use the Services log viewer to view the results of your chosen service logs.

You can select the services from the drop-down menu. The window has a split screen so that you can view two logs for two different services simultaneously. You

can also create additional tabs from where you can run additional service logs and you can move between tabs. There is also an advanced filter option which you can use to filter the results of a log.

The log view has the following options:

Window element	Description
New tab	Click this option to create new tabs to view additional service logs. For more information, see “Creating new tabs” on page 122.
Default tab	This tab displays automatically when you access the service log viewer.
Drop down menu	Use this option to select the service you want to run a log for.
Apply Filter	When the results of the server log display, type in a filter name then select the check box. You can use this option to filter the results. For more information about log viewer results, see “Service log viewer results.”
	Click to stop the log.
	Click to clear the log.
	Click to start the log again.

Service log viewer results

The log viewer displays information relating to the following features: date, time, policy name, and pool thread name.

The log viewer shows only the latest logging information. If there is an error in the service log the error message displays in red font. You can click the icon next to the error message to get more information about the error.

To refine the log results you want to view, use the **Filter** option. To use the filter type in a string or use a Java regular expression.

Important: The filter expression assumes the default settings in `java.util.regex` pattern. For example, the filter always assumes case-sensitive flag.

Example of a regular expression:

```
\bt[a-z]+\b
```

This expression matches any word starting with letter t and followed by one or more letters from (A to Z).

To apply the filter, select the **Apply filter** check box. The new log message that matches the filter expression is displayed.

You can view results of multiple services, the window has a split screen to view two service log results on the same tab.

You can also create more tabs to view additional service log results using the **New Tab** option. For more information about creating new tabs, see “Creating new tabs.”

Creating new tabs

You can create multiple tabs to view additional service logs.

Procedure

1. If you want to view more service logs, click the **New Tab** option to display the **Log name** dialog.
2. Type in the name of the new tab, click **OK** to create the new tab in the **Log viewer** window.
3. Populate the fields in the tab to run the service log, for more information see “Services log viewer” on page 120.
4. As you create more tabs and view results and you can move from one tab to the other by clicking the tab heading at the top of the window. For more information see “Service log viewer results” on page 121

Event mapping

Event mapping allows you to map incoming events to one or more specific policies.

In some services you can set events to trigger policies when they match a filter. You create each filter by entering the filter SQL and assigning a policy to it. The filter instructs the policy to run whenever an event matches the filter.

Creating event filters

Use this procedure to create an event filter.

Procedure

1. Click the **New Mapping** icon to open the Create a New Event Filter window.
2. Provide the required information to create the filter.
This filter specifies the type of event that maps to the policy. For information about the filter configuration options, see “Event filter configuration window” on page 123.
3. From the **Policy to Run** list, select the policy that you want to run for the event type.
4. Click **Active**.
5. Click **OK** and the service configuration window gets refreshed with the new filter showing in the table.

Event filter configuration window

Use this information to configure the event filter.

Table 57. Create New Event Filter Window

Window element	Description
Filter Expression	Type a filter expression. This filter specifies the type of event that maps to the policy. For example, you have created a policy to run when Netcool/Impact receives an event from an Oracle database with a Department table. You want the policy to run when the entry in the department location field (DepLoc) is London. You type: DepLoc = "London"
Policy to Run	Select the policy to assign to the filter and run for the event type.
Active	Selected to activate the filter or cleared to deactivate the filter.
Chain	When chaining policies, select the Chain option for each event mapping that associates a restriction filter with a policy name. See the <i>Policy Reference Guide</i> for more information.
Analyze Filter	Click to discover any conflicts with filter mappings that have been set for a service.

Event mapping table

Overview of the event mapping table fields.

Table 58. Event mapping table

Window element	Description
Select:	When you place your mouse over the words all or none , the words become underlined as links. <ul style="list-style-type: none">Click all to select all the rows of filters. You can then click Delete at the bottom of the list to delete all the previously defined filters.Click none to clear all the rows of filters
Restriction Filter	Contains the filter.
Policy Name	Contains the name of the policy that triggers when the event matches the restriction filter.
Active	Selected to activate the filter or cleared to deactivate the filter.
Chain	When chaining policies, select the Chain option for each event mapping that associates a restriction filter with a policy name. See the <i>Policy Reference Guide</i> for more information.
Move	Use the arrows to change the position of the filters in the table. The order of the filters is only important when you select to stop testing after the first match.
Edit	To edit a filter, click the Edit button next to it.

Editing filters

Use this information to edit a filter.

Procedure

1. Locate the filter in the table and click **Edit** to open the Edit Event Filter window.
2. Edit the filter text and select a policy to run, as necessary.
3. Click **OK** to save the information and close the window.

The filter in the table in the **Event Mapping** tab shows your edits.

Reordering filters

Whether the order of the filters is important depends on which Event Matching option you select.

- When you select the **Stop testing after first match** option Netcool/Impact checks an incoming event against the filters in the order they are listed in the table until it gets a single match. It does not continue checking after it finds the first match.
- When you select **Test event with all filters**, the order is not important.

Deleting filters

Use this procedure to edit a filter.

1. In the **Select:** column, select the filters that you want to delete. (Click the **All** link to select all the filters in the table.)
2. Click the Delete link.

Filter analysis

By analyzing the event mapping table you can check the syntax and scope of all your event filters.

To find any conflicts with filter mappings that have been set for a service, choose which filters you want to analyze by selecting either **Active filters** or **All filters** in the **Filters to analyze** menu.

The **Filter Syntax Analysis Result** section displays a lists of all syntax errors that were found in the filters, the position where these syntax errors occur, and a brief description of the error.

Filter range overlap analysis in the **Filter Range Overlap Analysis Result** section shows you which of your filters overlap and what is the scope of their overlap. You can analyze your filters against the active filters only, or against all your defined filters by selecting one of the options in the **Analyze the filter against** menu.

The **Consolidated filter expression** section displays an expression that corresponds to all your currently configured event filters. In other words, it is an expression that you would use in the **Filter Expression** field of the new event filter configuration window to create one filter incorporating all your filters.

Command execution manager service

The command execution manager is the service responsible for operating the command and response feature.

The service queues `JRExecAction` function calls to run external commands. The command execution manager only allows you to specify whether to print the service log to a file. There are no other configuration properties.

Command line manager service

Use the command-line manager service to access the Impact Server from the command line to configure services parameters as well as start and stop services.

When you configure this service, you specify the port to which you connect when you use the command line. You can also specify whether you want the service to start automatically when the Impact Server starts. The command-line manager is the service that manages the CLI. You can configure the port where the command-line service runs, and the startup and logging options for the service.

Command line manager service configuration window

Use this information to configure the command line manager service.

Table 59. Command Line Manager Service Configuration window

Window element	Description
Port	Select a port number where you want to run the service from the list or type the number. You telnet to this port when you use the CLI. The default is 2000.
Startup: Automatically when server starts	Select to automatically start the service when the server starts. You can also start and stop the service from the GUI.
Service log: Write to file	Select to write log information to a file.

Database event listener service

The database event listener service monitors an Oracle event source for new, updated, and deleted events.

This service works only with Oracle databases. When the service receives the data, it evaluates the event against filters and policies specified for the service and sends the event to the matching policies. The service listens asynchronously for events generated by an Oracle database server and then runs one or more policies in response.

You configure the service using the GUI. The configuration properties allow you to specify one or more policies that are to be run when the listener receives incoming events from the database server.

Database event listener service configuration window

You configure the database event listener service by setting events to trigger policies when they match a filter.

Table 60. Database Event Listener service configuration window

Window element	Description
Event Matching	

Table 60. Database Event Listener service configuration window (continued)

Window element	Description
Test events with all filters	Click this icon if, when an event matches more than one filter, you want to trigger all policies that match the filtering criteria.
Stop testing after first match	Click this icon if you want to trigger only the first matching policy. You can choose to test events with all filters and run any matching policies or to stop testing after the first matching policy.
New Mapping: New	Click this icon to create an event filter.
Analyze Event Mapping Table	Click this icon to view any conflicts with filter mappings that you have set for this service.
Startup: Automatically when server starts	Select to automatically start the service when the server starts. You can also start and stop the service from the GUI.
Service log: Write to file	Select to write log information to a file.

E-mail sender service

The e-mail sender is a service that sends e-mail through an external SMTP service.

The service lets you configure the local e-mail address information so that you can send e-mail notification to users and to other installations of Netcool/Impact. To configure the service, you provide the address for the local host and the originating e-mail address.

E-mail sender service configuration window

Use this information to configure the e-mail sender service.

Table 61. EmailSender Service Configuration window

Window element	Description
SMTP Host	Type the host name. Default is localhost.
From Address	Type the From address. Default value is Impact. Example of a valid address: root@flaco.tivoli.
Service log: Write to file	Select to write log information to a file.

Event processor service

The event processor manages the incoming event queue and is responsible for sending queued events to the policy engine for processing.

The event processor service sends events fetched from readers and listener to the policies. The service is responsible for managing events coming from the following event sources:

- OMNIbus event reader
- OMNIbus event listener
- Database event reader

- Database event listener
- JMS message listener
- WSNNotification listener

The event processor manages the incoming events queue and is responsible for sending queued events to the policy engine for processing.

The event processor is typically configured to start automatically when the Impact Server starts. On start-up, it runs with the minimum number of threads. It measures the performance on startup, increases the thread count, and compares the performance with the new thread configuration with the default configuration of minimum threads it started with. If there is an improvement in throughput, it runs with the new configuration and measures the performance again, until one of two events occurs:

- It reaches the limit set by the maximum number of threads
- It reaches a saturation point where increasing the number of threads further does not improve performance

Important: In a clustered environment, changes made to the event processor service using the GUI do not automatically propagate out from the primary Impact Server to the other servers in the cluster. To change configuration for all Impact Servers, log on to each server individually.

Event processor service configuration window

Use this information to configure the event processor service.

When you configure the maximum number of threads for the event processor service, for optimal performance the number of processing threads should be greater than, or equal to the size of the connection pool of the SQL data sources used in the policies being triggered. For information about viewing existing thread and connection pool information, see the information in the *Administration Guide* in the section *Command-Line tools, Event Processor commands*. See the **Select PoolConfig from Service where Name='EventProcessor'**;

Important: In a clustered environment, the event processor configuration is not replicated between servers. You must run the **Select PoolConfig from Service where Name='EventProcessor'**; command on the primary and the secondary servers.

Use the same considerations when configuring the maximum threads on a secondary server. The secondary server uses its own connection pool which is independent of the size of the connection pool in primary server. For example, a DB2 data source has connection pool size of 30. The DB2 data source is replicated between primary and secondary servers. There could potentially be $30+30 = 60$ connections made by Impact primary and secondary servers to the DB2 database. For optimal performance, the maximum number of threads with this setup of connection pool = 30 (should be at least 30 in each server of the cluster). The event processor configuration is not replicated between servers, so it must be set up manually in the secondary using CLI.

Table 62. Event processor window

Window element	Description
Minimum Number of Threads	Set the minimum number of processing threads that can run policies at one time.

Table 62. Event processor window (continued)

Window element	Description
Maximum Number of Threads	Set the maximum number of threads that can run policies at one time.
Processing Throughput: Maximize	If you set this property, the event processor tries to get the maximum performance out of the threads. This can result in high CPU usage. When you leave this field cleared, it runs conservatively at around 80% of peak performance.
Tuning configuration: Maintain on Restart	If you set this option, each time the event processor is started, it uses the same number of threads it had adjusted to in the earlier run. This feature is useful in cases where the environment where Netcool/Impact runs has not changed much from the previous run. The event processor can start with the maximum throughput immediately, rather than engaging in repeated tuning to reach the maximum.
Clear Queue: Clear	Click this icon to enable the event processor to delete unprocessed events that it has fetched from one or more event sources.
Service log: Write to file	Select to write log information to a file.

Hibernating policy activator service

The hibernating policy activator service monitors hibernating policies and awakens them at specified intervals.

You use the hibernating policy activator with X events in Y time solutions and similar solutions that require the use of hibernating policies. When you configure this service, you specify how often the service reactivates hibernating policies waiting to be activated. It can be a specific period or absolute time that you have defined.

Hibernating policy activator configuration

In the hibernation policy activator you can configure the wakeup interval, and the start up and logging options.

Hibernating policy activator configuration window

Use this information to configure the hibernating policy activator.

Table 63. Hibernating Policy Activator service configuration window

Window element	Description
Polling Interval	Select a polling time interval (in seconds) to establish how often you want the service to check hibernating policies to see whether they are due to be woken up. The default value is 3 seconds.
Process wakes up immediately	Select to run the policy immediately after wake-up. The wakeup interval is the interval in seconds at which the hibernating policy activator checks hibernating policies in the internal data repository to see if they are ready to be woken.

Table 63. Hibernating Policy Activator service configuration window (continued)

Window element	Description
Startup: Automatically when server starts	Select to automatically start the service when the server starts. You can also start and stop the service from the GUI.
Service log: Write to file	Select to write log information to a file.
Clear All Hibernations: Clear	Should it become necessary, click to clear all hibernating policies from the Impact Server.

Jabber service

The jabber service acts as a Jabber client and is responsible for logging on to the external Instant Messaging services and sending instant messages.

Jabber is a set of protocols and technologies that provides the means for two software entities to exchange streaming data over a network.

You use this service to use instant messaging systems to notify administrators, operators, and other users when certain events occur in your environment. Messages are sent during the execution of a policy when a call to the `SendInstantMessage` function is encountered. When the system processes the call, it passes the message text and the recipient to the jabber service, where it is routed to the specified recipient.

To send and receive instant messages, first set up the jabber service and then the jabber reader service.

Note: For this service to work, create a policy that uses instant messaging.

Adding resources to the Jabber ID

The Jabber service architecture allows a single user to log on to a jabber server (both public and private) multiple times simultaneously.

Resources that are appended to the jabber ID are used to enable the jabber components to distinguish between multiple connections. Each resource is unique for the user. Since you can configure only one jabber service per Impact Server, append only one user name and one resource per server.

The following example shows three separate IM addresses set up for one user. The user name is hamlet and the three resource names are castle, gate, and courtyard.

```
hamlet@jabber.denmark.org/castle  
hamlet@jabber.denmark.org/gate  
hamlet@jabber.denmark.org/courtyard
```

In a policy that uses the `InstantMessage` function, you could pass either:

hamlet@jabber.denmark.org in the **TO** field,

or add a resource name:

hamlet@jabber.denmark.org/castle,

Configuring Jabber service

Use this procedure to configure the Jabber service.

Procedure

1. Select JabberService to open the service configuration window and configure the jabber account.

In the configuration window, the **Jabber** tab is displayed by default. For information about configuration options, see “Jabber service configuration window - general settings.”

2. The four additional tabs enable you to set up your required transport account information.

For more information about transport accounts and their configuration, see “Jabber transport accounts” on page 131.

Jabber service configuration window - general settings

Use this information to configure the general settings for the Jabber service.

Table 64. JabberService Configuration window

Window element	Description
Jabber Server	Type the name of a private server (host name or IP address) or select a public server host name by clicking the link located under this field. When configuring the jabber service, you can use either a private or public jabber server. If you want to use a public server, enter the host name. For a private server enter either the host name or the IP address.
Jabber Port	Type the port number for the Jabber server. The default is 5222.
Username	Type a unique user name. Make the user name unique for use within Netcool/Impact only. You can use the same user name across Impact Servers. See “Adding resources to the Jabber ID” on page 129.
Password	Type a unique password for the impact user. Jabber service attempts to create an account at the time you configure the service if it does not exist. If you have not already registered the user names and password, you can enter them in the Username and Password fields. In most cases, they are registered automatically.
Resource	If you are using more than one user name for the same user on separate Impact Server, or for a separate Jabber client, type the resource names. The default resource name is Impact. See “Adding resources to the Jabber ID” on page 129.
Nickname	Type a nickname for this user. The default is impact.
SSL: Use an encrypted connection	Select to use an encrypted connection.
Startup: Automatically when server starts	Select to automatically start the service when the server starts. You can also start and stop the service from the GUI.
Service log: Write to file	Select to write log information to a file.

Jabber transport accounts

The Jabber service provides the ability to use transports.

