Redbooks
ibm.com/redbooks

IBM Netcool Operations Insight
A Scenarios Guide

Zane Bray
Rob Clark
Jeff Ditto
Manzoor Farid
Vasfi Gucer
Ahmed A Saleh
Maciej Olejniczak
Lanny Short
Steven Shuman

Analytics

IBM

Redbooks

International Technical Support Organization

IBM Netcool Operations Insight: A Scenarios Guide

July 2016

SG24-8352-00

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

First Edition (July 2016)

This edition applies to IBM Netcool Operations Insight Version 1.4.

Contents

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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, MD-NC119, Armonk, NY 10504-1785, US

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 jurisdictions 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

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

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

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

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises 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. The sample programs are provided “AS IS”, without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, 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 http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

DB2®
developerWorks®
IBM®
IBM SmartCloud™
Jazz™
Micromuse®
Netcool®
Redbooks®
Redbooks (logo)  ®
Tivoli®

The following terms are trademarks of other companies:

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

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

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

Find and read thousands of IBM Redbooks publications

- Search, bookmark, save and organize favorites
- Get personalized notifications of new content
- Link to the latest Redbooks blogs and videos

Get the latest version of the Redbooks Mobile App

iOS **Download Now** Android

Promote your business in an IBM Redbooks publication

Place a Sponsorship Promotion in an IBM® Redbooks® publication, featuring your business or solution with a link to your web site.

Qualified IBM Business Partners may place a full page promotion in the most popular Redbooks publications. Imagine the power of being seen by users who download millions of Redbooks publications each year!

ibm.com/Redbooks
About Redbooks → Business Partner Programs

THIS PAGE INTENTIONALLY LEFT BLANK

Preface

IBM® Netcool® Operations Insight empowers your IT operations to use real-time and historical analytics to identify, isolate, and resolve problems before they affect your business. Powered by IBM Tivoli® Netcool/OMNIbus and the transformative capabilities of cognitive analytics, Netcool Operations Insight consolidates millions of alerts from across local, cloud, and hybrid environments into a few actionable problems.

This IBM Redbooks® publication gives a broad understanding of Netcool Operations Insight and describes several scenarios that show the capabilities of this solution in a real-life environment. Each scenario features a different capability of Netcool Operations Insight. The scenarios are documented by using step-by-step figures with explanations to make them easier to implement in your own environment.

The scenarios in this book are broken into the following categories:

- Network Management-related scenarios
- Network Event and cognitive-related scenarios
- Network Event-related scenarios

The target audience of this book is network specialists, network administrators, and network operators.

Authors

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, Austin Center.

Zane Bray started at Micromuse® (pre-IBM acquisition) in 2001, working in Technical Support, Services, and then the Development lab. In this most recent role, he published successive versions of the Netcool Best Practices guides and documents that covered a range of topics that relate to the installation, upgrade, and solutions guides for Netcool products, including Netcool Operations Insight. He has also developed several out-of-the-box solutions and components including: Standard Multitier Architecture Configuration, X in Y for Netcool/OMNIbus, Netcool Health self-monitoring, and the new scope-based event grouping capability. He regularly speaks at user conferences that cover various related topics.

Rob Clark is a senior member of the IBM Tivoli Network Manager development team. He has 19 years of development and design experience at IBM in Network Management software, including product development, third-party integration, and customer consultation. His focus has been in the areas of time to value in customer deployments, best practices, software quality, and user experience. Most recently, Rob has been engaged in UX and IBM Design Thinking with customers and business partners.

Jeff Ditto is a Managing Consultant for the IBM Cloud - IT Service Management (ITSM) Lab Services Practice. Jeff brings a broad and significant experience base and perspective to the IBM ITSM practice with almost 20 years of experience in the Network Management Systems (NMS), Business Service Management (BSM), and Operation Support Systems (OSS). Having worked with the Netcool suite since 1997, Jeff has considerable development and operations experience across the breadth of the product line. Jeff joined IBM in November 2006 after working as a Network Management Systems (NMS) architect and development lead in the Federal government contracting arena. Jeff also has significant experience implementing and integrating vendor agnostic solutions in the NMS, BSM, and OSS space.

Manzoor Farid is an Architect for IBM's Cloud Middleware DevOps IT and leads the team in deploying and integrating IBM IT Service Management products in Cloud Software development environments. Manzoor is IT Infrastructure Library (ITIL) Foundation-certified and is a Certified Information Systems Security Professional (CISSP). He shares his team's experiences with implementing IT Service Management, Operations Analytics, and Security products at various customer demo engagements and at IBM Interconnect.

Vasfi Gucer is an IBM Redbooks Project Leader with the IBM International Technical Support Organization. He has more than 18 years of experience in the areas of systems management, networking hardware, and software. He writes extensively and teaches IBM classes worldwide about IBM products. His focus has been on cloud computing for the last three years. Vasfi is also an IBM Certified Senior IT Specialist, Project Management Professional (PMP), IT Infrastructure Library (ITIL) V2 Manager, and ITIL V3 Expert.

Maciej Olejniczak is a Cross-functional Software Support Team Leader in a collaborative environment. He works internationally with external and internal clients, IBM Business Partners, services, labs, and research teams. He is a dedicated account advocate for large customers in Poland. Maciej is an IBM Certified Expert in Actualizing IT Solutions: Software Enablement. He achieved a master level in implementing all activities that transform information technology from a vision to an actual working solution. Maciej is an Open Group Master Certified IT Specialist.

Ahmed A Saleh is an IBM certified Specialist. Ahmed joined IBM in 2005 as an application developer and is now part of Cairo Lab Services. Before joining IBM, Ahmed worked in developing Operation and Administration software for Class 5 switches and other software that is related to Telco industry. Ahmed is delivering services for Netcool suite of products: OMNIbus, Network Manager, and Configuration Manager.

Lanny Short is a member of the IBM Tivoli ISM Practice Team. He has 14 years of client consulting experience at IBM and 30 years totals experience with IBM. He is also experienced working with product development and product support for various IBM ISM solutions. His focus has been in the areas of time to value in customer deployments, best practices, software quality, and user experience. Most recently, Lanny has been engaged with customers who are deploying IBM's Application Performance Management solutions and with business partners who are working with customers to gain value from their Tivoli ISM product deployments.

Steve Shuman started out as an electrical and acoustics engineer. He now works to solve network configuration management issues at companies around the globe as the Netcool Configuration Manager architect, which is part of IBM Systems/Middleware Group. He has over 20 years of software, IT, and networking experience with expertise in network device operations and configuration. Steve has been associated with the Netcool Configuration Manager product and its customers from customer #1.

Thanks to the following people for their contributions to this project:

Jim Carey, John Griffith, Christopher Haynes, James Moore, Mark Simpson
IBM USA

Olivier Bonnet, Neil Bradbury, Jenny Li Kam Wa, Simon White
IBM UK

Shane O'Rourke, Tom Randles
IBM Ireland

Joakim Tenlen
IBM Sweden

Fred Harald Klein, Bert Holtwick
IBM Germany

Tobias Bautze
DICOS GmbH Kommunikationssysteme

Also, we want to express special thanks to **Simon Knights** and **Kliment Stephanov** from IBM UK for creating the following scenarios and providing technical guidance throughout the project:

- Known slow traffic between two points in a network
- Flood event detection
- Operator suspects an event is re-occurring and wants to investigate

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  ibm.com/redbooks
- Send your comments in an email to:

  redbooks@us.ibm.com
- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks
- Follow us on Twitter:

  http://twitter.com/ibmredbooks
- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806
- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm
- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Part 1

Introduction

In this part, the basic features and components of IBM Netcool Operations Insight are described.

1

IBM Netcool Operations Insight overview

In this chapter, we describe at a high level the features of the IBM Netcool Operations Insight (Netcool Operations Insight) solution.

This chapter includes the following topics:

1.1 Netcool Operations Insight at-a-glance

IBM Netcool Operations Insight uses real-time alarm and alert analytics, which are combined with broader historic data analytics. Netcool Operations Insight is powered by the fault management capabilities of IBM Tivoli Netcool/OMNIbus and IBM's leading big data technologies within IBM Operations Analytics - Log Analysis, which provides powerful event search and historical analysis in a single solution.

Netcool Operations Insight consists of a base solution for managing and analyzing application monitoring environments and an optional extension that is called Networks for Operations Insight, which widens the scope to include network discovery, visualization, event correlation, topology-based root-cause analysis, and configuration and compliance management. The Networks for Operations Insight capability is provided through the Network Manager and Netcool Configuration Manager products.

In addition, you can set up IBM Network Performance Insight as part of your Netcool Operations Insight solution to monitor network traffic performance. You also can integrate with other solutions, such as IBM Alert Notification and IBM Runbook Automation.

Netcool Operations Insight integrates infrastructure and operations management into a single coherent structure across business applications, virtualized servers, network devices and protocols, internet protocols, and security and storage devices.

Netcool Operations Insight includes the following capabilities:

- Event search

  Event search applies the search and analysis capabilities of Operations Analytics - Log Analysis to events that are monitored, recorded, and managed by Tivoli Netcool/OMNIbus. Events are transferred from the ObjectServer through the Gateway for Message Bus to Operations Analytics - Log Analysis, where they are brought into a data source and indexed for searching. After the events are indexed, you can search every occurrence of real-time and historical events.

  The Tivoli Netcool/OMNIbus Insight Pack is installed into Operations Analytics - Log Analysis and provides custom applications that search the events based on various criteria. The custom applications can generate dashboards that present event information to show how your monitoring environment is performing over time. Keyword searches and dynamic drill-down functions allow you to go deeper into the event data for detailed analysis. The applications can be run from the Operations Analytics - Log Analysis.

  Tooling can be installed into the web GUI that starts the applications from the right-click menus of the Event Viewer and the Active Event List. An "event reduction wizard" is also supplied that includes information and applications that can help you analyze and reduce volumes of events and minimize the "noise" in your monitored environment.

- Event analytics

  Event Analytics performs statistical analysis of Tivoli Netcool/OMNIbus historical event data. It can identify seasonal patterns, such as when and how frequently events occur. Seasonality analyses detects seasonal pattern in your events and shows them in reports and graphs. For example, an event that periodically occurs at an unscheduled specific time is highlighted. You can use the information from the seasonality reports to create suppression rules for network equipment or devices to reduce the number of events.

It can determine which events have a statistical tendency to occur together and output the results on a scheduled basis as event groups. You can deploy valid event groups as Netcool/Impact correlation rules. The rules act on the event data and show a single parent event from the event group, with all other events in the group as children. This ability reduces the number of events that are presented to operators.

Event Analytics is installed as two separate packages. One is installed into Netcool/Impact and the other is installed into the Netcool/OMNIbus web GUI. Both are required for Event Analytics to work.

- Networks for Operations Insight

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

- Topology search

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

- IBM Connections integration

  Netcool/Impact enables social collaboration through IBM Connections by automatically providing updates to key stake holders. It provides integration to IBM Connections by using a Netcool/Impact IBM Connections action function. The IBM Connections action function allows users to query forums and topics lists, create a forum or topic, and update topics.

  IBM Connections is a leading social software platform that can help your organization to engage the right people, accelerate innovation, and deliver results. This integrated, security-rich platform helps people engage with networks of experts in the context of critical business processes. Now, everyone can act with confidence and anticipate and respond to emerging opportunities.

- Network performance monitoring

  Network Performance Insight is a flow-based network traffic performance monitoring system. It provides comprehensive and scalable visibility on network traffic with visualization and reporting of network performance data for complex, multivendor, multi-technology networks.

- IBM Alert Notification

  IBM Alert Notification provides instant notification of alerts for any critical IT issues across multiple monitoring tools. It gives IT staff instant notification of alerts for any issues in your IT operations environment.

- IBM Runbook Automation

  IBM Runbook Automation empowers IT operations teams to resolve detected problems in a more efficient and effective way. Operators can focus their attention where it is needed and receive guidance about the best resolution with recommended actions and pre-filled context. With Runbook Automation, you can perform the following tasks:

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

Note: For more information about IBM Alert Notification and IBM Runbook Automation, see *Delivering Consistency and Automation with Operational Runbooks,* REDP-5347.

1.2 Netcool Operations Insight in IT Service Management context

Netcool Operations Insight provides end-to-end insight for smarter business decision and simplifies operations and reduces cost of operations. Netcool Operations Insight also provides services management personnel with improved visibility by introducing Network Health Dashboard, Network discovery, visualization, monitoring, and event correlation and root-cause analysis, which reduces mean-time to repair.

Netcool Operations Insight give IT Service Management personnel agile operations by providing the following features (see Figure 1-1 on page 7):

- Consolidated management:
  - Provide off-the-shelf integrations for rapid deployment
  - Correlate, enrich, and consolidate events into a single view
  - Scale from the smallest to largest environments
- Analytics:
  - Provide data-driven actionable insight from high-volume operations data
  - Identify frequent events
  - Recognize patterns
  - Recommend grouping and suppression
- Automation:
  - Streamline operations
  - Eliminate manual steps by automating actions
  - Automate routine processes
  - Take immediate action
- Built-in expertise:
  - Decades of IBM experience across thousands of customers
  - Reduction in actionable events ready for use

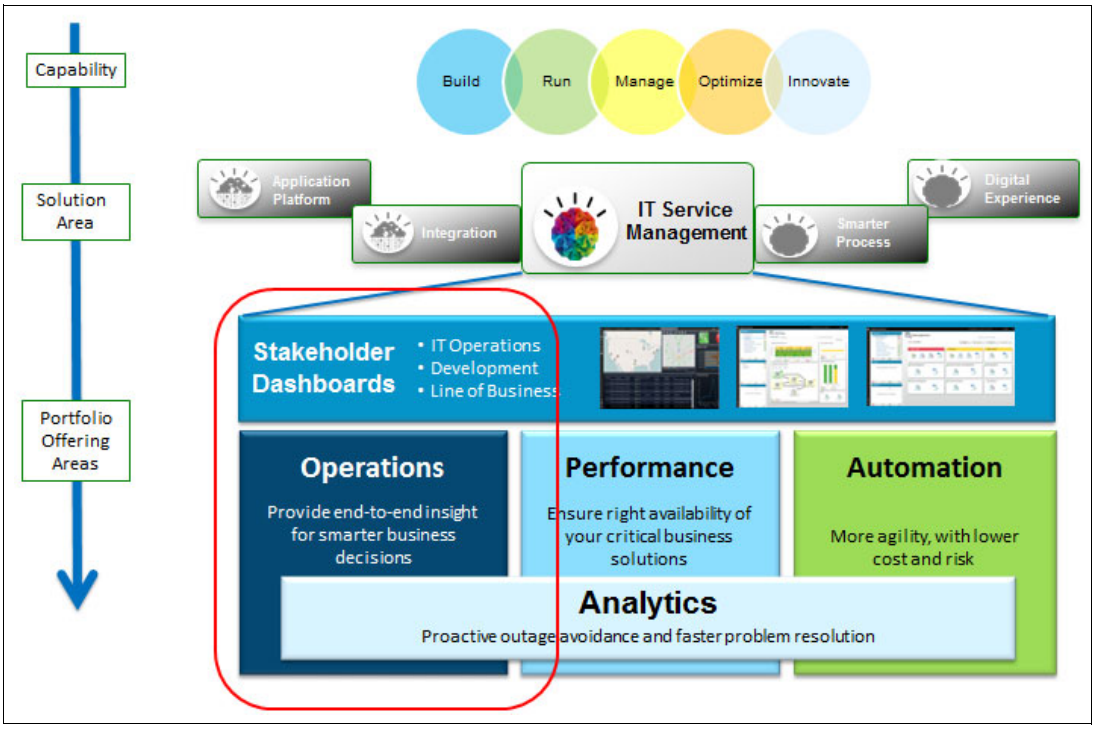

Figure 1-1 Netcool Operations Insight in IT Service Management context

1.3 Netcool Operations Insight Dashboards Services Hub

Netcool Operations Insight combines Netcool OMNIbus and Netcool Impact With the power of analytics and modern mobile dashboards to increase the effectiveness, efficiency, and reliability of Operations Management. It also simplifies or removes administrative tasks, which improves time to value and total cost of ownership.

IBM Dashboard Application Services Hub (DASH) service is a common web portal for IBM Netcool suite. When you log in to the DASH, you can access the Content Page, OMNIBus, Web GUI, Network Manager topology views, discovery configuration, and other integrations.

1.3.1 Navigation bar

By using the navigation bar, you can access the Search, Favorites, and Product pages navigation tools (see Figure 1-2 on page 8). Product pages provide access to DASH pages for each integrated product.

Figure 1-2 DASH Navigation bar

Click any folder's icon (a folder is a logical grouping for related pages and you can assign any icon to it) to browse to its content pages. We start by clicking the Default folder icon to see its contents. The folder comes empty, but as shown in Figure 1-3, it shows our custom sample page that named "austin map" under it.

Figure 1-3 Austin map

As shown in Figure 1-4, opening the Austin map page shows you a map for Austin that uses the Map Widget that is configured to use OpenStreetMap (OSM) service that is provided by a free tile server. (This feature requires internet connectivity if you do not have a replica for the server on-premises).

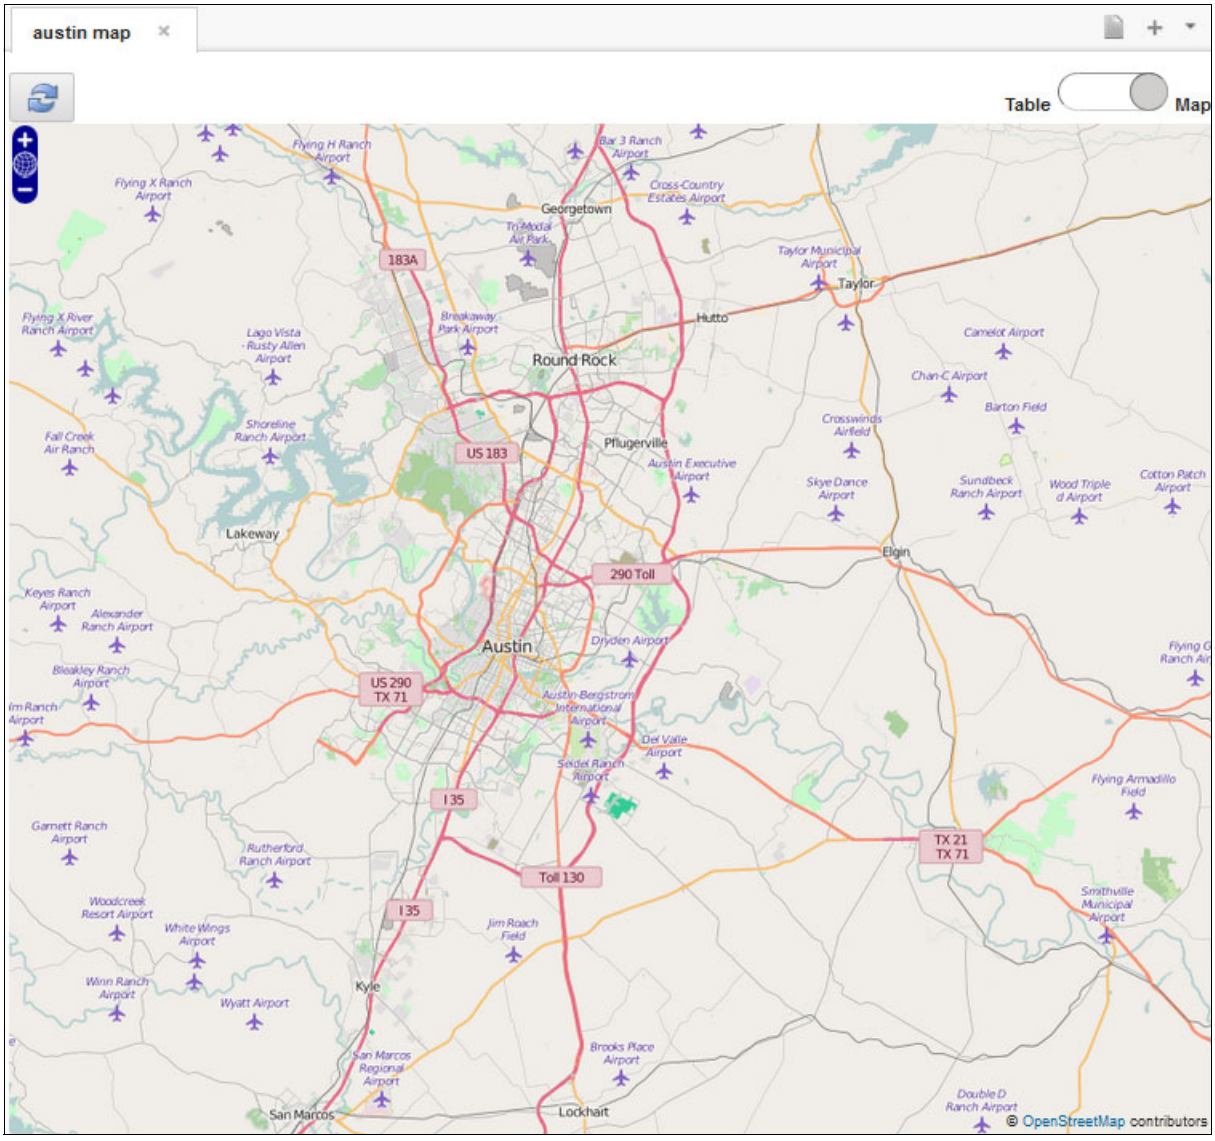

Figure 1-4 Map widget

This feature is helpful in visualizing topology over maps. Figure 1-5 shows devices that are laid over Google maps. For more information, see Chapter 3, “Geographic Discovery and Mapping” on page 67.

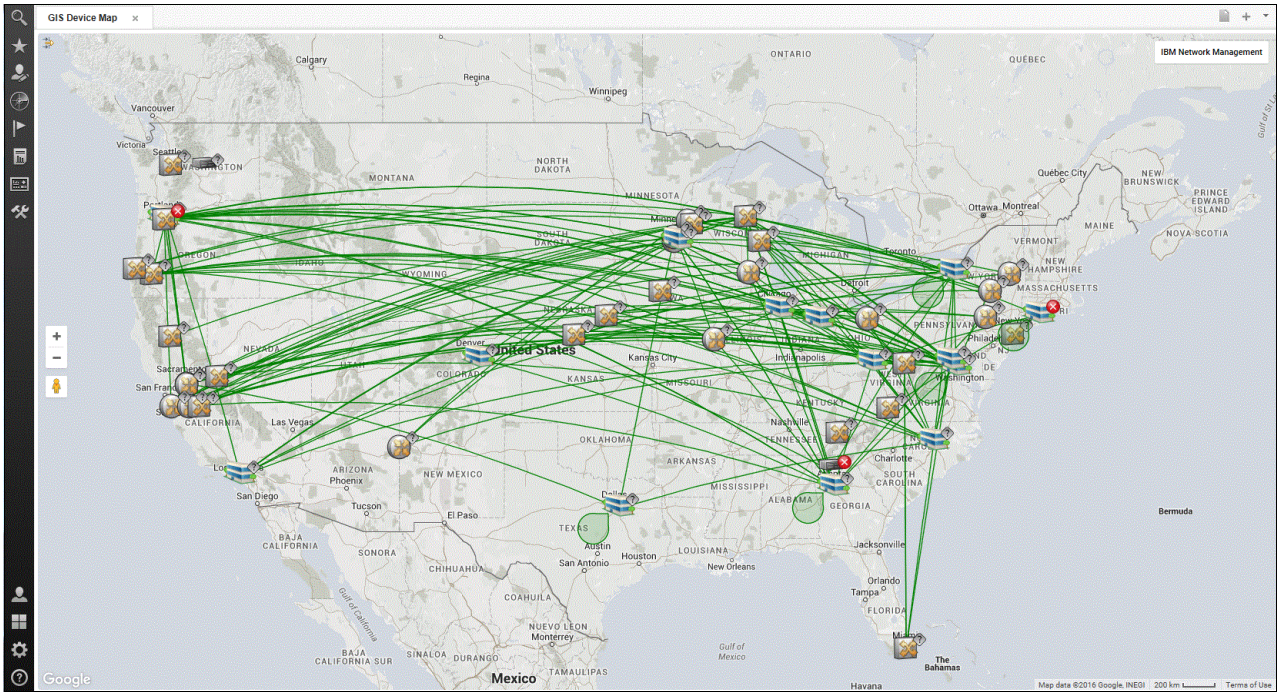

Figure 1-5 Topology visualization over map

1.3.2 Administration folder

The **Administration** folder contains the Netcool OMNIBus WebGUI and Network Manager administration GUI. Use the GUI in this folder to create OMNIBus filters, views, tool, menus, and Network Manager polling policies, as shown in Figure 1-6.

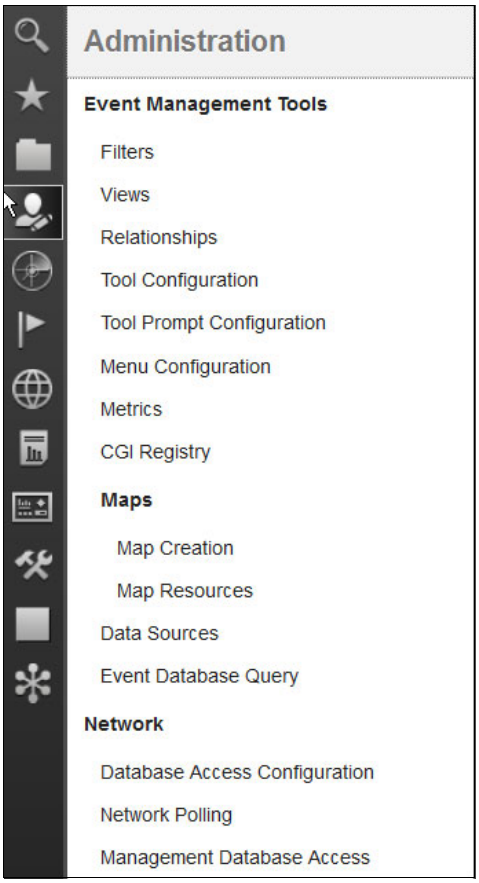

Figure 1-6 Administration folder.

1.3.3 Discovery folder

The Discovery folder opens links to the Network Discovery Status and Network Discovery Configuration pages, as shown in Figure 1-7.

Figure 1-7 Network Discovery status and configuration pages

Starting in version V1.4, some enhancements were added, such as the discovery of Cisco WiFi Access Point, which performs modeling for layer 2/3, SSID, 802.11 spec, channels, dependencies on DHCP. See Figure 1-8.

