




*Netcool Operations Insight 1.2*

---

**Example Guide**



Licensed Materials – Property of IBM



---

**Note:** Before using this information and the product it supports, read the information in “Notices” at the end of this document.

---

© Copyright IBM Corporation 2014.

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

## Contents

<b>About this publication .....</b>	<b>iv</b>
Intended audience .....	iv
What this publication contains.....	iv
Conventions used in this publication.....	iv
Typeface conventions .....	iv
Operating system-dependent variables and paths.....	v
<b>Introduction .....</b>	<b>1</b>
A fresh new look.....	1
Using the analytics tools reactively and proactively.....	2
<b>Example scenarios using the analytics tools .....</b>	<b>3</b>
Netcool Operations Insight as a reactive triage tool .....	3
Netcool Operations Insight as a proactive operational efficiency tool .....	5
<b>Building dashboards .....</b>	<b>9</b>
Using Netcool/Impact to feed a DASH widget.....	9
Check the connection to the IBM Netcool/Impact server .....	10
Create an Impact policy to fetch the data .....	11
Add and configure a DASH chart widget .....	13
Trademarks.....	17

## About this publication

This guide provides an overview of IBM Netcool Operations Insight 1.2 providing a worked example scenario taken from the IBM Pulse 2014 demo.

---

## Intended audience

This publication is intended for anyone interested in learning about IBM Netcool Operations Insight particularly around how it can be used in typical real world scenarios.

---

## What this publication contains

This publication contains the following sections:

- *Introduction* on page 2:  
Provides an overview of what IBM Netcool Operations Insight is and what this document contains.
  - *Example scenarios using the analytics tools* on page 3:  
Runs through some of the scenarios used in the Pulse 2014 demo to illustrate how the analytics tools can be used both as a triage tool for operational staff and also as a diagnosis tool for achieving operational efficiencies.
  - *Building dashboards* on page 9:  
Provides an introduction on how to populate DASH widgets with data from both the IBM Netcool/OMNIBUS Web GUI datasource as well as how to use IBM Netcool/Impact to retrieve data from a remote datasource such as a database and format the data for consumption by a DASH widget.
- 

## Conventions used in this publication

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

### 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. For example, on UNIX systems, the `$NCHOME` environment variable specifies the directory where the Network Manager core components are installed. On Windows systems, the same environment variable is `%NCHOME%`. The names of environment variables are not always the same in the Windows and UNIX environments. For example, `%TEMP%` in Windows environments is equivalent to `$TMPDIR` in UNIX environments.

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

## Introduction

IBM Netcool Operations Insight helps unify, reduce and prioritize your business service impacting events by using real-time and historical analytics. With the integration of these leading edge IBM analytics technologies, IBM Netcool Operations Insight powered by OMNIbus helps IT operations better understand their IT operations and infrastructure management:

- How can I reduce my event storm noise?
- How can I quickly identify and resolve problems from service impacting events?
- How can I move beyond reactive problem remediation to identify service impacting events before they occur?
- What part of my infrastructure can be more efficient and less costly?

IBM Netcool Operations Insight helps you:

- Improve operational agility with real time analytics;
- Accelerate problem identification and resolution with search analytics;
- Improve operational efficiency with historical analytics.

**Note:** For more information about IBM Netcool Operations Insight including prices and availability, please see the following page: <http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=AN&subtype=CA&htmlfid=897/ENUS214-071&appname=USN>

## A fresh new look

In addition to the new analytic capabilities, Netcool Operations Insight brings a new look and feel to the dashboard layer. This new look and feel is due to IBM's new Dashboard Application Services Hub (DASH) technology that powers the all new user interface.

As well as having a diverse array of new widgets, the integration possibilities are endless thanks to IBM Netcool/Impact - a part of Netcool Operations Insight - that can feed DASH widgets data from a wide variety of sources.



Netcool Operations Insight 1.2 Example Guide Introduction  
© Copyright IBM Corporation 2014.