Transports are extensions to the jabber service that allow connections to the following supported instant messaging services:

- AIM
- MSN
- Yahoo
- ICQ

You cannot use a transport without a configured jabber service. After you configure the jabber server information, you can set up transports if you need them.

To use a transport, set up an account for its use before you add it to the jabber service. Although jabber service attempts to create a jabber account at the time you configure the service, it does **NOT** attempt to create a transport account if it does not exist. You can go to the transport Web site to create the account, or you can download a client from the service provider and use the client to create the account.

The jabber service requires a dedicated service for each transport. For example, if you already have an AIM account with your personal sign-in name that you use to communicate with the standard AIM client, do not use that account with the jabber service.

Note: Before you use the jabber service, sign off your personal accounts everywhere they are active. The jabber service might attempt to sign you off from your client, but if it does fail in the attempt, you are not be able to sign on to the jabber service transport.

Do not attempt to use the same transport account with more than one running session of the jabber service. For example, if you have two Impact Servers running the Jabber service and both use the AIM transport, create a separate transport account for each running session. The only time separate accounts are **NOT** required is when the jabber services are running at separate times or when you are running different transports.

Jabber service configuration window - AIM transport settings

Use this information to configure the AIM transport account for the Jabber service.

Table 65. Jabber Service AIM window

Window element	Description
Service ID	Change the Service ID for the AIM service if the system-selected name is not correct. Attention: A Service ID is selected based on your entry in the Jabber Server field in the Jabber tab. Be sure to verify that the entry is correct so that your jabber service is usable.
Screen Name	Type the AIM screen name.
Password	Type the AIM password.

Jabber service configuration window - MSN transport settings

Use this information to configure the MSN transport account for the Jabber service.

Table 66. Jabber Service MSN window

Window element	Description
Service ID	Change the Service ID for the MSN service, if the system-selected ID is not correct. Attention: A Service ID is selected based on your entry in the Jabber Server field in the Jabber tab. Be sure to verify that the entry is correct so that your jabber service is usable.
Username	Type the Hotmail, Passport or MSN screen name. It must be in the format <i>username@hotmail.com</i> , <i>username@passport.com</i> or <i>username@msn.com</i> , depending on where you created the account.
Password	Type the Hotmail, Passport, or MSN password.

Jabber service configuration window - Yahoo transport settings

Use this information to configure the Yahoo transport account for the Jabber service.

Table 67. Jabber Service Yahoo window

Window element	Description
Service ID	Change the Service ID for the Yahoo service if the system-selected name is not correct. Attention: A Service ID is selected based on your entry in the Jabber Server field in the Jabber tab. Be sure to verify that the entry is correct so that your jabber Service is usable.
Username	Type the Yahoo screen name.
Password	Type the Yahoo password.

Jabber service configuration window - ICQ transport settings

Use this information to configure the ICQ transport account for the Jabber service.

Table 68. Jabber Service ICQ window

Window element	Description
Service ID	Change the Service ID for the ICQ service if the system-selected name is not correct. Attention: A Service ID is selected based on your entry in the Jabber Server field in the Jabber tab. Be sure to verify that the entry is correct so that your jabber service is usable.
UIN	Type the ICQ user identity number.
Password	Type the ICQ password.

Policy logger service

The policy logger service is responsible for managing the policy log.

The log is a text stream used to record messages generated during the runtime of a policy. The log contains both Netcool/Impact system messages and messages that you create when you write a policy. The policy logger service specifies an error-handling policy to activate when an error occurs during the execution of a policy. It also specifies the logging levels for debugging policies and which items must be logged. When you configure this service, you select a policy to handle the errors as they occur.

Policy logger configuration

You can configure the following properties of the policy logger.

- Error handling policy
- Highest log level
- Logging of SQL statements
- Logging of pre-execution function parameters
- Logging of post-execution function parameters
- Policy profiling
- Logging and reporting options

Policy logger service configuration window

Use this information to configure the policy logger service.

Table 69. Policy Logger Service configuration window

Window element	Description
Error-handling Policy	The error handling policy is the policy that is run by default when an error is not handled by an error handler within the policy where the error occurred.
Highest Log Level	<p>You can specify a log level for messages that you print to the policy log from within a policy using the Log function.</p> <p>When a <code>log()</code> statement in a policy is processed, the specified log level is evaluated against the number that you select for this field. If the level specified in this field is greater than or equal to the level specified in the policy <code>log()</code> statement, the message is recorded in the policy log.</p>
Log what	<p>Select what you want to appear in the log:</p> <ul style="list-style-type: none"> • All SQL statements. Select to print all the contents of all SQL statements made in calls to SQL database data sources. Logging SQL statements can help you debug a policy that uses external data sources. • Pre-execution Action Module Parameters. Select to print the values of all the parameters passed to a built-in action function before the function is called in a policy. These parameters include the values of built-in variables such as <code>DataItems</code> and <code>DataItem</code>. • Post-execution Action Module Parameters • All Action Module Parameters

Table 69. Policy Logger Service configuration window (continued)

Window element	Description
Policy Profiling: Enable	<p>Select to enable policy profiling. Policy profiling calculates the total time that it takes to run a policy and prints this time to the policy log</p> <p>You can use this feature to see how long it takes to process variable assignments and functions. You can also see how long it takes to process an entire function and the entire policy.</p>
Service log: Write to file	<p>Select to write log information to a file.</p> <p>You can also enable the collecting of report information through the service.</p>
Append Thread Name to Log File Name	Select this option to name the log file by appending the name of the thread to the default log file name.
Append Policy Name to Log File Name	Select this option to name the log file by appending the name of the policy to the default log file name.
Collect Reports: Enable	<p>Select to enable data collection for the Policy Reports.</p> <p>If you choose to enable the Collect Reports option, reporting related logs are written to the policy logger file only when the log level is set to 3.</p> <p>To see reporting related logs for a less detailed logging level for example, log level 1, the NCHOME/impact/etc/<servername>_policylogger.props file can be customized by completing the following steps:</p> <ol style="list-style-type: none"> 1. Add impact.policylogger.reportloglevel=1 to the NCHOME/impact/etc/<servername>_policylogger.props property. 2. Restart the Impact Server to implement the change.

Policy log files

You can use policy log files to provide a record of actions performed during the execution of a policy.

Multiple log files can be created as follows:

- 1 log file for each policy
- 1 log file for each thread in the event processor
- 1 log file for each policy for each thread

By default, a single policy log file is created.

Each log file is named by appending the name of the policy or the name of the thread to the default log file name. For example:

- If you were to run a policy named POLICY_01 and you selected to create log files on a per policy basis, the resulting log file would be named:

`servername_Policy_01_policylogger.log`

- If you selected to create log files on a per-thread basis, a possible log file name might be:

`servername_Policy_02HttpProcessor [5104] [2]_policylogger.log`

Where

HttpProcessor[5104] [2] is the name of the event processor thread where the policy is running on a Red Hat Linux system.

- If you selected to create log files on a per policy per thread basis, the log file name might be:

```
servername_Policy_02HttpProcessor [5104] [2]_policylogger.log
```

Enabling multiple policy log files

Use this procedure to enable multiple policy log files.

Procedure

1. In the PolicyLogger Service Configuration window, click the **Service Log: Write to File** option.
2. Select either the **Append Thread Name to Log File Name** or the **Append Policy Name to Log** file option, or both.

ITNM event listener service

The ITNM event listener service listens for events sent from ITNM.

After you install the ITNM DSA, you can optionally set up a ITNM event listener service. You only need to set up the listener service if you want to listen for events asynchronously from ITNM. For more information about ITNM TN or ITNM IP, see the guides for those products.

Configuring ITNM event listener service

Use this procedure to configure the ITNM listener service.

Procedure

1. Enter the required information in the service configuration window and save the configuration.
For information about the configuration options, see “ITNM event listener service configuration window.”
2. Before you start the event listener service, first stop all ITNM and rvd processes and enter the command:

```
$ITNM_HOME/bin/rvd -flavor  
116
```
3. Restart ITNM.
4. Make sure that the ITNM event listener service is started so that you can receive events from ITNM. (You have the option to have it start automatically when Netcool/Impact starts.)

ITNM event listener service configuration window

Use this information to configure the ITNM event listener service.

Table 70. ITNM Event Listener service configuration

Table element	Description
Listener Filter	Leave this field blank.
Policy to Execute	Select the policy to run when an event is received from the ITNM application. You can use the ITNMSampleListenerPolicy that was installed when you installed Netcool/Impact to help you understand the event listener functionality.
Name Service Host	

Table 70. ITNM Event Listener service configuration (continued)

Table element	Description
Name Service Port	
Name Service Context	
Name Service Object Name	
Direct Mode Class Name	Set this to: con.micromuse.dsa.precisiondsa.PrecisionEventFeedSource Note: Copy this class name exactly as it is written here, with no extra spaces.
Direct Mode Source Name	Type a unique name that identifies the data source, for example, ITNMServer.
Startup: Automatically when server starts	Select to automatically start the service when the server starts. You can also start and stop the service from the GUI.

Self monitoring service

The self monitoring service is used to send messages about the internal state of Impact Server to an ObjectServer.

If the event readers and listeners are running (OMNIBusEventReader, OMNIBusEventListener, DatabaseEventReader, DatabaseEventListener, JMSMessageListener), the self monitoring send status events regarding its event queue.

The self monitoring service provides the following types of monitoring:

Cluster monitoring

When Impact is running as a cluster, it provides information as to which server is the current primary and the current secondary. Impact also sends updates when the primary has gone down and one of the secondary servers assumes the primary role.

Data source monitoring

Provides information about the active data sources used by Netcool/Impact. It also gives information when the connection to the primary or backup host of the data source fails.

Memory and queue monitoring

Checks the heap utilization of the virtual machine used by Netcool/Impact and also the available system memory of the system where Impact is running at selected intervals and sends that information to ObjectServer as an event. Events warn users by severity level of conditions such as maximum heap utilization or insufficient system memory.

At intervals, Netcool/Impact checks to see whether it is approaching the maximum amount of available memory or whether the queue size is growing at a rate that exceeds a certain number. If so, the severity of the condition is determined and a corresponding event is sent to the ObjectServer. You can configure self monitoring to deduplicate the events, or send a new event to the ObjectServer every time a low memory or growing queue size condition occurs.

Self monitoring service configuration window

Use this information to configure the self monitoring service.

Table 71. Self Monitoring Service window

Window element	Description
ObjectServer Data Source	Select the ObjectServer that you want to use to send events.
Memory Status: Enable	Select to send status events regarding queue conditions of the event readers that are currently running.
Interval	Select or type how often the service must send status events to the ObjectServer.
Deduplication	Deduplication is enabled by default. See the Netcool/OMNIBus library for information about deduplication of events.
Queue Status: Enable	Select to enable the service to send events about the status of the event readers and listeners currently running.
Interval	Select or type (in seconds) how often the service must send queue status events.
Deduplication	Deduplication is enabled by default. See the Netcool/OMNIBus library for information about deduplication of events.
Cluster Status: Enable	Select to enable the service to send events about the status of the cluster to which it belongs. It sends events when: <ul style="list-style-type: none">• A Impact Server is started and joins the cluster• A server is stopped and removed from the cluster• A primary server is down and a secondary server becomes the new primary
Data Source Status: Enable	Select to enable the service to send the status when certain conditions occur with a data source. For example, the service sends a status message when a user tests a connection to a data source or when a connection cannot be established. Remember: Restart the service to apply the change.
Service Status: Enable	To enable service monitoring, select this check box and start the self-monitoring service. The self-monitoring service sends service status events to the ObjectServer. Remember: Restart the service to apply the change.
Startup: Automatically when server starts	Select to automatically start the service when the server starts. You can also start and stop the service from the GUI.
Service log: Write to file	Select to write log information to a file.

Database event reader service

The database event reader is a service that polls supported, external SQL data sources at regular intervals to get business events in real time.

It retrieves rows from a table, then converts the rows to event format, and passes them to Netcool/Impact for processing. The data source can be any of the supported SQL data sources. Conceptually, it is similar to the OMNIBus Event Reader, which polls the ObjectServer to get network fault events.

Configuring the database event reader service

Use this procedure to configure the database event reader service.

Procedure

1. Select the project for which you want to create the service.
2. From the **Service Type** list, select DatabaseEvent Reader to open the service configuration window.

The DatabaseEventReader Configuration window has two tabs, **General Settings** and **Event Mapping**.

3. Enter the required information in the General settings tab of the configuration window.

For information about general settings options, see “Database event reader configuration window - general settings.”

4. Enter the required information in the Event Mapping tab of the configuration window.

For information about general settings options, see “Event mapping” on page 122.

Database event reader configuration window - general settings

Use this information to configure the general settings of the database event reader.

Table 72. Database event reader configuration window - General Settings tab

Window element	Description
Service name	Type a unique name to identify the service.
Data Source	Select an external data source from the list. The data source must have a field that is guaranteed to be incremented every time a new record is added to avoid rereading the entire table every time the data source is accessed. If you want to use the GetUpdates function in a policy for this data source, the table also must have a time stamp field that is automatically populated when an insert or update occurs.
Data Type	After you select a data source, the system populates the data type field with a list of data types created in Netcool/Impact corresponding to that particular data source. Select a data type from the list.
Polling Interval	Select or enter a polling time interval to establish how often you want the service to poll the events in the event source. The polling time selections are in milliseconds and the default value is 3000 milliseconds
Restrict fields	Click Fields to access a selection list with all the fields that are available from the selected data source. You can reduce the size of the query by selecting only the fields that you need to access in your policy.
Startup: Automatically when server starts	Select to automatically start the service when the server starts. You can also start and stop the service from the GUI.
Service log: Write to file	Select to write log information to a file.

Table 72. Database event reader configuration window - General Settings tab (continued)

Window element	Description
Clear State	<p>When you click Clear, the internally stored value for the Key field and Timestamp field are reset to 0. This causes the event reader to retrieve all events in the data source at startup and place them in the event queue for processing.</p> <p>If the event reader is configured to get updated events, it requires the Timestamp field along with the Key field. The Timestamp field must point to a column in the table which is automatically populated with a timestamp when an insert or update occurs. The Key field must point to a column which uniquely identifies a row (it does not have to be an automatically incremented field).</p> <p>However, when the Actions Get updated events check box in the Event Mapping tab is not selected, you do not have to configure the Timestamp field. The Key field <i>MUST</i> in this case be an automatically incremented numeric field.</p> <p>You can only use Clear State to clear the event reader state when the service is stopped. Clicking Clear while the service is running does not change the state of the event reader.</p>
Clear Queue	Click Clear to enable the database event reader to delete unprocessed events that it has fetched from an SQL data source.

Database event reader configuration window - event mapping

Use this information to configure event mapping for the database event reader.

Table 73. DatabaseEvent Reader window - Event Mapping tab

Window element	Description
Event Matching:	
Test events with all filters	If an event matches more than one filter, trigger all policies that match the filtering criteria.
Stop testing after first match	Or select to trigger only the first matching policy.
Actions: Get updated events	Select to receive events that have been updated (all new events are automatically sent).
Time Stamp Field	When the event reader is configured to get updated events, it requires the TimeStamp field along with the Key field. The TimeStamp field must point to a column in the table that is automatically populated with a timestamp when an insert or update occurs. The Key field points to a column which uniquely identifies a row (it does not have to be an automatically incremented field). However, when Actions Get updated events is not selected, you do not have to configure the TimeStamp field, but the Key field must in this case be an automatically incremented numeric field.
Key Field	See Time Stamp Field.
New Mapping: New	Click to add a new filter.
Analyze Event Mapping Table	Click this icon to display any conflicts with filter mappings that you have set for this service.

Email reader service

The email reader service reads incoming email, and runs policies based on the contents of the email.

The email reader service polls a specified POP or IMAP host for email messages. The service reads email from a mailbox at intervals that you define when the service is created. The service is commonly used in escalation and notification policies to look for responses to email notifications that are sent out by Netcool/Impact.

If the number of emails waiting to be read from the email reader service is more than 25, the timeout value increases automatically. When the number of emails waiting to be read returns to less than 25. The timeout value is reset to the default value or the value specified in the service property file.

You can use this default service instead of creating your own, or in addition to creating your own.