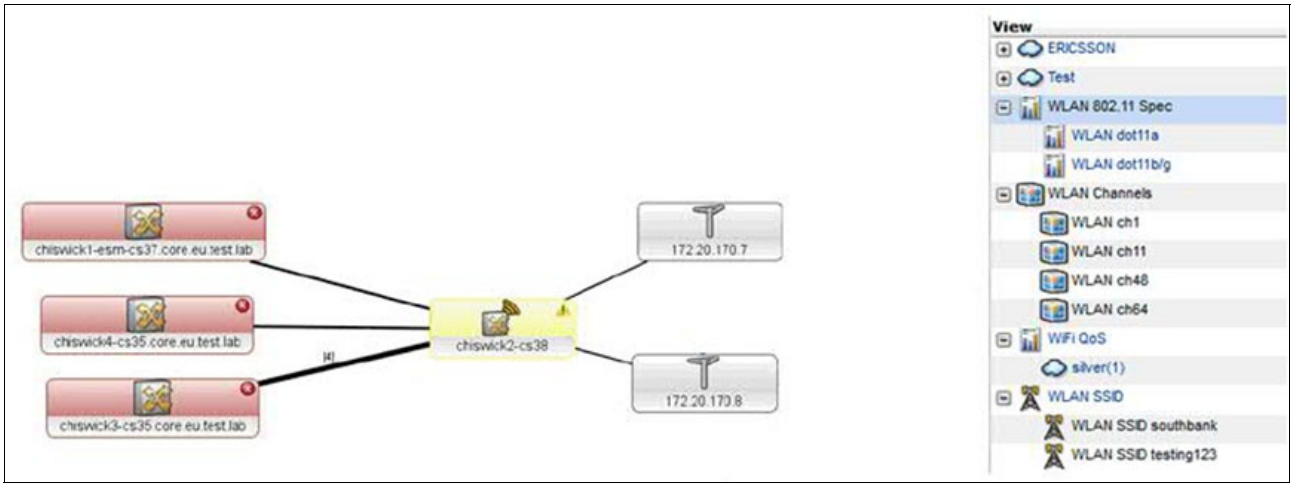

Figure 1-8 Cisco WiFi discovery

1.3.4 Incident folder

Clicking the Incident folder opens OMNIbus WebGUI-related views and Network views. You can browse the Event Dashboard, Event List, AEL, and work with OMNIBus events. You also can browse the Network views, Hop view, Health view, Fault-Finding view, SNMP MIB Browser, and real-time graphing, as shown in Figure 1-9 on page 13.

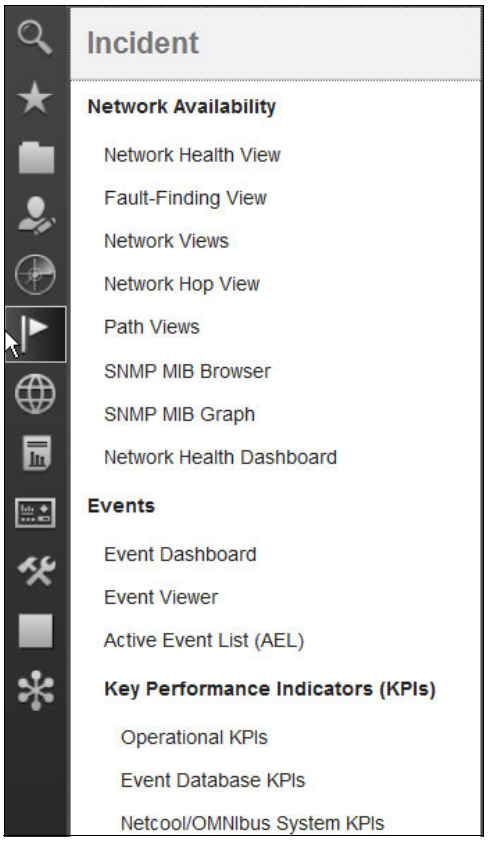

Figure 1-9 Events and Network Availability views

1.3.5 Network Health Dashboard

The Network Health Dashboard is a new feature in this release. As shown in Figure 1-10 on page 14, this feature answers the following questions at a glance:

- What devices or interfaces have been down longer than 1 hour, 24 hours?
- How is the availability level trending over the last 24 hours?
- What are the worst performers?
- Did any configuration changes coincide with this incident?
- What events are active?

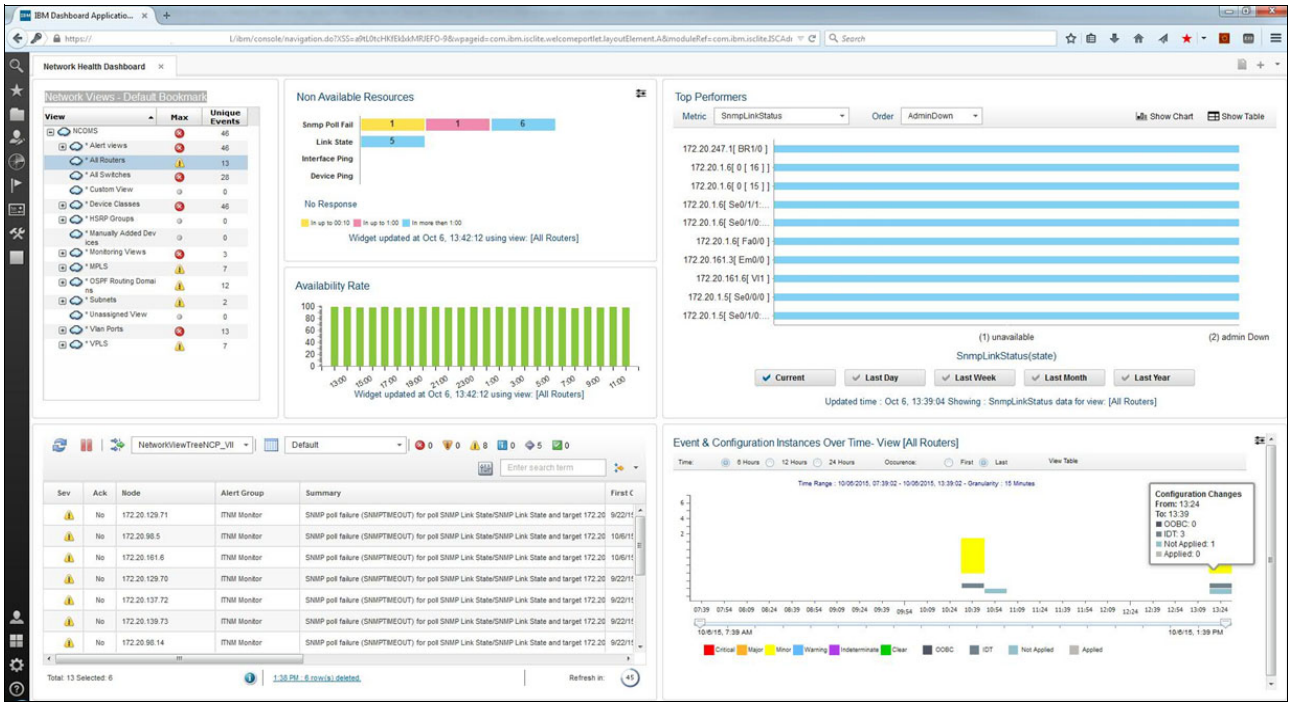

Figure 1-10 Network Health Dashboard

Unavailable Resources and Percentage Availability widgets

The Unavailable Resources and Percentage Availability widgets give more information about your network availability status in real time. They also gives you information about device pingability and SNMP polling failures, interface link state, and pingability.

You use the Network View bookmark to segment your network; for example, if you want to see router availability only, click the **All Routers** bookmark, as shown in Figure 1-11 on page 15.

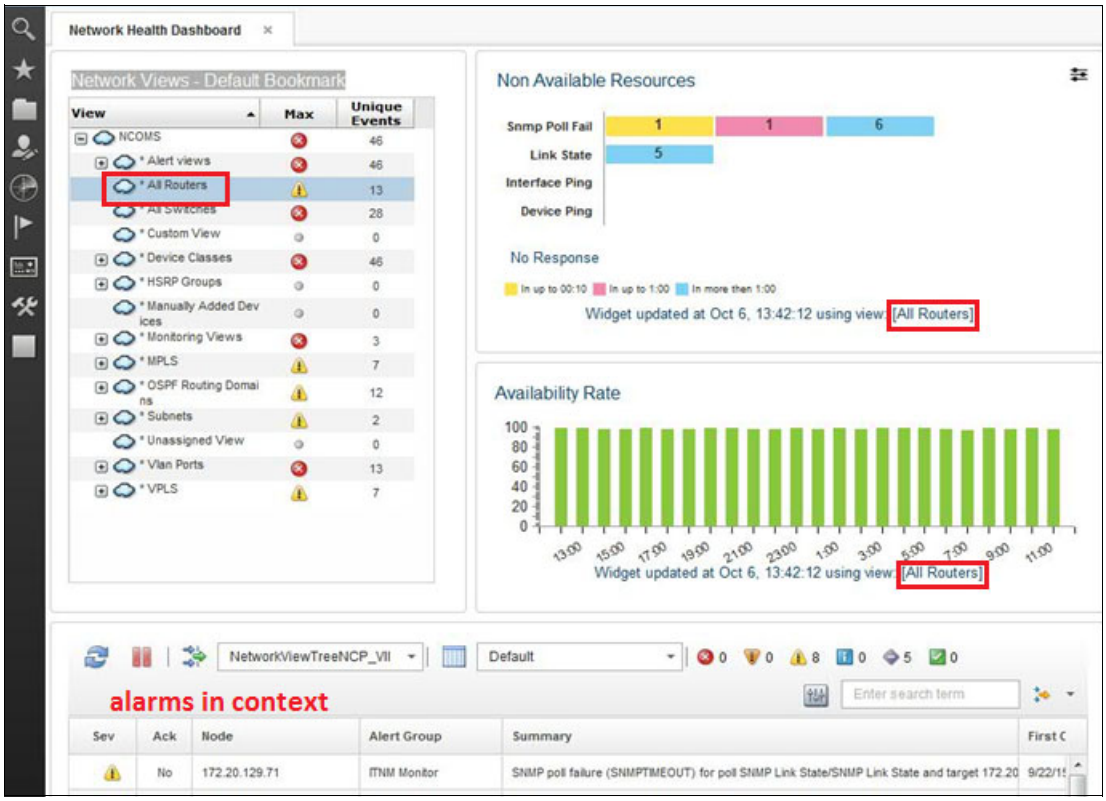

Figure 1-11 Selecting All Routers bookmark

1.3.6 Insights folder

Click the **Insights** folder to open the following analytics-related pages:

- Seasonal Events
- Related Events

Both pages are empty after installation and must be configured. It is in this folder where statistical analysis of Tivoli Netcool/OMNIbus historical events data is done. It can identify seasonal patterns, such as when and how frequently events occur, as shown in Figure 1-12.

Figure 1-12 Netcool Operations Insight Analytics

Seasonality analyses are output in reports and graphs so that you can find reoccurring event patterns, as shown in Figure 1-13.

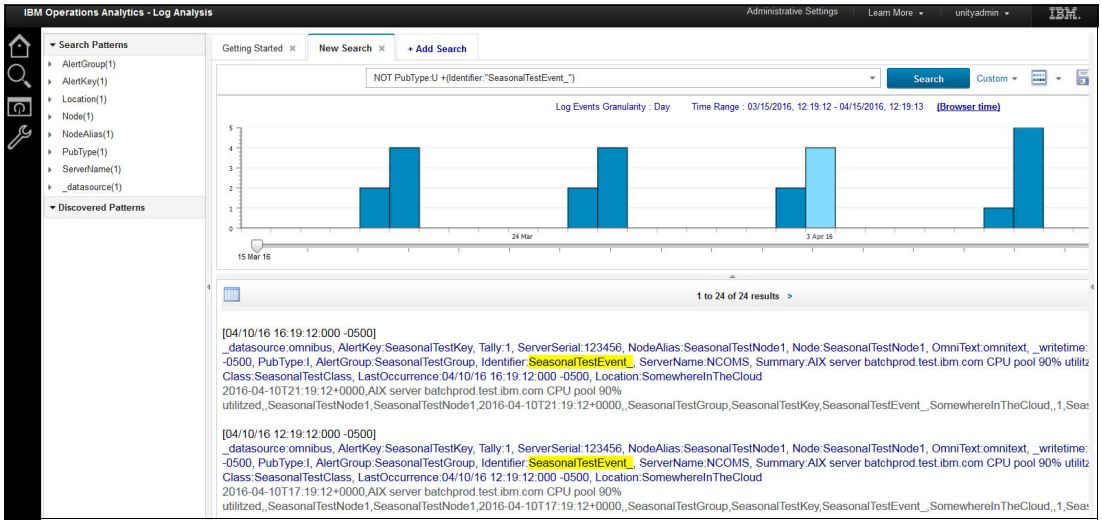

Figure 1-13 Seasonality analysis in Netcool Operations Insight

1.3.7 Reporting folder

Click the **Reporting** folder to start browsing through the out-of-box reports that are installed and included with Netcool Operations Insight components. Reports are built on top of IBM Tivoli Common Reporting engine (see Figure 1-14).

Figure 1-14 Reporting link

You get OMNIBus historical reports as part of OMNIBus installation, which is used by the analytics engine. You get approximately 55 reports for Tivoli Network Manager if you install the product. For Tivoli Netcool Configuration Manager, you get 15 different reports about compliance and security. Moreover, a customized report is also possible by using the Report Studio tool.

The installed package reports are shown in Figure 1-15.

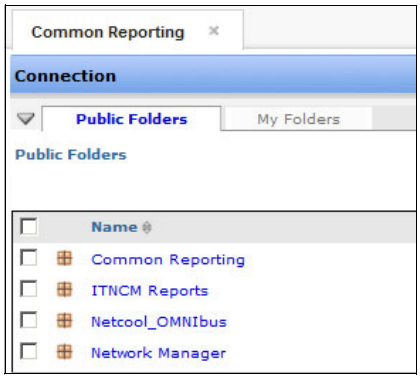

Figure 1-15 Out-of-box reporting

1.3.8 Configurations folder

Click the **Configurations** folder icon to open Netcool Configuration Manager ITNCM base and compliance GUI, as shown in Figure 1-16. The JNLP application is started when you click any one of those applications.

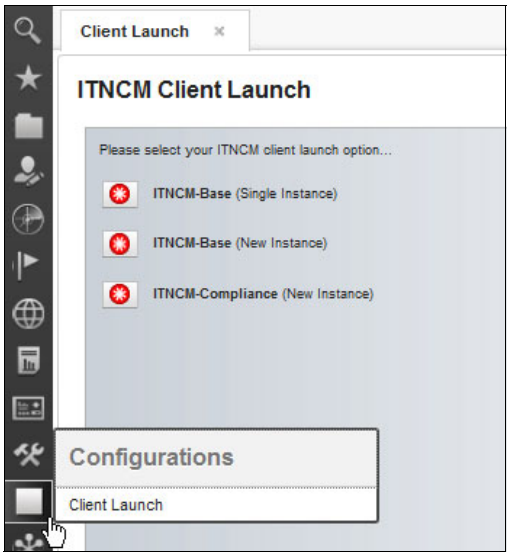

Figure 1-16 Netcool Configuration Manager client launch

Click the **Console Integrations** icon to open any integrated product. Based on your installation, you might have Netcool Impact only, or have Netcool Impact and NPI (see Figure 1-17).

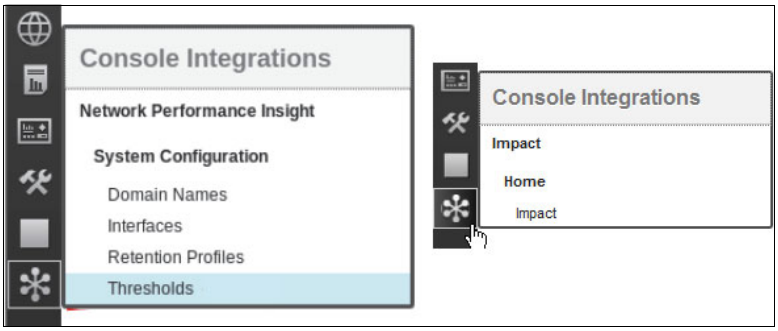

Figure 1-17 Console Integrations

Click **Impact** to open the Netcool Impact GUI, as shown in Figure 1-18. Impact is a core component for Netcool Operations Insight and all enrichments and event analytics depends on it.

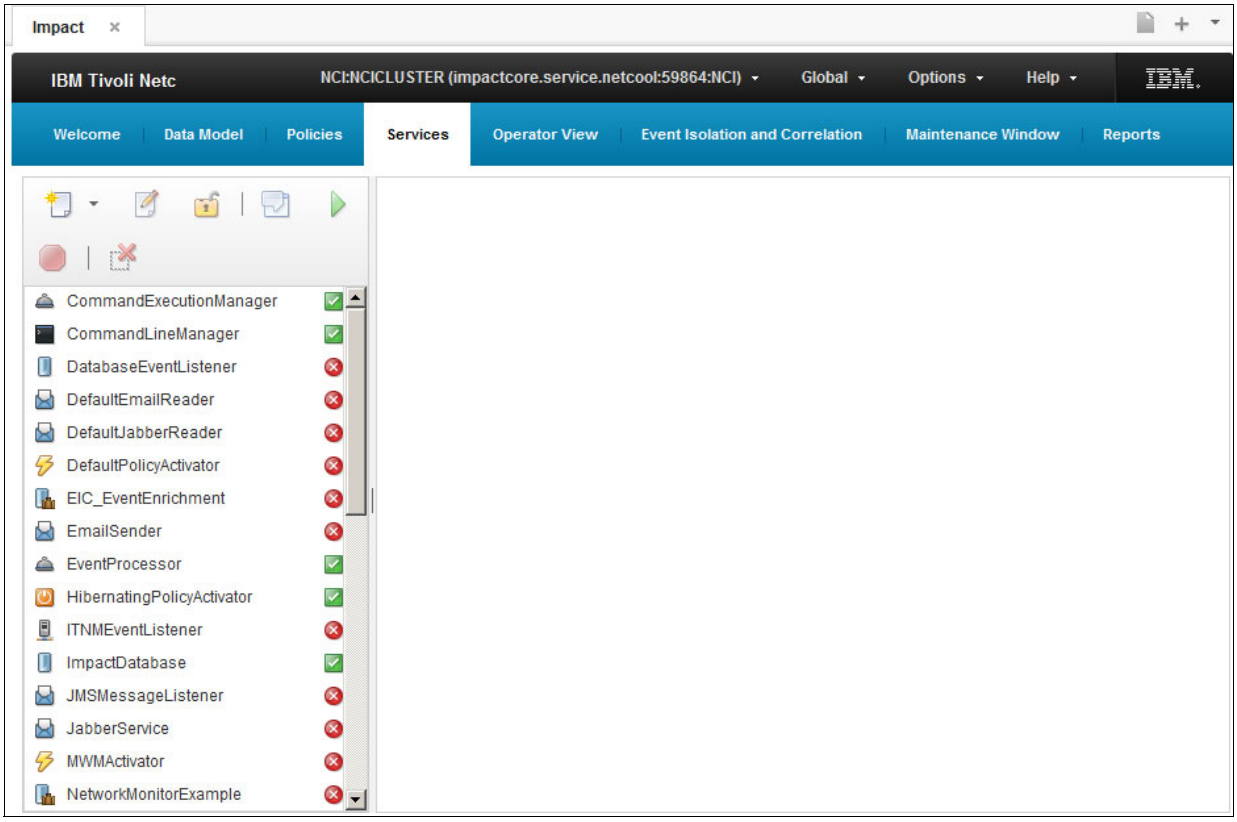

Figure 1-18 Impact GUI

1.4 Our environment for the scenarios

Most of the scenarios that are described in this publication were implemented by using the environment that is described in this section. This environment is the same environment that we created for *IBM Netcool Operations Insight Version 1.4 Deployment Guide*, SG24-8365. This environment is a typical Netcool Operations Insight environment that is configured for high availability (HA).

We used a separate environment for some of the scenarios. Those environments were similar to the one that is used here, but with different host names. For more information, see “Scenario Topology” section of each scenario.

1.4.1 High-level architecture

Figure 1-19 shows the high-level architecture that was used to deploy a multi-tiered Netcool Operations Insight environment with HA. The following Netcool Operations Insight components are shown in Figure 1-19:

- IBM Tivoli Netcool/OMNIbus (OMNI)
- IBM Jazz™ for Service Management (JazzSM)
- IBM DB2®
- IBM Tivoli Netcool Impact (NCI)
- IBM Tivoli Network Manager (NM)
- IBM Tivoli Netcool Configuration Manager (NCM)
- IBM Operations Analytics Log Analysis (IOALA)

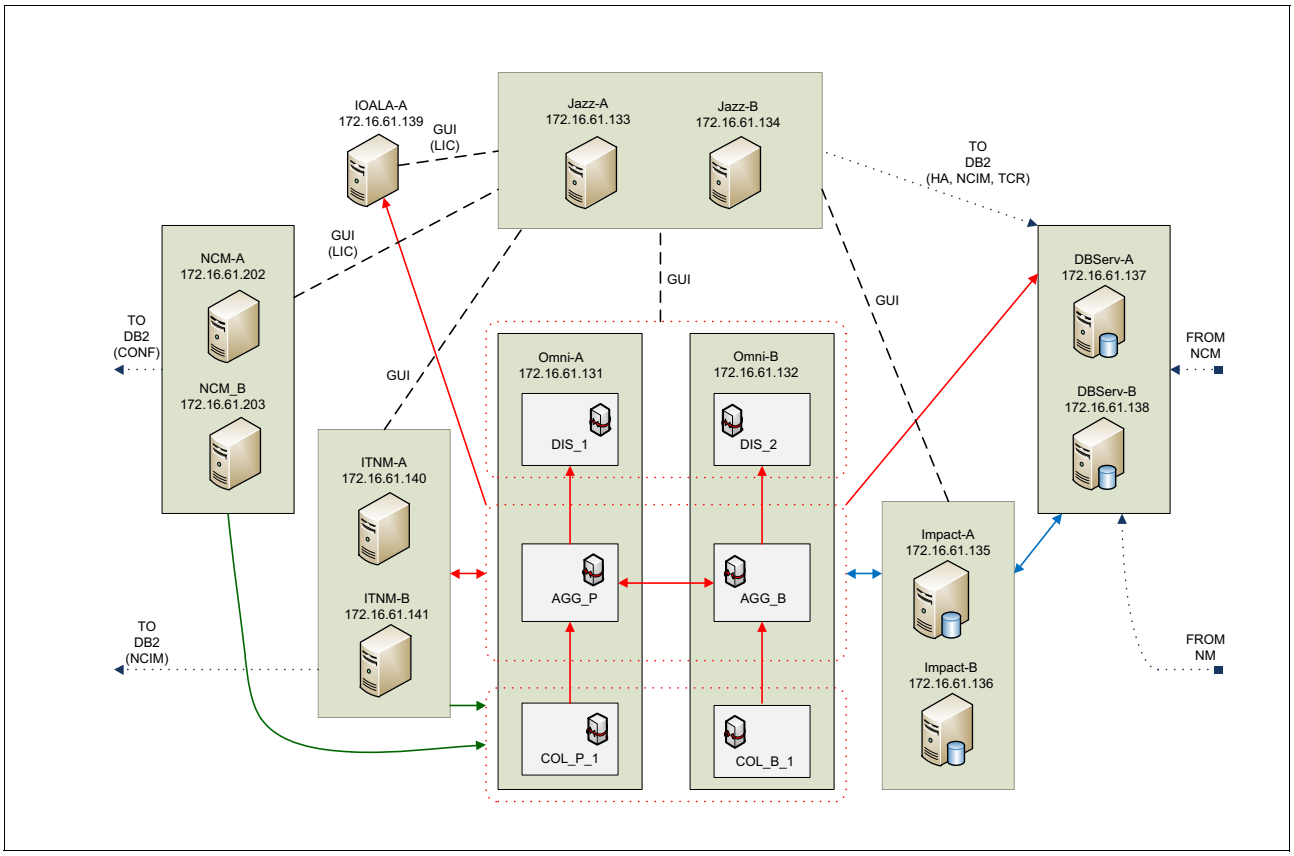

Figure 1-19 High-level architecture

As shown in Figure 1-19, JazzSM is deployed with HA and it is collecting data from Network Manager servers, Netcool Configuration Manager, Netcool Impact, DB2, IBM Operations Analytics - Log Analysis, and Netcool/OMNIbus.

Note: For more information about installing Netcool Operations Insight, see *IBM Netcool Operations Insight Version 1.4 Deployment Guide*, SG24-8365.

1.4.2 Environment database and connections

Figure 1-20 shows the Netcool Operations Insight environment from a database perspective. More specifically, it shows the relationships of each Netcool Operations Insight component and the DB2 instance to which they are connected.

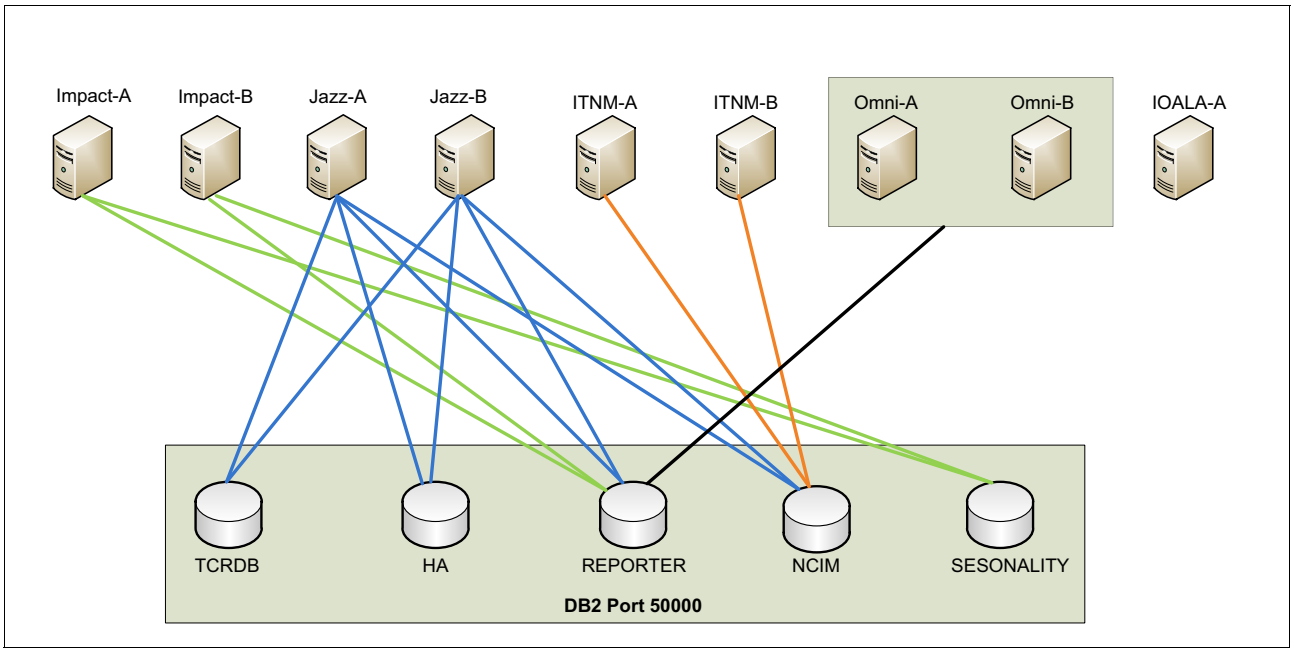

Figure 1-20 Database connection in Netcool Operations Insight

1.4.3 Ports used

The main ports that are needed during the Netcool Operations Insight deployment are listed in Table 1-1.

Table 1-1 Main ports that are used for Netcool Operations Insight

Netcool Operations Insight component	Port used
Omnibus - Aggregation Layer	4100
Omnibus - Collection Layer	4101
Omnibus - Display Layer	4102
JazzSM	16310 - 16316
Impact	16311
Log Analysis	9987
Network Manager	7968

The ports that are used in our environment for the connections between all the Netcool Operations Insight components are shown in Figure 1-21.