---

## Using the analytics tools reactively and proactively

As this guide will illustrate, IBM Netcool Operations Insight can be used reactively as a triage tool to improve operational agility - to "*find the needle in the haystack*" - and also proactively for improving operational efficiency.

In the former capacity as a reactive triage tool, IBM Netcool Operations Insight comes with preconfigured tools that can be used to perform keyword searches from pertinent events on the entire event archive for the specified time windows. As the examples will show, this is extremely powerful and can rapidly bring key information to light.

In the latter capacity as a proactive analysis tool, IBM Netcool Operations Insight can be used to identify "low hanging fruit" - for example: identifying sets of events that create the most "noise" and yet add the least value. Such events are prime candidates for ObjectServer housekeeping initiatives. By reducing event noise from operator screens, operators will be able to pinpoint the real issues sooner and not waste time dealing with events that don't require operator involvement. This in turn can save a great deal of time for operations teams, which can be directly translated into cost savings.

## Example scenarios using the analytics tools

This section outlines two main scenarios in the context of IBM Netcool Operations Insight that highlight some of the capabilities and value that the analytics capabilities bring to the table. Both scenarios are based on real world activities that were carried out on customer sites and yielded tremendous results - like the ones contained herein.

### Netcool Operations Insight as a reactive triage tool

A common issue that operations face is when there is a failure on a core component that has a large downstream affect in terms of affected components and services. Operators will typically receive a huge flurry of events - often referred to as an "event storm" - and often trying to find the root cause of the outage can be likened to trying to find the proverbial "needle in a haystack".

One part of IBM Netcool Operations Insight is as an analysis tool for doing just that. All events that are received into Netcool/OMNIbus are sent to the Smart Cloud Analytics - Log Analysis (SCA-LA) engine where the historical event history can then be analysed.

#### EXAMPLE

Widgetcom have a number of large data centres in the United States. An automated process makes a change to the configuration of one of the ports on a core switch causing it to malfunction. As a result, a large number of monitored systems and services suddenly go offline. This results in around 2,000 critical outage events hitting the operator screens.



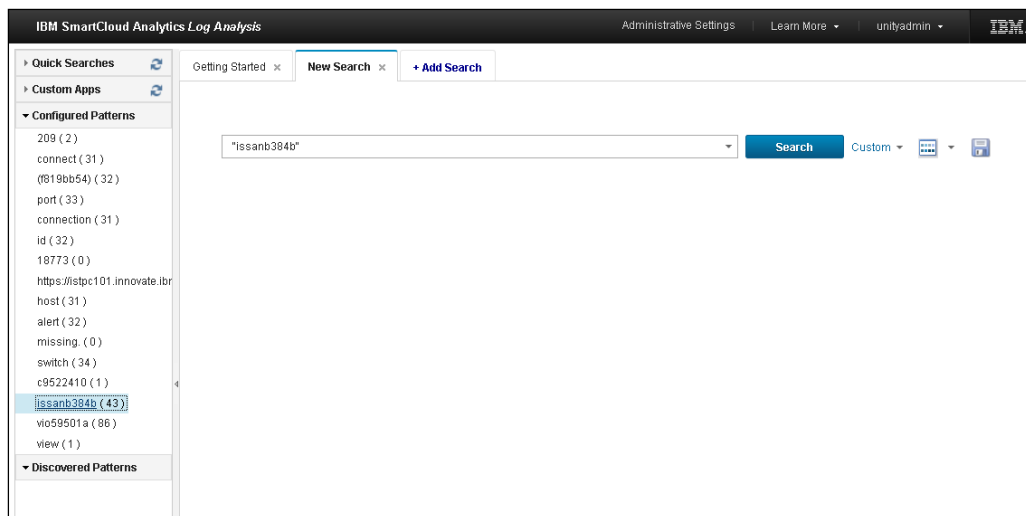
The operator seeing this in their Event Viewer decides to use the Log Analysis tool to perform a keyword search on a number of the events. A keyword search is one of the out-of-the-box tools that comes with IBM Netcool Operations Insight and allows a user to select an