The email reader tries to process the body of the email as name and value pairs. This means that the service is looking for the body of the email to be in IPL syntax. To change the default behavior to ignore the body of the email, insert the following property `impact.<EmailReaderName>.ignorebody=true` into the file `$IMPACT_HOME/etc/<ImpactServerName>_<EmailReaderName>.props` and restart the service.

Email reader service configuration window

Use this information to configure the email reader service.

Table 74. Create New email Reader Service Configuration window

Window element	Description
Service name	Type a unique name to identify the service.
Host:	Type the mail host name.
Protocol:	Select one of the following options from the drop-down menu: POP3 or IMAP.
Port:	Select the port to connect to the mail server. The default POP3 port is 110. The default IMAP port is 143.
Log in As:	Type a login name. The default value is the one you use to log on to Netcool/Impact.
Password:	Type your password. The letters you type are replaced with asterisks.
Polling Interval:	Select how often (in seconds) the service polls the POP or IMAP host for new email messages.
Policy Name:	Select a policy to run for this event.
Startup: Automatically when server starts	Select to automatically start the service when the server starts. You can also start and stop the service from the GUI.
Service log: Write to file	Select to write log information to a file.

Event listener service

Event listeners monitor non-ObjectServer event source events.

Event listener services typically work with DSAs that allow bi-directional communication with a data source. If you need to configure the Event Listener service to work your DSA, for a configuration procedure refer to the documentation for that DSA.

Event listener service configuration window

Use this information to configure the event listener service.

Table 75. EventListener service configuration window

Table element	Description
Service name	Type a unique name to identify the service.
Listener Filter	Leave this field blank.
Policy to Execute	Select the policy to run when an event is received from the database server.
Name Service Host	Type in the name of the service host.
Name Service Port	Provide the port over which the name service host is accessed.
Name Service Context	Type in the name service context.
Name Service Object Name	Type in the name of the service object.
Direct Mode Class Name	Type in the direct mode class name.
Direct Mode Source Name	Provide a unique name that identifies the data source.
Startup: Automatically when server starts	Select to automatically start the service when the server starts. You can also start and stop the service from the GUI.

JMS message listener

The JMS message listener service runs a policy in response to incoming messages sent by external JMS message providers.

The message provider can be any other system or application that can send JMS messages. Each JMS message listener listens to a single JMS topic or queue. There is one default JMS message listener named `JMSMessageListener`. You can create as many listener services as you need, each of which listens to a different topic or queue.

A JMS message listener is only required when you want Netcool/Impact to listen passively for incoming messages that originate with JMS message producers in your environment. You can actively send and retrieve messages from within a policy without using a JMS message listener.

JMS message listener service configuration properties

Use this information to configure the JMS listener service.

Table 76. JMSMessageListener Service configuration window

Window element	Description
Service name	Type a unique name to identify the service.
Policy To Execute	Select the policy that you created to run in response to incoming messages from the JMS service.

Table 76. *JMSMessageListener Service configuration window (continued)*

Window element	Description
JMS Data Source	JMS data source to use with the service. You need an existing and valid JMS data source for the JMS Message Listener service to establish a connection with the JMS implementation and to receive messages. For more information about creating JMS data sources, see “JMS data source configuration properties” on page 58.
Message Selector	The message selector is a filter string that defines which messages cause Netcool/Impact to run the policy specified in the service configuration. This string must be specified using the JMS message selector syntax. Message selector strings are similar in syntax to the contents of an SQL WHERE clause, where message properties replace the field names that you might use in an SQL statement. The content of the message selector is dependent on the types and content of messages that you anticipate receiving with the JMS message listener. For more information about message selectors, see the JMS specification or the documentation distributed with your JMS implementation. The message selector is an optional property.
Durable Subscription: Enable	You can configure the JMS message listener service to use durable subscriptions for topics which allow the service to receive messages when it does not have an active connection to the JMS implementation. A durable subscription can have only one active subscriber at a time. Only a JMS topic can have durable subscriptions.
Client ID	Client ID for durable subscription. It defines the client identifier value for the connection. It must be unique in the JMS Implementation.
Subscription Name	Subscription Name for durable subscription. Uniquely identifies the subscription made from the JMS message listener to the JMS Implementation. If this property is not set, the name of JMS DSA listener service itself is used as its durable subscription name, which is JMSMessageListener by default.
Clear Queue: Clear	Clear the message waiting in the JMSMessageListener queue that has not yet been picked by the EventProcessor service. It is recommended not to do this while the Service is running.
Service: Automatically when server starts	Select to automatically start the service when the server starts. You can also start and stop the service from the GUI.
Service log: Write to file	Select to write log information to a file.

Jabber reader service

The jabber reader service is the instant message listener service.

The jabber reader service listens for new messages from instant messaging accounts. When it receives a message, it creates a new EventContainer and populates it with the contents of the message. It then starts the policy specified in its configuration settings and passes the EventContainer to it. The policy is then processed.

This is a default service that you can use instead of creating your own, or in addition to creating your own. You can create multiple readers to associate with the single jabber service. The jabber readers need the jabber service to be running. This interdependency between the jabber service and the jabber reader is automated by selecting the **Startup Automatically when JabberService starts** check box when configuring the reader.

Jabber reader service configuration window

Use this information to configure the Jabber reader service.

Table 77. New JabberReader Service configuration window

Window element	Description
Service name	Type a unique name to identify the service.
Allow only specified users to activate policy	If you want to create a list of specific users that are <i>permitted</i> to activate the policy selected in the Policy field, select this button.
Prevent specified users from activating policy	If you want to create a list of users that are <i>blocked</i> from activating the policy selected in the Policy field, select this button.
Add User to list	Click the Add User to List button to open a window to create a list of users that are either to be blocked or allowed, depending on your radio button selection.
User Address	Enter the address of the user you want to add to the list and click Apply . Repeat for every user that you want to add to the list.
Tell blocked users that they are blocked	If you have created a blocked or allowed list, select this option to notify the non-allowed users. This notification is sent when the jabber reader receives a message from either type of prevented user.
Policy	The policy selected here is activated by the jabber reader whenever an instant message event occurs. You must create a policy to use with the jabber reader service.
Startup: Automatically when server starts	Select to automatically start the service when the server starts. You can also start and stop the service from the GUI.
Service log: Write to file	Select to write log information to a file.

OMNibus event listener service

The OMNibus event listener service is used to integrate with Netcool/OMNibus and receive immediate notifications of fast track events.

The OMNIBus event listener is used to get fast track notifications from OMNIBus using the Accelerated Event Notification feature of OMNIBus. It receives notifications through the Insert, Delete, Update, or Control (IDUC) channel. To set up the Netcool/OMNIBus event listener, you must set its configuration properties using the GUI. The configuration properties allow you to specify one or more policies that are to be run when the OMNIBus event listener receives incoming events from Netcool/OMNIBus. For more information about OMNIBus triggers and accelerated event notification, see the *OMNIBus Administration Guide*.

Important:

- The OMNIBus event listener service works only with OMNIBus 7.3 to monitor ObjectServer events.
- If the Impact Server and OMNIBus server are located in different network domains, for the OMNIBus event listener service to work correctly, you must set the **Iduc.ListeningHostname** property in the OMNIBus server. This property must contain the IP address or fully qualified hostname of the the OMNIBus server. For more information about this property, refer to the OMNIBus documentation.

Setting up the OMNIBus event listener service

Use this procedure to create the OMNIBus event listener service.

Procedure

1. In the Tivoli Integrated Portal, in the navigation tree, click **System Configuration > Event Automation > Services**, to open the **Services** tab.
2. If required, select a cluster from the **Cluster** list.
3. Click the **Create New Service** icon in the toolbar and select **OMNIBusEventListener** to open the configuration window.
4. Enter the required information in the configuration window.
5. Click the **Save** icon in the toolbar to create the service.
6. Start the service to establish a connection to the ObjectServer and subscribe to the IDUC channel to get notifications for inserts, updates, and deletes.

OMNIBus event listener service configuration window

Use this information to configure the OMNIBus event listener service.

Table 78. OMNIBusEventListener Service configuration window

Table Element	Description
Service name	Type a unique name to identify the service.
Data Source	Select an OMNIBus version 7.2 ObjectServer data source. Make sure your data source has a configured, and valid connection to an ObjectServer. You can use the default ObjectServer data source that is created during the installation, defaultobjectserver.
Startup: Automatically when server starts	Select to automatically start the service when the server starts. You can also start and stop the service from the GUI.
Service log: Write to file	Select to write log information to a file.
Event Matching	
Test events with all filters	Select this option, to test events with all filters and run any matching policies. When an event matches more than one filter, all policies that match the filtering criteria will be triggered.

Table 78. OMNIBusEventListener Service configuration window (continued)

Table Element	Description
Stop testing after first match	Select this option, if you want to trigger only the first matching policy.
New Mapping: New	Click to create a new event filter. This filter specifies the type of event that maps to the policy.
Analyze Event Mapping Table	Click to check the validity of event filters.

OMNIBus event reader service

OMNIBus event readers are services that monitor a Netcool/OMNIBus ObjectServer event source for new, updated, and deleted alerts and then runs policies when the alert information matches filter conditions that you define.

The event reader service uses the host and port information of a specified ObjectServer data source so that it can connect to an Objectserver to poll for new and updated events and store them in a queue. The event processor service requests events from the event reader. When an event reader discovers new, updated, or deleted alerts in the ObjectServer, it retrieves the alert and sends it to an event queue. Here, the event waits to be handled by the event processor.

You configure this service by defining a number of restriction filters that match the incoming events, and passing the matching events to the appropriate policies. The service can contain multiple restriction filters, each one triggering a different policy from the same event stream, or it can trigger a single policy.

You can configure an event reader service to chain multiple policies together to be run sequentially when triggered by an event from the event reader.

Important: Before you create an OMNIBus event reader service, you must have a valid ObjectServer data source to which the event reader will connect to poll for new and updated events.

OMNIBus event reader configuration

You can configure the following properties of an OMNIBus event reader.

- Event reader name
- ObjectServer event source you want the event reader to monitor
- Interval at which you want the event reader to poll the ObjectServer
- Event fields you want to retrieve from the ObjectServer
- Event mapping
- Event locking
- Order in which the event reader retrieves events from the ObjectServer
- Start up, service log, and reporting options

OMNIBus event reader service General Settings tab

Use this information to configure the general settings of the OMNIBus event reader service.

Table 79. EventReader service - general settings tab

Table Element	Description
Service name	Type a unique name to identify the service.
Data Source	Select an OMNIBusObjectServer data source. The ObjectServer data source represents the instance of the Netcool/OMNIBus ObjectServer that you want to monitor using this service. You can use the default ObjectServer data source that is created during the installation, defaultobjectserver.
Polling Interval	<p>The polling interval is the interval in milliseconds at which the event reader polls the ObjectServer for new or updated events.</p> <p>Select or type how often you want the service to poll the events in the event source. If you leave this field empty, the event reader polls the ObjectServer every 3 seconds (3000 milliseconds).</p>
Restrict Fields: Fields	<p>You can complete this step when you have saved the OMNIBusEventReader service. You can specify which event fields you want to retrieve from the ObjectServer. By default, all fields are retrieved in the alerts. To improve OMNIBus event reader performance and reduce the performance impact on the ObjectServer, configure the event reader to retrieve only those fields that are used in the corresponding policies.</p> <p>Click the Fields button to access a list of all the fields available from the selected ObjectServer data source.</p> <p>You can reduce the size of the query by selecting only the fields that you need to access in your policy. Click the Optimize List button to implement the changes. The Optimize List button becomes enabled only when the OMNIBusEventReader service has been saved.</p>
Startup: Automatically when server starts	Select to automatically start the service when the server starts. You can also start and stop the service from the GUI.
Service log: Write to file	Select to write log information to a file.
Collect Reports: Enable	Select to enable data collection for the Policy Reports.
Clear State: Clear	<p>When you click the Clear State button, the Serial and StateChange information stored for the event reader is reset to 0. The event reader retrieves all events in the ObjectServer at startup and places them in the event queue for processing. If the event reader is configured to get updated events, it queries the ObjectServer for all events where StateChange >= 0. Otherwise, it queries the ObjectServer for events where Serial > 0.</p> <p>You can use the Clear State button only to clear the event reader state when the service is stopped. Clicking the button while the service is running does not change the state of the event reader.</p>
Clear Queue: Clear	Click to clear unprocessed events.

OMNIBus event reader service Event Mapping tab

In the Event Mapping tab, you set events to trigger policies when they match a filter.

Table 80. Event Mapping tab

Window element	Description
Event Matching	
Test events with all filters	<p>Select this option to test events with all filters and run any matching policies.</p> <p>If an event matches more than one filter, all policies that match the filtering criteria will be triggered.</p>
Stop testing after first match	Select this option to stop testing after the first matching policy, and trigger only the first matching policy.
Actions	
Get updated events	Select to receive updated events as well as new events from the ObjectServer (All new events are automatically sent). See also the description of the Order By field below for more information.
Get status events	Select to receive the status events that the Self Monitoring service inserts into the ObjectServer.
Run policy on deletes	Select if you want the event reader to receive notification when alerts are deleted from the ObjectServer. Then, select the policy you want to run when notification occurs from the Policy list.
Policy	Select a policy to run when events are cleared from the ObjectServer.
Event Locking: enable	<p>Select if you want to use event order locking and type the locking expression in the Expression field.</p> <p>Event locking is a feature that allows a multi-threaded event processor to categorize incoming alerts based on the values of specified alert fields and then to process them within a category one at a time.</p> <p>With event locking enabled, if more than one event exists with a certain lock value, then these events are not processed at the same time. These events are processed in a specific order in the queue.</p> <p>You use event locking in situations where you want to prevent a multi-threaded event processor from attempting to access a single resource from more than one instance of a policy running simultaneously.</p>
Expression	<p>The locking expression consists of one or more alert field names.</p> <p>To lock on a single field, specify the field name, for example: Node</p> <p>To lock more than one field, concatenate them with the + sign, for example: Node+Severity</p> <p>If the value of that field is the same in both events, then one event is locked and the second thread must wait until the first one is finished.</p>
New Mapping	Click to add a new filter.

Table 80. Event Mapping tab (continued)

Window element	Description
Order by	<p>If you want to order incoming events retrieved from the ObjectServer, type the name of an alert field or a comma-separated list of fields. The event reader will sort incoming events in ascending order by the contents of this field.</p> <p>This field or list of fields is identical to the contents of an ORDER BY clause in an SQL statement. If you specify a single field, the event reader sorts incoming events by the specified field value. If you specify multiple fields, the events are grouped by the contents of the first field and then sorted within each group by the contents of the second field, and so on.</p> <p>For example, to sort incoming events by the contents of the Node field, type Node.</p> <p>To sort events first by the contents of the Node field and then by the contents of the Summary field, type Node, Summary.</p> <p>You can also specify that the sort order is ascending or descending using the ASC or DESC key words.</p> <p>For example, to sort incoming events by the contents of the Node field in ascending order, type the following Node ASC.</p> <p>Note that all events retrieved from the ObjectServer are initially sorted by either the Serial or StateChange field before any additional sorting operations are performed. If you select the Get updated events option (see the Actions check box in the Event Mapping section of the window), the events are sorted by the StateChange field. If this option is not specified, incoming events are sorted by the Serial field.</p>
Analyze Event Mapping Table	Click to analyze the filters in the Event Mapping table.

OMNIbus Event Reader event locking examples

The following examples explain how Event Locking works.

The Event Processing service receives events in blocks from the Event Reader service and places them in a queue. These events are picked up as threads sequentially and sent to the respective policies for processing.

Example of event locking on single field

In this example, event locking is enabled with the event locking expression set to Severity and then configured with four threads. With event locking set on Severity, only one event with the same value of Severity can be processed at any instant.

The Event Processor receives from the Event Reader events with the following severities:

3 4 3 5 4 4 2 3 5
F L
F: First Element in the Queue
L: Last Element in the Queue

Since the Event Processor has four threads configured, the first thread receives the first event with Severity=3 from the queue and sends it to a policy for processing. The second thread receives the event with Severity=4 and sends it to a policy for processing. Although two remaining threads are available for processing, the next event Severity=3 cannot be processed because an event with Severity=3 is already being processed (the first event in the queue). Until the processing of the first event is complete, the other threads cannot begin, since they would violate the locking criteria.