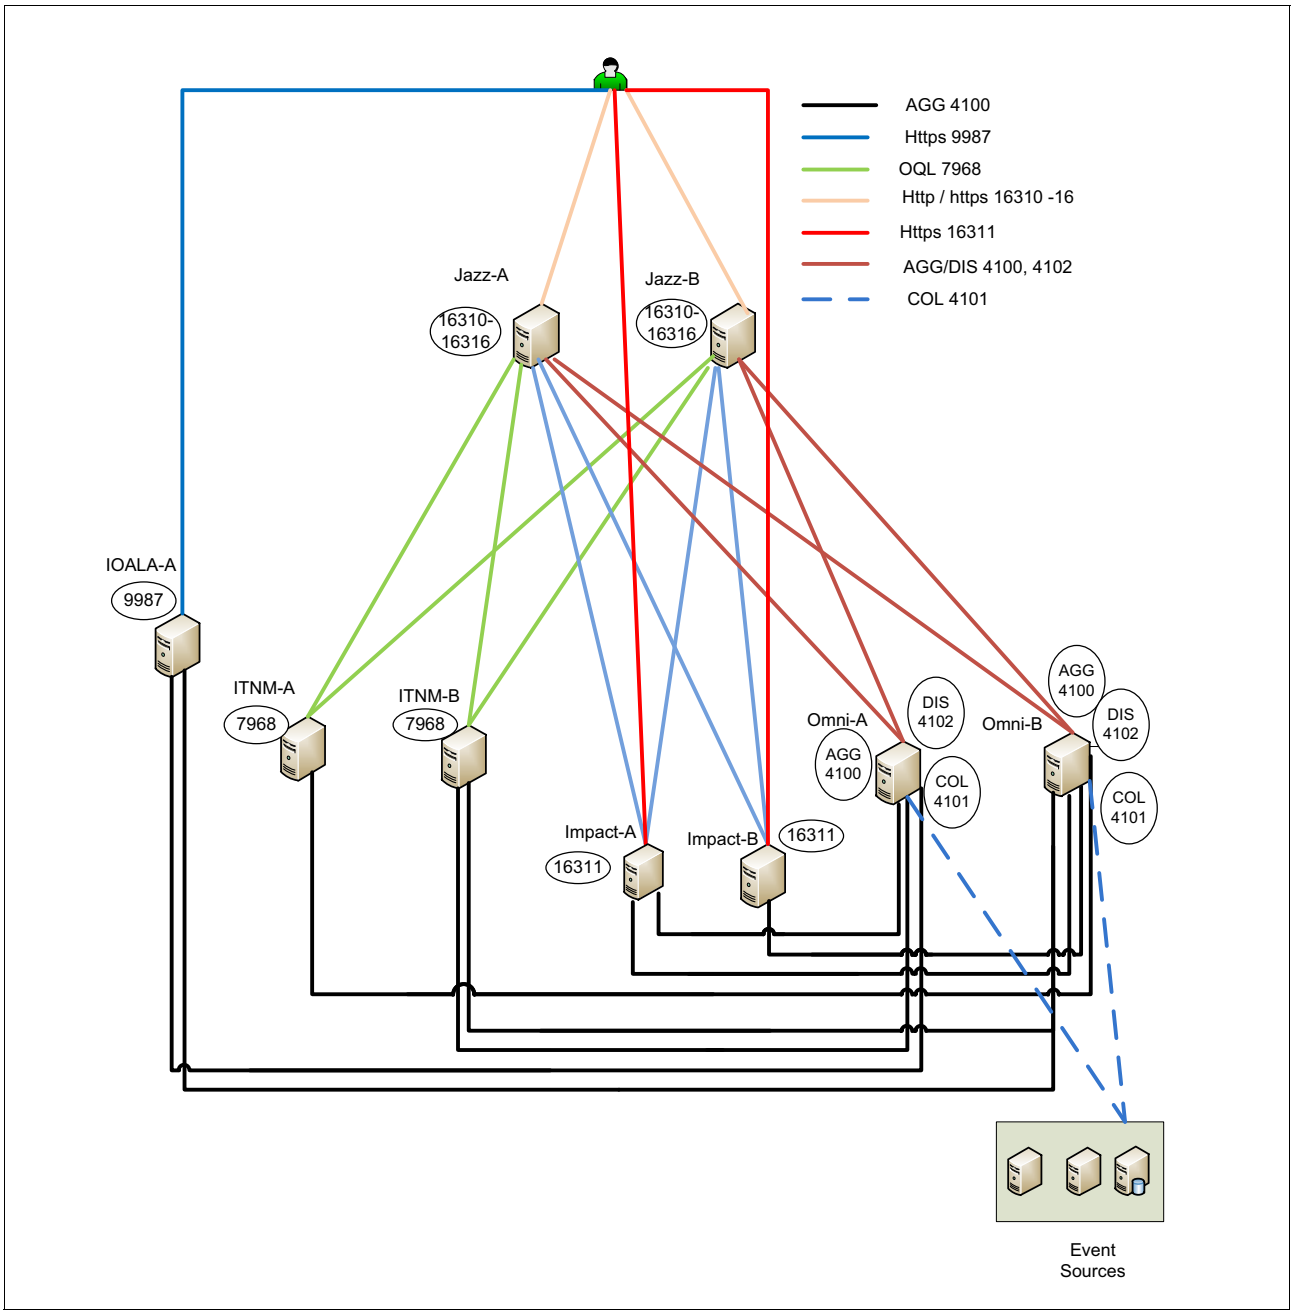

Figure 1-21 Ports that are used to connect the Netcool Operations Insight components

1.5 Summary

Netcool Operations Insight integrates across all IBM and non-IBM Management Solutions for a total IT solution serves enterprises demand. It combines all data by using web dashboards, mobile dashboards, and simple data integration solutions to extend Operations Analytics for prediction and search.

Part 2

Network management-related scenarios

In this part, the network management-related scenarios are described.

2

Networks for Operations Insight and the Network Health Dashboard

This chapter describes scenarios that are based on the Networks for Operations Insight feature, including integrated visualization of Network infrastructure with the Network Health Dashboard.

This chapter includes the following topics:

2.1 Networks for Operations Insight and Network Health Dashboard overview

IBM Networks for Operations Insight is an optional feature that can be added to a deployment of the base IBM Netcool Operations Insight solution. It provides service assurance in dynamic network infrastructures.

Networks for Operations Insight features the following capabilities:

- Network discovery
- Visualization
- Event correlation and root-cause analysis
- Configuration and Compliance Management

These capabilities provide service assurance in dynamic network infrastructures. The Networks for Operations Insight capability is provided by setting up the following products in Netcool Operations Insight:

- IBM Tivoli Network Manager
- IBM Tivoli Netcool Configuration Manager
- IBM Network Performance Insight

Network Health Dashboard is one of the main features that is provided by the Networks for Operations Insight solution. The Network Health Dashboard is available only with Network Manager as part of Netcool Operations Insight.

2.2 Scenario description

For more information about system components and default settings in the test environment, see Chapter 1, “IBM Netcool Operations Insight overview” on page 3.

Blue Bank & OmniFinance are large banking companies with an extensive network. They use several different vendors’ equipment across their widely distributed infrastructure.

By using the Network Health Dashboard, IT Operators or Subject Matter Experts (SME) can identify and troubleshoot network outages fast and resolve it quickly.

2.2.1 Business value

By using the Network Health Dashboard, IT Practitioners can perform the following tasks:

- View the state of network infrastructure in a single pane of glass
- View real-time monitoring and availability data
- Observe real-time Performance of best and worst performing devices
- Have an aggregated view of device configurations that were applied or not applied
- See real-time events in context
- Identify and troubleshoot network outages and issues and quickly resolve them

2.3 Scenario topology

For this scenario, we used the environment that is described in 1.4, “Our environment for the scenarios” on page 18.

2.4 Scenario steps

This section describes steps to use the Network Health Dashboard to quickly identify and resolve network outages or issues.

2.4.1 Administering the Network Health Dashboard

An IT administrator can configure how data is displayed and which data is displayed in the Network Health Dashboard. Device configuration change data can be displayed in the Configuration and Event Timeline if the integration with Netcool Configuration Manager is set up.

Tip: To fit the quantity of widgets onto a single window, a minimum resolution of 1536 x 864 is needed.

If the resolution is less than this minimum, scroll bars are seen on one or more of the widgets in the Network Health Dashboard.

More information: For more information about the integration with Netcool Configuration Manager, see the following topic in the Network Manager Knowledge Center:

`https://ibm.biz/BdrxsD`

For more information about the Integration of NPI into NOI, see this website:

`https://ibm.biz/BdrxsX`

Changing dashboard layout

The layout of the dashboard can be changed to the IT Administrator’s discretion. For example, widgets can be repositioned or resized.

More information: For more information about changing the dashboard content or layout, see the following topic in the Network Manager Knowledge Center:

`http://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/isc/edit_page.html`

Changing the refresh period for widgets on the Network Health Dashboard

The Network Manager widgets within the Network Health Dashboard update by default every 20 seconds. This update frequency can be changed by performing the following steps:

1. Edit the following configuration file:

   `$NMGUI_HOME/profile/etc/tnm//nethealth.properties.`

2. Find the following line and update the refresh period to the wanted value in seconds:

   `nethealth.refresh.period=60`

3. Save the file.
4. Close and reopen the Network Health Dashboard tab to affect the changes.

Note: The Event Viewer widget updates every 60 seconds by default.

Changing the colors that are associated with event severity values that are used in the Configuration and Event Timeline

The colors that are associated with event severity values that are used in the Configuration and Event Timeline can be changed by performing the following steps:

1. Edit the following configuration file:

   `$NMGUI_HOME/profile/etc/tnm/status.properties`.
2. Find the properties `status.color.background.severity_number`, where `severity_number` corresponds to the severity number. For example, 5 corresponds to Critical severity.
3. Change the RGB values for the severity values, as wanted.
4. Save the file.

Disabling start of the Network View tab when selecting a network view in the Network Health Dashboard

When a user selects a network view in the Network Health Dashboard (called *Network View*) a second tab is opened by default, as shown in Figure 2-1. This tab contains a dashboard that consists of the Network Views GUI, Event Viewer, and Structure Browser. It also displays the selected network view.

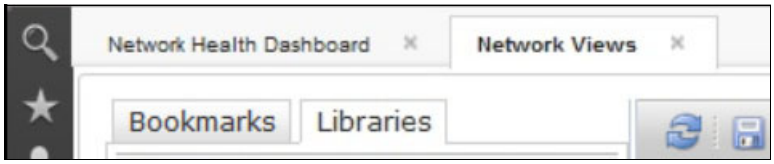

Figure 2-1 Network View Tab

If the network views are large, displaying this second tab can affect system performance. To avoid this performance affect, starting the Network View tab can be disabled by performing the following steps:

1. Edit the following configuration file:

   `$NMGUI_HOME/profile/etc/tnm/topoviz.properties.`
2. Find the lines that are shown in Example 2-1.

Example 2-1 Excerpt from the $NMGUI_HOME/profile/etc/tnm/topoviz.properties file

```
# Defines whether the dashboard network view tree fires a launchPage event when
the user clicks a view in the tree
topoviz.networkview.dashboardTree.launchpage.enabled=true
```

3. Set the property `topoviz.networkview.dashboardTree.launchpage.enabled` to `false`.
4. Save the file.

More information: For more information about troubleshooting the Network Health Dashboard, see the following topic in the IBM Knowledge Center:

https://ibm.biz/BdrxsN

2.4.2 Custom dashboards

Pages can be created that act as dashboards for displaying information about the status of parts of a network or edit dashboards, such as the Network Health Dashboard. Many widgets can be used for creating or modifying dashboards. The widgets are provided with Network Manager, Tivoli Netcool/OMNIbus Web GUI, and from other products that are deployed in the Dashboard Application Services Hub environment.

For more information: For more information about creating Custom Dashboards, see the following topic on the Networks for Operations Insight Knowledge Center:

https://ibm.biz/Bdrxix

For more information about creating and editing pages in the Dashboard Application Services Hub, see the following topic on the Network Manager Knowledge Center:

http://www.ibm.com/support/knowledgecenter/SSSHRK_4.2.0/isc/edit_page.html

For more information about Dashboard Application Services Hub, see the Jazz for Service Management information center at this website:

http://www.ibm.com/support/knowledgecenter/SSEKCU/welcome

2.4.3 Monitoring the Network Health Dashboard

The Network Health Dashboard can be monitored by selecting a network view within an area of responsibility, such as a geographical area, or a specific network service, and reviewing the data that appears in the other widgets on the dashboard.

If users set up a default network view bookmark that contains the network devices within an area of responsibility, these views appear in their network view tree within the dashboard.

Monitoring the Network Health Dashboard includes the following key steps:

- “Displaying device and interface availability in a network view” on page 29
- “Displaying overall network view availability” on page 30
- “Displaying highest and lowest performers in a network view” on page 32
- “Displaying the Configuration and Event Timeline” on page 37

Displaying device and interface availability in a network view

By using the Unavailable Resources widget, users can monitor within a selected network view the number of device and interface availability alerts that are open for more than a configurable amount of time. By default, this widget charts the number of device and interface availability alerts that were open for up to 10 minutes, for more than 10 minutes but less than one hour, and for more than one hour.

Complete the following steps to monitor the number of open device and interface availability alerts within a selected network view:

1. Click the Incident icon and select **Network Availability** → **Network Health Dashboard**.

2. In the Network Health Dashboard, select a network view from the network view tree in the **Network Views** window at the top left. The other widgets update to show information based on the network view that you selected.

   The Unavailable Resources widget updates to show device and interface availability in the selected network view. A second tab, called *Network View*, opens. This tab contains a dashboard comprised of the Network Views GUI, Event Viewer, and Structure Browser, and it displays the selected network view. You can use this second tab to explore the topology of the network view that you are displaying in the Network Health Dashboard.

3. In the Unavailable Resources widget, to determine the number of unavailable devices and interface alerts, use the following sections of the chart and note the colors of the stacked bar segments and the number inside each segment. Click any one of the bars that are shown in Figure 2-2 to show the corresponding alerts for the devices and interfaces in the Event Viewer at the bottom of the Network Health Dashboard:

   – SNMP Poll Fail

     Uses color-coded stacked bars to display the number of SNMP Poll Fail alerts within the specified time frame.

   – SNMP Link-State

     Uses color-coded stacked bars to display the number of SNMP Link State alerts within the specified time frame.

   – Interface Ping

     Uses color-coded stacked bars to display the number of Interface Ping alerts within the specified time frame.

   – Device Ping

     Uses color-coded stacked bars to display the number of Device Ping alerts within the specified time frame.

Tips: Consider the following points:

- By default, all of the bars that are described in this section are configured to display. However, the Unavailable Resources widget can be configured to display specific bars only. For example, if the widget is configured to display only the Device Ping and the Interface Ping bars, only those bars are displayed in the widget.
- By default, the data in the Unavailable Resources widget is updated every 20 seconds.

Color coding of the stacked bars is shown in Figure 2-2.

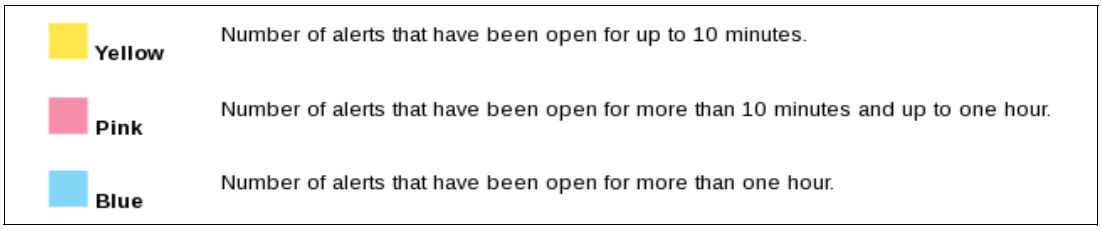

Figure 2-2 Stacked bar color codes

Displaying overall network view availability

IT Operators and SMEs can monitor overall availability of chassis devices within a selected network view by using the Percentage Availability widget.

Complete the following steps to display the overall availability of chassis devices within a selected network view:

1. Click the **Incident** icon and select **Network Availability** → **Network Health Dashboard**.
2. In the Network Health Dashboard, select a network view from the network view tree in the Network Views at the top left. The other widgets update to show information that is based on the network view that you selected.

   In particular, the Percentage Availability widget updates to show overall availability of chassis devices in network view. A second tab, called Network View, opens. This tab contains a dashboard that consists of the Network Views GUI, Event Viewer, and Structure Browser. It displays the selected network view. This second tab can also be used to explore the topology of the network view that is being displayed in the Network Health Dashboard.

   The Percentage Availability widget (as shown in Figure 2-3) displays 24 individual hour bars. Each bar displays a value, which is an exponentially weighted moving average of ping results in the past hour (the bar appears on the completion of the hour only). The bar value represents a percentage availability rate rather than a total count within that hour. The color of the bar indicates the following percentage availability:

   – Green: 80% or more
   – Orange: Between 50% and 80%
   – Red: Less than 50%

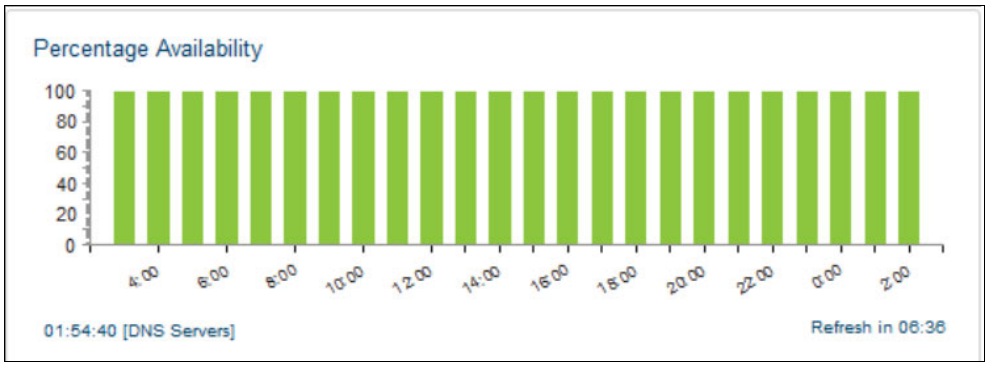

Figure 2-3 Percentage Availability widget

Tip: If the Percentage Availability widget is taking a long time to refresh, one possible solution is to increase the number of threads available for this widget. This solution is most suitable for customers with large networks. You can increase the number of threads available by performing the following steps:

1. Edit the following configuration file:

   ```
   $NMGUI_HOME/profile/etc/tnm//nethealth.properties
   ```

2. Find the following lines:

   ```
   ## Widget thread count for availability widget
   nethealth.threads.availability=5
   ```

3. Increase the value of the nethealth.threads.availability property. The maximum possible value is 10.
4. Save the file.

Displaying highest and lowest performers in a network view

The highest and lowest poll data metrics across all devices and interfaces within a selected network view can be monitored by using the Top Performers widget. Complete the following steps:

1. Click the **Incident** icon and select **Network Availability** → **Network Health Dashboard**.
2. In the Network Health Dashboard, select a network view from the network view tree in the Network Views at the top left. The other widgets update to show information that is based on the network view that is selected.
3. In the Top Performers widget (Figure 2-6 on page 34), select from the following controls to display chart, table, or trace data in the Top Performers widget, and how to display it:
   - Metric

     Select this drop-down list to display a selected set of poll data metrics. The metrics that are displayed in the drop-down list depend on which poll policies are enabled for the selected network view. Select one of these metrics to display associated data in the main part of the window. Figure 2-4 shows an example of metric drop-down choices.

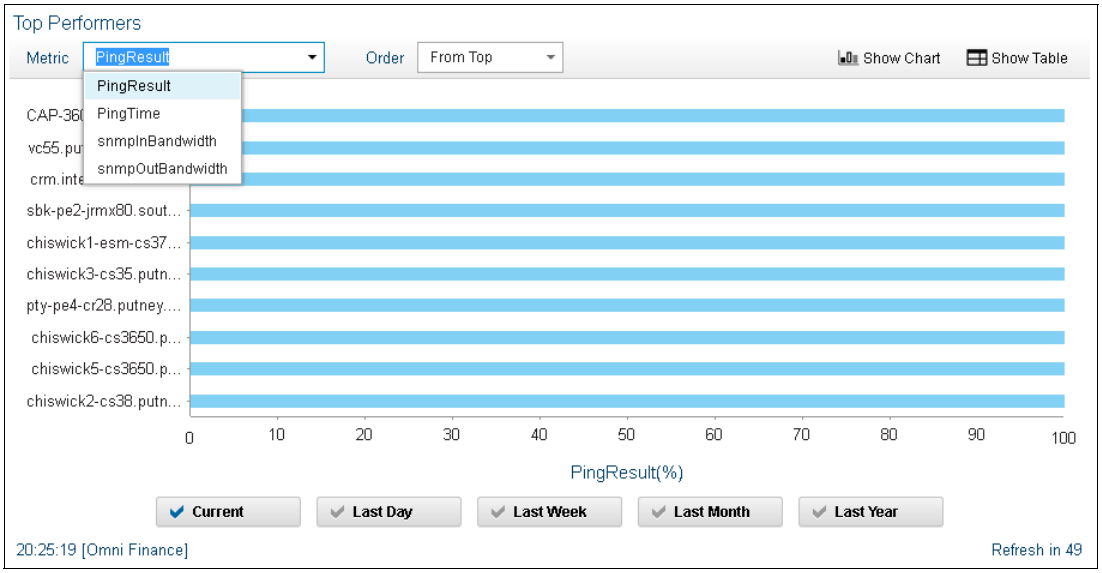

Figure 2-4 An example of metric drop-down choices

   - Order

     Select this drop-down list to display the statistic to apply to the selected poll data metric. The following statistics are available for all metrics (except the SnmpLinkStatus metric):

     - From Top: Displays a bar chart or table that shows the 10 highest values for the selected metric. The devices or interfaces with these maximum values are listed in the bar chart or table.
     - From Bottom: Displays a bar chart or table that shows the 10 lowest values for the selected metric. The devices or interfaces with these minimum values are listed in the bar chart or table.

     The following statistics are available for the SnmpLinkStatus metric. In each case, a bar chart or table displays and shows devices for the selected statistic:

     - Unavailable: This statistic displays by default. Devices with this statistic are problematic.

- Admin Down: Devices with this statistic are not problematic as Administrators change devices to this state.
- Available: Devices with this statistic are not problematic.

The widget lists devices or interfaces depending on which of the following metrics was selected:

- If the metric that was selected applies to a device, such as memoryUtilization, the top 10 list contains devices.
- If the metric that is selected applies to an interface, such as ifInDiscards, the top 10 list contains interfaces.

Figure 2-5 shows an example of Top Performer view with top 10 statistics within current time frame.

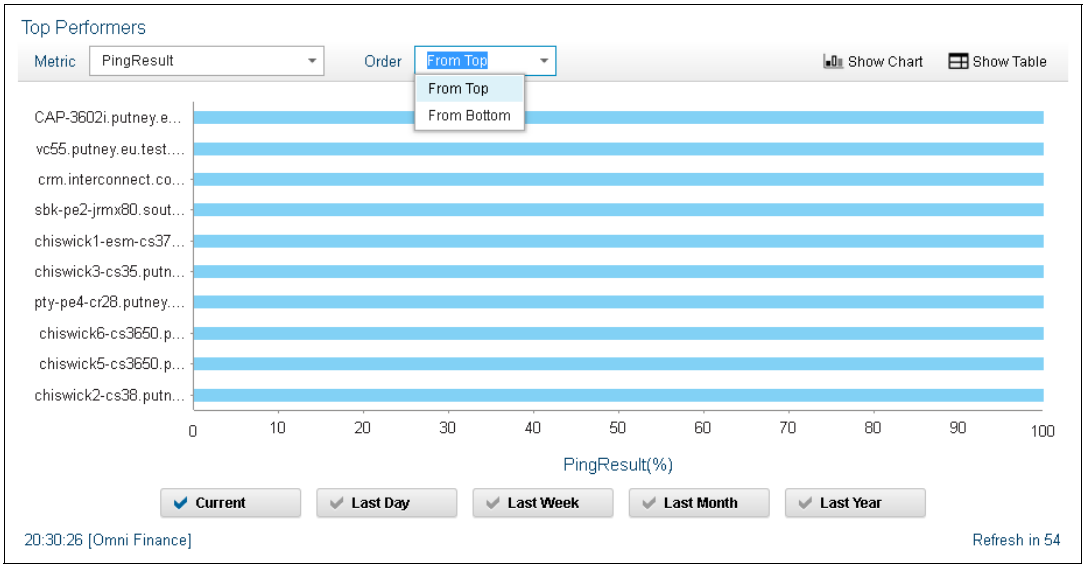

Figure 2-5 Example of Top Performer view with top 10 statistics within current time frame

– Show Chart

Displays a bar chart with up to the 10 highest or lowest values. Show Chart is the display option when the widget is first opened.

– Show Table

  Displays a table of data that is associated with up to the 10 highest or lowest values. Figure 2-6 shows the Top Performers widget view.

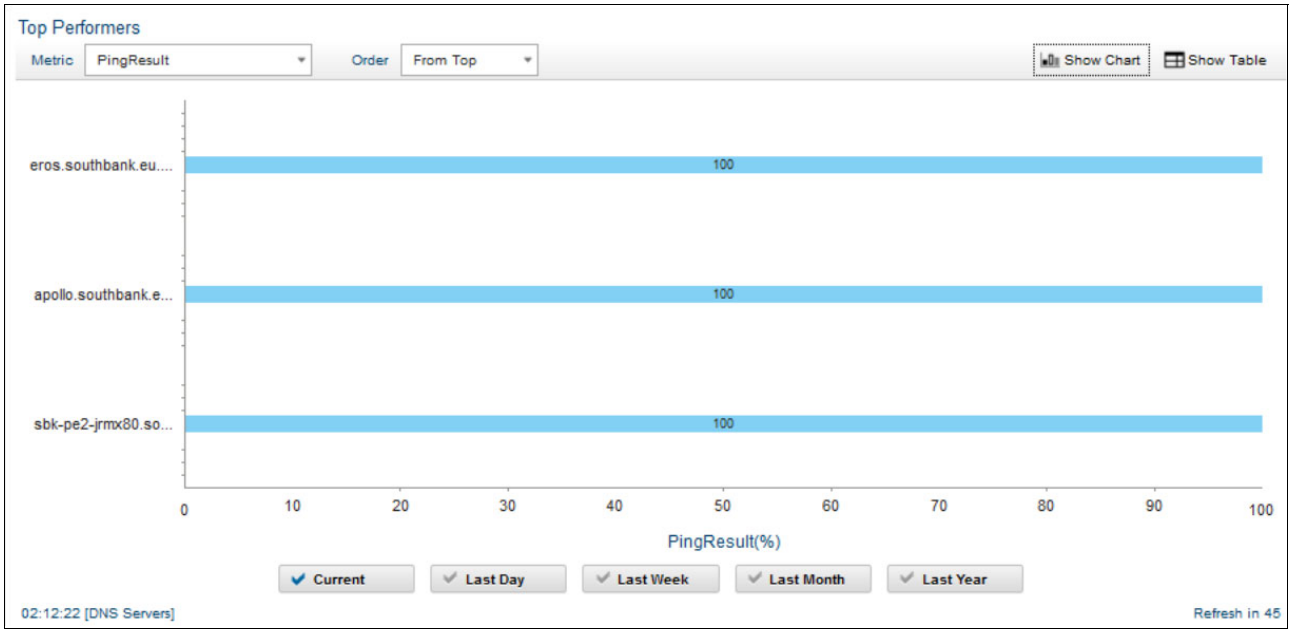

Figure 2-6 Example of a Top Performers widget view

– Define Filter

  This button appears in Show Table mode only. Click this button to define a filter to apply to the Top Performers table data.

The main part of the window contains the data in one of the following formats:

– Chart

  Bar chart with the 10 highest or lowest values. Click any bar in the chart to show a time trace for the corresponding device or interface.

– Table

  Table of data that is associated with the 10 highest or lowest values. The table contains the following columns:

  - Entity Name: Name of the device or interface.
  - Show Trace: Click a link in one of the rows to show a time trace for the corresponding device or interface.
  - Last Poll Time: Last time this entity was polled.
  - Value: Value of the metric the last time this entity was polled.

An example of Table format view is shown in Figure 2-7.