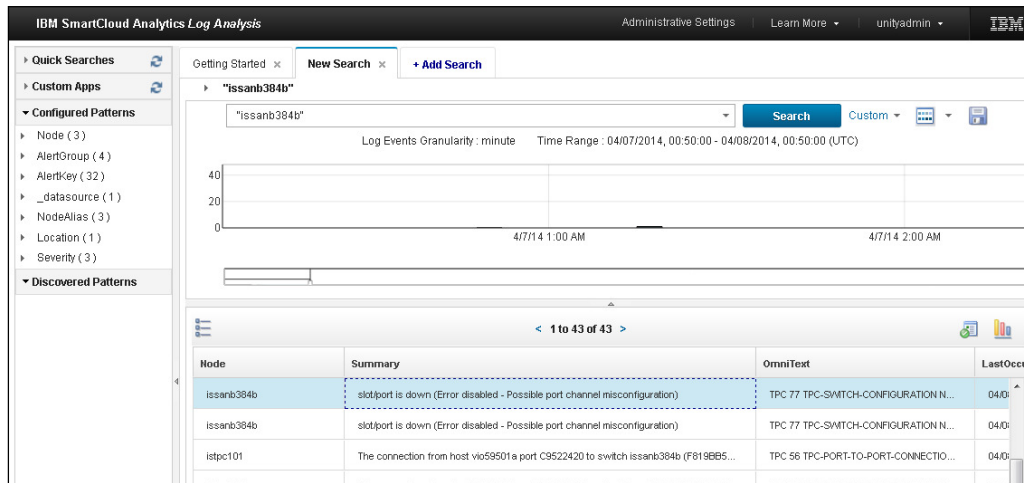
event or number of events and then compare those events with the entire event history for the specified time window. It does this by taking the words contained in the Node, AlertGroup and Summary fields and then seeing how many times each of those words occurs across the event history for the specified time window.



On running the Log Analysis tool, the operator sees a number of hits. Some of the keyword hits are not so useful - such as the keywords "connect", "connection" or "host" - but, by applying some tribal knowledge, the operator notices the hostname of one of their core routers "issanb384b" has occurred many times.



The operator can then select the core router hostname and further drill down to see the events that contain that hostname. By carrying out this further search, the operator is immediately presented with events showing that there is a port channel misconfiguration on the core router concerned.



It is immediately clear to the operator at this point that such a misconfiguration is a major problem. The operator therefore immediately dispatches an engineer to correct this problem. With the misconfiguration corrected, the monitored devices and services all auto-clear and the event storm is halted.

--

This example demonstrates how an operator who would typically have taken hours to manually go through many hundreds of events has instead quickly found the root cause of the major outage within a few mouse clicks. The power of being able to perform in-context keyword searches of the entire event history within a specified time window using the Log Analysis tool allows an operator to distil down, summarise and make sense of vast quantities of event data.

## Netcool Operations Insight as a proactive operational efficiency tool

In addition to being a powerful triage tool to improve operation agility, Netcool Operations Insight also provides an analyst the ability to identify ways to make operational efficiencies.