If the thread that picked the second event in the queue (with Severity=4) finishes processing before the first event, it waits along with the other two threads until the first event has finished processing. When the thread that picked up the first event in the queue is finished, three threads pick up the third, fourth, and fifth events from the queue, since they have different Severity values (3, 5, 4).

At this point, the remaining thread cannot pick up the next event (sixth in the queue) from the queue because an event with the same Severity level (4) is already processing (fifth in the queue).

Example of event locking on multiple fields

In the example above, locking is on a single field, Severity. You can also lock on more than one field by concatenating them with the plus (+) operator. If you lock, for example, on the Node and Severity fields, you can use one of the following event locking expressions:

Node+Severity

or:

Severity+Node

Event locking on multiple fields works in the same way that locking on a single field except that in this instance, two events with the same combination of fields cannot be processed at the same instant. In other words, if two events have the values for Node as abc and xyz and both have the value for Severity as 5, then they can be processed simultaneously. The only case when the two events cannot be processed together is when the combination of Node and Severity is the same for the events. In other words, if there are two events with the Node as abc and Severity as 5, then they cannot be processed together.

Policy activator service

The policy activator service activates policies at startup or at the intervals you specify for each selected policy.

This is a default service that you can use instead of creating your own, or in addition to creating your own.

Policy activator configuration

In a policy activator you can configure the policy activator name, the activation interval, the policy you want to run at intervals, and the start up and logging options.

Policy activator service configuration window

Use this information to configure the policy activator service.

Table 81. Create New Policy Activator Service configuration window

Window element	Description
Service name	Type a unique name to identify the service.
Activation Interval	Select how often (in seconds) the service must activate the policy. The minimum value is 0; the default value is 10.
Policy	Select the policy you want the policy activator to run.
Startup: Automatically when server starts	Select to automatically start the service when the server starts. You can also start and stop the service from the GUI.
Service log: Write to file	Select to write log information to a file.

Web Services Notification Listener service

Netcool/Impact provides the implementation for the consumer part of the WS-notification standard through the Web Services Notification Listener service.

Through the WS-Notification standard the monitoring environment can generate asynchronous events. These events can be consumed by a consumer. Netcool/Impact provides support for the consumer component of the WS-Notification model. The following producer applications are currently supported by the Impact consumer:

- Apache Muse WSN producer
- Websphere 6.1 WSN producer
- Websphere 7.0 WSN producer

The subscription manager, notification broker, or the producer support are not provided.

Important: The Impact consumer supports simple text, XML, and UTF-8 strings. If an event sent to the consumer contains NLV characters, they must be encoded in UTF-8 format so that the Log Viewer can interpret them correctly as national language version characters.

IBM Websphere Application Server V7.0 must be upgraded to support Websphere 7.0 WSN producer. Install the patches in the following order:

1. WebSphere Application Server V7.0 Fix Pack 5 (7.0.0.5)
2. IF PK93527

These patches must be installed in this order. Do not install them at the same time.

Note: SOAP version 1.1. is not supported. You must use SOAP version 1.2.

Web Services Notification Listener service configuration window

Use this information to configure the Web Services Notification listener service.

Table 82. Web Services Notification Listener Service window fields

Tab	Description
Service name	Type a unique name to identify the service.
Enter Subscription or Producer End point Reference	End point reference where the subscription needs to be sent to. Depending on the environment, this could be your producer or notification broker. Identify the EPR and give the value here. Here is an example of the EPR format: <code>http://<host>:<port>/<WSN Application Context URI></code>
Subscribe for all topics for the End Point Reference	Use this option to select the policy to which all the topics will be associated. Note: You can either subscribe to all topics and select a policy for all topics or you can specify individual topics and a policy for each such a topic. If you select to subscribe to all topics, you cant specify individual topics.
Select Policy associated with all the topics	From the list select a policy that will be associated with all the topics. This option is not active if the Subscribe for all topics for the End Point Reference: option is not selected.
Add a new Topic to Policy Mapping	Use this option to add the topics that you are interested in. After you select this option you are asked to enter a qualified name and select the policy associated with the name.
Startup: Automatically when server starts	Select to automatically start the service when the server starts. You can also start and stop the service from the GUI.
Service log: Write to file	Select to write log information to a file.

Chapter 11. Operator views

An operator view is a custom web-based tool that you use to view events and data in real time and to run policies that are based on that data.

Viewing operator views

Procedure

1. To view the basic and advanced operator views that are currently defined in IBM Tivoli Netcool/Impact log on to the GUI.
2. In the navigation tree, expand **System Configuration > Event Automation**, click **Operator Views** to open the **Operator Views** tab.
3. From the **Cluster** list, select the cluster you want to use.
4. From the **Project** list, select the project you want to use.
5. Double click the operator view to see the details or right click the operator view and click **Edit**.

Results

You can also list all the operator views that are currently defined by opening the following URL in a web browser:

`http://hostname:port/opview:`

Operator views overview

An operator view is a custom Web-based tool that you use to view events and data in real time and to run policies that are based on that data.

The simplest operator views present a basic display of event and business data. More complex operator views can function as individual GUIs that you use to view and interact with event and business data in a wide variety of ways. You can use this kind of GUI to extensively customize an implementation of Netcool/Impact products and other Tivoli Monitoring applications.

Typically, you create operator views to:

- Accept incoming event data from Netcool OMNIBus or another application.
- Run a policy that correlates the event data with business data that is stored in your environment.
- Display the correlated business data to a user.
- Run one or more policies based on the event or business data.
- Start another operator view based on the event or business data.

One common way to use an operator view is to configure it to be started from within the Netcool OMNIBus event list. Netcool/Impact operators can view related business data for an event by right-clicking the event in the event list and viewing the data as displayed in the view. The business data might include service, system, or device information related to the event, or contact information for administrators and customers that are affected by it.

Operator views are not limited to use as Netcool OMNIBus tools. You can use the operator view feature to create a wide variety of tools that display event and business data to users.

Operator view types

Basic and advanced operator views are supported.

- Basic operator views that you use to display data in a preformatted Web page. For more information about basic operator views, see “Basic operator views.”
- Advanced operator views that you use to display data using any HTML formatting that you choose. For more information about advanced operator views, see “Advanced operator views.”

Basic operator views

You use basic operator views to view real-time data in a preformatted web page and to run policies based on that data.

A basic operator view has the following display elements:

Event panel

Displays incoming event information from Netcool OMNIBus or information from another application that can be expressed in name/value pairs.

Actions panel

You use it to run one or more policies from within the operator view.

Information groups panel

Displays sets of data that you define when you create the view, or when you manually edit the operator view policy.

You create basic operator views using the GUI. The GUI automatically creates the corresponding display page and operator view policy.

If you need to customize the appearance of the view or the type of information displayed in the information group panel, you can manually edit the display page using a text editor. You can edit the operator view policy using the GUI.

Advanced operator views

You use advanced operator views to view real-time data in an HTML-formatted Web page and to run policies based on that data.

Unlike basic operator views, which must use the provided preformatted page design, advanced operator views have no restrictions on the appearance of the resulting Web page.

You can use any type of HTML formatting to specify how an advanced operator view is displayed and you can display data in an advanced view in any format that is viewable using a Web browser. You can also further customize advanced operator views using cascading styles sheets (CSS) and browser scripting languages.

For detailed information about how to create and view advanced operator views, see the *Operator View Guide*.

Operator view components

An overview of the operator view components.

Display page

Text file that contains HTML content and special instructions called smart tags that determine what data to display and how to display it.


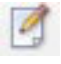


Operator view policy

Policy that contains the logic required to retrieve and manipulate the data displayed in the view.

Operator views panel controls

An overview of the icons and controls used in the operator view panel.

Table 83. Operator views panel controls

Control	Description
	Click this icon to create a basic operator view.
	Select an operator view and use this icon to edit it. Alternatively, you can edit an operator view by right clicking its name and selecting Edit in the menu.
	Click this icon to view the operator view display for the selected operator view. Alternatively, right click an operator view and select View .
	Select an operator view from the list and click this icon to delete it. Alternatively, right click an operator view and select Delete .

Layout options

When you create a basic operator view using the GUI, you can use the layout options and the associated preview feature to specify how different parts of the tool are arranged on the resulting web page.

The following table shows the display panels in a basic operator view:

Table 84. Operator view display panels

Display Panel	Description
Event panel	<p>Displays information, if any, passed from Netcool OMNIBus or another application to the operator view. This information can be fields in a Netcool OMNIBus event, or any other information that can be expressed as a set of name/value pairs.</p> <p>You can configure the layout so that the event panel is displayed on the top or the bottom of the operator view, or not at all.</p>
Action panel	<p>Contains a list of policies associated with this view. You can configure the layout so that the action panel is displayed on the top, the bottom, the left or the right of the display, or not at all.</p>
Information group panel	<p>Displays sets of information retrieved from data types. This data is often business data that is related to event information passed to the view from Netcool OMNIBus or another application.</p>

Action panel policies

You can use the action panel editor in the GUI to specify one or more policies that are displayed in the action panel of a basic operator view.

The action panel presents a list of policies that the user can start from within the view. This is an optional part of the operator view function. You use the action panel to start policies only, you cannot use it to display data that is returned by a policy. An advanced operator view, however, does provide the capability to display this data.

Information groups

An information group is a set of dynamic data that is displayed when you open the view.

This is often business data that is related to event information that is passed to the view from Netcool OMNIbus or another application. The data that is displayed in an information group is obtained by a query to a data source either by filter or by key.

When you create a basic operator view using the GUI, you can specify one or more information groups that are to be displayed by the view.

The following table shows the properties that you specify when you create an information group:

Table 85. Information group configuration properties

Property	Description
Group name	Unique name that identifies the information group.
Type	Type of query to a data source. Available options are: <ul style="list-style-type: none">• By Key: Key expression that specifies which data to retrieve from the data type.• The filter syntax is similar to the contents of the WHERE clause in an SQL select statement.• By Filter: SQL filter string that specifies which data to retrieve from the data type.
Data type	Data type that contains the data that you want to display.
Value	Add a value.
Style	Layout style for data items in the resulting information group. Options are Tabbed and Table.

You can customize the information that is displayed in the information groups by editing the operator view policy.

Creating a basic operator view

About this task

Complete the following steps to create basic operator views:

Procedure

1. Log on to the GUI.

2. In the navigation tree, expand **System Configuration > Event Automation** click **Operator Views** to open the **Operator Views** tab.
3. From the **Cluster** list, select the cluster you want to use.
4. From the **Project** list, select the project you want to use.
5. Click the **New Operator View** icon to open the **New Operator View**.
6. In the **Operator View Name** field, enter a unique name for the operator view. You cannot edit the name once the operator view is saved.
7. In the **Layout Options** area, specify the position of the event panel and action panel in the operator view. You can preview the appearance of the operator view using the images available in the **Preview** area.
8. Click the **Action Panel** link, select one or more action policies that the user can open from within the operator view.
9. Click the **Information Groups** link. Use the following steps to create one or more information groups:
 - a. Click the **New Information Group** icon to insert a new row into the information groups table.
 - b. In the **Group Name** field, type a unique name for the group.
 - c. From the **Type** list, select **By Filter** or **By Key** to specify whether the information group retrieves data from a data type by filter or by key.
 - d. From the **Data Type** list, select the data type that contains the information you want to display.
 - e. In the **Value** field, enter a filter or key expression. If the **Type** is **By Filter** adding a value is optional. If the **Type** is **By Key** then the value is mandatory.
 - f. In the **Style** list, select **Tabbed** or **Table** to specify how the operator view shows the resulting data.
 - g. Press Enter on your keyboard to confirm the value you are adding to the information group (or press Escape on your keyboard to cancel the edit).
 - h. Repeat these steps to create multiple information groups for any operator view.
 - i. To edit an information group, click the item you want to edit and change the value.
 - j. To delete one or more information groups, multiselect the rows groups using the Ctrl and shift keys on the keyboard, then click **Delete**.
 - To sort rows up or down, select a row or multiple rows to activate the **Move Up** and **Move Down** arrows on the toolbar. Click the required icon to move the rows up or down by one row.
10. Click the **Save** icon on the main editor toolbar to implement the changes.

Editing operator views

About this task

You can only modify the policy that is associated with an advanced operator view using the GUI or an external text editor. If you modify the policy using an external text editor, you must import the policy manually after you make your changes. You are not required to stop and restart the Impact Server or GUI Server after modify an existing operator view policy. Any changes that you make to the policy are immediately recognized by the system.

You can modify the display page that is associated with an advanced operator view using an external text editor only. Do not attempt to modify the display page using the tools provided by the GUI for use with basic operator views. If you modify the display page using the GUI, the changes that you make override the HTML content and smart tags in the existing HTML file.

Procedure

1. To modify a basic operator view, log on to the GUI.
2. In the navigation tree, expand **System Configuration > Event Automation**, click **Operator Views** to open the **Operator Views** tab.
3. From the **Cluster** list, select the cluster you want to use.
4. From the **Project** list, select the project you want to use.
5. Double click on the operator view you want to modify, or click the **Edit Operator View** icon.
6. Modify the operator view configuration properties as required. You cannot modify the **Operator View Name**.
7. Click the **Save** to implement the changes.

Deleting operator views

Procedure

1. To delete a basic or advanced operator view log on to the GUI.
2. In the navigation tree, expand **System Configuration > Event Automation**, click **Operator Views** to open the **Operator Views** tab.
3. From the **Cluster** list, select the cluster you want to use.
4. From the **Project** list, select the project you want to use.
5. Select the operator view you want to delete and click the **Delete** icon on the toolbar, or right click the operator view and click **Delete**.
6. A confirmation message displays.
7. Click **OK** to delete the operator view.

The operator view is removed from the GUI display and the operator view display file and policy are removed from the system.

Displaying operator views in TIP

Follow this procedure to create a page in Tivoli Integrated Portal (TIP), and display an operator view in it.

Before you begin

Make sure your TIP user rights allow you to create pages.

Procedure

1. Log on to TIP.
2. In the TIP navigation pane, click the **Settings > Pages** node.
3. Click the **New Page** button, in the **Pages** window.
4. In the **Page Settings** window, provide the following information:

Page name

The name of the new page.

Page location

The position of the new page in the navigation pane.

Page layout

Choose the “Classic”, or the “Freeform” layout for the new page.

Optional

Click **OK** to continue.

5. In the **Choose a Portlet** window, add a portlet to the new page.
 - a. Select the **Operator View** portlet to add it to your page.
 - b. Click **OK** to continue.

You can disregard the error message, and click **OK**. It is only to inform you that your new operator view widget requires configuration.

6. Click the **Save** button in the upper right hand corner of the page.
7. To configure the new operator view widget, select **Personalize**, or **Edit Shared Settings** from the pull down menu available in the upper right hand corner of the widget.

If you select **Personalize**, the new configuration will apply only to the user, who set the preferences. The **Edit Shared Settings** option allows the administrator to set the preferences for all users. The latter, is the recommended choice for operator view widgets.

8. In the **Operator View** window, provide the following information:

Portlet Title

Type in the portlet title. Optional.

Cluster

Select a cluster from the menu. Required.

Operator View Name

Type in the filename of the existing operator view, without the extension. For example, EIC_configure. Required.

Click **OK** to update the operator view with the new settings.

Note: Clicking the **Restore Defaults** button, clears all entered information.

Chapter 12. Event Isolation and Correlation

Event Isolation and Correlation is provided as an additional component of the Netcool/Impact product. Event Isolation and Correlation is developed using the operator view technology in Netcool/Impact. You can set up Event Isolation and Correlation to isolate an event that has caused a problem. You can also view the events dependent on the isolated event.

Overview

Netcool/Impact has a predefined project, **EventIsolationAndCorrelation** that contains predefined data sources, data types, policies, and operator views. When all the required databases and schemas are installed and configured you must set up the data sources. Then, you can create the event rules using the objectserver sql in the Event Isolation and Correlation configuration view from the Tivoli Integrated Portal. You can view the event analysis in the operator view, **EIC_Analyze**.

To set up and run the Event Isolation and Correlation feature the following steps need to be completed.