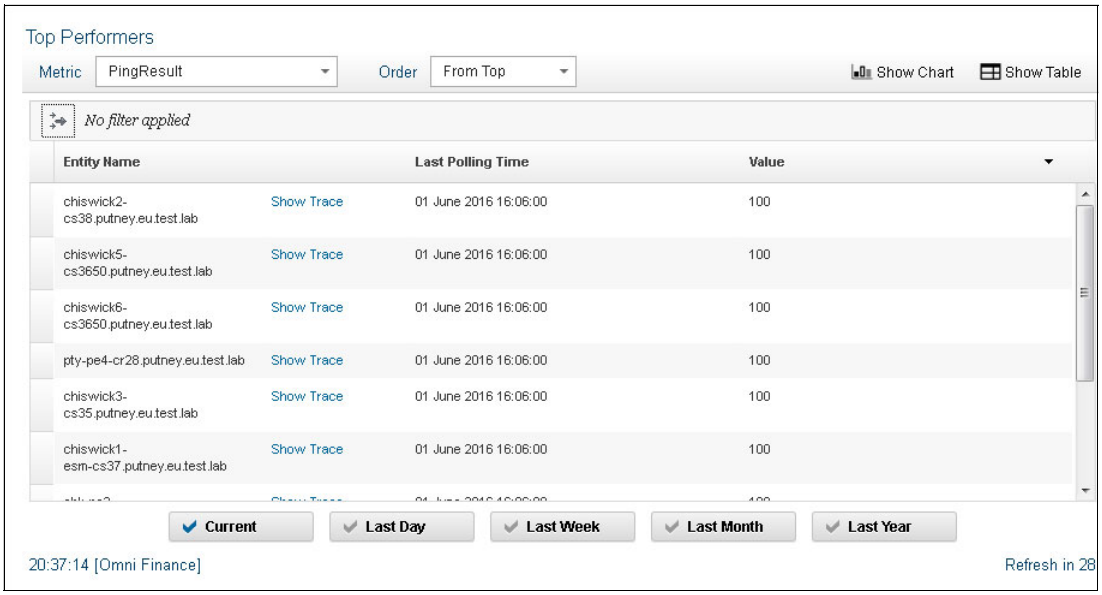

Figure 2-7 Top Performers view with Show Table mode

– Trace

Time trace of the data for a single device or interface. Navigate within this trace by performing the following operations:

i. Zoom into the trace by moving your mouse wheel forward.

ii. Zoom out of the trace by moving your mouse wheel backward.

iii. Double-click to restore the normal zoom level.

iv. Click within the trace area for a movable vertical line that displays the exact value at any point in time.

Click one of the following buttons to specify the current or historical poll data to display in the main part of the window. This button updates the data regardless of which mode is presented: bar chart, table, or time trace.

Tip: If your administrator opted not to store poll data for any of the poll data metrics in the Metric drop-down list, historical poll data is not available when you click any of the following buttons:

- Last Day
- Last Week
- Last Month
- Last Year

– Current

Click this button to display current raw poll data. When in time trace mode, the time trace shows anything up to two hours of data, depending on the frequency of polling of the associated poll policy.

– Last Day

Click this button to show data based on a regularly calculated daily average.

In bar chart or table mode, the top 10 highest or lowest values are shown based on a daily exponentially weighted moving average (EWMA).

In time trace mode, a time trace of the last 24 hours is shown that is based on the average values.

In the Last Day section of the widget, EWMA values are calculated by default every 15 minutes and are based on the previous 15 minutes of raw poll data. The data that is presented in this section of the widget is then updated with the latest EWMA value every 15 minutes.

Last day or last 24 hours?: In the GUI, this metric is referred to as Last Day, but it is actually the last 24 hours. Similar logic is used for Last Week, Last Month, and Last Year, which is described next.

– Last Week

  Click this button to show data based on a regularly calculated weekly average.

  In bar chart or table mode, the top 10 highest or lowest values are shown based on a weekly EWMA.

  In time trace mode, a time trace of the last seven days is shown, based on the average values.

  In the Last Week section of the widget, EWMA values are calculated by default every 30 minutes and are based on the previous 30 minutes of raw poll data. The data that is presented in this section of the widget is then updated with the latest EWMA value every 30 minutes.

– Last Month

  Click this button to show data that is based on a regularly calculated monthly average.

  In bar chart or table mode, the top 10 highest or lowest values are shown based on a monthly EWMA.

  In time trace mode, a time trace of the last 30 days is shown, based on the average values.

  In the Last Month section of the widget EWMA values are calculated by default every two hours and are based on the previous two hours of raw poll data. The data that is presented in this section of the widget is then updated with the latest EWMA value every two hours.

– Last Year

  Click this button to show data based on a regularly calculated yearly average.

  In bar chart or table mode, the top 10 highest or lowest values are shown based on a yearly EWMA.

  In time trace mode, a time trace of the last 365 days is shown, based on the average values.

  In the Last Year section of the widget EWMA values are calculated by default every day and are based on the previous 24 hours of raw poll data. The data that is presented in this section of the widget is then updated with the latest EWMA value every day.

  The Historical Polling Data for Top Performers with the Current option that is selected is shown in Figure 2-8 on page 37.

Figure 2-8 View current or Historical Polling Data for Top Performers

Displaying the Configuration and Event Timeline

A timeline can be displayed for all devices in a selected network view with device configuration changes and network alert data over a period of up to 24 hours by using the Configuration and Event Timeline widget. Correlation between device configuration changes and network alerts on this timeline can help identify where configuration changes might lead to network issues.

To display a timeline that shows device configuration changes and network alert data for all devices in a selected network view, complete the following steps:

1. Click the **Incident** icon and select **Network Availability** → **Network Health Dashboard**.
2. In the Network Health Dashboard, select a network view from the network view tree in the Network Views at the top left. The other widgets update to show information that is based on the network view that you selected. In particular, the Configuration and Event Timeline updates to show configuration change and event data for the selected network view.

   In the Configuration and Event Timeline widget, configuration changes that are displayed in the Configuration and Event Timeline can be any of the changes that are described in “Changes that are managed by Netcool Configuration Manager” on page 37.

> **Tip:** If Netcool Configuration Manager is not installed, no configuration data is displayed in the timeline.

Changes that are managed by Netcool Configuration Manager

The following changes are made under full Netcool Configuration Manager control. The timeline differentiates between scheduled or policy-based changes, which can be successful (Applied) or unsuccessful (Not Applied), and one-time changes made by using the IDT Audited terminal facility within Netcool Configuration Manager:

- Applied

  A successful scheduled or policy-based set of device configuration changes that were made under the control of Netcool Configuration Manager.

- Not Applied

  An unsuccessful scheduled or policy-based set of device configuration changes that were made under the control of Netcool Configuration Manager.

- IDT

  Device configuration changes made by using the audited terminal facility within Netcool Configuration Manager that allows one-time command-line based configuration changes to devices.

- Unmanaged changes:

  – OOBC

    Out-of-band-change. Manual configuration change that was made to device where that change is outside of the control of Netcool Configuration Manager.

    Events are displayed in the timeline as stacked bars, in which the color of each element in the stacked bar indicates the severity of the corresponding events. Move your mouse over the stacked bars to view a tooltip that lists the number of events at each severity level. The X-axis granularity for events and configuration changes varies depending on the time range that you select for the timeline.

    The X-axis granularity in the Configuration and Event Timelines is listed in Table 2-1.

Table 2-1 X-axis granularity in the Configuration and Event Timelines

If you select this time range ...	Then, the X-axis granularity is ...
6 hours	15 minutes
12 hours	30 minutes
24 hours	1 hour

More information: For more information about the different types of configuration change, see the Netcool Configuration Manager knowledge center at this website:

`http://www-01.ibm.com/support/knowledgecenter/SS7UH9/welcome`

Select from the following controls to define what data to display in the Configuration and Event Timeline:

– Time

  Select the duration of the timeline:

  - 6 Hours: Click to set a timeline duration of 6 hours.
  - 12 Hours: Click to set a timeline duration of 12 hours.
  - 24 Hours: Click to set a timeline duration of 24 hours.

– Events by Occurrence:

  - First Occurrence: Click to display events on the timeline based on the first occurrence time of the events.
  - Last Occurrence: Click to display events on the timeline based on the last occurrence time of the events.

– Show Table

  Displays the configuration change data in tabular form. The table contains the following columns:

Note: If Netcool Configuration Manager is not installed, this button is not displayed.

- Number: Serial value that indicates the row number.
- Device: Host name or IP address of the affected device.
- Unit of Work (UoW): In the case of automated Netcool Configuration Manager configuration changes, the Netcool Configuration Manager UoW under which this configuration change was processed.
- Result: Indicates whether the change was successful.
- Start Time: The time at which the configuration change began.
- End Time: The time at which the configuration change completed.
- User: The user who applied the change.
- Description: Textual description that is associated with this change.

– Show Chart

Click this button to switch back to the default graph view.

Note: If Netcool Configuration Manager is not installed, this button is not displayed.

Use the sliders under the timeline to zoom in and out of the timeline. The legend under the timeline shows the colors that are used in the timeline to display the following items:

- Event severity values
- Configuration change types

If the integration with Netcool Configuration Manager was set up but there is a problem with data retrieval from Netcool Configuration Manager, the configuration change types that are shown in the legend are marked with an icon that indicates that the integration with Netcool Configuration Manager is set up but no configuration management data is available.

The Configuration and Event Timeline Filter creation window is shown in Figure 2-9 on page 40.

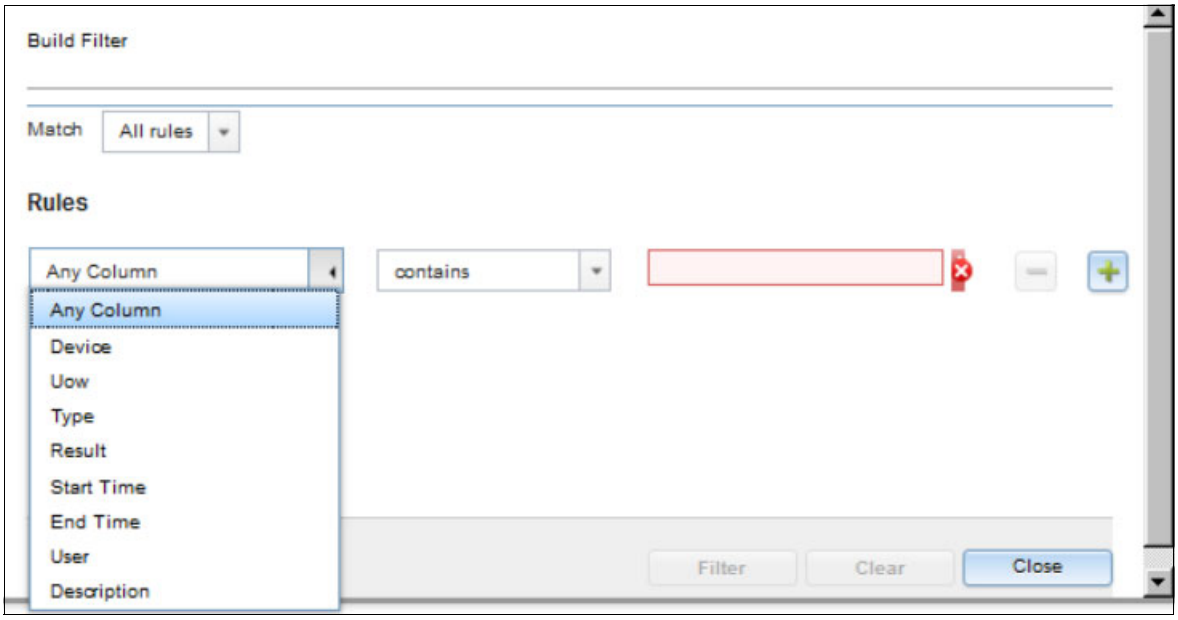

Figure 2-9 The Configuration and Event Timeline Filter creation window

2.4.4 Configuring the Network Health Dashboard for users

A user can configure the Network Health Dashboard to display the data as they want to see it.

This section describes the steps for the following processes:

- "Configuring the network view tree to display in the Network Health Dashboard"
- "Configuring the Unavailable Resources widget" on page 42
- "Configuring the Configuration and Event Timeline" on page 44

Configuring the network view tree to display in the Network Health Dashboard

Users of the Network Health Dashboard can configure a default bookmark to limit the data that is displayed in the Network Health Dashboard to the network views within their area of responsibility.

The network views tree in the Network Health Dashboard automatically displays the network views in a user's default network view bookmark. If there are no network views in the default bookmark, a message is displayed with a link to the Network Views GUI in which the user can add network views to the default bookmark. The network views that can be added to the default bookmark are displayed in the network tree within the Network Health Dashboard.

Complete the following steps:

1. Within the displayed message, click the link that is provided. The Network Views GUI opens in a second tab.
2. To add network views to a bookmark:
   a. Click the **Incident** icon and select **Network Availability** → **Network Views** → **Libraries**.
   b. Select a single network view:

i. From the network view library drop-down list that is above the network view tree, select the **network view library** that contains the network view to add to the bookmark.

ii. In the network view tree, browse to the wanted network view.

c. Right-click the wanted network view and click **Add to Bookmark**. An example is shown in Figure 2-10.

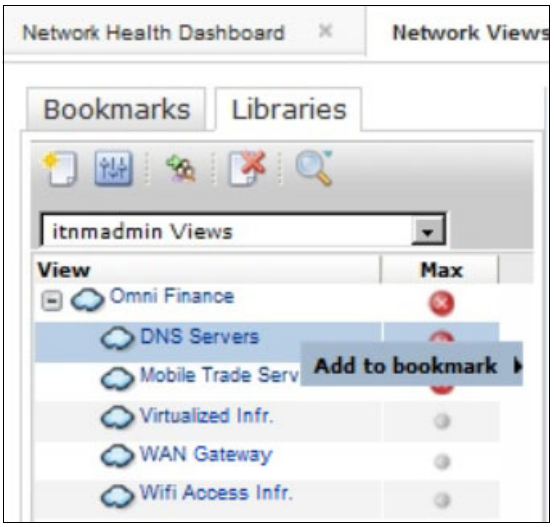

Figure 2-10 Adding Network Views to bookmarks for networks for Omnifinance

d. In the submenu, click the name of the bookmark to which to add the network view. Only the bookmarks that the user has permission to read and write to appear in the list. For example, to add the selected network views to a bookmark named `Bookmark1`, click **Add to Bookmark** → **Bookmark1**.

You can add the Network views to the bookmark by using one of the following methods:

- If the parent network view is selected, the parent network view and all of its child network views are added to the bookmark. If the parent network view is a dynamic network view, the views are added to or removed from the bookmark automatically as child network views are added to or removed from the parent following network discovery.
- If one or more child network views are selected (whether of a standard or dynamic parent network view), only the network views that are selected are added to the bookmark. In addition, individually selected child network views of dynamic parent network views are automatically removed from the bookmark if the child network view is removed following network discovery.

Note: Whichever view is added to a bookmark, the system automatically adds its parent, the parent's parent, and so on, until it gets to the top of the tree. Manually added network views appear without an asterisk to the left of the view name. System-added network views appear with an asterisk to the left of the view name. For example, if a parent network view is added to the bookmark, the parent network view and all of its child network views are added to the bookmark.

If the parent network views were manually added, it appears without an asterisk. If the child network views were added by the system they appear with an asterisk.

The network views tree in the Network Health Dashboard displays the network views in your newly configured default bookmark.

3. To remove network views from a bookmark:
   a. Click the **Incident** icon and select **Network Availability** → **Network Views** → **Bookmarks**.
   b. Select a single network view:
      iii. From the network view bookmarks drop-down list that is above the network view tree, select the network view bookmark that contains the network view to remove from the bookmark.
      iv. In the network view bookmarks tree, browse to the wanted network view.

      Note: Only network views from the bookmark that were manually added to the bookmark earlier can be removed. Consider the following points:

      - The network views that can be removed appear without an asterisk to the left of the view name. These views are manually added views.
      - The network views that cannot be removed appear with an asterisk to the left of the view name. These views were added by the system and are parents or children of manually added views.

   c. Right-click the wanted network view and click **Remove view from bookmark**.

      Note: A view can be removed only from a bookmark with read-write permission on that bookmark, or the user has the ncp_bookmark_admin role.

      The selected view is removed from the bookmark.

      Note: If a view is removed that is part of a hierarchy of other views that were previously added, that view is not physically removed from the tree. Instead, that view is marked with an asterisk (*).

Configuring the Unavailable Resources widget

Availability data can be displayed in the Network Health Dashboard by using the Unavailable Resources widget. IT Operators or SMEs can configure the widget to display availability data that is based on ping polls only, and not based on SNMP polls. They can also configure the time duration thresholds to apply to availability data displayed in this widget.

For example, by default the widget charts the number of device and interface availability alerts that were open for up to 10 minutes, more than 10 minutes, and more than one hour. These thresholds can be modified as needed.

To configure which availability data is displayed by the Unavailable Resources widget, complete the following steps:

1. Click the **Incident** icon and select **Network Availability** → **Network Health Dashboard**.
2. In the Unavailable Resources widget, click **User Preferences**, as shown in Figure 2-11 on page 43.

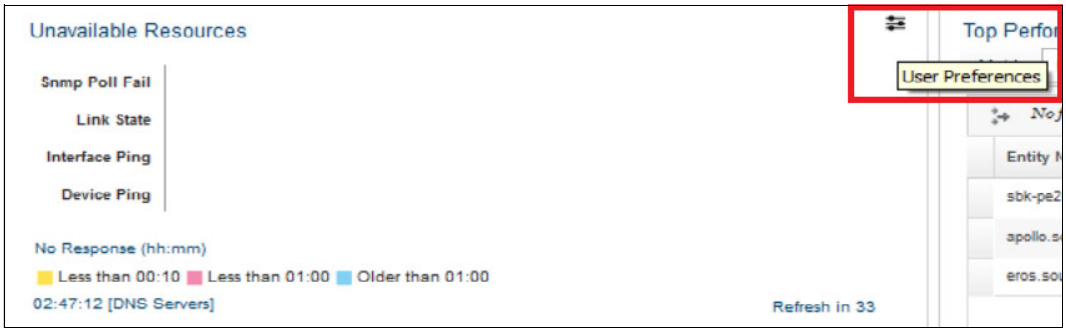

Figure 2-11 Unavailable Resources widget selection

To configure the Unavailable Resources widget, select the following options and the upper and lower field arrows (see the red box in Figure 2-13 on page 44) as applicable:

- Device

  Configure which device alerts to monitor in the Unavailable Resources widget to retrieve information about device availability. By default, all of these options are selected.

- Device Ping

  Select the option to monitor Default Chassis Ping alerts. Selecting this option causes the Unavailable Resources widget to provide an indication of the number of open device ICMP (ping) polling alerts.

- SNMP Poll Fail

  Select this option to monitor SNMP Poll Fail alerts. Selecting this option causes the Unavailable Resources widget to provide an indication of the number of open SNMP Poll Fail alerts.

- Interface

  Configure which interface alerts to monitor in the Unavailable Resources widget to retrieve information about interface availability. By default, all of these options are selected.

- Interface Ping

  Select this option to monitor Default Interface Ping alerts. Selecting this option causes the Unavailable Resources widget to provide an indication of the number of open interface ICMP (ping) polling alerts.

- Link State

  Select this option to monitor SNMP Link State alerts. Selecting this option causes the Unavailable Resources widget to provide an indication of the number of open SNMP Link State alerts.

- Thresholds:

  - Upper

    Specify an upper threshold in hours and minutes. By default, the upper threshold is set to one hour. This threshold causes the chart in the Unavailable Resources widget to update. When the amount of time that any availability alert in the selected network view remains open exceeds the one hour threshold, the relevant bar in the Unavailable Resources chart updates to show this unavailability as a blue color-coded bar section.

- Lower

  Specify a lower threshold in hours and minutes. By default, the lower threshold is set to 10 minutes. This threshold causes the chart in the Unavailable Resources widget to update. When the amount of time that any availability alert in the selected network view remains open exceeds the 10-minute threshold, the relevant bar in the Unavailable Resources chart updates to show this unavailability as a as a pink color-coded bar section.

  Figure 2-12 shows the Unavailable Resources Widget.

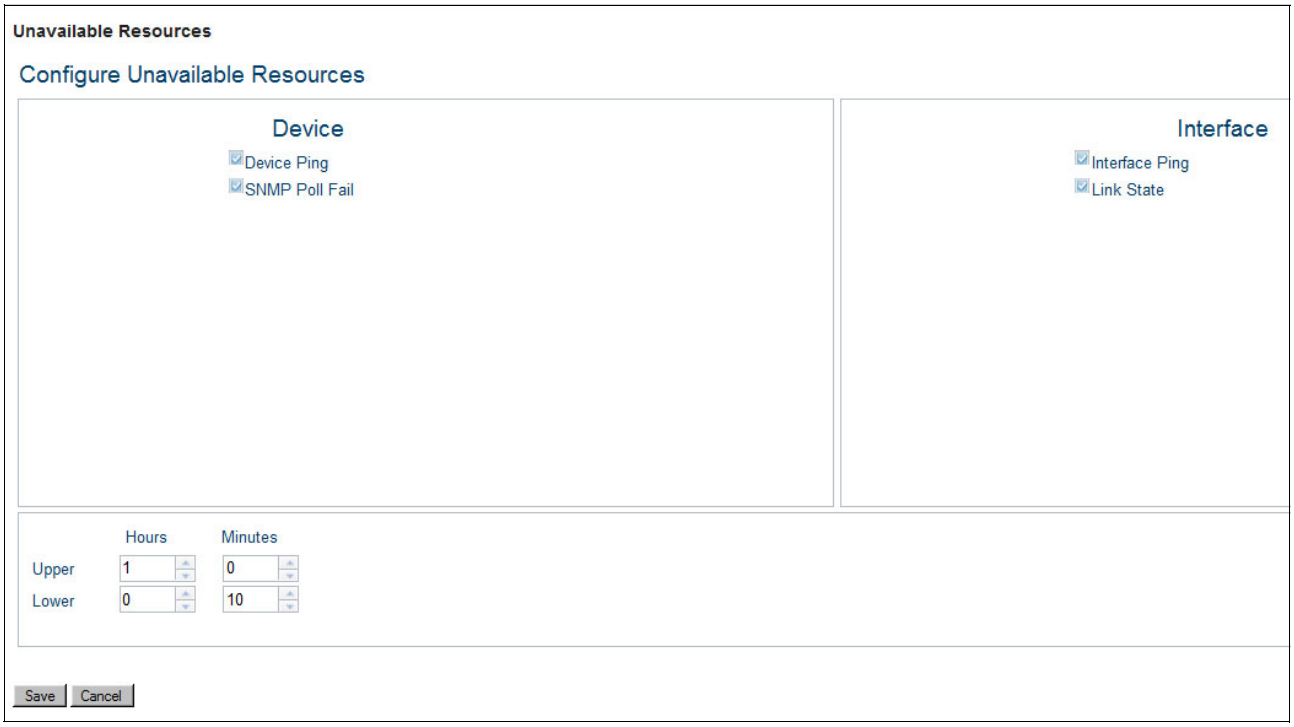

Figure 2-12 Configuring the Unavailable Resources Widget

Configuring the Configuration and Event Timeline

Event severity values can be configured to display on the Configuration and Event Timeline.

Complete the following steps to configure which event severity values to display on the Configuration and Event Timeline:

1. Click the Incident icon and select **Network Availability** → **Network Health Dashboard**.
2. In the Configuration and Event Timeline widget, click the **User Preferences** button, as shown in Figure 2-13.

Figure 2-13 Selecting the User Preferences button

3. To configure the Configuration and Event Timeline, use the following lists:
   - Available Severities

     By default, lists all event severity values. These event severity values are all displayed in the Configuration and Event Timeline.

To remove an item from this list, select the item and click the right-pointing arrow. You can select and move multiple values at the same time.

– Selected Severities

By default, no event severity values are displayed in this list. Move items from the Available Severities list to this list to show only those values in the Configuration and Event Timeline. For example, to show only Critical and Major in the Configuration and Event Timeline, move the Critical and Major items from the Available Severities list to the Selected Severities list.

To remove an item from this list, select the item and click the left-pointing arrow. You can select and move multiple values at the same time (see Figure 2-14).

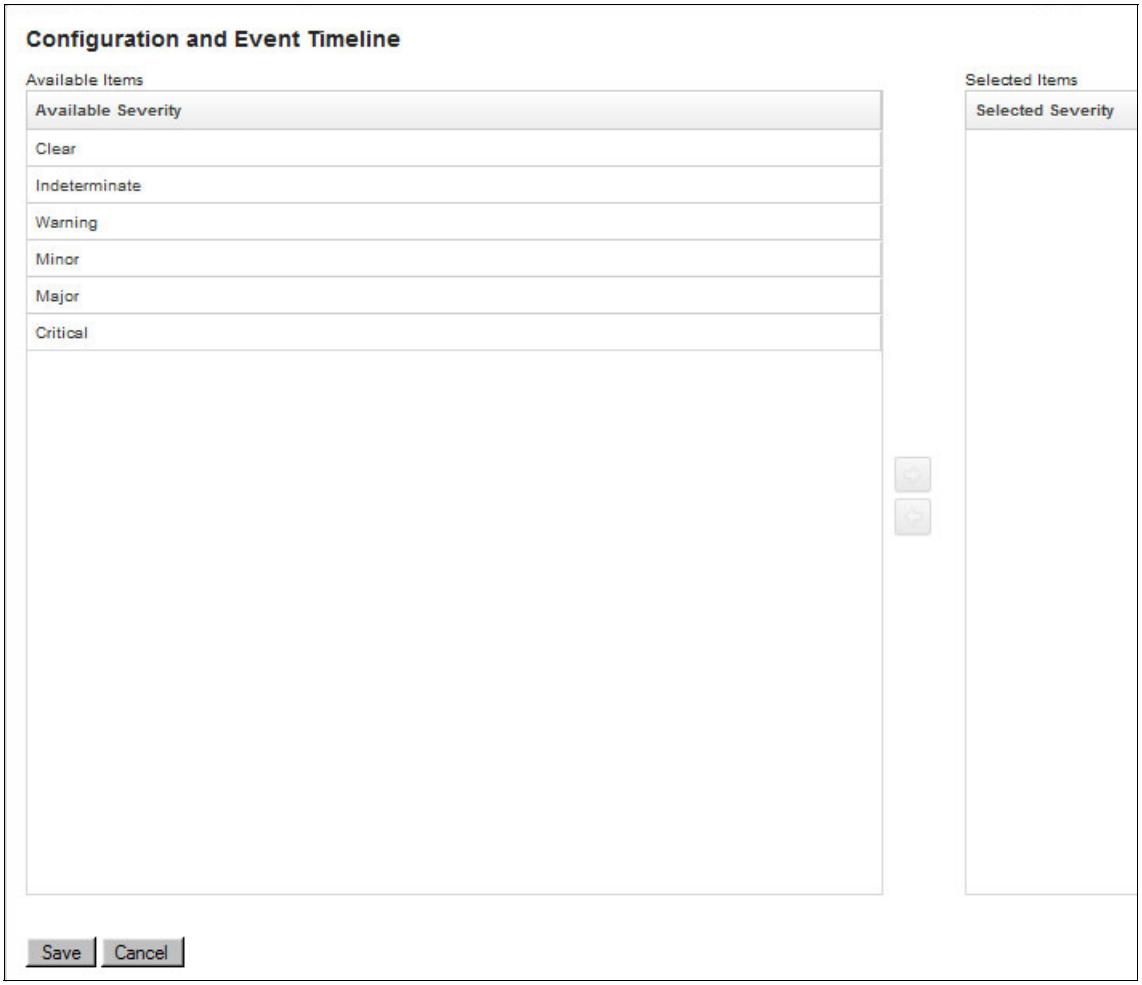

Figure 2-14 Detailed Configuration and Event Timeline view

2.4.5 Scenario personas

The following personas are featured in this scenario:

- Annette: Annette is an IT Operator. In her role, she wants to quickly identify outages or incidents and resolve them fast.
- Brock: Brock is an IT Subject Matter Expert (SME). As a SME, he wants to increase efficiency and agility in the Operations Center (OC) by reducing noise, identifying problems faster, and issuing resolution with the power of automation and actionable insights.