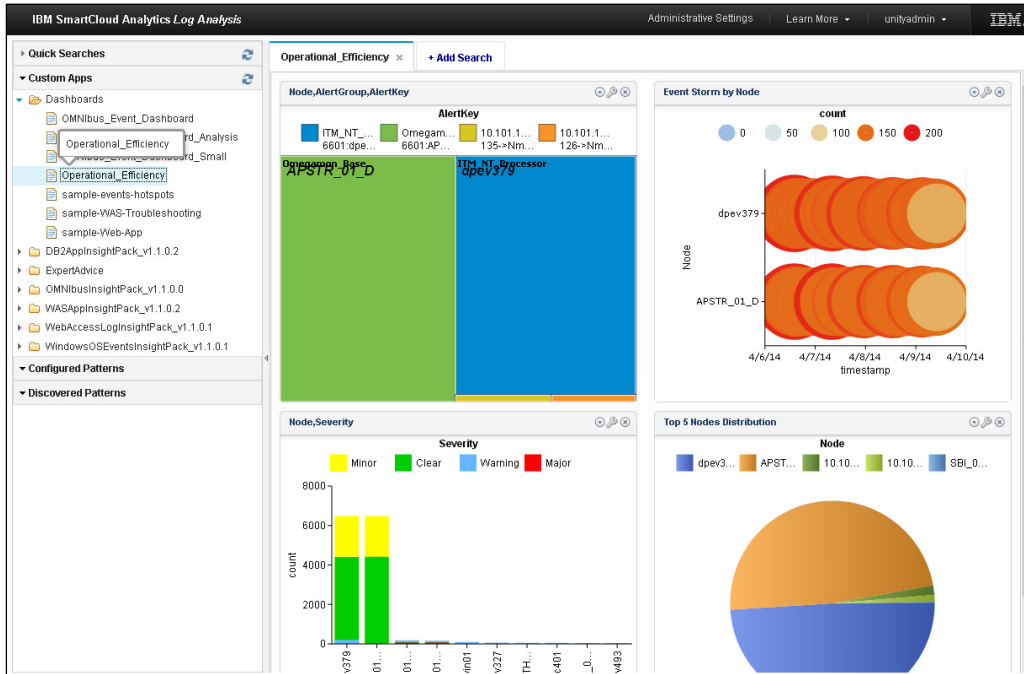
First, the Log Analysis tool allows a user to quickly construct dashboards that summarise the entire dataset in different ways. Second, the Seasonality functionality can be used to identify chronic issues within the dataset - that is, events that occur with some degree of regularity.

The very fact that events are seasonal in conjunction to the nature of the seasonality can quickly guide an analyst to the root cause of the problem. In addition to this, an operator solely working from an Event List would likely not notice that any given event occurs with any sort of discernable pattern since they are only viewing the current event set which is not in the context of the event history. Only a tool that has visibility to the entire data set would be able to spot such patterns. An operator would likely not be able to make sense of such vast quantities of data needed to spot such patterns.

**EXAMPLE**

Widgetcom's analyst is trying to reduce the overall event numbers and has been tasked with trying to identify potential candidate sets of events that could be dropped from the Netcool/OMNIBus ObjectServer.

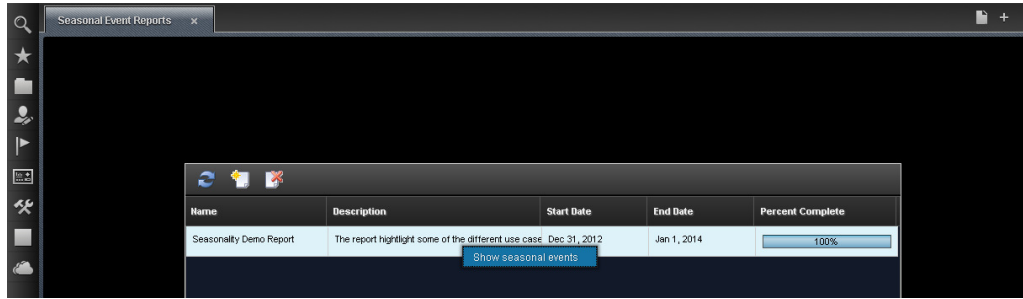
The first place to look is in the Log Analysis tool. The analyst runs a preconfigured *Operational Efficiency* dashboard showing different views of the historical event data to see what information can be gleaned from there.



The first thing the analyst sees is that the large majority of events are coming from just two nodes: *APSTR\_01\_D* and *dpev379*. Second, the large majority of those events are either clear or minor. This being the case, the analyst decides to investigate further the types of events that are being received from these two nodes. Either the minor issues being experienced by these two nodes could simply be fixed - or these events could be suppressed until such time as the issues are fixed.