1. Install Netcool/Impact.
2. Install DB2 or use an existing DB2 installation.
3. Configure the DB2 database with the DB2 Schema in the Netcool/Impact launchpad.
4. Install Discovery Library toolkit from the Netcool/Impact launchpad.
If you already have a Tivoli® Application Dependency Discovery Manager (TADDM) installation, configure the discovery library toolkit to consume the relationship data from TADDM. You can also consume the data through the loading of Identity Markup Language (IDML) books. For additional information about the discovery library toolkit, see the *Tivoli Business Service Manager Administrator's Guide* and the *Tivoli Business Service Manager Customization Guide*. These guides are available in the Tivoli Business Service Manager 6.1.0.1 information center available from the following url, <https://www.ibm.com/developerworks/wikis/display/tivolidoccentral/Tivoli+Business+Service+Manager>.
5. In the Tivoli Integrated Portal, configure the data sources and data types in the **EventIsolationAndCorrelation** project to use with the Impact Server.
6. Create the event rules in the UI to connect to the Impact Server.
7. Configure WebGUI to add a new launch point.

Detailed information about setting up and configuring Event Isolation and Correlation, is in the *Netcool/Impact Solutions Guide*.

Event Isolation and Correlation policies

The **EventIsolationAndCorrelation** project has a list of predefined policies that are specific to Event Isolation and Correlation.

The following policies are in the **EventIsolationAndCorrelation** project and support the Event Isolation and Correlation feature and must not be modified:

- **EIC_IsolateAndCorrelate**

- **EIC_eventrule_config**
- **EIC_utils**
- **Opview_EIC_Analyze**
- **Opview_EIC_confSubmit**
- **Opview_EIC_configure**
- **Opview_EIC_requestHandler**

Event Isolation and Correlation operator views

The **EventIsolationAndCorrelation** project has a list of predefined operator views that are specific to Event Isolation and Correlation.

- **EIC_Analyze** shows the analysis of an event query.
- **EIC_confSubmit** supports the configuration of Event Isolation and Configuration.
- **EIC_configure** configures the event rules for Event Isolation and Configuration.
- **EIC_requestHandler** supports the configuration of Event Isolation and Configuration.

Configuring Event Isolation and Correlation data sources

All the Event Isolation and Correlation-related features are associated with the project, **EventIsolationAndCorrelation**. Configure the necessary data sources, data types, and data items for the event isolation and correlation.

Procedure

1. From the Tivoli Integrated Portal, click **System Configuration > Event Automation > Data Model**.
2. From the project list, select the project **EventIsolationAndCorrelation**. A list of data sources specific to the **EventIsolationAndCorrelation** feature display.
 - **EIC_alertsdb**
 - **SCR_DB**
 - **EventrulesDB**
3. For each data source, update the connection information, user ID, and password and save it.
4. Configure **EIC_alertsdb** to the object server where the events are to be correlated and isolated.
5. Configure **SCR_DB** to the Services Component Registry database.

Note: When configuring the Services Component Registry (SCR) data sources, you must point the data sources to what is commonly called the SCR. The SCR is a schema within the TBSM database that is created when you run the DB2 schema configuration step. The schema is called **TBSMSCR**. The database has a default name of **TBSM**.

6. Configure **EventRulesDB** to the Services Component Registry database.

Configuring Event Isolation and Correlation data types

The **EventIsolationAndCorrelation** project has a list of predefined data types that are specific to Event Isolation and Correlation. Except for the data type **EIC_alertquery** which you must configure, the remaining data types are preconfigured and operate correctly once the parent data sources are configured.

About this task

The following list shows the Event Isolation and Correlation data sources and their data types:

- **EIC_alertsdb**
 - **EIC_alertquery**
- **SCR_DB**

The following data types are used to retrieve relationship information from the Services Component Registry.

- **bsmidentities**
- **getDependents**
- **getRscInfo**

- **EventRulesDB**

The following data types used by the database contain the end user configuration for Event Isolation and Correlation.

- **EVENTRULES**
- **EIC_PARAMETERS**

Procedure

1. To configure the **EIC_alertquery** data type, right click on the data type and select **Edit**.
2. The **Data Type Name** and **Data Source Name** are prepopulated.
3. The **State** check box is automatically selected as **Enabled** to activate the data type so that it is available for use in policies.
4. **Base Table**: Specifies the underlying database and table where the data in the data type is stored.
5. Click **Refresh** to populate the table. The table columns are displayed as fields in a table. To make database access as efficient as possible, delete any fields that are not used in policies. For information about adding and removing fields from the data type see "SQL data type configuration window - Table Description tab" on page 71.
6. Click **Save** to implement the changes.

Creating, editing, and deleting event rules

How to create, edit, and delete an event rule for Event Isolation and Correlation.

Procedure

1. Select **System Configuration > Event Automation > Event Isolation and Correlation** to open the Event Isolation and Correlation page tab.
2. Click the **Create New Rule** icon to create an Event Rule. While creating this item the configure page has empty values for various properties.
3. Click the **Edit the Selected Rule** icon to edit the existing event rules.
4. Click the **Delete the Selected Rule** icon to delete an event rule from the system and the list.

Creating an event rule

Complete the following fields to create an event rule.

Procedure

1. **Event Rule Name:** Specify the event rule name. The event rule name must be unique across this system. When you select **Edit** or **New** if you specify an existing event rule name, the existing event rule is updated. When you edit an event rule and change the event rule name, a new event rule is created with the new name.
2. **Primary Event:** Enter the SQL to be executed against the objectserver configured in the data source **EIC_alerts db**. The primary event is the event selected for analysis.

The primary event filter is used to identify if the event that was selected for analysis has a rule associated with it. The primary event filter is also used to identify the object in the Services Component Registry database that has the event associated with it. The object may or may not have dependent entities. During analysis, the event isolation and correlation feature finds all the dependent entities and there associated events.

For example, the primary event has 3 dependent or child entities and each of these entities has 3 events has associated with it. In total there are 9 dependent events. Any of these secondary events could be the cause of the primary event. This list of events is what is termed the list of secondary events. The secondary event filter is used to isolate one or more of these events to be the root cause of the issue.

3. **Test SQL:** Click **Test SQL** to test the SQL syntax specified in the primary event. Modify the query so that only one row is returned. If there are multiple rows, you can still configure the rule. However, during analysis only the first row from the query is used to do the analysis.
4. **Secondary Events:** The text area is for the SQL to identify the dependent events. When you specify the dependent events, you can specify variables or parameters which can be substituted from the primary event information. The variables are specified with the @ sign. For example, if the variable name is *dbname*, it must be specified as *@dbname@*. An example is Identifier = 'BusSys Level 1.2.4.4' and Serial = @ser@. The variables are replaced during the analysis step. The information is retrieved from the primary event based on the configuration in the parameters table and displays in the **Variables Assignment** section of the page.
5. **Extract parameters:** Click **Extract Parameters** to extract the variable name between @ and populate the parameter table. Once the variable information is extracted into the table, you can edit each column.
 - a. Select the field against the regular expression you want to execute, and a substitution value is extracted.
 - b. Enter the regular expression in the regular expression column. The regular expression follows the IPL Syntax and is executed using the RExtract function.
 - c. When the regular expression is specified, click **Refresh** to validate the regular expression and check that the correct value is extracted. The table contains the parameters.
6. **Limit Analysis results to related configuration items in the Service Component Registry:** Select this check box if the analysis is to be limited to related configuration items only. If the check box is not selected, the dependent query will be returned.
7. **Primary Event is a root cause event:** Select this check box to identify whether the primary event is the cause event and rest of events, are symptom only events.

8. **Event Field:** Identifies the field in the event which contains the resource identifier in the Services Component Registry. Select the field from the drop-down menu that holds the resource identifier in the event.
9. **Time window in seconds to correlate events:** Add the time period the event is to analyze. The default value is 600 seconds. The events that occurred 600 seconds prior to the primary event are analyzed.
10. Click **Save Configuration** to add the configuration to the backend database.
11. Now the event rules are configured, the event is ready to be analyzed. You can view the event analysis in the in the **EIC_Analyze** page.

Configuring WebGUI to add a new launch point

Configure the WebGUI with a launch out context to launch the analysis page.

About this task

WebGUI can be configured to launch the analysis page. Refer to the procedure for launch out integration described in the following URL, http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/index.jsp?topic=/com.ibm.netcool_OMNIBus.doc_7.3.1/webtop/wip/task/web_con_integrating.html.

The URL you need for Event Isolation and Correlation is `<TIPHOSTNAME>:<TIPPORT>/opview/displays/NCICLUSTER-EIC_Analyze.html`. Pass the serial number of the selected row for the event.

Note: NCICLUSTER is the name of the cluster configured during the installation of Netcool/Impact. You must use the name of your cluster whatever it is, in the URL. For example, in Tivoli Business Service Manager the default cluster name is TBSMCLUSTER. To launch from Tivoli Business Service Manager, you would need to use the following html file, TBSMCLUSTER-EIC_Analyze.html.

Launching the Event Isolation and Correlation analysis page

How to launch the Event Isolation and Correlation analysis page.

About this task

There are two ways to launch the Event Isolation and Correlation analysis page.

- Manually by using the webpage and Event Serial number.
- Using the launch out functionality on Active Event List (AEL) or Lightweight Event List (LEL) from WebGUI in the Tivoli Enterprise Portal.

Procedure

Open a browser on Netcool/Impact. Use one of the following options:

- Point to `<TIPServer>:<TIPPort>/opview/displays/NCICLUSTER-EIC_Analyze.html?serialNum=<EventSerialNumber>`. Where `<TIPServer>` and `<TIPPort>` are the Netcool/Impact GUI Server and port and `EventSerialNumber` is the serial number of the event you want to analyze. To launch the analysis page outside of the AEL (Action Event List), you can add `serialNum=<Serial Number>` as the parameter.
- The Event Isolation and Correlation analysis page can be configured to launch from the Active Event List (AEL) or LEL (Lightweight Event List) within

WebGUI. For more information see, “Configuring WebGUI to add a new launch point” on page 165. When you create the tool you have to specify only `<TIPSERVER>:port/opview/displays/NCICLSTER-EIC_Analyze.html`. You do not have to specify **SerialNum** as the parameter, the parameter is added by the AEL tool.

Viewing the Event Analysis

View the analysis of an Event query in the **EIC_Analyze** page.

About this task

The input for the **EIC_IsolateAndCorrelate** policy is the serial number of the event through the `serialNum` variable. The policy looks up the primary event to retrieve the resource identifier. The policy then looks up the dependent events based on the configuration. The dependent events are further filtered using the related resources, if the user has chosen to limit the analysis to the related resources. Once the serial number has been passed as the parameter in WebGUI, you can view the event from the AEL or LEL and launch the Analyze page.

Procedure

Select the event from the AEL or LEL and launch the Analyze page. The **EIC_Analyze** page contains three sections:

- **Primary Event Information:** shows the information on the selected event. This is the event on which the event isolation and correlation analysis takes place.
- **Correlated Events:** shows information about the dependent events identified by the tool. Dependant events are identified as the events that are associated with the dependant child resources of the device or object that is associated with the primary event. These events are displayed in the context of dependent resources that were identified from the Services Component Registry.
- **Event Rule process:** shows the rule which was identified and processed when this primary event was analyzed.

Chapter 13. Reporting tools

The reports provide information about your network and network operators and help you to assess the efficiency of your configuration.

Accessing reports

Use this procedure to access the reports.

Procedure

1. From the Tivoli Integrated Portal navigation tree, expand **Reporting > Event Automation**.
2. Select the report you want to run, the tab for the specified report opens.
The following reports are available:
 - Action Efficiency Report
 - Action Error Report
 - Policy Efficiency Report
 - Policy Error Report
 - Node Efficiency Report
 - Operator Efficiency Report
 - Impact ROI Efficiency Report
 - Impact Profile Report
3. In the tab menu, select the date and time ranges. Select the view option you want, either **Chart View** or **Tabular view** then run the report. The time range displays in local time. For more information see “Viewing Reports” and “Reports toolbar” on page 168.

Viewing Reports

The reports present their data in graphical and tabular format. Use the chart view tab and the tabular view tab to switch between these two formats.

Chart view

The chart view presents the report data in graphical format. The legend shows the color code for each action. The descending order in the legend reflects the order from left to right in the chart.

Tabular view

The tabular view presents the report data in a table. To get more detail for a particular row of the table, select the row, then click the **DrillDown** icon on the toolbar above the table. The table refreshes automatically and loads the information for the row. To return to the main report view click the **Drillup** arrow icon on the toolbar.

If you are viewing a multi-page report, use the **Page** and **Row** controls at the bottom of the table. In the **Page** field, click the arrows to get to the page you want to view. In the **Row** field, use the arrows to adjust the number of rows that display

per page. The minimum number of rows is three and the maximum is 50 per page. The total number of rows that display on a page is shown on the lower right corner of the table.

Multi-page reports have **Previous** and **Next** links so that you can move from page to page. You can also click the individual page numbers to move to specific pages.

Note: When viewing the contents of reports, Netcool/Impact loads the data from the HSQL database as long as the number of items is within the threshold limit. The default threshold limit is 10000. The threshold limit is set in `NCHOME/impact/etc/server.props` using the property, `impact.dataitems.threshold`. To view data exceeding the threshold limit, the `impact.dataitems.threshold` property would need to be modified and the server restarted. Note that the higher the value is set, the more memory is consumed.

Reports toolbar

You use the report toolbar to perform a basic report configuration.

You can find some toolbar controls, for example, the time of the report, selection fields, or the refresh report icon, can be found in all reports. Other controls can be found only in specific reports.

Selecting the time range

Use the **Start** and **End** fields in the report toolbar to configure the date and time range for a report. Click a field to activate the date and time menus. The time range displays in local time.

The default parameter for the date is four weeks, with the last day being the current date and time. The time can be set at 15 minute intervals.

Time Range: Start	6/9/2011	11:18 AM	End	7/7/2011	11:18 AM
--------------------------	----------	----------	-----	----------	----------

Report icons

This table explains the function of the icons that you can find in the reports.

Table 86. Report icons









Icon	Description
	Click to refresh the report data after changing the report parameters.
	Open a window that you can use to change the report parameters. You can find this icon only in the Impact Profile report and Impact ROI Efficiency report.
	Open a window where you associate the business processes with a policy. You can find this icon only in the Impact ROI Efficiency report.
	Clear all Impact Profile Report data. You can find this icon only in the Impact Profile report.

Table 86. Report icons (continued)

Icon	Description
	Click this arrow to generate a report.
	Stop collecting data for this report. This icon can be found only in the Impact Profile report.
	In the report tabular view, you can drill down to view more detailed information about a row, by selecting a row, and then clicking this icon. This icon is only enabled after you select a table row.
	Click this icon to return to the main table view of a report, after you drill down for more detail.

Action Efficiency report

The Action Efficiency report shows the total number of actions that were processed over a selected time range.

Using this report you can learn how many actions the Impact Server performed and which actions you are using the most.

The chart view shows how many times each action has been performed for the Impact Server.

The tabular view contains a table that shows how many times an action was run and the average time it took to process the action.

The detail view shows the action name, the name of the policy executed, the time it took to process the action in seconds, and the time it was processed.

To use this report, enable reporting in the Policy Logger service configuration window. See “Policy logger service configuration window” on page 133.

Action Error report

The Action Error report shows the actions that generated errors each time a policy executed.

The report shows you how many action errors occurred in Netcool/Impact over a time period that you selected.

The **Chart View** reports how many times each action failed.

The **Tabular View** tab contains a table that shows the number of errors for each action.

You can drill down to see the policies where the errors occurred, the time the errors occurred, and the error messages that resulted.

The detailed view shows the following details:

- Type of action
- Policy it belongs to

- Time the policy executed
- Error message it generated.

To use this report, enable reporting in the Policy Logger service configuration window. See “Policy logger service configuration window” on page 133.

Impact Profile report

The Impact Profile report provides information about the efficiency of the Impact Server.

The detailed view shows the following details about Netcool/Impact configuration:

- SQL query
- Policy that issued this query
- Type of action
- Data source queried
- Metric

Configuring Impact Profile report

Use this procedure to configure the Impact Profile report.

Procedure

1. From the Tivoli Integrated Portal, expand **Reporting > Event Automation** select **Impact Profile Report**.
2. From **Impact Profile Report** toolbar, click **Open Configuration** to open the Impact Profile Rules Editor window.

Use this window to set the parameters for the report. For more details about the available parameters, see “Impact Profile Report rules editor” on page 171.