2.4.6 Viewing the Network Health Dashboard

For the scenarios in the next sections, complete these initial steps:

1. Log in to the Dashboard Application Services Hub, as shown in Figure 2-15.

Figure 2-15 Log in prompt into Dashboard

2. To access the Network Health Dashboard, browse to **Incident** → **Network Health Dashboard**, as shown in Figure 2-16.

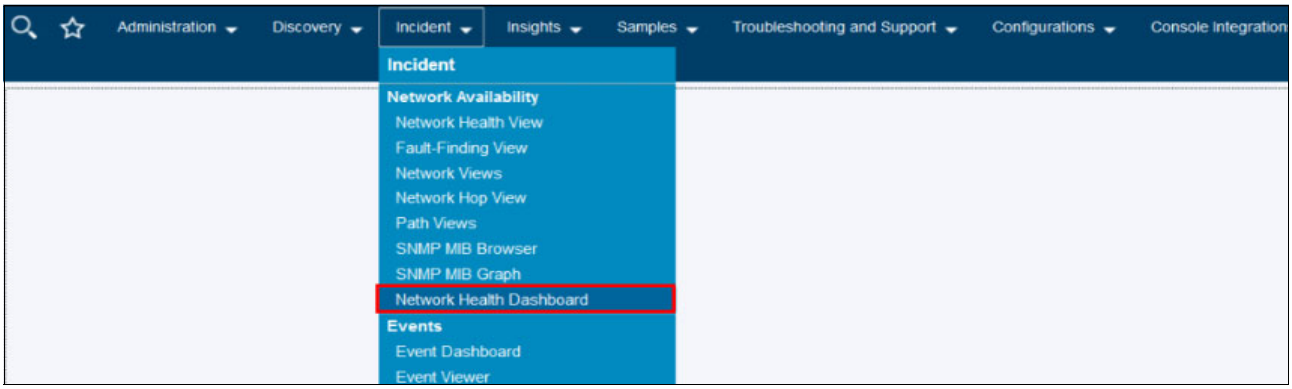

Figure 2-16 Accessing the Network Health Dashboard.

2.4.7 Scenario 1: IT Operator uses Network Health Dashboard and Runbook Automation and Alert Notification

In this scenario, an IT Operator uses Network Health Dashboard and Runbook Automation to be quickly notified of an outage and takes steps to resolve the outage.

This scenario features the following steps:

1. Annette, the IT Operator, receives alert notification of an outage. She receives the alert on her mobile device and then follows the link that is sent with the alert to login and view the alert details.

   She can see the details of the incident, as shown in Figure 2-17.

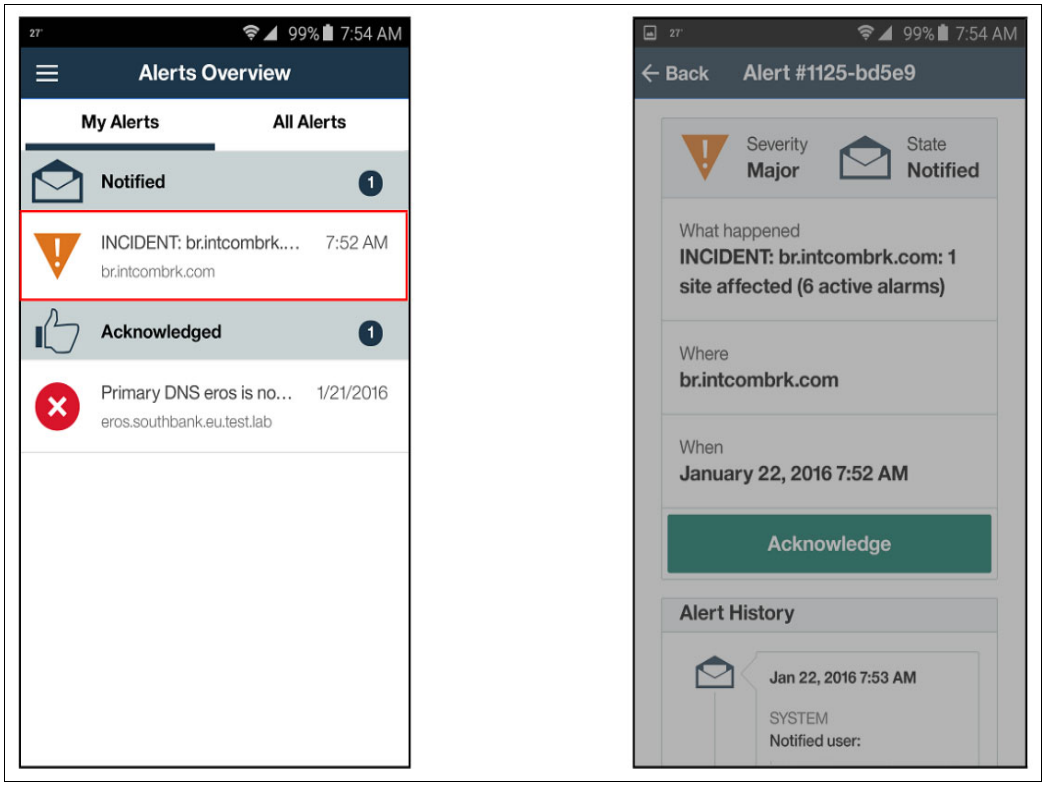

Figure 2-17 Alert notifications received

2. Annette sees that multiple clients cannot access the Mobile Trade server. She can isolate the events and network details that are related to this Service by clicking **Mobile Trade Server** under the Network Views, as shown in Figure 2-18.

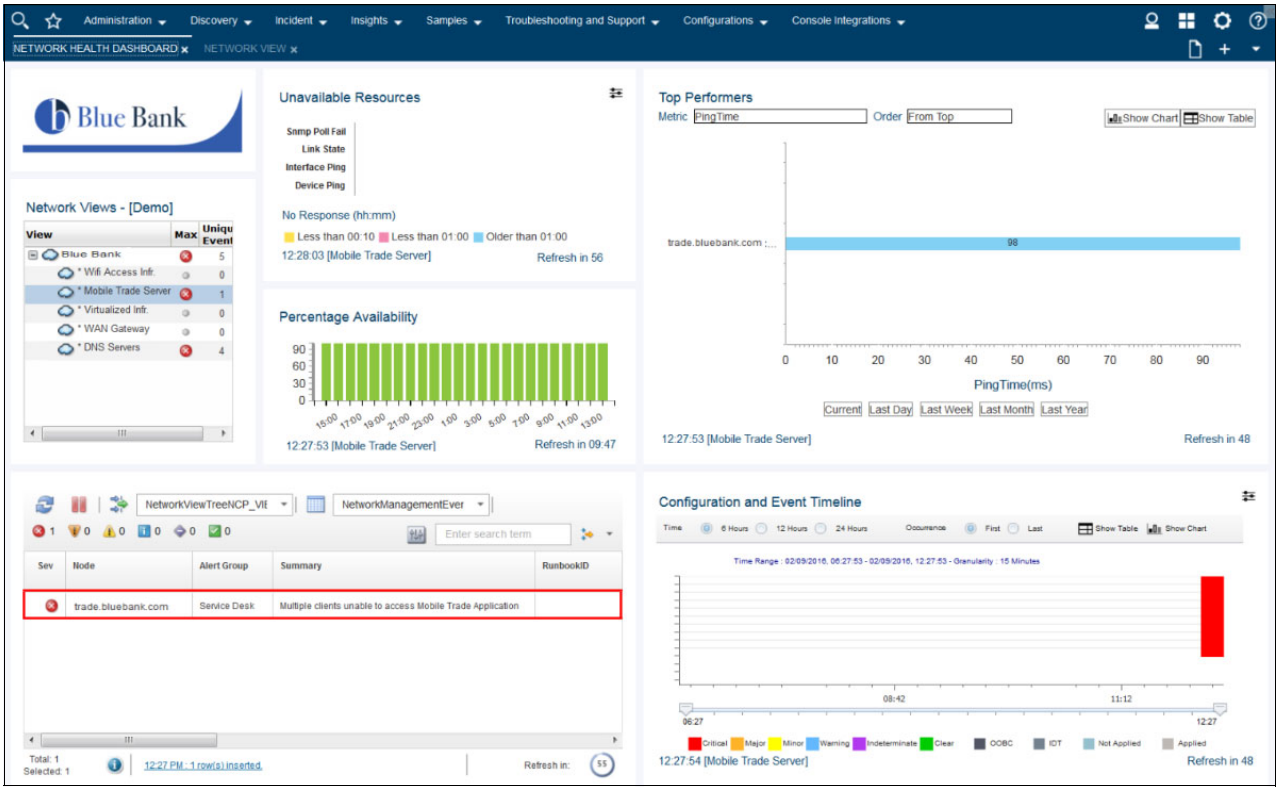

Figure 2-18 View Network events that are related to the Mobile Trade Server

3. Annette also finds a BIND (DNS) failure event on primary DNS Server.

   In addition, Netcool Performance Insight (NPI) data shows a DNS traffic loss event on the Primary DNS server and a traffic increase event on Secondary DNS server, as shown in Figure 2-19.

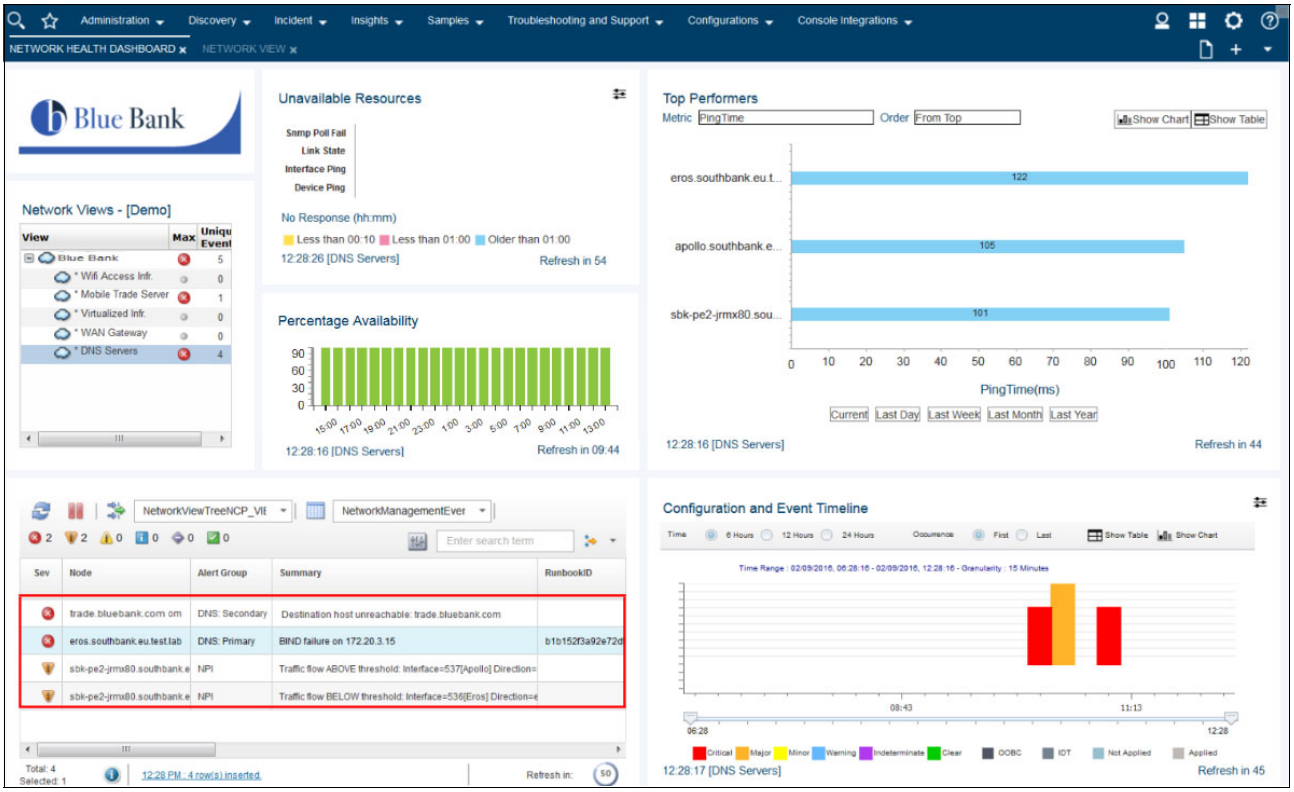

Figure 2-19 View events that are associated with the DNS Servers

4. Annette checks Wifi Access Infrastructure, Virtualized Infrastructure Utilization, WAN Gateway, and Service availability. All appear to be OK, as shown in Figure 2-20.

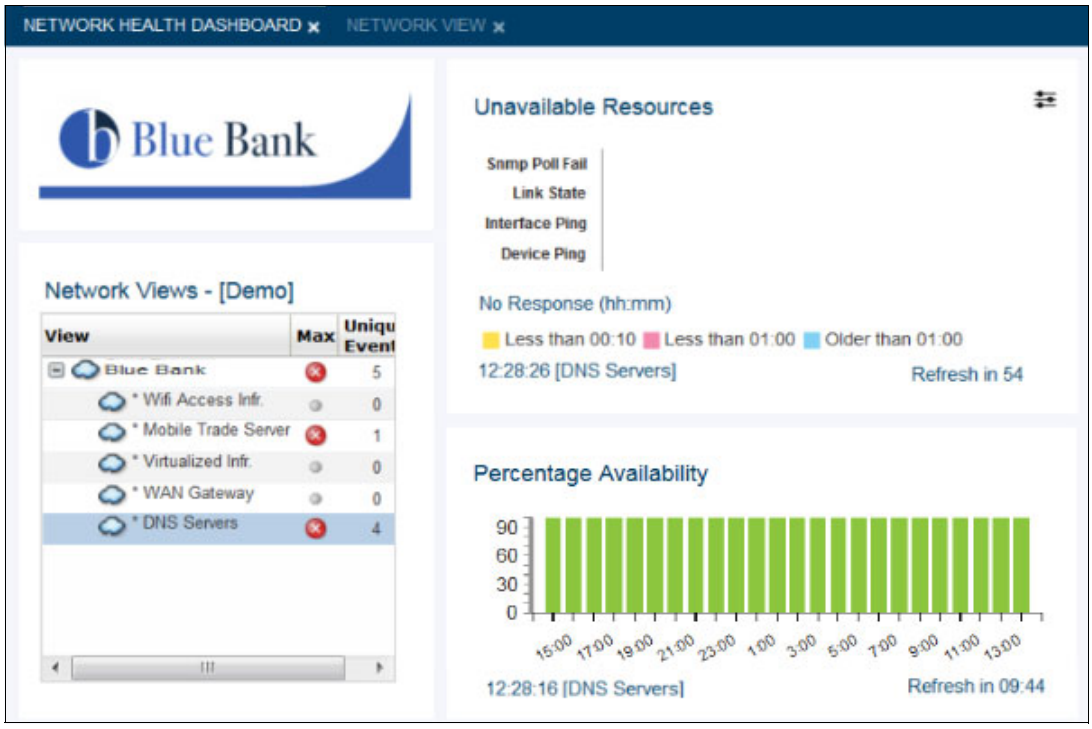

Figure 2-20 Verify network services that might be problematic

5. A Runbook ID that is associated with the DNS Service is shown automatically by the alert viewer, as shown in Figure 2-21. Annette concludes that the issue is DNS-related based on confirmation that there are no issues with Wifi Access infrastructure, Virtualized infrastructure Utilization, WAN Gateway, and Service Availability, coupled with Netcool Operations Insight having identified a BIND (DNS) Failure in Step 3.

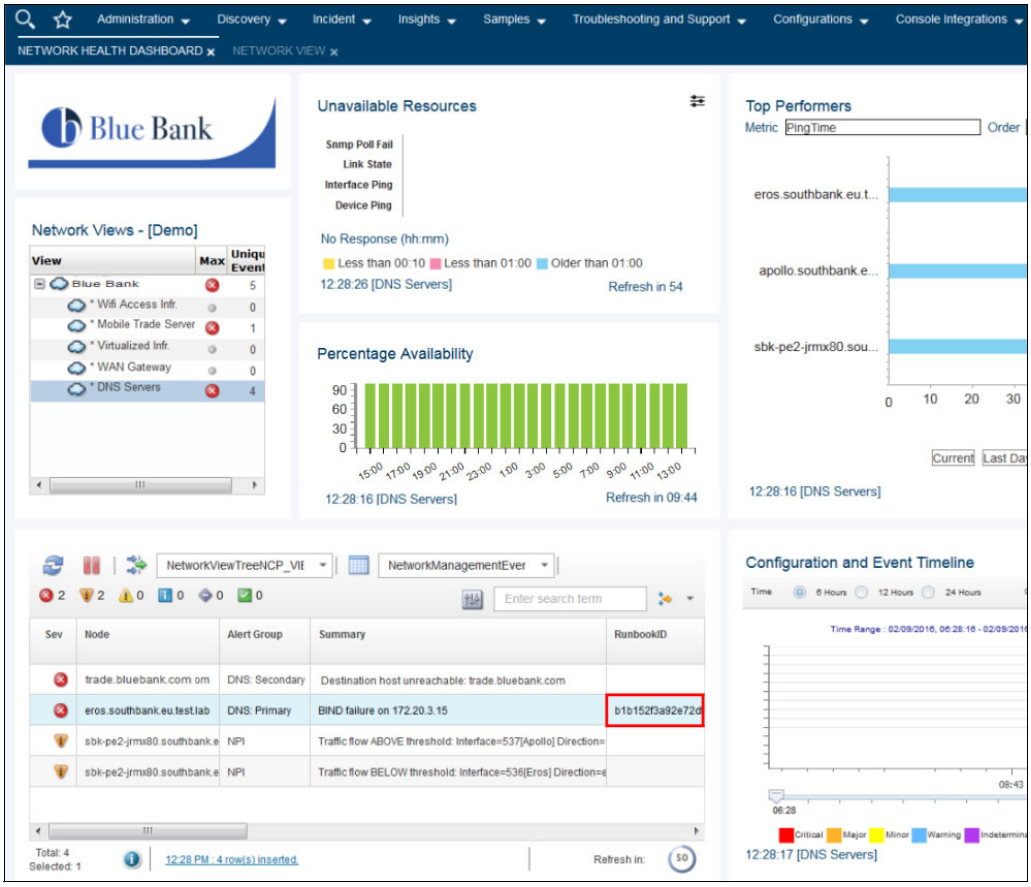

Figure 2-21 Network Health Dashboard Events list show Runbook IDs that are associated with events

6. Annette starts the Runbook to restart Primary DNS, as shown in Figure 2-22.

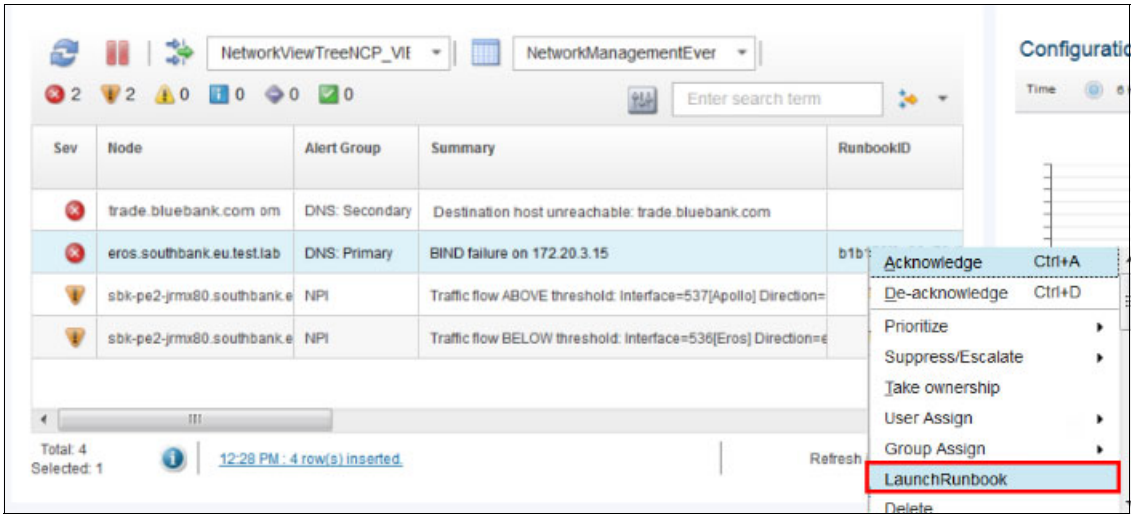

Figure 2-22 Annette can start a Runbook that is associated with a particular event type

7. Annette runs the Runbook by following the steps to Restart the Primary DNS server. Because it was successful, she records that the Runbook works (see Figure 2-23).

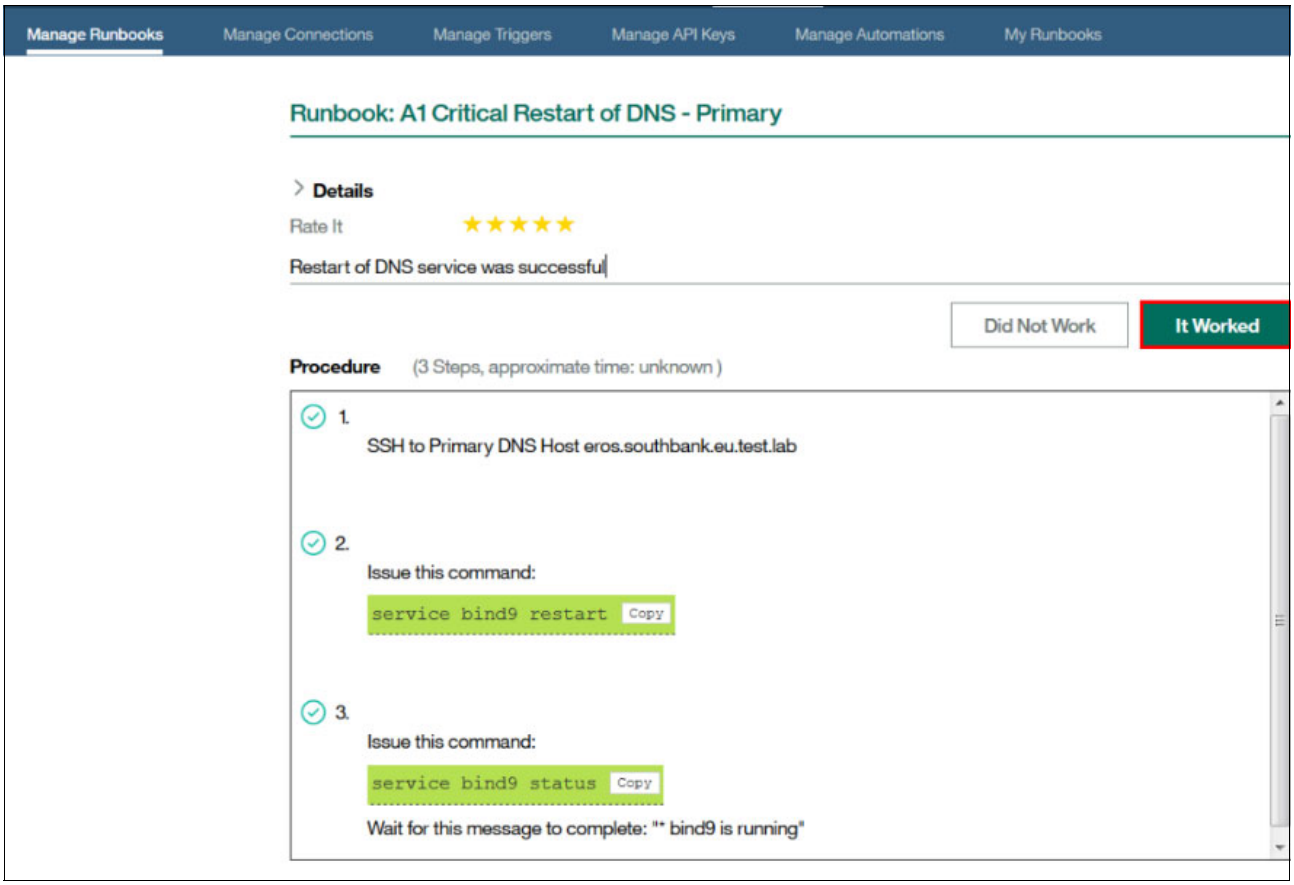

Figure 2-23 Runbook to restart the Primary DNS Service executed successfully

More information: For more information about Runbook Automation, see this website:
https://ibm.biz/Bdrx2d

8. Annette observes that the Secondary DNS server cannot resolve the IP of `trade.bluebank.com`, as seen in Figure 2-24.

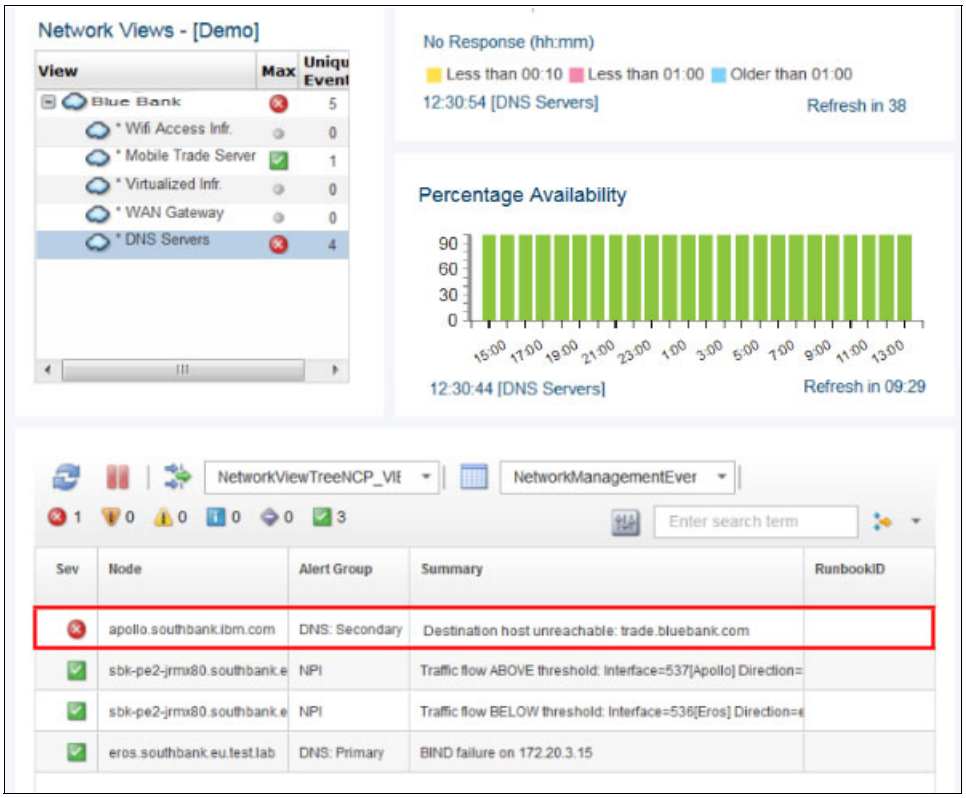

Figure 2-24 Secondary server cannot resolve the address of the Mobile Trade Server

9. Annette escalates the issue to Brock (the SME) by right-clicking the Event that is associated with the Secondary DNS server, and clicking **Suppress/Escalate** → **Escalated-Level 3**, as shown in Figure 2-25.