By taking this action, the analyst was able to drastically reduce the event set which, in turn, dramatically cleaned up the operator screens due to the reduced event numbers being displayed.

Next the analyst turns their attention to the seasonality reports.



The analyst runs a seasonality report for the event set for the last month and discovers around a dozen events that the seasonality algorithm has identified as being 100% seasonal. Having a 100% seasonality quotient means that some aspect of the recurrent nature of the event occurs with predictable regularity.

The first event in the report showing 100% seasonality is a storage event whereby a storage array is reporting low storage availability.



By drilling down into this event, the analyst sees that this event always occurs at 3:41pm.



This leads the analyst to suspect an automated process is running at this time and is causing the issue. Since the event is a storage event, the analyst first looks at the automated backup system. Sure enough, there is a backup job that runs every day at 3:41pm. By simply provisioning more space for this backup job, this problem stops happening. This in turn prevents events being generated, trouble tickets being raised - saving operations time and money.

By using the traditional event based working methods, by the time investigations were made into any given problem, it may have stopped occurring. This means that the root cause of the problem is sometimes never deduced. In this example, the very fact that the event was seasonal was the biggest clue to its root cause.

This example shows how the seasonality analytics tool within Netcool Operations Insight enabled the analyst to identify the chronic recurring problems within *Widgetcom's* environment - enabling the organisation to make operational efficiencies, saving time and money.

## Building dashboards

As touched on in the introductory section, DASH boasts an impressive array of powerful and versatile widgets from which sophisticated dashboards can be built - such as those used in the Pulse 2014 demo.



The data that feeds DASH widgets comes from data providers. A data provider is, as the name suggests, a source of data that a DASH widget can be configured to load data from. The IBM Netcool/OMNIBus Web GUI product documentation lists the data sets from Web GUI that can be used by each type of DASH widget including tables, lists, charts, volume bars and gauges.

[http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.netcool\\_OMNIBus.doc.7.4.0/webtop/wip/reference/web\\_con\\_dash\\_compatibilitywithtwlwidgets.html](http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.netcool_OMNIBus.doc.7.4.0/webtop/wip/reference/web_con_dash_compatibilitywithtwlwidgets.html)

## Using Netcool/Impact to feed a DASH widget

IBM Netcool Operations Insight also comes with IBM Netcool/Impact. IBM Netcool/Impact is a powerful data processing engine that can access data from a myriad of different sources and then (among other things) be used to pull data from remote datasources, process the data and then send the data to DASH widgets in a consumable format.

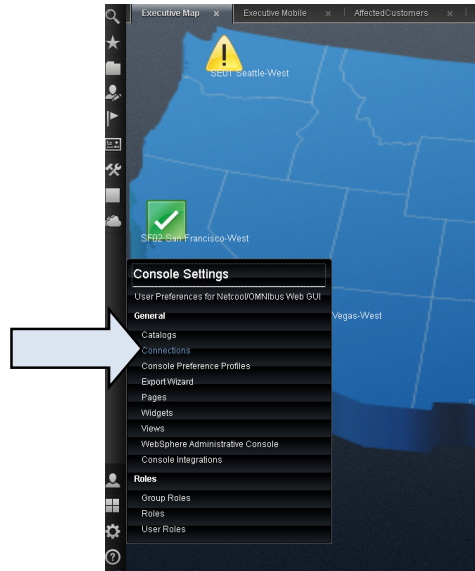
The data format DASH widgets are expecting to receive are typically a series of name/value pairs. Hence you can use IBM Netcool/Impact to read data from the source data source, store those values in an array, and then return the array to the DASH widget. This section outlines how this is done.