Impact Profile Report data

The Impact Profile Report Rules editor lists all the queries you can use to generate an Impact Profile Report.

Table 87. Impact Profile Report Parameters

Impact Profile	Description	Impact Profile Rule
Queries sent to same data source by same policy more than n times in n seconds	The number of "hotspot" queries sent to the same data source by the same policy in more than a specified number of seconds.	SQL Query XinY Rules
Queries done more than n times in n seconds that are taking more than n milliseconds	Measures the number of queries made in a specified number of seconds that take more than a specified number of milliseconds.	SQL Hotspot Rules
Queries made more than n times in n seconds that return more than n rows	Counts the number of queries made in a specified number of seconds that return more than a specified number of rows.	SQL Hotspot Rules

Table 87. Impact Profile Report Parameters (continued)

Impact Profile	Description	Impact Profile Rule
Inserts into any types more than n times in n seconds that are taking more than n milliseconds.	Measures the number of SQL inserts into any type of data type in a specified time window that take more than a specified number of milliseconds.	SQL Hotspot Rules
Internal types written more than n times in n seconds	Measures the number of internal data types that are accessed more than a specified number of times in a specified number of seconds.	Internal Type Rules
Same identifier updated by ReturnEvent more than n times in n seconds	Measures the number of return events that update events using the same identifier as the source event.	Return Event Rules
Same identifier inserted into the same Object Server that events are read from.	Measures the number of new events that were sent to the ObjectServer that use the same identifier that they read the event from.	Return Event Rules
JRExec calls done more than n times in n seconds that are taking more than n seconds	Measures the number of "troublesome" JRExec calls in more than a specified number of times in a specified time period.	JRExecAction Rules
Hibernations built up in memory more than n (true/false)	Measures whether the number of hibernations that have built up in memory is more than a specified number over the lifetime of the server.	Hibernation Rules

Impact Profile Report rules editor

Use the following rules and settings to edit the Impact Profile queries.

Procedure

1. Select the rule in the form that is associated with the query you want to edit.
2. **SQL Query XinY Rules**

Use this option to change the settings for the following query:

Queries sent to same data source by same policy more than n times in n seconds

- Select the **Count Threshold** to set the number of SQL queries to be run.
- Select the **Count Time Window** to set the time window the measurement is to be based on.

3. **SQL Hotspot Rules**

Use this option to change the settings for the following queries:

Queries done $>n$ times in n seconds that are taking more than n milliseconds

Queries made $>n$ times in n seconds that return $>n$ rows

Inserts into any types $>n$ times in n seconds that are taking $>n$ milliseconds

- Select the **Insert Execution Time Threshold** to set the time threshold for the SQL inserts.
 - Select the **Query Execution Time Threshold** to set the time threshold for query execution.
 - Select the **Query Return Row Threshold** to set the threshold for the number of queries to be retrieved.
 - Select the **Count Threshold** to set the threshold for the number of SQL statements to be run.
 - Select the **Count Time Window** to set the time window the measurement is to be based on.
4. **JRExecAction Rules**
- Use this option to change the setting for the following query:
JRExec calls done more than n times in n seconds that are taking > n seconds
- Select the **Count Threshold** to set the threshold for the number of JRExecActions to be run.
 - Select the **Execution Time Threshold** to set the threshold for how long the JRExecActions must take.
 - Select the **Time Window** to set the time window the measurement is to be based on.
5. **Internal Type Rules**
- Use this option to change the settings for the following query:
Internal types written more than n times in n seconds.
- Select the **Count Threshold** to set the number of times internal data types are written to.
 - Select the **Time Window** to set the length of time the profile is based on.
6. **ReturnEvent Rules**
- Use this option to change the settings for the following queries:
Same identifier updated by ReturnEvent > n in n seconds
Same identifier inserted into the same Objectserver that Netcool/Impact reads events from
- Select the **Count Threshold** to set the count threshold for the number of returned events.
 - Select the **Time Window** to set the length of time the profile is based on.
7. **Hibernation Rules**
- Use this option to change the settings for the following queries:
Set the number of hibernations to be held in memory
Hibernations built up in memory > n (true/false)
- Select the **Hibernation in Memory Threshold** to set the number of hibernations to be held in memory.
8. Click **OK** to accept the parameter changes. Click **Refresh Report** to update the parameters in the Impact Profile Rules Editor.

Impact ROI Efficiency report

The Impact Return on Investment (ROI) Report shows operator time saved as a result of a Netcool/Impact deployment compared to the time it would take an operator to solve the identical problem manually.

The manual times defaults, provided with the Netcool/Impact installation, are calculated from industry statistics for common tasks. You associate the relevant policies with these calculations before you turn on report data collection in the Policy Logger service. In order for each calculation to work, you must associate at least one policy with it. The saved time is based on how many times the corresponding policies are executed against the manual process time of the ROI business process during a specified period of time. After you associate relevant policies with the calculations, you turn on Impact ROI Efficiency Reporting.

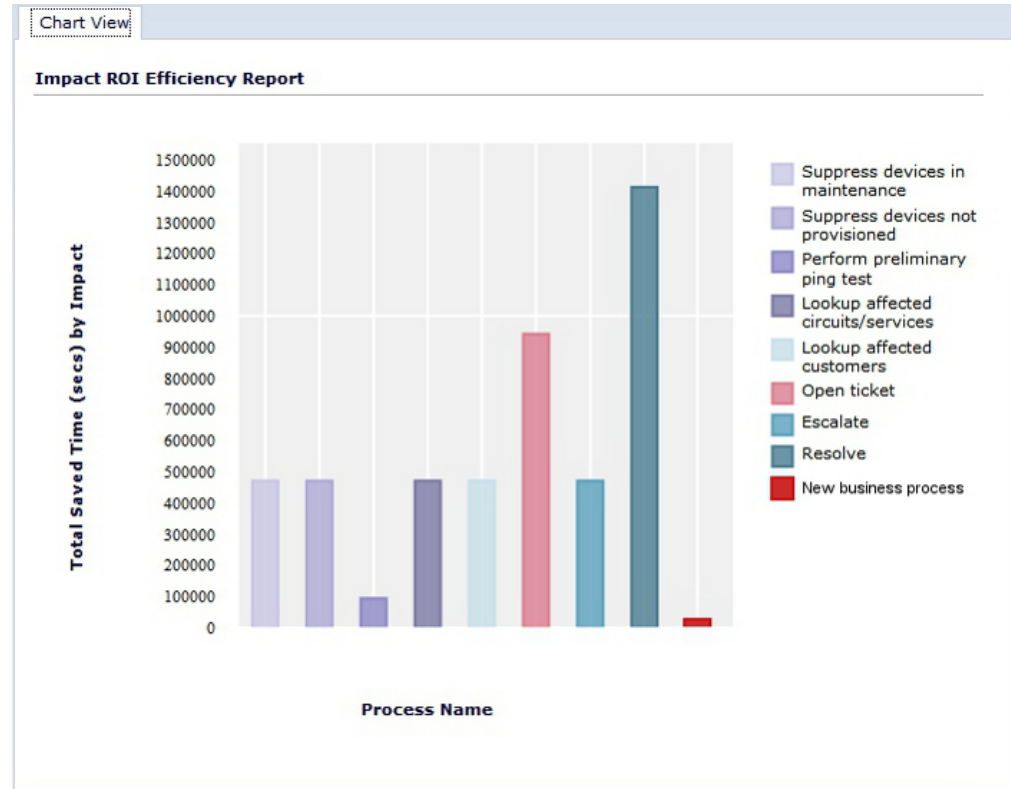


Figure 2. Impact ROI Efficiency report

Important: Before you configure this report, enable report data collection in the Policy Logger service. For more information, see “Policy logger service configuration window” on page 133.

Report views

The chart view presents the report data in graphical format. The legend shows the color code for each process. You can hover the mouse cursor over a process in the chart view to highlight it, and see the total time saved in seconds after automating the process.

The tabular view shows the following details:

- The process time
- The time it would take an operator to perform the task manually
- The time saved in seconds by automating the process

Impact ROI Efficiency report business processes

A business process is an action that is typically performed manually by an operator.

The Impact ROI Efficiency report is installed with eight default business processes:

- Suppress devices in maintenance
- Suppress devices not provisioned
- Perform preliminary ping test
- Lookup affected circuits/services
- Lookup affected customers
- Open ticket
- Escalate
- Resolve

These business processes are provided as examples only. To use one of them, you need to associate it with a relevant policy. You can also add your own business processes, as necessary.

Creating a sample Impact ROI Efficiency report

To configure your Impact ROI Efficiency report, you need to associate the relevant policies and business processes.

Procedure

1. From the Tivoli Integrated Portal navigation tree, select **Reporting > Event Automation > Impact ROI Efficiency Report**.
2. Click the **Configuration** icon and select the **Configure Business Process** option, to add a business process.

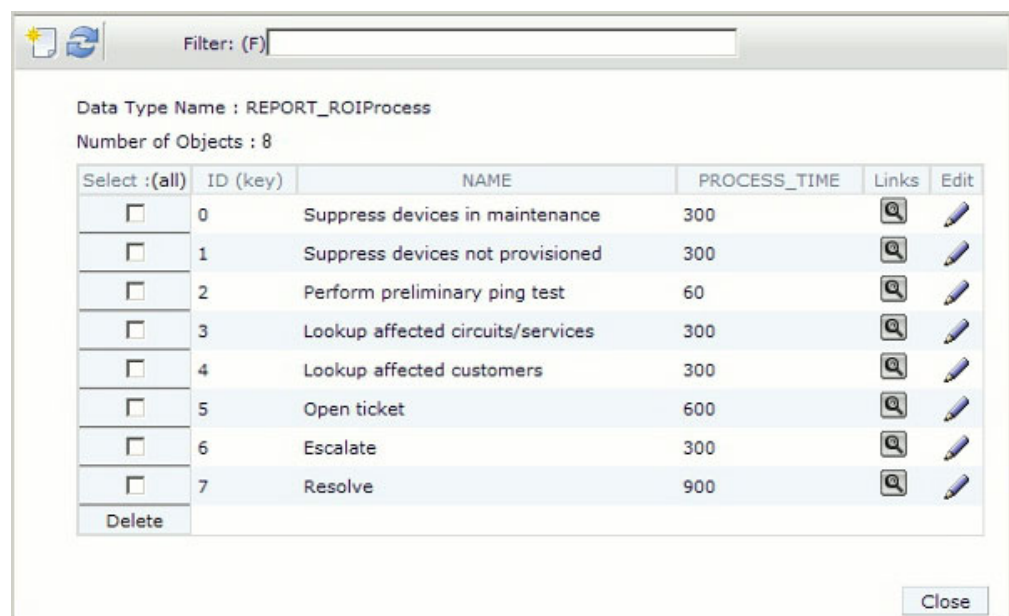


Figure 3. Configure Business Process window

3. In the editor toolbar, click the **New Data Item** icon.
 - a. Type or select an ID for the data item.

- b. Type a name for the business process.
 - c. Type the manual time for this process.
4. Click **OK** and then repeat for each new process you want to add.
5. Click **Configure** and select **Configure Policy and Report Mapping** to associate the processes with a policy.
 - a. Select the policy that you want associate with a process.
 - b. Select the processes you want to map to from the **Available Processes** list.
 - c. Click **Add** to move them to the **Assigned Processes** list.
 - d. Optional: If you decide you do not want to associate a process to this policy, select it and click **Remove** to move it back to the **Available Processes** list.
 - e. Click **Apply** and close the window.
6. Using the **Time Range** controls select the time period for which you want to run the report.
7. Click **Refresh Report**.
The configured report displays in the editor.

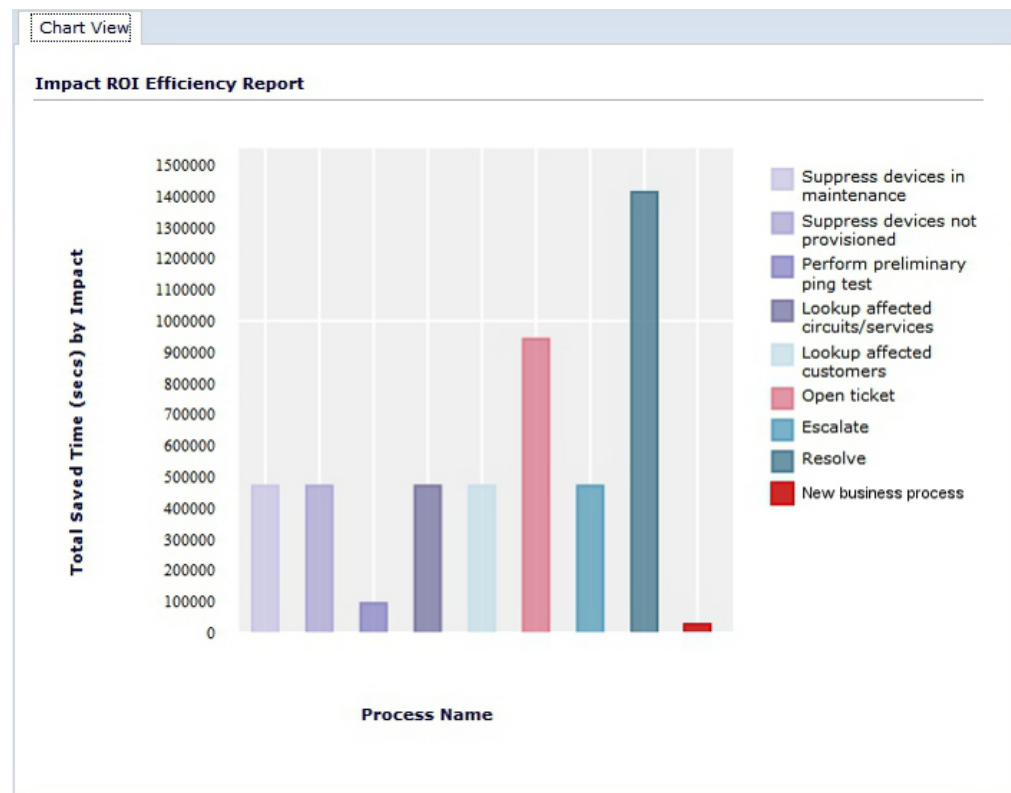


Figure 4. Impact ROI Efficiency Report: chart view

The legend on the left shows the color code for each process. The descending order of the legend reflects the order from left to right in the chart.

8. Click the **Tabular View** tab.

Tabular View		
Impact ROI Efficiency		
Action	Average Time	Execution Count
Suppress devices in maintenance	300	17100
Suppress devices not provisioned	600	17100
Perform preliminary ping test	60	3420
Lookup affected circuits/services	300	17100
Lookup affected customers	300	17100
Open Ticket	600	34200
Escalate	300	17100
Resolve	900	51300
New business process	300	1100
Pages 1 of 1 Rows 25 Total: 9		

Figure 5. Impact ROI Efficiency Report: tabular view

The new business process saved an operator 1100 seconds.

Node Efficiency report

The Node Efficiency Report records the number of alerts generated by a node.

The chart view shows the unique event count for each node.

The tabular view shows the node name and the unique event count.

The detail view shows the following information for each node:

- Node name
- Severity level
- Information recorded in the Objectserver Summary field
- Location of the node
- Whether the event has been acknowledged
- Unique Event Name

To use this report, enable reporting in the Event Reader Service Configuration window. See “OMNIBus event reader service **General Settings** tab” on page 145

Operator Efficiency report

The Operator Efficiency reporting tool records how quickly operators respond to events.

For each operator, the report records the following information:

- Operator name
- The average time between when the event first occurs and the operator acknowledgment of the event

The chart view displays the average acknowledgment time for each operator.

The tabular view shows the following information:

- Operator name
- Average event Acknowledgment time in seconds
- Acknowledgment count

The detail view shows the following information for each operator:

- The operator name
- Each unique event
- The entry in the event list Summary field
- Acknowledgment count
- The severity level assigned to the event

Policy Efficiency report

The Impact Policy Efficiency Report records historical information about the performance of all your policies.

Each time a policy runs, the time taken to run it is recorded. When reporting is switched on, you can see a table of all policies and the average execution time and count for each one.

The chart view shows the average time in seconds each policy took to run.