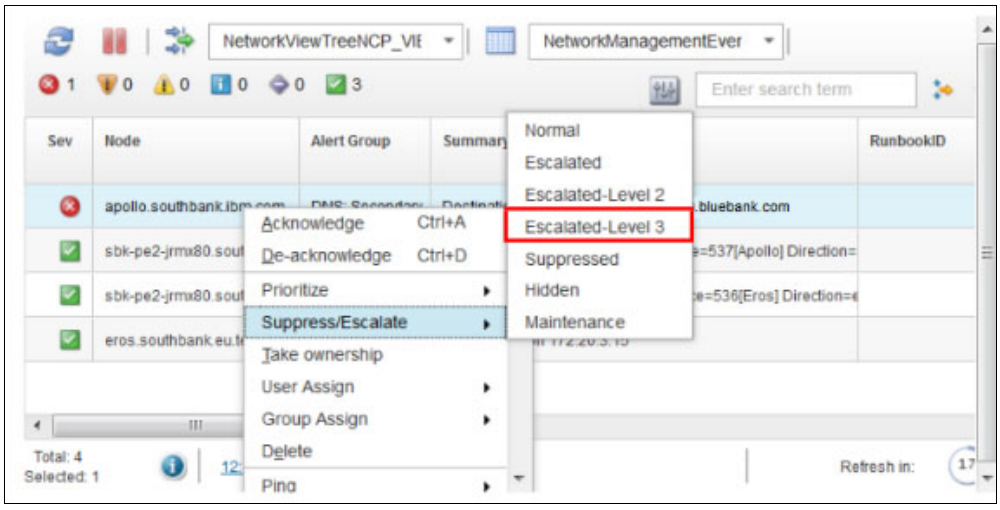

Figure 2-25 Escalating events by using the Network Health Dashboard

2.4.8 Scenario 2: IT SME using Network Health Dashboard

In this scenario, IT SME Brock uses the Network Health Dashboard to quickly resolve an escalated Event. Brock receives an escalated Alert Notification that Primary DNS is down from Annette (see 2.4.5, “Scenario personas” on page 45) and continues the problem determination process as described in the following process:

1. Brock logs in and brings up the Network Health Dashboard by clicking **Incident** → **Network Health Dashboard**, as shown in Figure 2-26.

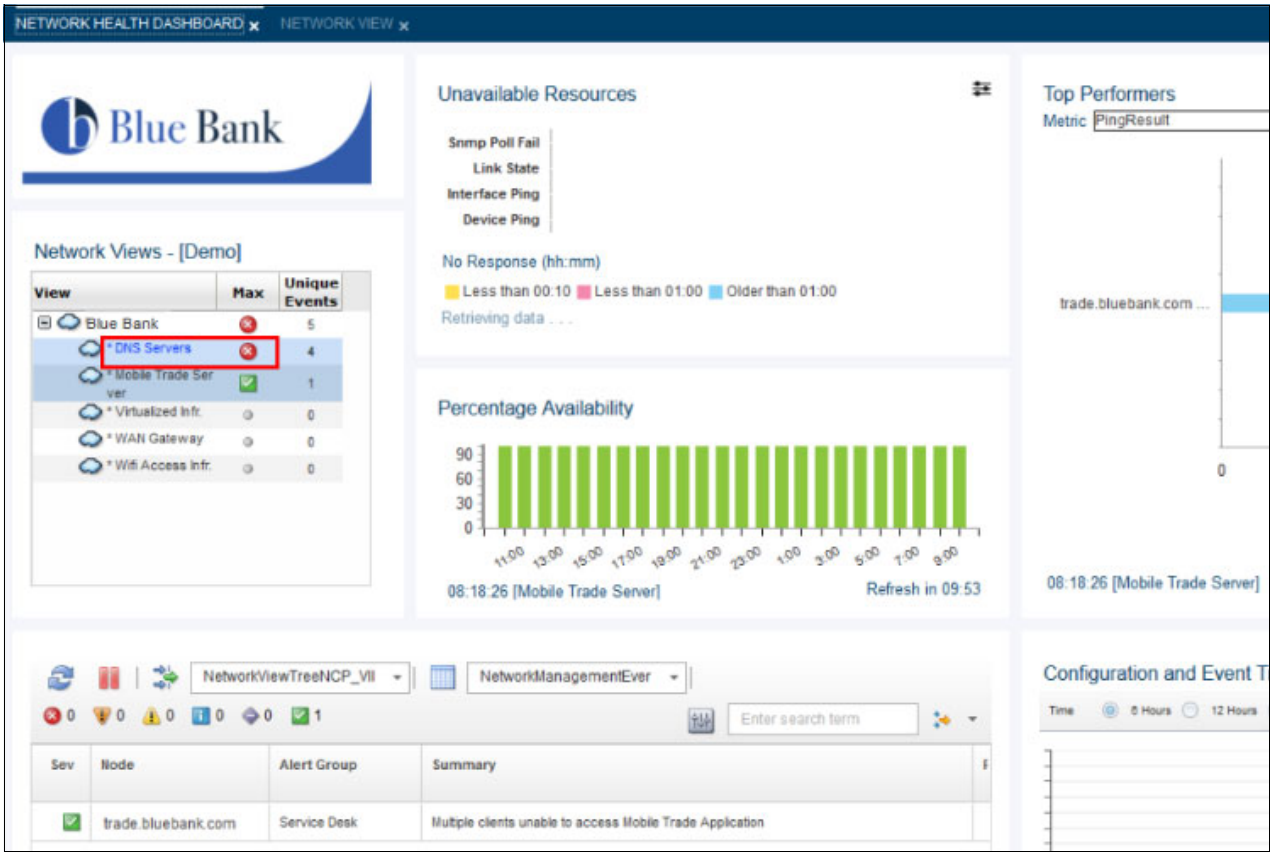

Figure 2-26 Network View shows issues with the DNS Servers

Brock observes that the Secondary DNS server cannot resolve the internal web service.

2. Brock reviews the Network topology that is associated with the DNS Servers by double-clicking the DNS Servers line in the Network Views box (see Figure 2-27).

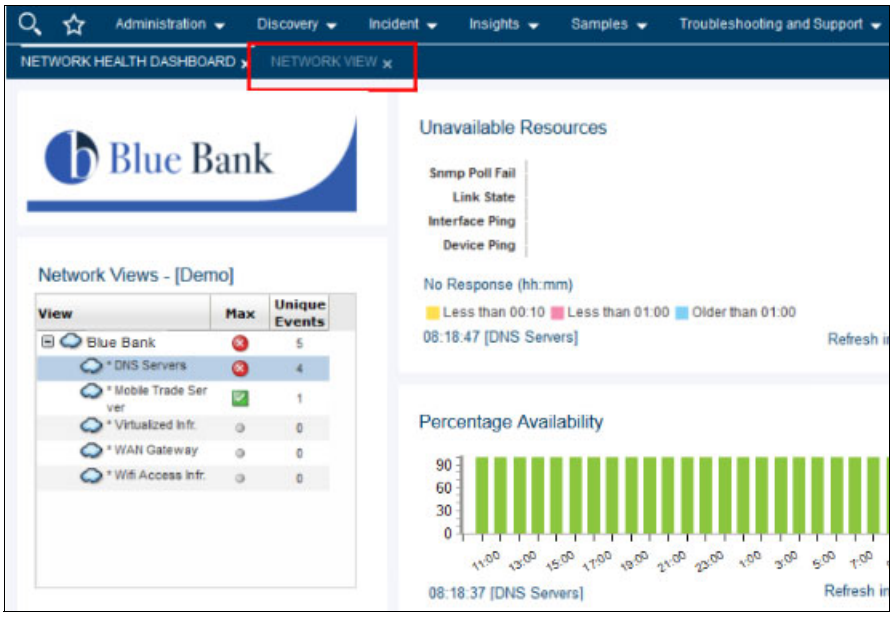

Figure 2-27 Network Views box

3. Brock clicks the Network View tab for the DNS network topology, as seen in Figure 2-28.

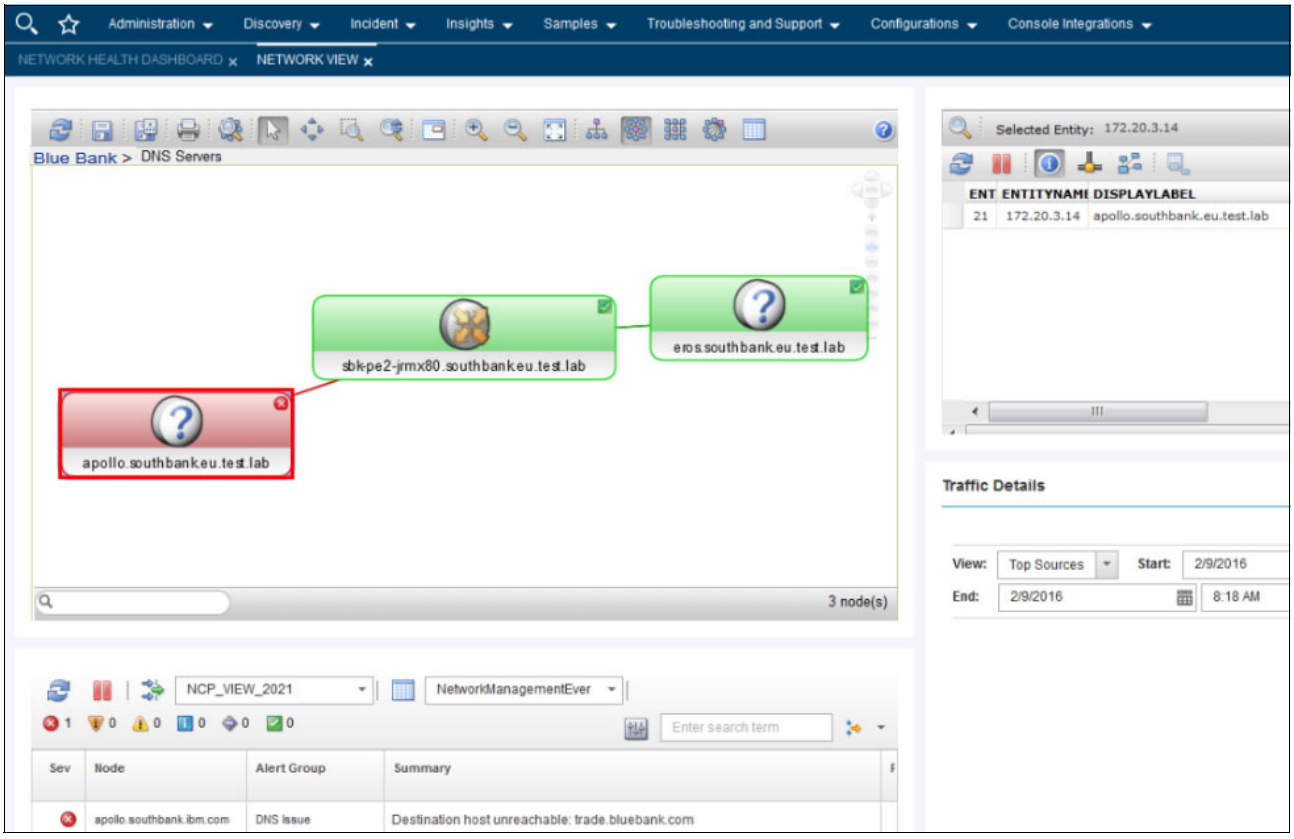

Figure 2-28 Network view of the Primary and Secondary DNS servers

4. Brock uses the `nslookup` tool to check the DNS configuration. He selects **Webtools** → **Launch WebTools GUI**, as shown in Figure 2-29.

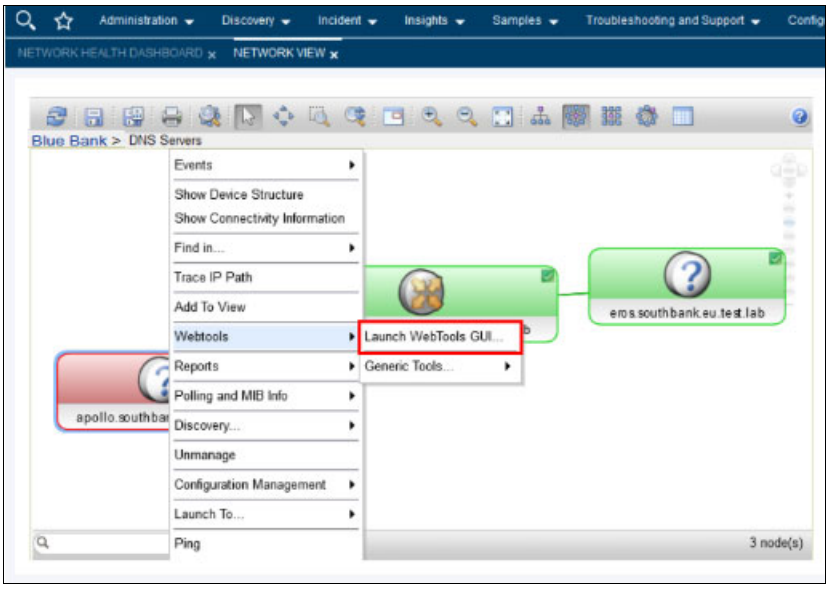

Figure 2-29 Starting the WebTools GUI by using Network View

5. Brock starts the `nslookup` tool from the WebTools GUI tab, as shown in Figure 2-30.

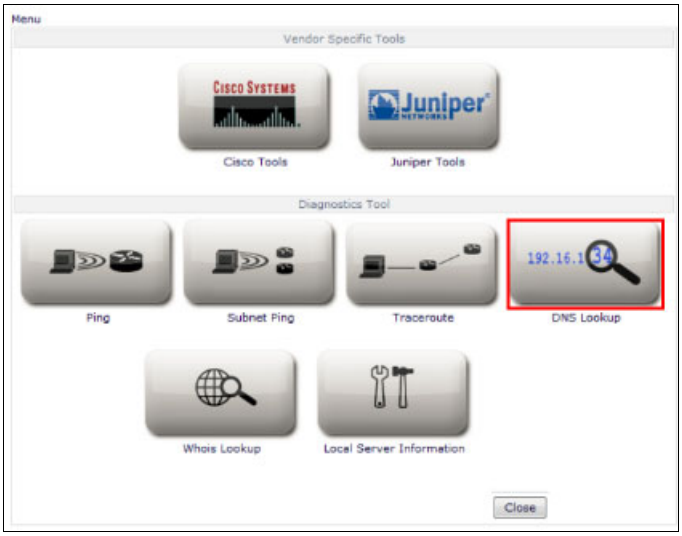

Figure 2-30 Network Manager Webtools GUI that is used to start DNS lookups

6. Brock finds that the DNS configuration is out of synch between the primary and backup servers, as shown in Figure 2-31.

Figure 2-31 Primary and Secondary DNS servers are resolving the IP address

7. Brock reviews the network switch that is connected to the Secondary DNS server in the Network View to investigate the connectivity statistics of that server (see Figure 2-32).

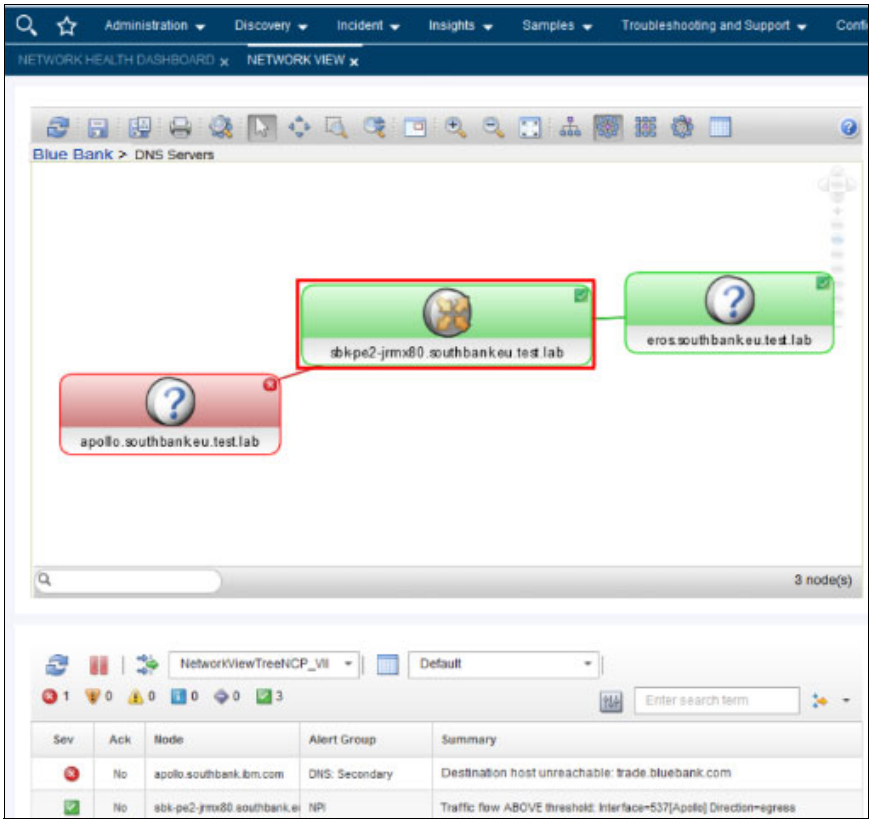

Figure 2-32 Reviewing the switch statistics of the Secondary DNS server port

8. Brock drills down to the Port where the Secondary DNS server is connected by clicking the **Show Interfaces** icon within the Structure Browser, as shown in Figure 2-33.

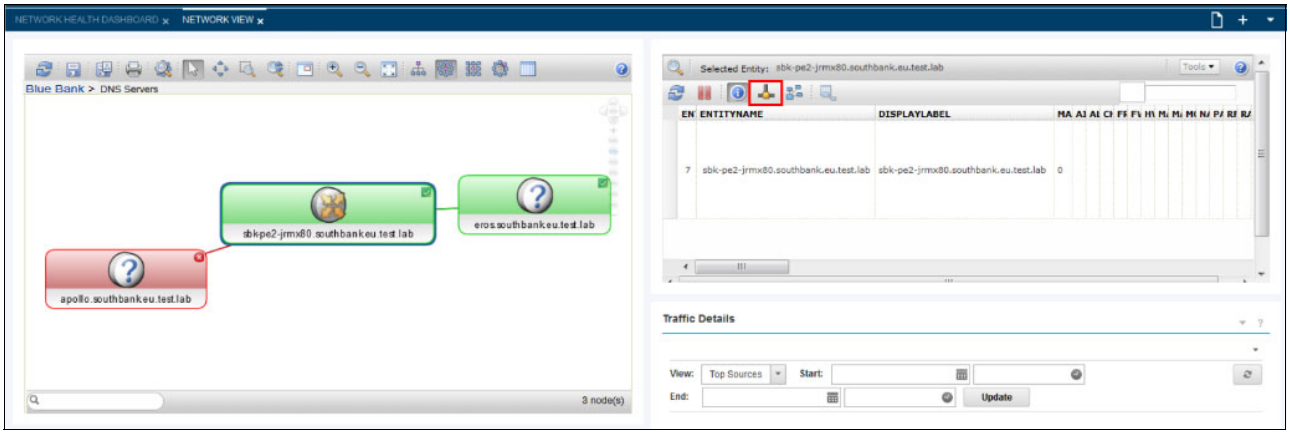

Figure 2-33 Selecting the Show Interfaces icon within the Structure Browser for the secondary DNS server's switch

9. Brock identifies the specific switch port that the secondary DNS server is connected to and views the traffic for that port as provided from Netcool Performance Insight (NPI), as shown in Figure 2-34.

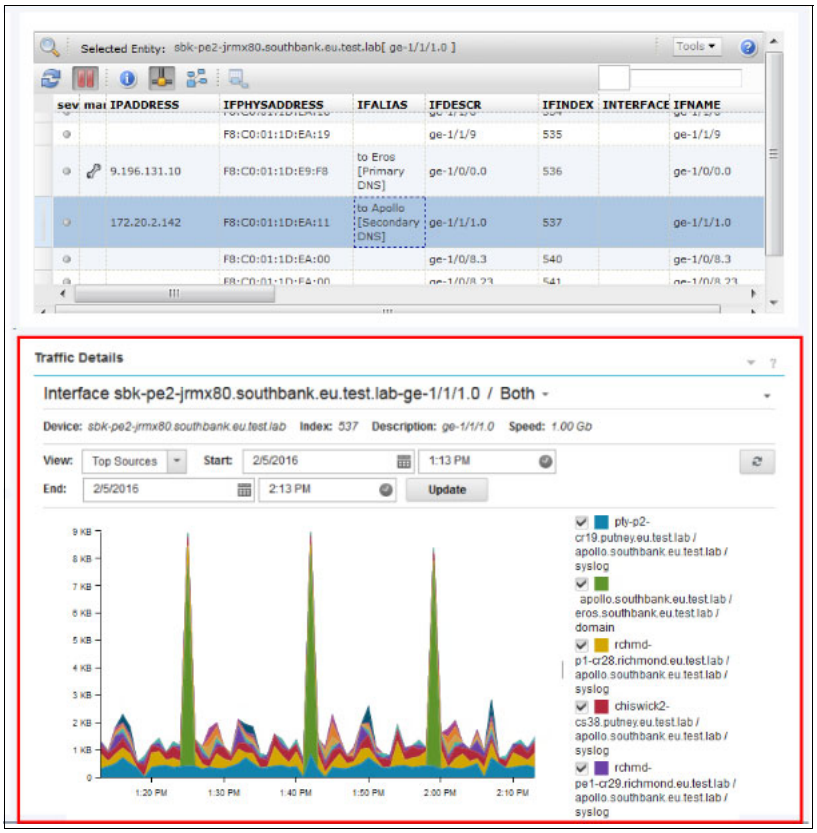

Figure 2-34 Port connecting the DNS server is identified by using the Structure Browser

10. By using Netcool Performance Insight (NPI) flow, Brock confirms that there is no DNS zone traffic to keep IP addresses synchronized between the Primary and Secondary DNS servers now. Instead, he sees traffic from the Secondary to Primary DNS but not from Primary to Secondary, as shown in Figure 2-35.

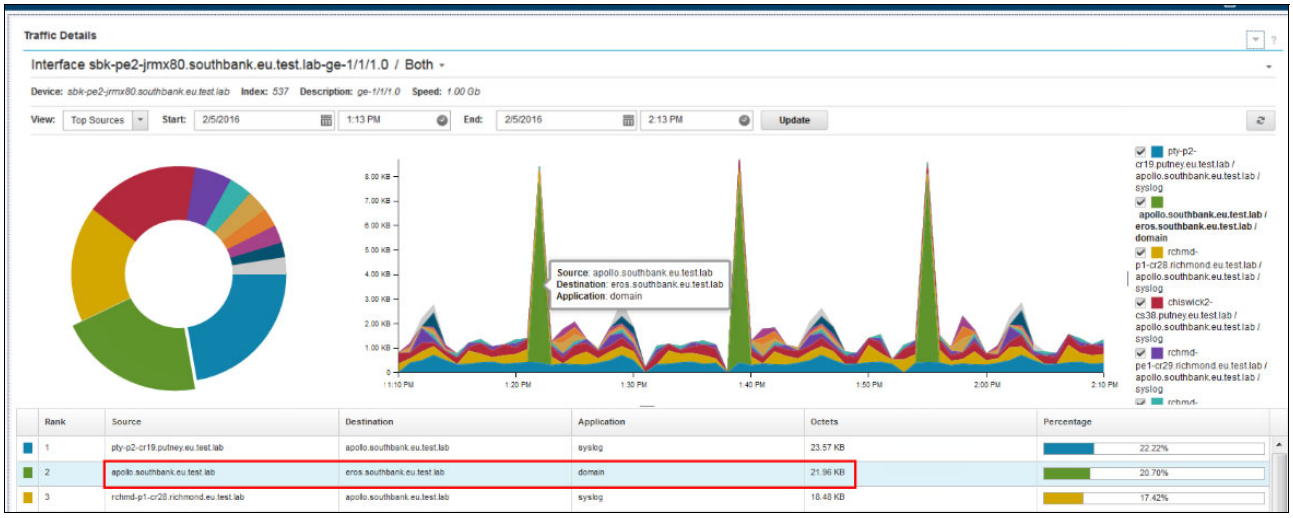

Figure 2-35 No DNS Zone traffic between Primary and Secondary DNS servers

11. Brock reviews traffic from one week ago and observes that DNS Zone traffic was flowing between Primary and Secondary DNS servers, as shown in Figure 2-36.

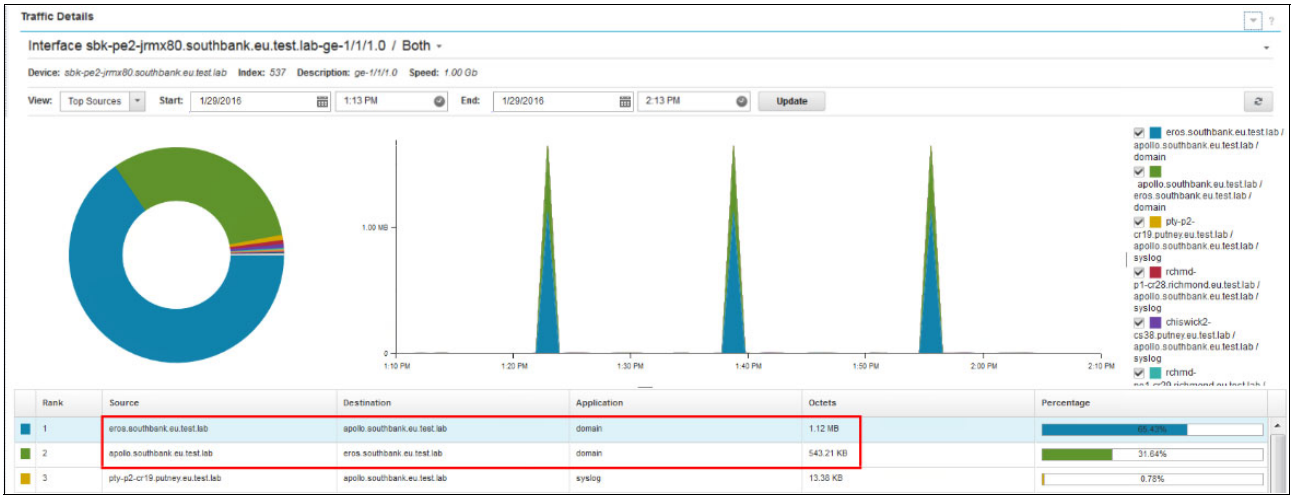

Figure 2-36 Reviewing traffic from one week previous

12. Within the Network View tab, Brock starts the Netcool Configuration Manager in the context of the Secondary DNS Server, as shown in Figure 2-37.