### EXAMPLE

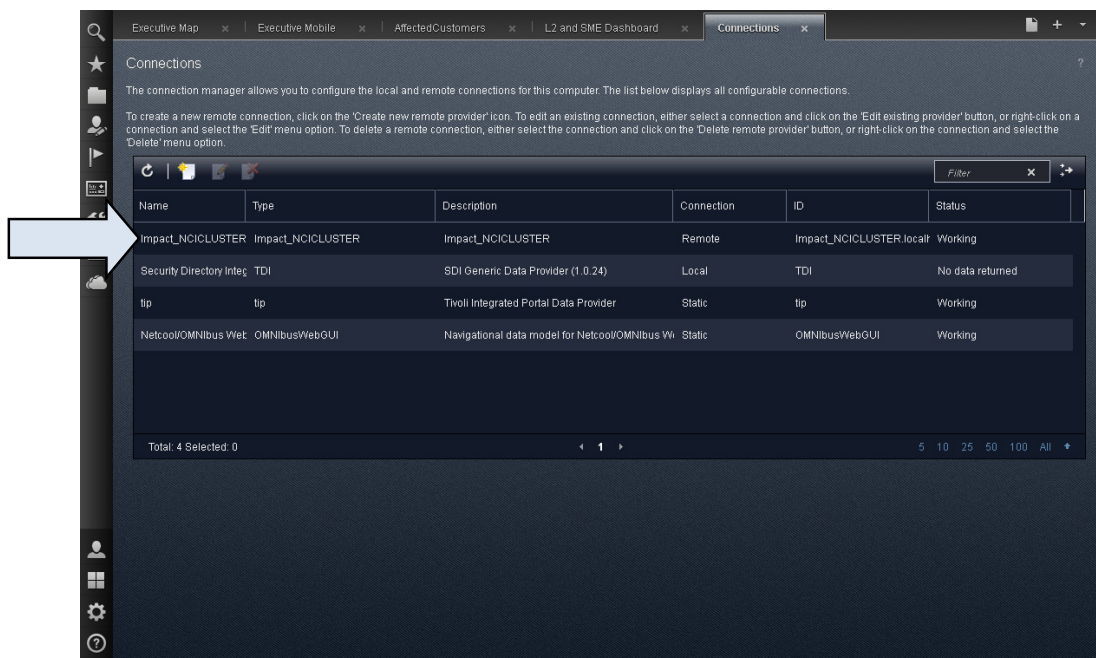
In this example, we will use IBM Netcool/Impact to pull data from a database table, store the pertinent data in an array and then configure DASH to connect to IBM Netcool/Impact, pull the data and then use it to populate a DASH chart widget.

## Check the connection to the IBM Netcool/Impact server

The first step in getting IBM Netcool/Impact to feed data to a DASH widget is to ensure the connection between DASH and Impact is correctly configured and operational. This is done by clicking on the "Connections" option under the "Console Settings" menu.



The Connections dialogue box lists the connections that the DASH server has to. The connection to the IBM Netcool/Impact server should be shown in the list and the Status should be "Working" as in the following screen shot.



If this connection is not present, you will need to add a new connection to your IBM Netcool/Impact server. This can be done by clicking on the "Create new remote provider" button and entering the host/port details for the IBM Netcool/Impact server.

If the connection is present but not working, it may be that the IBM Netcool/Impact server is down or that the host/port information is not correct.

## Create an Impact policy to fetch the data

The next step is to construct an IBM Netcool/Impact policy to fetch the data from the remote data source. The following example is taken from the Pulse 2014 demo. It fetches a set of data from a database called "KVMDCNETS" and stores the fields "NETWORK", "OS" and "NOV" into the array fields "Name", "Status" and "NOV" respectively from each returned row. The array is called "dataCenterNetworks".

**Note:** You can store as many fields as you like into the array even if they will not all be used by the target DASH widget. You will specify which fields to use when you configure the DASH widget. In this way a single policy could potentially be used to populate multiple different DASH widgets all using different parts of the data available in the array.