The **Tabular View** tab shows the policy name, the average time it took in seconds to run it, and how many times it ran in the specified time range.

The detail view shows the following information:

- The name of the policy
- The policy run time
- The time the policy ran

To use this report, enable reporting in the Policy Logger service configuration window. See “Policy logger service configuration window” on page 133

Policy Error report

The Policy Error Report gives you a list of the policies that generated errors along with how many times each policy was run.

The chart view shows the error count for each policy.

The tabular view shows the failure count for each policy within the specified date range.

This detail view shows the following detail:

- The policy name
- The times the policy executed
- The error message generated.

To use this report, enable reporting in the Policy Logger service configuration window. See “Policy logger service configuration window” on page 133

Chapter 14. Maintenance Window Management

Maintenance Window Management (MWM) is an add-on for managing Netcool OMNIbus maintenance windows.

MWM can be used with OMNIbus versions 7.1, 7.2, and 7.3. A maintenance time window is a prescheduled period of downtime for a particular asset. Faults and alarms, also known as events, are often generated by assets undergoing maintenance, but these events can be ignored by operations. MWM creates maintenance time windows and ties them to Netcool OMNIbus events that are based on OMNIbus fields values such as **Node** or **Location**. Netcool/Impact watches the Netcool OMNIbus event stream and puts these events into maintenance according to the maintenance time windows. The Netcool/Impact **MWMActivator** service located in the **System Configuration > Event Automation > Services** in the **MWM** project must be running to use this feature. For more information about maintenance windows, see “About MWM maintenance windows.”

About MWM maintenance windows

Use the Maintenance Window Management (MWM) web interface to create maintenance time windows and associate them with Netcool OMNIbus events.

Netcool OMNIbus events are based on OMNIbus field values such as **Node** or **Location**. The Netcool OMNIbus events are then put into maintenance according to these maintenance time windows. If events occur during a maintenance window, MWM flags them as being in maintenance by changing the value of the OMNIbus field, integer field, SuppressEsc1 to 6 in the alerts.status table.

A maintenance time window is prescheduled downtime for a particular asset. Faults and alarms (events) are often generated by assets that are undergoing maintenance, but these events can be ignored by operations. MWM tags OMNIbus events in maintenance so that operations know not to focus on them. You can use MWM to enter one time and recurring maintenance time windows.

- **One time windows** are maintenance time windows that run once and do not recur. **One Time Windows** can be used for emergency maintenance situations that fall outside regularly scheduled maintenance periods. You can use them all the time if you do not have a regular maintenance schedule.
- **Recurring time windows** are maintenance time windows that occur at regular intervals. MWM supports three types of recurring time windows:
 - **Recurring Day of Week**
 - **Recurring Date of Month**
 - **Every nth Weekday**

Maintenance time windows must be linked to OMNIbus events in order for MWM to mark events as being in maintenance. When you configure a time window, you also define which events are to be associated with the time window. The MWM supports the use of **Node**, **AlertGroup**, **AlertKey**, and **Location** fields for linking events to time windows.

Logging on to Maintenance Window Management

Use the Tivoli Integrated Portal to access Maintenance Window Management (MWM).

Procedure

1. In the Tivoli Integrated Portal, expand **Troubleshooting and Support > Event Automation**.
2. Click **Maintenance Window Management** to open MWM. The main menu options are **Add One Time**, **Add Recurring**, and **View Windows**. There is also a **Time Zone** menu for setting your time zone. For more information about using these options, see “About MWM maintenance windows” on page 179.

Creating a one time maintenance window

Create a one time maintenance time window for a particular asset.

Procedure

1. Click the **Add One Time** link to view the form to create a one time maintenance window.
2. Enter the appropriate values in the fields **Node**, **AlertGroup**, **AlertKey**, and **Location**.
Select the **Equals** or **Like** options next to each field.
3. Click the calendar icon to select the **Start Time** and **End Time** for the maintenance time window.
4. Click **Add Window** to create the window.
5. Click **View Windows** to see the configured window.

Creating a recurring maintenance window

Create a recurring maintenance time window for a particular asset.

Procedure

1. Click the **Add Recurring** link to view the form for creating the different types of recurring time windows.
2. Enter the appropriate values in the fields **Node**, **AlertGroup**, **AlertKey**, and **Location**.
Select the **Equals** or **Like** options next to each field.
3. Select the **Start Time** and **End Time** for the maintenance time window.
4. Select the type of recurring window and complete the details.
 - **Recurring Day of Week** These windows occur every week on the same day and at the same time of day. For example, you can set the window to every Saturday from 5 p.m. to 12 a.m. Or you can set the window for multiple days such as Saturday, Sunday, and Monday from 5 p.m. to 12 a.m.
 - **Recurring Day of Month** These windows occur every month on the same date at the same time of day. For example, you can set the window to every month on the 15th from 7 a.m. to 8 a.m. Or you can set the window for multiple months.
 - **Every nth Weekday** These windows occur every month on the same day of the week at the same time. For example, you can set the window to the first and third Saturday of the month from 5 p.m. to 12 a.m.

5. Click **Add Window** to create the window.
6. Click **View Windows** to verify that your time window has been added.

Viewing maintenance windows

Click the **View Windows** link to view a toolbar which contains links to the different types of windows. Your viewing options are:

- **One Time**
- **Day of Week**
- **Day of Month**
- **nth Weekday**
- **Active Windows**

If maintenance windows are defined in any of these window categories, click a link to view a list of defined maintenance windows.

The color of the status icon indicates whether the window is active (green), expired (red), or has not started yet (blue, future).


You can use the delete icon to delete a maintenance window.

Chapter 15. Configuration documenter

The **Configuration Documenter** is used to view the configuration of a Netcool/Impact installation.

Configuration documenter overview

The **Configuration Documenter** is an integral part of any Netcool/Impact installation.



IBM Tivoli Netcool/Impact		
Configuration Documenter		Server Name: NCI_044 Version: Netcool/Impact 6.1 Generated on: Wed Jul 13 05:44:17 EDT 2011
Cluster Status Server Status Data Sources Data Types Policies Services		
Cluster Status for NCICLUSTER_044		
Primary Server	Host	
NCI_044 (Current Instance)	kiev	
Server Status		
Memory Status	Current Usage (in MB)	Max Limit (in MB)
Heap Memory Utilization	256	1200
Service Name	Number of Events in Queue	Event Source
EventProcessor	0	

Figure 6. An example of the configuration documenter screen

With the **Configuration Documenter**, you can view the detailed information about the system components:

- Cluster status: the name and the host where primary and secondary servers are running. Also, which server is currently the primary.
- Server status: information about the running services and the memory status for the server.
- Data sources: each defined data source
- All data types, including predefined data types such as **Doc**, **Schedule**, and **LinkType**; user-defined internal and external data types; and report data types.
- Policies
- Services


Opening the configuration documenter

Use this procedure to open the **Configuration Documenter** for a selected cluster.

Procedure

1. Log on to Tivoli Integrated Portal.
2. Select any component of Netcool/Impact to open it in a tab, in the workspace on the right.

For example, select **System configuration > Event Automation > Services**.

3. In the cluster selection menu, select a cluster for which you want to open the **Configuration Documenter**.
4. Click the **Configuration Documenter** icon, , next to the **Cluster** menu, to open the configuration documenter for the selected cluster.

The configuration documenter opens in a new browser window. Use the table of contents to view your cluster status, server status, data sources, data types, policies, and services.

Viewing the cluster status

Use this procedure to view the information about the current cluster status in the configuration documenter.

Procedure

1. Open the configuration documenter.
2. Select **Cluster Status** in the table of contents at the top of the page.

Depending on the status of the current server in the cluster, you can view the following information:

 - The current server is the primary server
 - The name and host where the primary server is running.
 - The name and host of each secondary server.
 - The current server is a secondary server
 - The name and host where the primary server is running.
 - Startup replication status, whether it was successful, and also how long it took for it to happen.

Important: Click the link in the secondary server name to open the documenter page for this server.

Viewing the server status

Use this procedure to view the information about the current server status in the configuration documenter.

Procedure

1. Open the configuration documenter.
2. Select **Server Status** in the table of contents at the top of the page.

The **Server Status** section contains the following information:

Memory status

Shows the maximum heap size and the current heap size in MB that the Java Virtual Machine, where Netcool/Impact is running, can use.

Service status

Shows the number of events available in the event queues for the various event-related services like readers, listeners, and **EventProcessor**. It does not provide information about all the services that are currently running, only the status for event-related services. For each of these services, you can see from where the service is reading events. For example, for **OMNIBusEventReader** that would include the name of the datasource, whether events are being read from the primary, or backup source of that datasource, and additional

connection-related information like the host, port, and the username that is used to connect to the datasource.

Remember: In the case of the primary server, you can view the queue status for readers or listeners and **EventProcessor**. For a secondary server, you can view only the queue status for **EventProcessor**, because the readers or listeners run only on the primary server.

Viewing data sources

Use this procedure to view the data source details in the configuration documenter.

Procedure

1. Open the configuration documenter.
2. Choose **Data Sources** in the table of contents at the top of the page.
A list of defined data sources displays showing the data source names and data source types.
3. Choose the data source you want to view.
The data source details list displays showing host, port, and database information.

Viewing data types

Use this procedure to view data type details in the configuration documenter.

Procedure

1. Open the configuration documenter.
2. Choose **Data Types** in the table of contents at the top of the page.
3. Choose a data type from the data type list.
You can view the following details about a data type:
 - Field Name
 - Display Name
 - Data source name (for external data types). By clicking the data source name, you can display the connection information.
 - Configuration information for each of the fields in the data type, including the **Field Name**, **Display Name**, **Key field**, **Alias**, **Default Expression**, and **Choices**.
 - Dynamic links associated with the data type
4. To see the connection information for an external data type, click the data source name.

Viewing policies

You can use the configuration documenter to view the policy details.

Procedure

1. Open the configuration documenter.
2. Choose **Policies** in the table of contents at the top of the page.
3. Choose a policy from the Policy list.

Viewing services

Use this procedure to view service details in the configuration documenter.

Procedure

1. Open the configuration documenter.
2. Choose **Services** in the table of contents at the top of the page.
3. Choose the service you want to view from the Services list.

You can use the configuration documenter to view the following information about a service:

- Name
- Class Name
- Run status (running or not running)
- Auto start configuration
- Logging configuration
- Configuration properties

4. Select the associated policy link to see it displayed in the documenter.

Appendix A. Accessibility

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully. These are the major accessibility features you can use with *Netcool/Impact* when accessing it on the *IBM Personal Communications* terminal emulator:

- You can operate all features using the keyboard instead of the mouse.
- You can read text through interaction with assistive technology.
- You can use system settings for font, size, and color for all user interface controls.
- You can magnify what is displayed on your screen.

For more information about viewing PDFs from Adobe, go to the following web site: <http://www.adobe.com/enterprise/accessibility/main.html>

Appendix B. Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement might not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758 U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to

IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information in softcopy form, the photographs and color illustrations might not be displayed.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe, Acrobat, PostScript and all Adobe-based trademarks are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.



Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other product and service names might be trademarks of IBM or other companies.

Glossary

This glossary includes terms and definitions for Netcool/Impact.

The following cross-references are used in this glossary:

- See refers you from a term to a preferred synonym, or from an acronym or abbreviation to the defined full form.
- See also refers you to a related or contrasting term.

To view glossaries for other IBM products, go to www.ibm.com/software/globalization/terminology (opens in new window).

A

assignment operator

An operator that sets or resets a value to a variable. See also operator.

B

Boolean operator

A built-in function that specifies a logical operation of AND, OR or NOT when sets of operations are evaluated. The Boolean operators are &&, || and !. See also operator.

C

command execution manager

The service that manages remote command execution through a function in the policies.

command line manager

The service that manages the command-line interface.

Common Object Request Broker Architecture (CORBA)

An architecture and a specification for distributed object-oriented computing that separates client and server programs with a formal interface definition.

comparison operator

A built-in function that is used to compare two values. The comparison operators are ==, !=, <, >, <= and >=. See also operator.

control structure

A statement block in the policy that is executed when the terms of the control condition are satisfied.

CORBA

See Common Object Request Broker Architecture.

D

database (DB)

A collection of interrelated or independent data items that are stored together to serve one or more applications. See also database server.

database event listener

A service that listens for incoming messages from an SQL database data source and then triggers policies based on the incoming message data.

database event reader

An event reader that monitors an SQL database event source for new and modified events and triggers policies based on the event information. See also event reader.

database server

A software program that uses a database manager to provide database services to other software programs or computers. See also database.

data item

A unit of information to be processed.

data model

An abstract representation of the business data and metadata used in an installation. A data model contains data sources, data types, links, and event sources.

data source

A repository of data to which a federated server can connect and then retrieve data by using wrappers. A data source can contain relational databases, XML files, Excel spreadsheets, table-structured files, or other objects. In a federated system, data sources seem to be a single collective database.

data source adapter (DSA)

A component that allows the application to access data stored in an external source.

data type

An element of a data model that represents a set of data stored in a data source, for example, a table or view in a relational database.

DB See database.

DSA See data source adapter.

dynamic link

An element of a data model that represents a dynamic relationship between data items in data types. See also link.

E

email reader

A service that polls a Post Office Protocol (POP) mail server at intervals for incoming email and then triggers policies based on the incoming email data.

email sender

A service that sends email through an Simple Mail Transfer Protocol (SMTP) mail server.

event An occurrence of significance to a task or system. Events can include completion or failure of an operation, a user action, or the change in state of a process.

event processor

The service responsible for managing events through event reader, event

listener and email reader services. The event processor manages the incoming event queue and is responsible for sending queued events to the policy engine for processing.

event reader

A service that monitors an event source for new, updated, and deleted events, and triggers policies based on the event data. See also database event reader, standard event reader.

event source

A data source that stores and manages events.

exception

A condition or event that cannot be handled by a normal process.

F

field A set of one or more adjacent characters comprising a unit of data in an event or data item.

filter A device or program that separates data, signals, or material in accordance with specified criteria. See also LDAP filter, SQL filter.

function

Any instruction or set of related instructions that performs a specific operation. See also user-defined function.

G

generic event listener

A service that listens to an external data source for incoming events and triggers policies based on the event data.

graphical user interface (GUI)

A computer interface that presents a visual metaphor of a real-world scene, often of a desktop, by combining high-resolution graphics, pointing devices, menu bars and other menus, overlapping windows, icons and the object-action relationship. See also graphical user interface server.

graphical user interface server (GUI server)

A component that serves the web-based graphical user interface to web browsers through HTTP. See also graphical user interface.

GUI See graphical user interface.

GUI server

See graphical user interface server.

H

hibernating policy activator

A service that is responsible for waking hibernating policies.

I

instant messaging reader

A service that listens to external instant messaging servers for messages and triggers policies based on the incoming message data.

instant messaging service

A service that sends instant messages to instant messaging clients through a Jabber server.

IPL See Netcool/Impact policy language.

J

Java Database Connectivity (JDBC)

An industry standard for database-independent connectivity between the Java platform and a wide range of databases. The JDBC interface provides a call level interface for SQL-based and XQuery-based database access.

Java Message Service (JMS)

An application programming interface that provides Java language functions for handling messages.

JDBC See Java Database Connectivity.

JMS See Java Message Service.

JMS data source adapter (JMS DSA)

A data source adapter that sends and receives Java Message Service (JMS) messages.

JMS DSA

See JMS data source adapter.

K

key expression

An expression that specifies the value that one or more key fields in a data item must have in order to be retrieved in the IPL.

key field

A field that uniquely identifies a data item in a data type.

L

LDAP See Lightweight Directory Access Protocol.

LDAP data source adapter (LDAP DSA)

A data source adapter that reads directory data managed by an LDAP server. See also Lightweight Directory Access Protocol.

LDAP DSA

See LDAP data source adapter.

LDAP filter

An expression that is used to select data elements located at a point in an LDAP directory tree. See also filter.

Lightweight Directory Access Protocol (LDAP)

An open protocol that uses TCP/IP to provide access to directories that support an X.500 model and that does not incur the resource requirements of the more complex X.500 Directory Access Protocol (DAP). For example, LDAP can be used to locate people, organizations, and other resources in an Internet or intranet directory. See also LDAP data source adapter.

link An element of a data model that defines a relationship between data types and data items. See also dynamic link, static link.