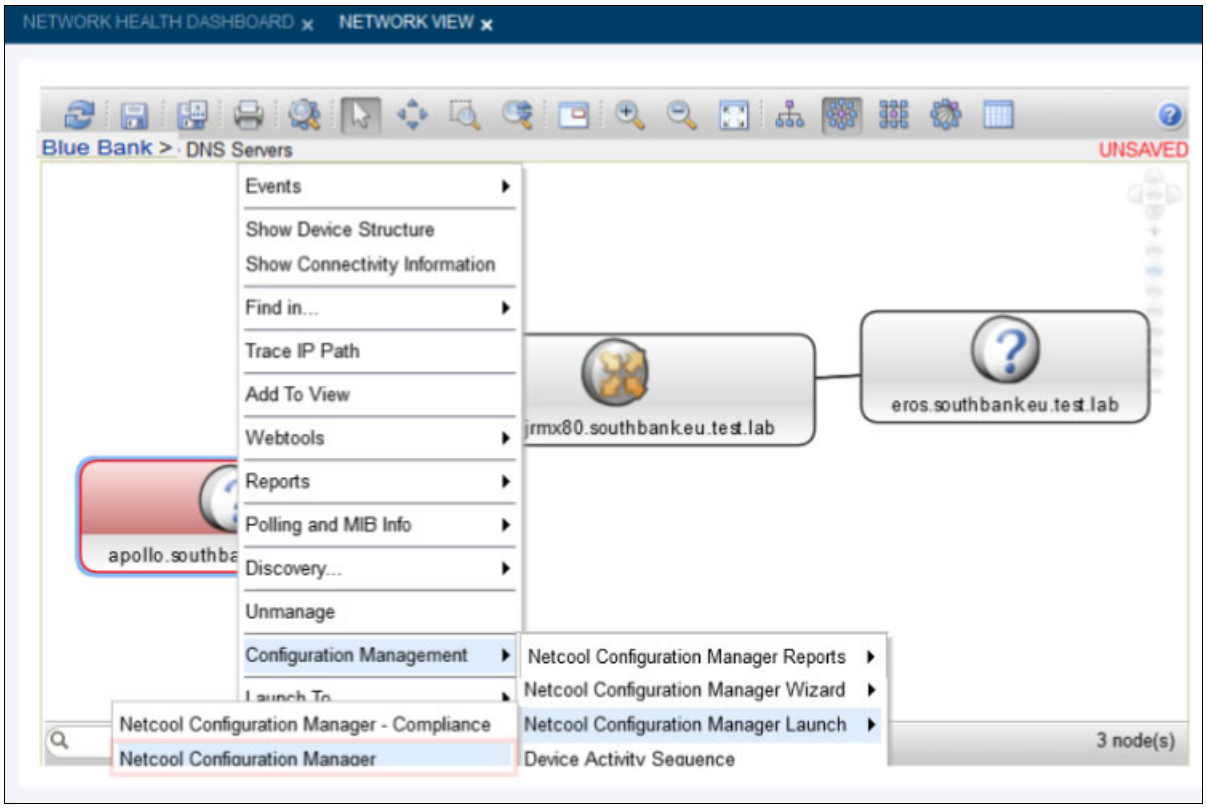

Figure 2-37 Starting Netcool Configuration Manager for the Secondary DNS Server

13. The Netcool Configuration Manager starts and Brock selects the **Secondary DNS Server** to review any changes that were made on it, as shown in Figure 2-38.

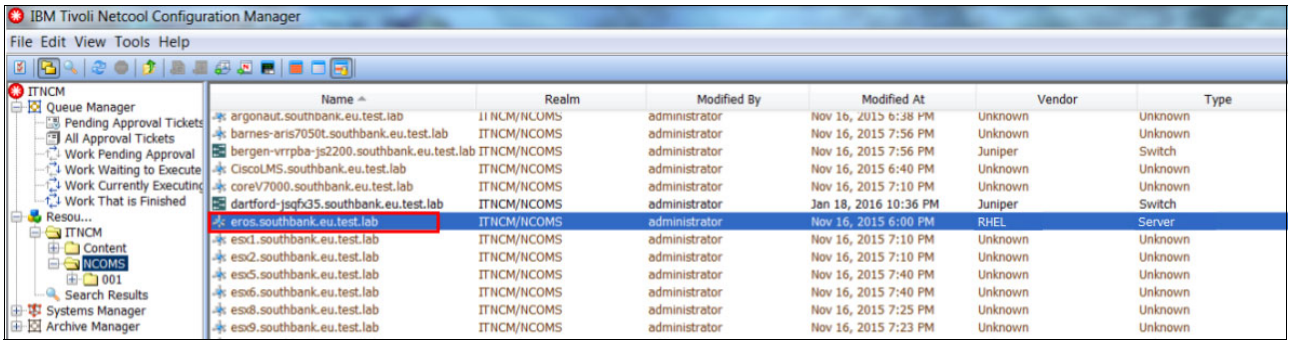

Figure 2-38 Changes that were made on the Secondary DNS Server

14. Brock verifies that the TCP port 53 is blocked on the server firewall. Then, by using the Netcool Configuration Manager, he starts a manual Change to reopen the port (see Figure 2-39).

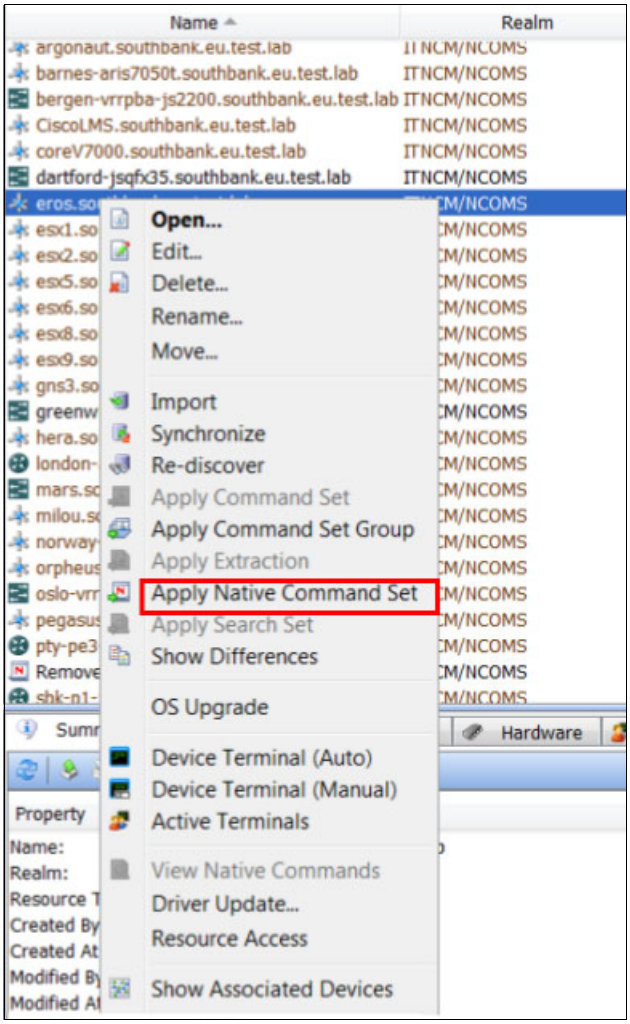

Figure 2-39 Applying a manual change in Netcool Configuration Manager

15. Brock selects the Native Command to remove the DNS (TCP port 53) block (see Figure 2-40).

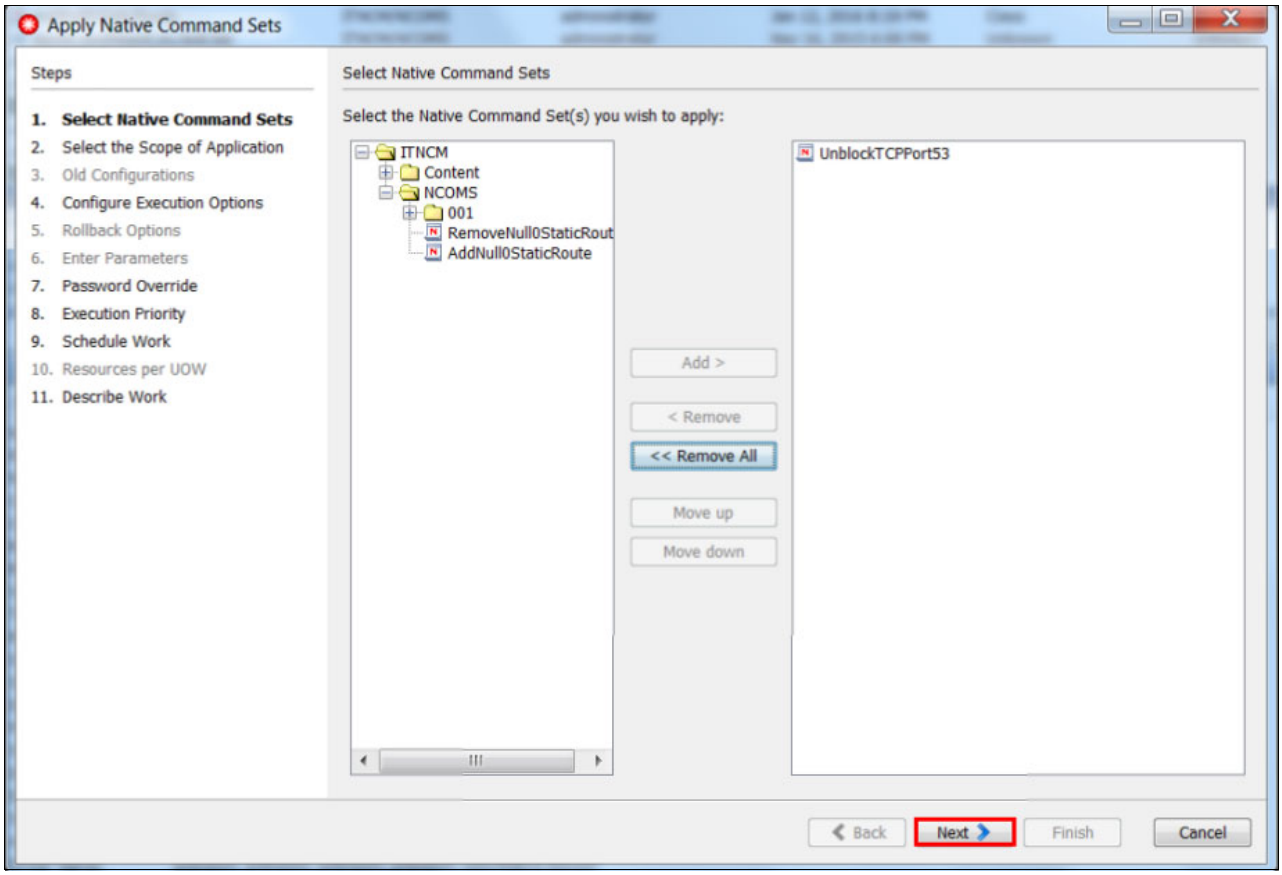

Figure 2-40 Implementing the change to unblock TCP port 53 on the Secondary DNS Server

16. Brock submits this UoW in the Netcool Configuration Manager. He receives confirmation that it was successfully run, as shown in Figure 2-41.

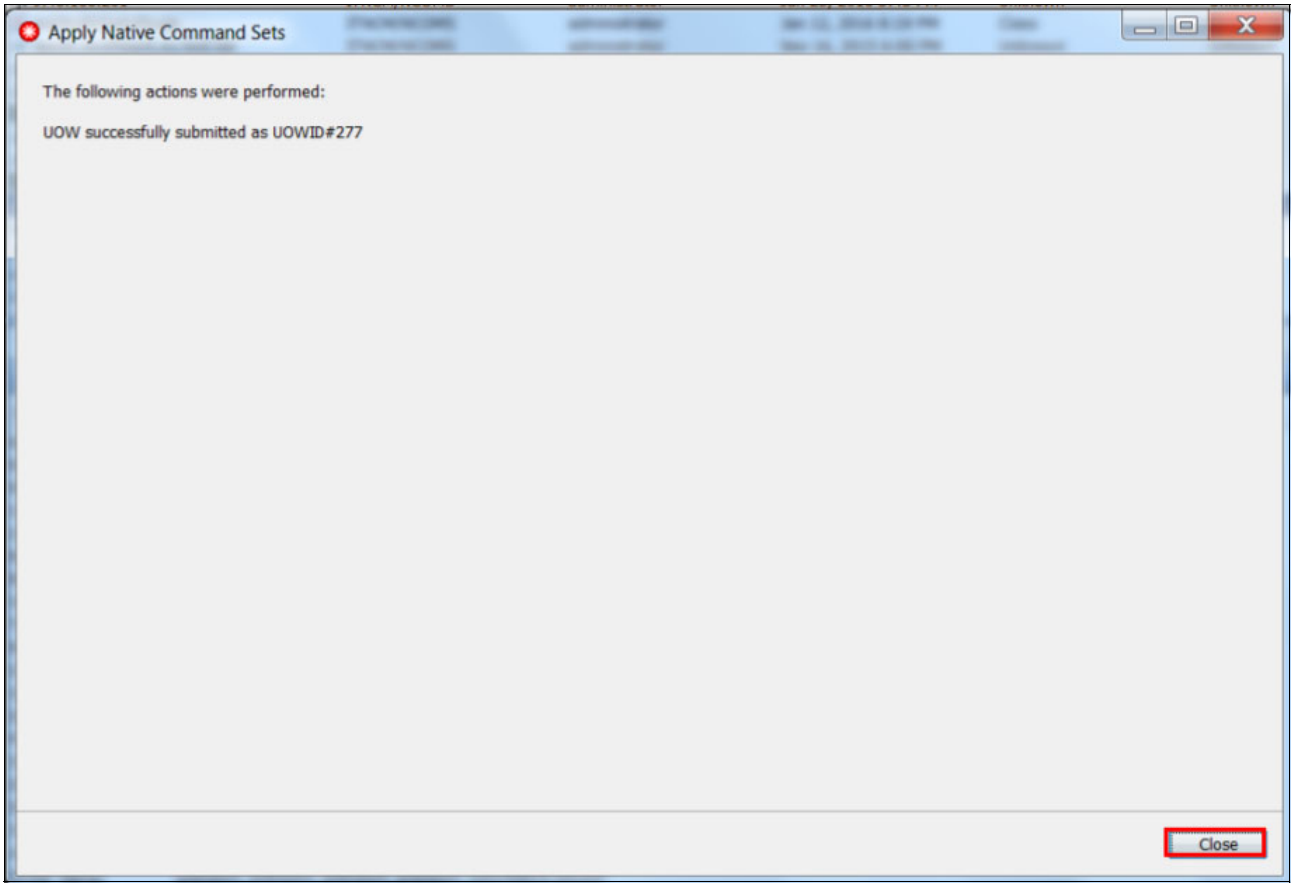

Figure 2-41 Netcool Configuration Manager Command Set applied successfully

17. By using the Network Health Dashboard, Brock verifies that DNS Network Operations are back to normal, as shown in Figure 2-42.

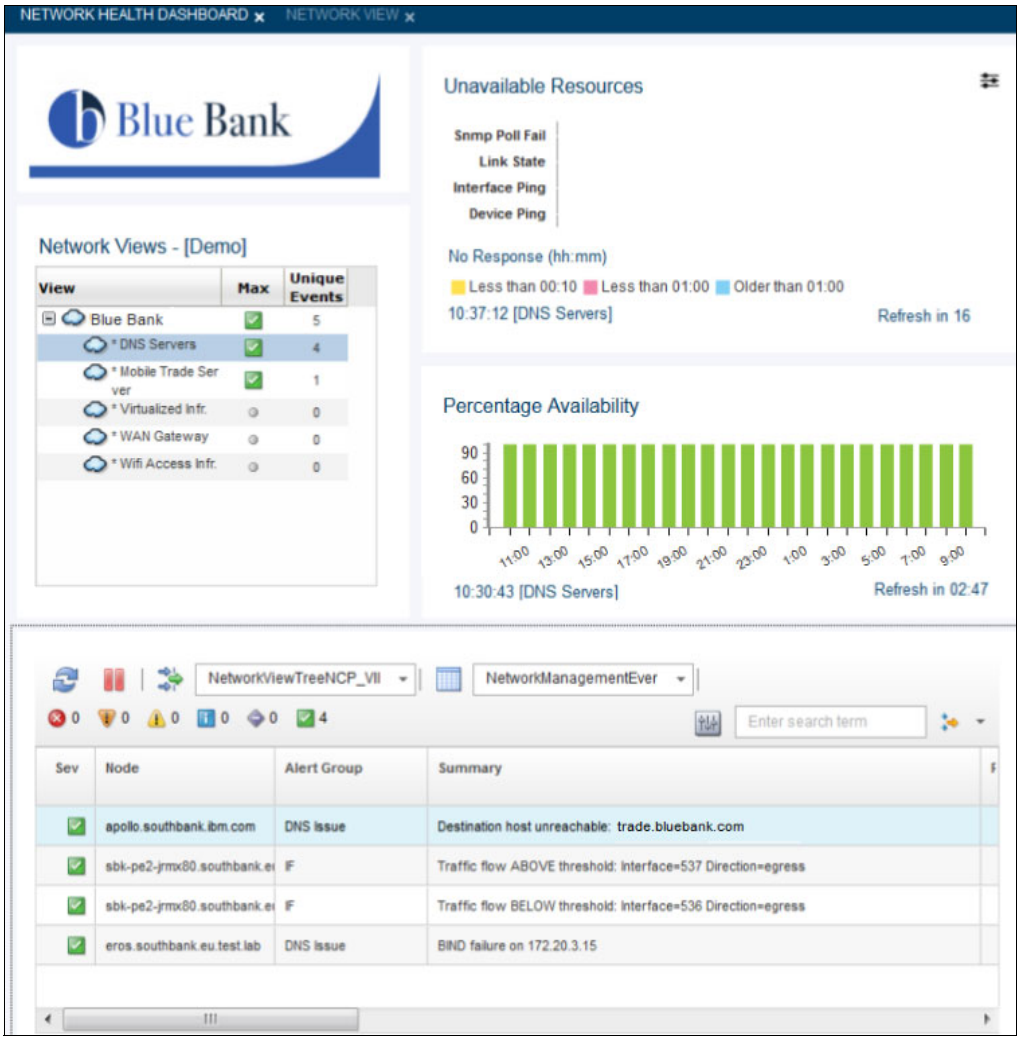

Figure 2-42 Confirmation that DNS Service is restored in a single view with Network Health Dashboard

2.5 Summary

As described in our scenarios, the Network Health Dashboard is a significant and beneficial new feature of Networks for Operations Insight. By using it, the IT Practitioner can have a single, consolidated view of IBM Tivoli Network Manager, IBM Tivoli Netcool Configuration Manager, IBM Network Performance Insight, and OMNIbus events.

When fully integrated, it can identify and troubleshoot network outages fast and resolve them quickly.

Note: For more information about the Network Health Dashboard configuration, see this website:

https://ibm.biz/BdrxQh

3

Geographic Discovery and Mapping

This chapter features a scenario that describes how to automatically include and use Geographic Information System (GIS) data for network devices that are discovered by the Network Manager. Also described is how Geographic mapping can help operators and other users visually orient problems that are based on location.

This early-release capability is available for download and is fully supported by the Development team. For more information about the link to download the distribution package, see 3.5, “Summary” on page 81.

This chapter includes the following topics:

3.1 Scenario description

This scenario describes how you can use the Geographic maps to gain an at-a-glance understanding of where the devices are and how they are connected, including the active status of the devices and links. It is also an alternative mechanism that can be used to drill down for more information about devices and links.

3.1.1 Business value

Geographic mapping shortens time to triage events or diagnose problems by understanding the geographic layout around the issue intuitively. The geographic orientation helps operators align outside knowledge with the network problems they are seeing. It is helpful not only for location sensitive devices on cell towers, but also to see status across the network backbone and core networks around the world. Operators can quickly see all the device details for alarms, inventory, contacts, and connections at their fingertips. Administrators can set up flexible views to match areas of responsibility that are tied to operators based on location (latitude-longitude dimensions), device types, services, and so on.

3.2 Geographic map basics

Before describing the scenario steps, we describe how to use the Geographic map in Network Manager.

3.2.1 Use of the Geographic map

The Geographic map displays the individual devices according to their latitude and longitude values. Those devices that are closely collocated are aggregated into a *Building*.

Complete the following steps to use the Geographic map:

1. Right-click the **Building site** icon to see the list of all devices there, including their status, event counts, connectivity, and details.

2. Use the right-click menu on any row to use the standard context tools for browsing to other topology views, structure browser, Top Performers, or the diagnostic tools, such as Ping and Trace route, as shown in Figure 3-1.

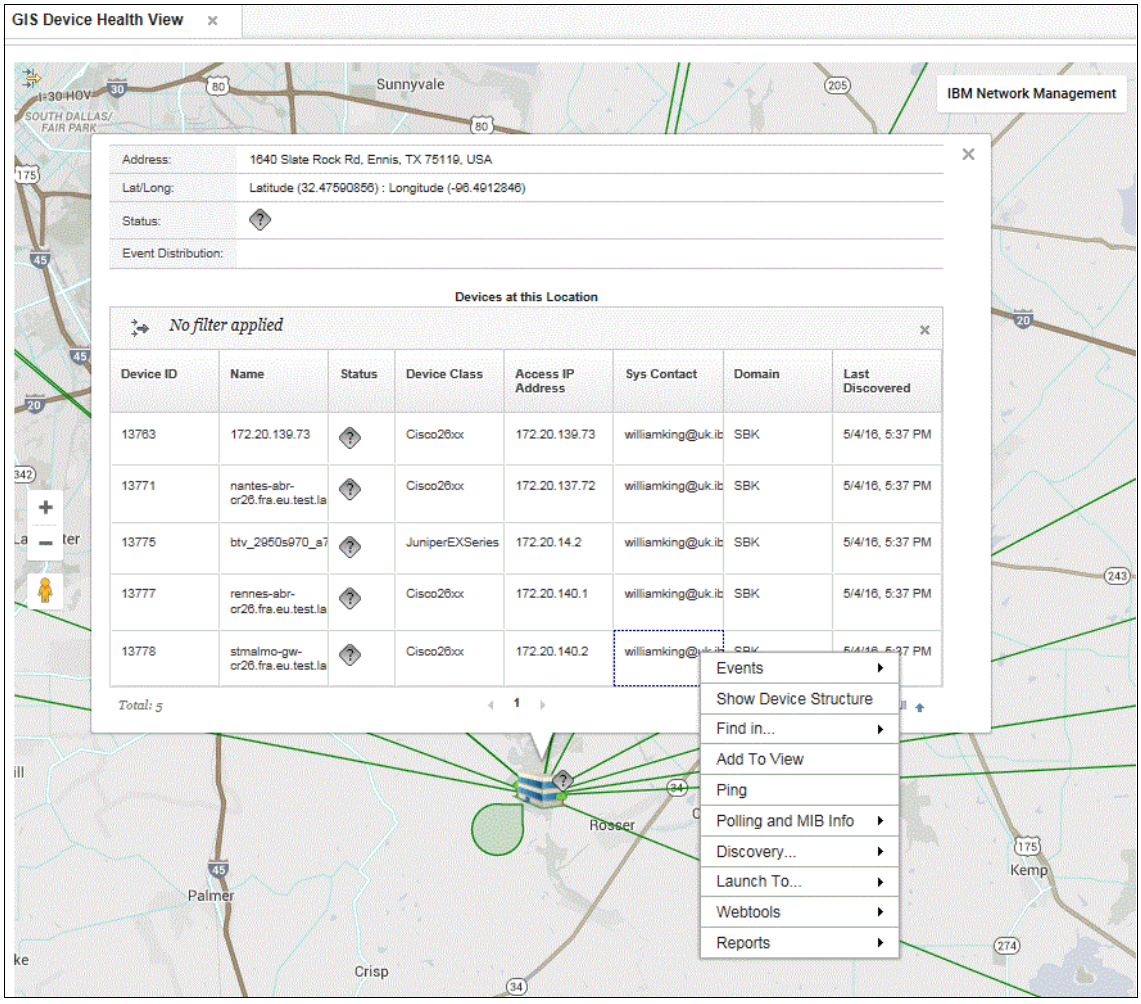

Figure 3-1 Right-clicking building site link

3. You can also right-click a link to show all of the connections it represents. Right-click a single row to perform actions on that link, as shown in Figure 3-2.

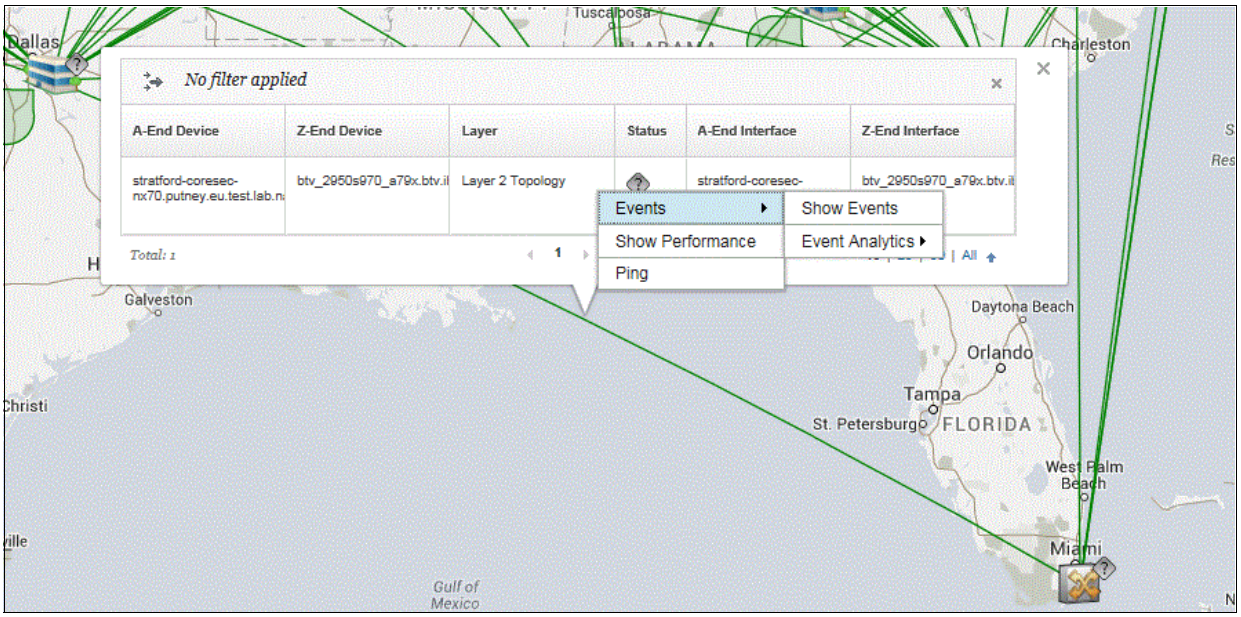

Figure 3-2 Right-click a single row

4. The device status is taken from the active events and aggregated in the case of the Building site. Use the Map Configuration icon at top left to control what event severities to include.
5. Connectivity is also aggregated from Buildings and by right-clicking the link, a table of all the links and their status is shown. Use the Map Configuration icon to control what connectivity layers to display.

   The Map Configuration icon can also be used to dynamically filter the devices on display by Network Manager domain and device type.

3.2.2 Use of the Health view

Access this page by click in the DASH task menu **Incident** → **GIS Device Health View**, as shown in Figure 3-3.