The screenshot shows the IBM Tivoli Netcool/Impact 7.1 interface. The 'Policies' tab is active, and a list of policies is shown on the left. The 'DataCenterNetworks' policy is selected. The right pane displays the source code for this policy, which is a JavaScript script designed to fetch data from a database and store it in an array.

```

Log("DataCenterNetworks: Reporting network usage for datacenter");
// CREATE ARRAY TO STORE NAME VALUE PAIRS
dataCenterNetworks = {};

// FETCH DATA FROM DATABASE
result=GetByFilter("KVMDCNETS", "DATACENTER='DC01-DC'", false);

// ITERATE OVER RETRIEVED ROWS AND STORE IN ARRAY
count = 0;
while (count < Length(result)) {
    network = NewObject();
    network.Name = result[count].NETWORK;
    network.Status = result[count].OS;
    network.NOV = Int(result[count].NOV);
    dataCenterNetworks = dataCenterNetworks + network;
    count = count + 1;
}


Log(dataCenterNetworks);
Log("DataCenterNetworks: exiting policy");
  
```

By the time the policy has completed its execution, the array "dataCenterNetworks" will contain the collection of name/value pairs needed to populate the DASH chart widget.

Before exiting the policy, an output parameter needs to be defined in order to pass the array back to the component that has called the policy - in this case the DASH chart widget.



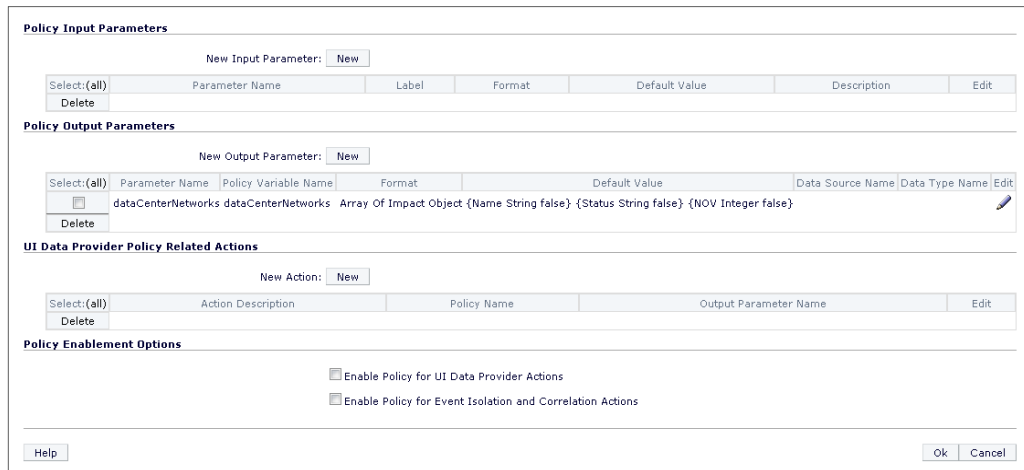
To create an output parameter, click on the "Configure Policy Settings" button as highlighted below in red:



The resulting popup dialogue box gives you the option to add output parameters. In this example, we have created an output parameter by the name "dataCenterNetworks" which will map to the policy variable of the same name.

**Note:** It is essential that the variable name is correct otherwise it will return a null value.

The format you need to configure for the output parameter is "Array Of Impact Object" since the variable in the policy is an array. An array of Impact objects is essentially an array of name/value pairs.




Click on "Ok" to save the parameters and then save your IBM Netcool/Impact policy.


We have now configured the policy to retrieve the data from the database, store the name/value pairs into an array and then pass the array back to whatever calls it within the variable named "dataCenterNetworks".

## Add and configure a DASH chart widget

The last step is to add a DASH chart widget to your DASH page and configure the data set it will use. Drag the Line Chart widget onto your page and then select "Edit" from the pull down menu on the newly added widget. The next step is to define a dataset for the widget to use. Enter the name of your newly created IBM Netcool/Impact policy - in this case: "DataCenterNetworks" - and click "Search".