M

mathematic operator

A built-in function that performs a mathematic operation on two values. The mathematic operators are +, -, *, / and %. See also operator.

mediator DSA

A type of data source adaptor that allows data provided by third-party systems, devices, and applications to be accessed.

N

Netcool/Impact policy language (IPL)

A programming language used to write policies.

O

operator

A built-in function that assigns a value to a variable, performs an operation on a value, or specifies how two values are to be compared in a policy. See also assignment operator, Boolean operator, comparison operator, mathematic operator, string operator.

P

policy A set of rules and actions that are required to be performed when certain events or status conditions occur in an environment.

policy activator

A service that runs a specified policy at intervals that the user defines.

policy engine

A feature that automates the tasks that the user specifies in the policy scripting language.

policy logger

The service that writes messages to the policy log.

POP See Post Office Protocol.

Post Office Protocol (POP)

A protocol that is used for exchanging network mail and accessing mailboxes.

precision event listener

A service that listens to the application for incoming messages and triggers policies based on the message data.

S

security manager

A component that is responsible for authenticating user logins.

self-monitoring service

A service that monitors memory and other status conditions and reports them as events.

server A component that is responsible for maintaining the data model, managing services, and running policies.

service

A runnable sub-component that the user controls from within the graphical user interface (GUI).

Simple Mail Transfer Protocol (SMTP)

An Internet application protocol for transferring mail among users of the Internet.

Simple Network Management Protocol (SNMP)

A set of protocols for monitoring systems and devices in complex networks. Information about managed devices is defined and stored in a Management Information Base (MIB). See also SNMP data source adapter.

SMTP See Simple Mail Transfer Protocol.

SNMP

See Simple Network Management Protocol.

SNMP data source adapter (SNMP DSA)

A data source adapter that allows management information stored by SNMP agents to be set and retrieved. It also allows SNMP traps and notifications to be sent to SNMP managers. See also Simple Network Management Protocol.

SNMP DSA

See SNMP data source adapter.

socket DSA

A data source adaptor that allows information to be exchanged with external applications using a socket server as the brokering agent.

SQL database DSA

A data source adaptor that retrieves information from relational databases and other data sources that provide a public interface through Java Database Connectivity (JDBC). SQL database DSAs also add, modify and delete information stored in these data sources.

SQL filter

An expression that is used to select rows in a database table. The syntax for the filter is similar to the contents of an SQL WHERE clause. See also filter.

standard event reader

A service that monitors a database for new, updated, and deleted events and triggers policies based on the event data. See also event reader.

static link

An element of a data model that defines a static relationship between data items in internal data types. See also link.

string concatenation

In REXX, an operation that joins two characters or strings in the order specified, forming one string whose length is equal to the sum of the lengths of the two characters or strings.

string operator

A built-in function that performs an operation on two strings. See also operator.

U

user-defined function

A custom function that can be used to organize code in a policy. See also function.

V

variable

A representation of a changeable value.

W

web services DSA

A data source adapter that exchanges information with external applications that provide a web services application programming interface (API).

X

XML data source adapter

A data source adapter that reads XML data from strings and files, and reads XML data from web servers over HTTP.

Index

A

- About Netcool/Impact page 4
- absolute time ranges
 - adding 68
- accessibility viii, 187
 - entry fields 5
- accessing reports 167
- Action Efficiency report 169
- Action Error report 169
- action functions 104
- action panel
 - policies 156
- add-ons
 - Maintenance Window Management 179, 181
- auto-saved policy 99
- automated project deployment 20

B

- basic operator view
 - action panel policies 156
 - creating 156
 - information groups 156
 - layout options 155
- books
 - see publications vii, viii
- browser requirements 2

C

- Cache Settings tab
 - External Data Types editor 75
- changing default font
 - See Policy Editor
- character encoding 1
- clear version control file locking 21
- Composite data types 85, 86
- configuration documenter 183
 - cluster status 184
 - opening 183
 - overview 183
 - server status 184
- Configuring a linked field on a composite data type 86
- configuring data sources 162
- configuring data types 163
- conventions
 - typeface xii
- CORBAMediator DSA data sources 55
- creating 93
- Creating an event rule 164
- Creating composite data types 85
- Creating editing and deleting an event rule 163
- creating linked fields 86
- Custom Fields tab
 - internal data types editor 63
- custom policy
 - See policy

- customer support x

D

- daily time ranges
 - adding 67
- data caching 75
- data items 87
 - adding 87
 - deleting 88
 - editing 88
 - overview 31
 - viewing 87
- Data Items editor
 - filtering the view 88
- data model 23
 - components 23
 - task pane icons 24
- data models
 - setting up 23
- data sources
 - categories 25
 - CORBAMediator DSA 55
 - creating 28
 - DB2 34
 - deleting 28
 - DirectMediator DSA 56
 - editing 28
 - Flat File 36
 - GenericSQL 51
 - HSQldb 52
 - Informix 37
 - JMS 57, 58
 - LDAP 54
 - Mediator DSA 55, 56
 - MS_SQL 39
 - MYSQL 40
 - ObjectServer 42
 - ODBC 43
 - Oracle 45, 47
 - overview 25, 33
 - PostgreSQL 48
 - predefined 27
 - SNMPDirectMediator 56
 - v1 and v2 56
 - SQL data source
 - Informix 37
 - SQL database 34
 - Sybase 49
 - testing connections to 28
 - user defined 26
- data type
 - LDAP 76
 - Packed OID 80
 - performance statistics 61
 - SNMP 79
 - table 81
- data type caching 75
- data types 61, 73
 - caching 62
 - caching types 62

- data types (*continued*)
 - categories 29
 - configuring LDAP 77
 - configuring Packed OID SNMPDirectMediator 80
 - configuring SQL 71
 - configuring SQL data types
 - Table Description tab 71
 - configuring table data types for SNMPDirect Mediator 82
 - deleting 30
 - Doc 84
 - editing 30
 - external 29
 - configuring 64
 - deleting a table row 65
 - Flat File 76
 - internal 29, 63
 - internal data types editor 63
 - Mediator DSA 78
 - overview 29
 - predefined 29, 65
 - configuring time range groups 66
 - schedules 68
 - time range groups and schedules
 - overview 66
 - time range groups specifications and combinations 66
 - SNMP 79
 - SQL 70
 - viewing 30
 - viewing performance statistics 61
- DB2 data sources
 - creating 34
- DeployProject
 - parameters 21
- DeployProject policy 21
- directory names
 - notation xii
- disability 187
- Doc data types
 - adding a field 84
 - adding data items 84
- dynamic links 89
 - browsing 93
 - creating 90
 - deleting 92
 - editing 92
 - link by key 91
 - link by policy 92
 - linking methods 90
 - links by filter 90

E

- education
 - See Tivoli technical training
- entry field accessibility
 - using keystroke combinations 5
- environment variables
 - notation xii

- event filter
 - configuration options 123
- Event Isolation and Correlation 161, 162, 163, 164
- Event Isolation and Correlation operator
 - views 162
- Event Isolation and Correlation
 - policies 161
- event listener
 - adding filters 122
 - service 122
- event mapping 122
- event mapping table 123
- event readers
 - configuration 145
- external data type
 - editor 71
- external data types
 - configuring 64
 - configuring SQL 71
 - editor 65, 73
 - LDAP 77
 - Mediator DSA 78
 - Pack OID SNMPDirectMediator 80
 - table DirectMediator 82

F

- FailedEvent
 - overview of data types 84
 - viewing data items 84
- failover 33
 - configurations 33
- filter
 - for event listener services 122
- filters
 - analysis 124
 - delete 124
 - editing 124
 - reordering 124
- fixes
 - obtaining ix
- Flat File
 - creating data type 76
- Flat File data sources
 - creating data sources 36
- functions
 - action 104

G

- GenericSQL data sources
 - creating 51
- getting started 6
- Global project
 - editing and deleting items 16
- global repository
 - adding an item to 17
 - clearing version control locking 21
 - deleting an item from 17
 - overview 17
 - viewing data 17
- globalization 1
- glossary 193
- Graphical User Interface
 - overview 1

- GUI
 - See Graphical User Interface

H

- hibernating policy activator
 - configuration 128
- Hibernation data types
 - overview 85
- HSQldb data sources
 - creating 52

I

- Impact Profile report 170
- Impact ROI Efficiency report 173
 - business processes 174
 - scenario 174
- Informix data sources
 - creating 37
- internal data types
 - editor
 - Custom Fields tab 63
- IPL functions 104
- IPv6 2
- ITNM DSA data type 69

J

- jabber service
 - configuring AIM transport
 - account 131
 - configuring MSN transport
 - account 132
- Jabber service
 - adding resources to the Jabber
 - ID 129
 - configuring 130
 - configuring ICQ transport
 - account 132
 - configuring transport accounts 131
 - configuring Yahoo transport
 - account 132
- JMS
 - data source 57, 58

K

- key expressions 91

L

- Launching the Event Isolation and Correlation analysis page 165
- LDAP 47
- LDAP data sources
 - creating 54
- LDAP External Data Type editor
 - LDAP Info tab 77
- LDAP external data types 77
- Link by Key 91
- Link Editor 93
- links 89
 - categories 31
 - dynamic 89, 90

- links (*continued*)
 - overview 31
 - static 89, 93
- LinkType data items
 - configuring 84
- LinkType Data Type
 - overview 83
- logging on
 - prerequisites 2
- logon 2

M

- main tabs 4
- maintenance schedules 66
- manuals
 - see publications vii, viii
- Mediator DSA
 - CORBA data sources 55
 - data sources 55, 56
 - data types 78
 - Direct Mediator data sources 56
 - SNMPDirectMediator data
 - sources 56
 - viewing data types 78
- MS-SQL Server data sources
 - creating 39
- multiple policy logs 135
- MWM
 - See Maintenance Window Management
- MySQL data sources
 - creating 40

N

- navigation 4
- navigation panel
 - selecting clusters 6
- nci_policy script 99
- negative time range groups 66
- Netcool/Impact components 4
- Node Efficiency report 176
- notation
 - environment variables xii
 - path names xii
 - typeface xii

O

- ObjectServer data sources
 - creating 42
- ODBC data sources
 - creating 43
- online publications
 - accessing viii
- Operator Efficiency report 176
- operator view
 - advanced 154
 - basic 154
 - components 155
 - controls 155
 - deleting 158
 - modifying 157
 - opening in TIP 158
 - types 154

- operator view (*continued*)
 - viewing 153
- operator view EIC_Analyze 166
- operator views 153
 - overview 153
- Oracle data source
 - connecting over LDAP 47
 - creating 45, 47
 - integration with RAC cluster 47
- ordering publications viii
- override time range group 66
- overview 1
- Overview 85, 161

P

- Packed OID SNMPDirectMediator data types
 - configuring 80
- path names
 - notation xii
- Performance Statistics report
 - for data types 61
- personalizing 119
- policies
 - accessing 95
 - working with 95
- Policies 98
- policies overview 95
- policy
 - auto saved 99
 - custom 98
 - deleting 98
 - DeployProject 21
 - developing custom policy 96
 - editing 98
 - graphical view 102
 - log files 134, 135
 - optimizing policy 102
 - predefined 112
 - recovering 99
 - syntax checking 101
 - task pane icons 96
 - uploading 111
 - version control interface 111
 - viewing 95
 - wizard 98
 - wizards 96
 - writing 96
- policy activators
 - configuration 150
- policy editor 99
 - personalizing 110
- Policy Editor
 - browsing data types 103
 - changing default font 111
 - graphical view 102
 - optimizing policy 102
 - run policy option 102
 - setting runtime parameters 103
 - toolbar controls 99
- Policy Efficiency report 177
- Policy Error report 177
- policy logger
 - configuration 133
- policy runtime parameter
 - attributes 104
- policy syntax highlighter 101
- positive time range groups 66
- PostgreSQL data sources
 - creating 48
- predefined data items
 - adding absolute time range groups 68
 - adding daily time range groups 67
 - adding weekly time range groups 67
- predefined data types
 - configuring time range groups 66
 - Doc 84
 - FailedEvent overview 84
 - Hibernation 85
 - Linktype 83
 - LinkType data items
 - configuring 84
 - overview 29, 65
 - schedules 68
 - configuring schedules 68
 - time range groups
 - specifications and combinations 66
 - time range groups and schedules
 - overview 66
 - viewing FailedEvent data items 84
- predefined policy
 - See* policy
- problem determination and resolution xi
- projects
 - automated project deployment 20
 - cluster 6
 - components 16
 - creating 18
 - deleting 20
 - DeployProject policy 21
 - editing 19
 - editing and removing 16
 - editor configuration window 19
 - overview 15
 - project menu icons 18
 - viewing project members 19
 - working with 15
- publications vii
 - accessing online viii
 - ordering viii

Q

- query caching 75

R

- RAC Cluster Support 47
- recovering
 - auto-saved policy 99
- report
 - Impact Profile 170
- reports 167
 - Action efficiency 169
 - Action Error 169
 - Impact Profile 170
 - Impact ROI Efficiency 173
 - navigating 167
 - Node Efficiency 176
 - Operator Efficiency 176

- reports (*continued*)
 - Policy Efficiency 177
 - Policy Error 177
 - toolbar 168
 - viewing 167
- run policy option
 - See* Policy Editor
- runtime parameters 103

S

- schedules
 - adding 68
 - configuring 68
 - overview 68
- selecting projects
 - overview 6
- service
 - command execution manager 125
 - command line manager 125
 - database event listener 125
 - database event reader 137, 138, 139
 - e-mail sender 126
 - event listener 141
 - event processor 126, 127
 - hibernating policy activator 128
 - ITNM event listener 135
 - Jabber reader 143
 - jabber service 129
 - JMS message listener 141
 - OMNIBus event listener 144
 - OMNIBus event reader 145, 146, 147, 148
 - policy activator 149, 150
 - policy logger 133
 - self monitoring 136, 137
 - Web Services Notification Listener 150, 151
- service log 121
- service log viewer 120
 - creating new tabs 122
 - results 121
- Service Status panel
 - service icons 116
 - status icons 116
- services 119
 - accessing 115
 - configuring 119
 - displaying log files 120
 - e-mail reader 140
 - list 117
 - overview 115
 - starting 120
 - stopping 120
 - working with 115
- setting runtime parameters
 - See* Policy Editor
- SNMP data types
 - configuring 79
- SNMP DSA
 - data sources 34
- SNMPDirect Mediator v3 data sources 56
- SNMPDirectMediator
 - data sources 56
 - v1 and v2 56
- SNMPDirectMediator data sources 56

- SNMPDirectMediator data types
 - packed OID 80
 - table 82
- socket DSA
 - data source 34
- Software Support
 - contacting x
 - overview ix
 - receiving weekly updates ix
- SQL data sources
 - DB2 34
 - flat file 36
 - GenericSQL 51
 - HSQldb 52
 - MS-SQL Server 39
 - MySQL 40
 - ObjectServer 42
 - ODBC 43
 - Oracle 45, 47
 - PostgreSQL 48
 - Sybase 49
- SQL data types
 - adding a field to the table 73
 - configuring 71
 - deleting a table row 65
 - flat file 76
- SQL database DSAs
 - failover 33
 - failover configurations 33
- static links 89, 93
 - creating an Internal data type 93
 - Link editor 93
- Sybase data sources
 - creating 49
- Sybase data types
 - Setting the Exclude this field option 73

T

- Table Description tab
 - SQL External Data Types editor 71
- table OID SNMPDirectMediator data types
 - configuring 82
- TBSM specific projects 16
- time range groups 66
 - absolute 68
 - configuring 66
 - daily 67
 - specifications and combinations 66
 - weekly 67
- TIP
 - opening operator view 158
- Tivoli Information Center viii
- Tivoli technical training viii
- training
 - Tivoli technical viii
- typeface conventions xii
- types browser
 - browsing data types
 - See Policy Editor
 - Policy Editor 103

U

- user-defined services
 - creating 119
 - deleting 120

V

- variables
 - notation for xii
- version control
 - file locking 21
- version control interface
 - See policy
- viewing
 - data sources 185
 - data types 185
 - policies 185
 - services 186
- Viewing Event Isolation and Correlation results 165, 166

W

- WebGUI 165
- weekly time ranges
 - adding 67



Printed in USA

SC23-8830-04