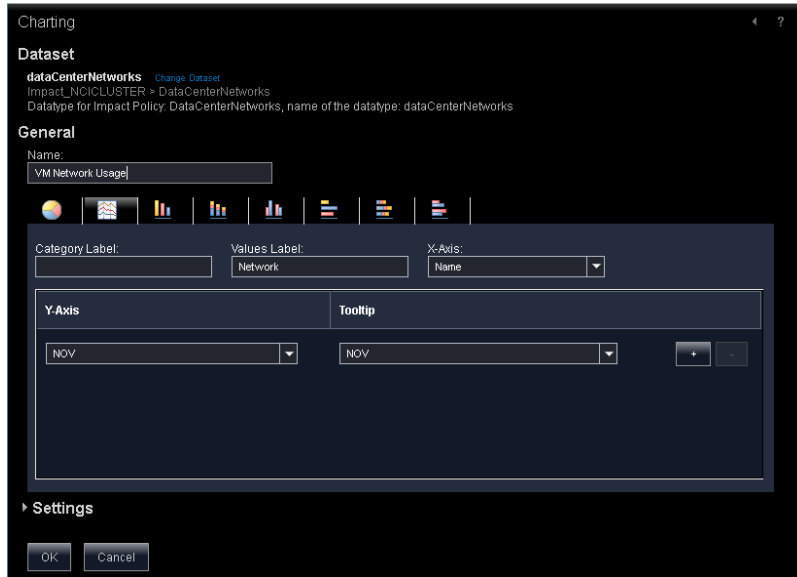


The search will return the found datasource "DataCenterNetworks" which is your IBM Netcool/Impact policy. The datasource information shows that the available datatype for this datasource is called: "dataCenterNetworks". This is the output parameter that was created within the policy - which points to the array variable of the same name.



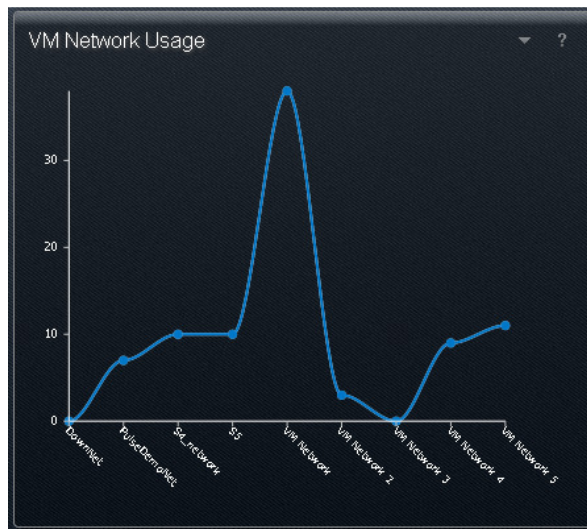
After having selected the datatype "dataCenterNetworks", you will be prompted by the DASH Line Chart widget configuration dialogue box to enter the parameters for the chart.

In this example, we will give the chart the name "VM Network Usage", use the "Name" field for the X axis and the corresponding values in the "NOV" field for the Y axis.



Once the values have been added, click "OK" to save the configuration, and then "Save and Exit" to save the page configuration.

When the page loads, the DASH chart widget will be populated with the data IBM Netcool/Impact has pulled from the database.



By applying this example to the other DASH widgets, you can use IBM Netcool/Impact to create rich and diverse dashboards within the DASH framework.

## Notices

This information was developed for products and services offered in the U.S.A.

IBM® may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation  
Licensing  
2-31 Roppongi 3-chome, Minato-ku  
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation  
958/NH04  
IBM Centre, St Leonards  
601 Pacific Hwy  
St Leonards, NSW, 2069  
Australia

IBM Corporation  
896471/H128B  
76 Upper Ground  
London  
SE1 9PZ  
United Kingdom

IBM Corporation  
JBF1/SOM1 294  
Route 100  
Somers, NY, 10589-0100  
United States of America

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

---

## Trademarks

IBM, the IBM logo and [ibm.com](http://ibm.com) are 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 [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.



Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Adobe, Acrobat, Portable Document Format (PDF), PostScript, and all Adobe-based trademarks are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.



Licensed Materials – Property of IBM