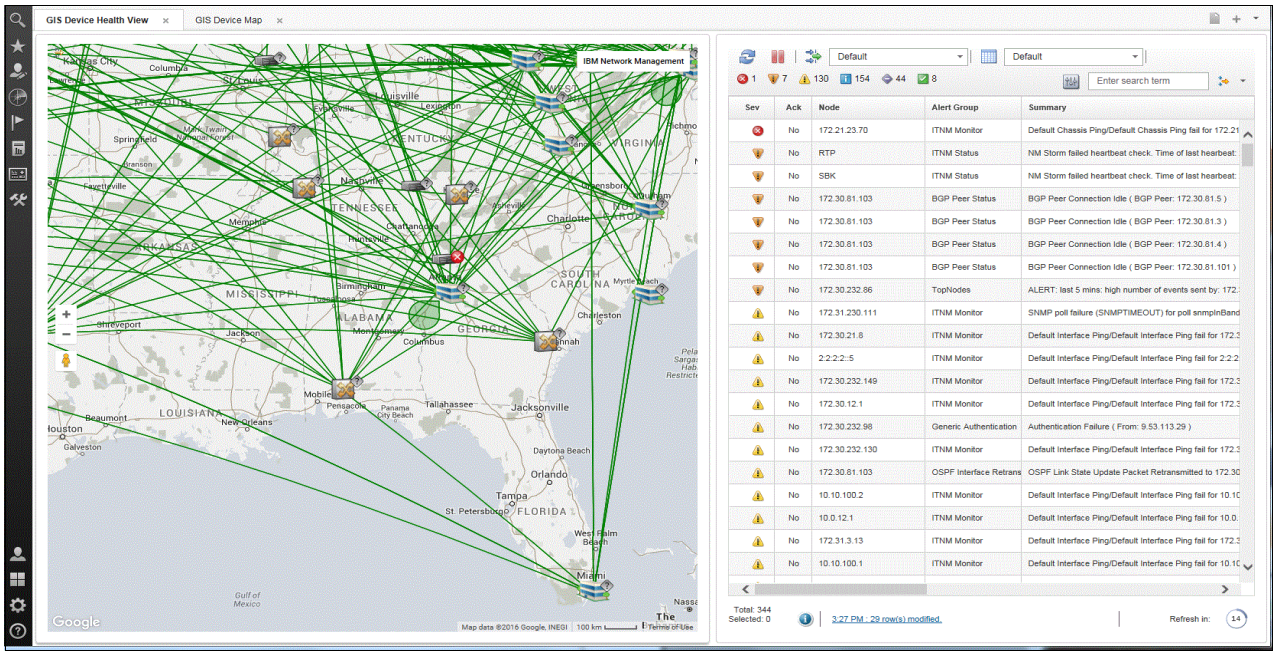

Figure 3-3 GIS Device Health view

This page contains the Geographic map and an Event Viewer, which shows events in context for the device or Building site that is clicked. By default, the page layout features a horizontal split that shows more event columns, but with a narrower map view. If you prefer, you can edit the DASH page and drag the windows into a vertical split (see Figure 3-3) for a larger map area on wide-screen monitors.

3.2.3 Browsing to the Geographic map

If you are looking for geospatial information about a device from one of the topology views, you can browse to the Geographic map from the Hop View, Path View, or a Network View. Right-click the device and click **Find in** → **Geographical Map**. A new browser page opens with the Geographic map, as shown in Figure 3-4.

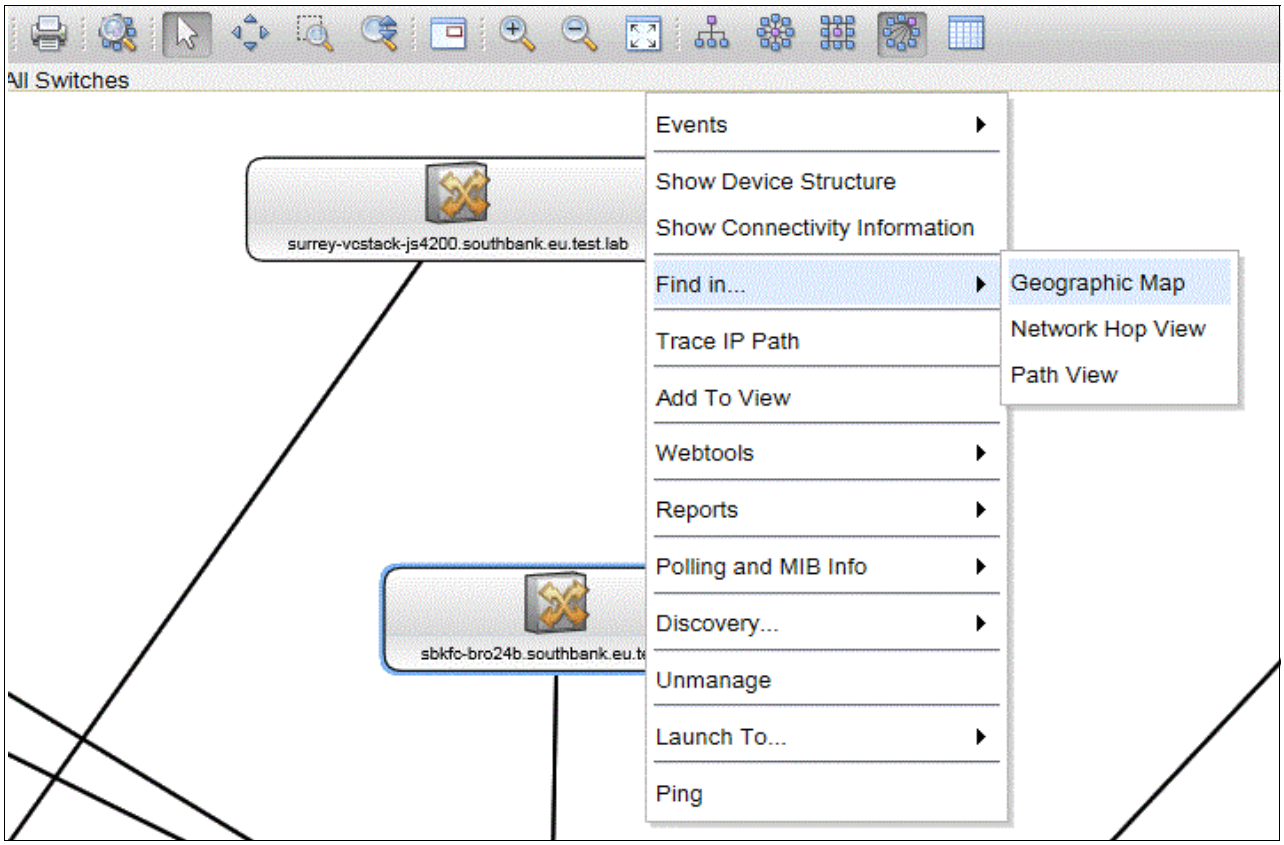

Figure 3-4 Browsing the to Geographic map

3.2.4 Use of Network Views for the Geographic map

You can use the inter-portlet events to control the Geographic map from the same Network Views tree that is used by the Network Health Dashboard.

This view is similar to the Network Health page, but uses the Geographic map instead of the network view topology. In this scenario, an operator can select a network view to show areas of responsibility or other grouping. As with the Network Health Dashboard, the network view tree works off the default bookmark for each user, which makes it easier for each user to create a custom set of useful network views.

Click a view to change the context of the geographical map and all events for devices in that view. Click a device or Building site icon to see events for that device or site, only (see Figure 3-5 on page 73).

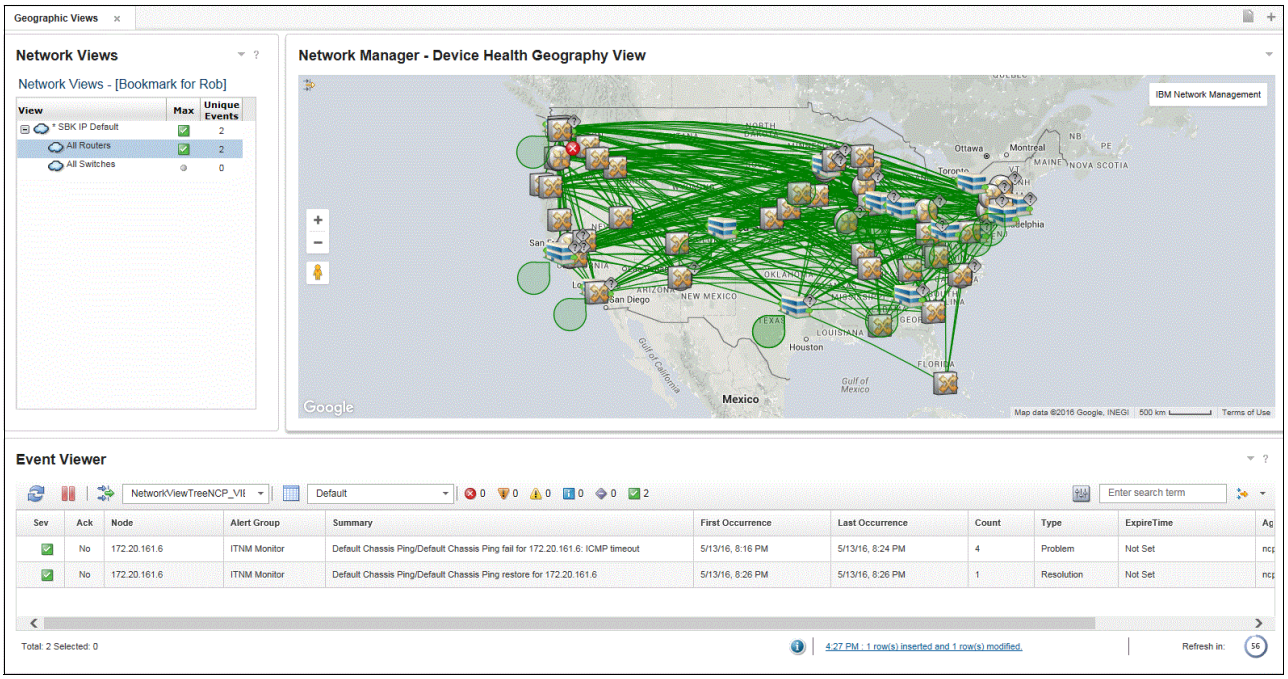

Figure 3-5 Custom Geographic Views page

Building the Geographic Views page

Complete the following steps to build the custom page for the Geographic Views:

1. Click **Console Settings** → **Pages**.
2. Click **New Page** and enter a name for the page. In this example, the name *Geographic Views* is used.
3. Click **Location** and drag the Geographic Views menu item that is in the `console/Incident/Network Geography` directory, as shown in Figure 3-6.

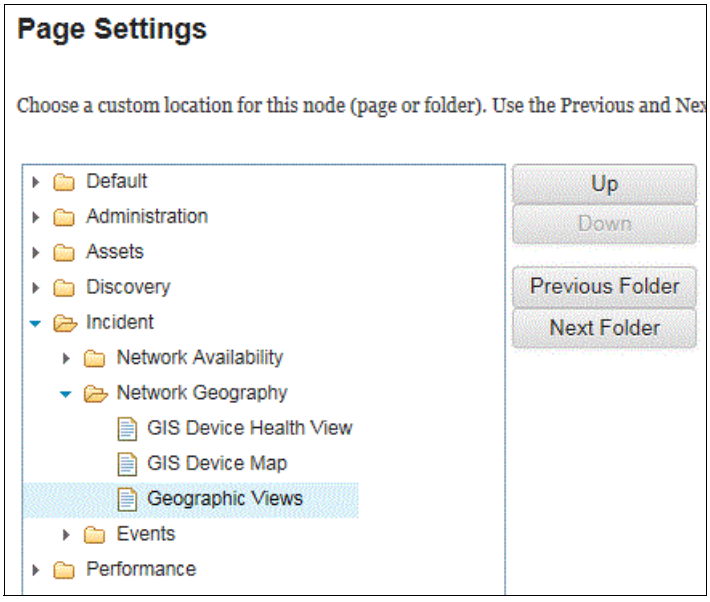

Figure 3-6 Menu location for the new Geographic Views

4. Select **Proportional Layout** and then, click **OK**.

   We select each of the three widgets for this page in turn and drag and expand each widget into position.

5. Click **Network Manager Widgets** and drag the Dashboard Network Views onto the workspace and expand it (as shown in Figure 3-5 on page 73) or to your own design. Then, by using the menu at the top left, select **Home** to return to the widget Home level.

6. From the Home level, click **Network Manager (Experimental) Widgets** and drag the GIS Device Map widget into place and expand, as appropriate. Return to the widget Home.

7. Click **Netcool/OMNIbus Widgets** and drag the Event Viewer widget into place.

   Your page is complete now. The final step for to set the wiring events so that the contextual changes work as wanted.

8. Select the **Event Viewer** widget and then, from the Widget pull-down menu on the tool bar at the top of the page, select the **Events**, as shown in Figure 3-7.

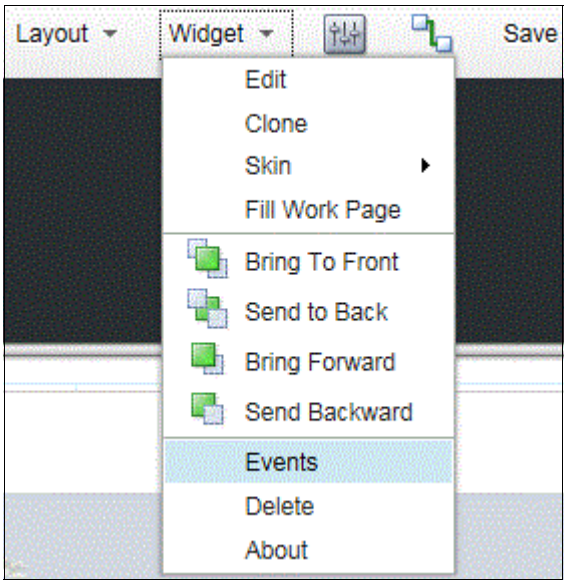

Figure 3-7 Set Wiring Events for widgets

9. Adjust the selected published and subscribed events to match the settings that are listed in Table 3-1. No changes are necessary for the other two widgets.

Table 3-1 Configuration table

Widget	Published Events	Subscribed Events
Event Viewer	(Uncheck NodeClickedOn)	showEvents NodeClickedOn
GIS Device	showEvents	NodeClickedOn
Dashboard Network Views	NodeClickedOn	-

10. Click **Save and Exit**. The widget now can be tested.

Note: The Network View uses bookmarks only, so you must click **Incident** → **Network Views** first to set up a default bookmark and add the views from the Library. If completed this task for the Network Health Dashboard, it does not need to be repeated. For more information, see this website:

https://ibm.biz/BdrDXv

Later for production, you might want to edit this page to set appropriate security roles.

3.3 Scenario topology

The solution that is used in this scenario includes system components that are installed on the following host names:

- hostname1:
  - OMNIbus WebGUI 8.1
  - Network Manager GUI 4.2
  - Network Health Dashboard 4.2 (optional)
- hostname2: ITNM 4.2 core
- hostname3: DB2 10.5

The capability that is described in this chapter is not included with Network Manager. For more information about where to download, install, and set up the distribution package, see 3.4, “Scenario steps”.

For this scenario, we use Google Maps API.

For more information about system components and default settings that were used in the test environment, see Chapter 1, “IBM Netcool Operations Insight overview” on page 3.

3.4 Scenario steps

This section describes that steps that are used to deploy and configure Geographic maps that show devices and sites that are overlaid and aggregated on a world map.

This early-release capability is available for download and is fully supported by the Development team. For more information about downloading the distribution package, see 3.5, “Summary” on page 81.

For convenience, start by making the distribution package available to the Network Manager core server, DB2 server, and the GUI server.

From the distribution package, we use the files that are shown in Example 3-1 on page 76 and any of their supporting files. The files that are shown in Example 3-1 on page 76 might change in future distribution updates; therefore, refer to the documentation and modify these steps accordingly.

Example 3-1 Files in the distribution package

```
ncp_gis.war
nm_rest.war
bin/db2/modifyGISTables.sql
bin/db2/update_geo_location.sh
bin/install.sh
samples/GeoByLookup/data/core_lat_long.csv
samples/GeoByLookup/stitchers/ACMEDeviceLocationEnrich.stch
samples/GeoByLookup/stitchers/PopulateDNCIM_CustomGeography.stch
etc/tnm/tools/findInGeoMapByDevice.xml
etc/tnm/tools/findInGeoMapByView.xml
etc/tnm/menus/ncp_findIn_submenu.xml
etc/tnm/menus/ncp_gis_device_menu.xml
etc/tnm/menus/ncp_gis_link_menu.xml
```

3.4.1 Updating network connectivity and inventory model database

Update the NCIM schema for the standard `ncim.geographicLocation` table that is included with Network Manager 4.2 GA.

NCIM topology database: The Network Connectivity and Inventory Model (NCIM) topology database is a relational database that Network Manager uses to consolidate topology data about the physical and logical composition of devices, layer 1, layer 2, and 3 connectivity, routing protocols, and network technologies, such as OSPF, BGP, and MPLS Layer 3 VPNs.

Run the following command on the db2 server:

```
bin/db2/update_geo_location.sh NCIM db2inst1 <password>
```

Where:

- `NCIM` is the database name
- `db2inst1` is a database instance user with privileges to modify the schema
- `<password>` is the `db2inst1` password

3.4.2 Installing the Geographic maps on DASH

This task is performed on the GUI server where the Network Manager GUI is installed.

Installing the GIS Device Map widget

Before installing the widget, you must edit the defaults at the top of the `bin/install.sh` file to match your environment.

For example, the text in **bold** might need to be changed:

```
JAZZSM_PROFILE=/opt/IBM/netcool/JazzSM/profile
JAZZSM_CELL=JazzSMNode01Cell
SERVER_NAME=server1
CMD=RESTART
GIS_CTX_ROOT=/ibm/console/ncp_gis
REST_CTX_ROOT=/ibm/console/nm_rest
USERNAME=smadmin
PASSWORD=netc00l
```

Run the **bin/install.sh -d** command. You must run this command from the top-level directory of the distribution kit.

Use `-d` to deploy for the first time and `-r` to redeploy or upgrade the war file.

When it finishes, log in to DASH and ensure that you have a new menu section under `Incident` that is called `Network Geography` with the following two items:

- `GIS Device Health View`
- `GIS Device Map`

Geographic map context menus

Copy the files into place that add the navigation to the Geographic map from the context menus.

The Geographic map context menus are defined in the following files:

- `ncp_gis_device_menu.xml`
- `ncp_gis_link_menu.xml`
- `ncp_findIn_submenu.xml`

From the distribution package, find and copy the following files:

```
[root@ncrobitm tnm]# cp menus/ncp_gis*
/opt/IBM/netcool/gui/precision_gui/profile/etc/tnm/menus/

[root@ncrobitm tnm]# cp menus/ncp_findIn_submenu.xml
/opt/IBM/netcool/gui/precision_gui/profile/etc/tnm/menus/
```

Copy the following supporting tools files:

```
tools/findInGeoMapByDevice.xml
tools/findInGeoMapByView.xml

[root@ncrobitm tnm]# cp tools/findInGeo*
/opt/IBM/netcool/gui/precision_gui/profile/etc/tnm/tools/
```

3.4.3 Custom enrichment for the geographical data

The task that is described in this section is performed on the Network Manager core server.

There are various ways to customize Network Manager's discovery process to enrich devices with the geographical location data. We chose a method that starts with a flat file as the source of the data (`core_lat_long.csv`). We import this file manually into a database custom staging table. During discovery, the two custom stitchers that are provided merge this data into the standard NCIM tables in Network Manager 4.2.

We use the `core_lat_long.csv` file as the original source of the location data per IP address for this scenario.

Tip: This method is convenient and helps you to get started quickly. For production, you might make changes to automatically maintain this data in the staging table from your source data. Also, you might consider the use of the EntityName instead of the IP address to index the device.

Creating the staging SQL table

Create an SQL file (for example, name it `createAcmeGeoLocation.sql`), with the following command:

```
CREATE TABLE NCIM.ACMEGEOLOCATION (
                                    IP VARCHAR(255) NOT NULL,
                                    ADDRESS VARCHAR(255),

                                    CITY VARCHAR(255),
                                    STATE VARCHAR(255),
                                    COUNTRY VARCHAR(255),
                                    LATITUDE DECIMAL(10 , 8) NOT NULL DEFAULT 0,
                                    LONGITUDE DECIMAL(11 , 8) NOT NULL DEFAULT 0
                            );
```

Use the following command to apply this file:

```
db2batch -d NCIM -a db2inst1/<password> -f createAcmeGeoLocation.sql
```

Where:

- `NCIM` is the database name
- `db2inst1` is a database instance user with privileges to create the table
- `<password>` is the `db2inst1` password

Importing the data to the staging table

Import the data in `core_lat_long.csv` into the `NCIM.ACMEGEOLOCATION` table by running the commands that use the same database account, as shown in Example 3-2.

Example 3-2 Importing the data to the staging table

```
db2 => connect to NCIM

   Database Connection Information
   Database server        = DB2/LINUXX8664 10.5.3
    SQL authorization ID   = DB2INST1
    Local database alias   = NCIM

db2 => import from core_lat_long.csv of del insert into
ncim.ACMEGEOLOCATION(ip,address,city,state,country,latitude,longitude)
   SQL3109N  The utility is beginning to load data from file "core_lat_long.csv".

   SQL3110N  The utility has completed processing.  "155" rows were read from the
   input file.

   SQL3221W  ...Begin COMMIT WORK. Input Record Count = "155".

   SQL3222W  ...COMMIT of any database changes was successful.

   SQL3149N  "155" rows were processed from the input file.  "154" rows were
   successfully inserted into the table.

   Number of rows read         = 155
   Number of rows skipped      = 0
   Number of rows inserted     = 154
   Number of rows updated      = 0
   Number of rows rejected     = 1
```

```
   Number of rows committed    = 155

db2 => quit
```

Tip: As shown in Example 3-2 on page 78, the `core_lat_long.csv` file contains 154 locations and a header line, which is correctly rejected by the **db2** command.

Deploying the custom geographic stitchers

Complete the following steps to deploy the custom geographic stitchers:

1. Find the following stitchers in `samples/GeoByLookup/stitchers/`:
   - `ACMEDeviceLocationEnrich.stch`
   - `PopulateDNCIM_CustomGeograpghy.stch`
2. Copy the stitchers to `$NCHOME/precision/disco/stitchers/DNCIM/`.
3. Edit `$NCHOME/precision/disco/stitchers/DNCIM/InferDNCIMObjects.stch` to add a line to run the custom `ACMEDeviceLocationEnrich.stch` stitcher:

```
// Build the WLAN model
if(wlanEnabled <> NULL)
{
   ExecuteStitcher('PopulateDNCIM_WLAN', domainId, isRediscovery,
   dynamicDiscoNode);
}
delete(wlanEnabled);

// CUSTOM (5/4/2016): Enrich GEO Location data
ExecuteStitcher('ACMEDeviceLocationEnrich', domainId, isRediscovery,
dynamicDiscoNode );
```

Important: If multiple domains are used but the geographic data is imported to only one of the domains, create a domain-specific version of this stitcher; for example, `InferDNCIMObjects.SBK.stch`, assuming SBK is your domain name.

Running a discovery

You are now ready to run a discovery. After the discovery is run, check whether the `ncim.geographicLocation` table was populated with the location data for devices that match those devices in `core_lat_long.csv`. This table is a standard NCIM table and you might decide not to populate all of the columns.

3.4.4 Registering for the Google Map Key and Client ID

This widget is based on an integration with the Google Maps API. As a deployer and user of this solution, you must ensure that you are compliant with the Terms of Use for Google Maps. Complete the following steps:

1. For Google, obtain a valid key from this website:

   https://developers.google.com/maps/

2. Search for the Google Maps JavaScript API and register for a Key/Client ID. This file uses the following format:

   `AIzaSyBPSOoWeQkjVS1ogd3X6GRh8MPjqIyRKCq`

3. Add this key to the `config.json` file at:

   `/opt/IBM/netcool/JazzSM/profile/installedApps/JazzSMNode01Cell/isc.ear/ncp_gis.war/resources/config.json`

   The default `config.json` file without the license key is shown in Example 3-3.

Example 3-3 Default config.json file without the license key

```
"mapProvider": {
        "baseLayers": [
            {
                "baseLayerName": "Google",
                "baseLayerURL":
"https://maps.google.com/maps/api/js?libraries=visualization,drawing,geometry&v
=3.exp"
```

The `config.json` file after the key is applied (note the '&') is shown in Example 3-4.

Example 3-4 After the key is applied

```
"mapProvider": {
        "baseLayers": [
            {
                "baseLayerName": "Google",
                "baseLayerURL":
"https://maps.google.com/maps/api/js?libraries=visualization,drawing,geometry&v=3.
exp&AIzaSyBPSOoWeQkjVS1ogd3X6GRh8MPjqIyRKCq"
```

As distributed, there is no default license applied. When the user starts the Geographic map in DASH, the warning that is shown in Example 3-8 is displayed.

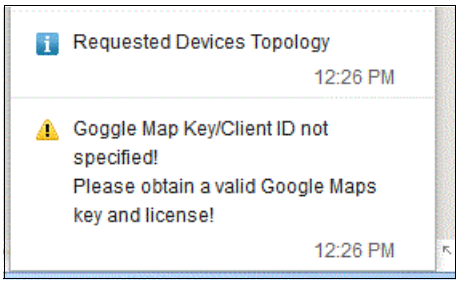

Figure 3-8 Invalid license

3.4.5 No known problems status

By default, the status is “unknown” if there are no events for a device and it is shown with a question mark. You can change this status so that no active events means “no known problems” and marked as green.

Complete the following steps:

1. Edit the `config.json` file that is in `/opt/IBM/netcool/JazzSM/profile/installedApps/JazzSMNode01Cell/isc.ear/ncp_gis.war/resources/config.json`
2. Find “sevColours” and change the name pair **“**`unknown`**” : “**`grey`**”** to the value “green”, as shown in Example 3-5.

Example 3-5 Find “sevColours” and set “unknown” to “green”

```
"sevColours" : {
                        "unknown" : "green",
                        "clear" : "green",
```

3. Change the unknown icon to be the same as the clear icon:
   a. Go to the directory: `/opt/IBM/netcool/JazzSM/profile/installedApps/JazzSMNode01Cell/isc.ear/ncp_gis.war/resources/common_assets/status_icons`
   b. Run the following commands:

```
cp unknown.png unknown.png.orig
cp clear.png unknown.png
```

3.5 Summary

This chapter described how to set up the Geographic mapping scenario for ITNM 4.2 with DASH in a test environment. Various use cases also were described to show the value for operators and others, including a big wall status map.

This early-release capability is available for download. At the time of this writing, the Sprint 1 version was used. For more information about the latest distribution packet, see this website:

https://ibm.biz/BdrDHA

4

Golden configuration and dynamic compliance

This scenario describes how users can easily create compliance definitions by using all or part of a device's configuration.

A *golden configuration* is a configuration version that is used in compliance management as an ideal configuration against which configurations from similar devices can be compared. Any differences that are found are recorded as compliance evaluation failures. You can use this feature to quickly compare many thousands of details without having to manually create any evaluations.

In cases where a range of values is acceptable, you can edit the golden configuration to provide a regular expression and evaluations are treated as failures only if the values that are found in the compared configurations do not satisfy the regular expression in the golden configuration.

This chapter includes the following topics:

4.1 Scenario description

As device configurations and network complexity increase, providing a simplified method of ensuring that the as-built network conforms as much as possible to the as-designed network becomes increasingly crucial to efficient network operations. This section describes how a configuration snippet can be imported as a “real” device and then used to create a Netcool Configuration Manager Compliance Policy.

The first example uses direct xpath evaluations and the second example uses a context xpath for an interface.

Note: XML Path Language (xpath) is an expression that is used to browse to a specific location within an XML document.

The reader should be familiar with Netcool Configuration Manager Configuration and Compliance features and functionality.

4.1.1 Business value

Rapidly evolving network architectures require the implementation discipline to keep them aligned to the design intent. The use of golden configuration or dynamic compliance allows the network compliance function to be implemented in a timely and efficient manner by using configurations or partial configurations that are based on network engineering best practices. Adherence to network compliance can be implemented contemporaneously with the feature rollout.

4.2 Scenario topology

For this scenario, we used the environment that is described in 1.4, “Our environment for the scenarios” on page 18.

4.3 Scenario steps

Our scenario features the following high-level steps:

1. You start with a text file of native configuration settings that are derived from a real device configuration. Ideally, this configuration is based on an actual device that was used to develop the engineering best practice. You then add regular expressions to make the text file more generic.
2. Create a virtual Network Resource, a File Transfer resource, and an Authentication resource within Netcool Configuration Manager. You use these resources to import the configuration file into the network resource.
3. Edit the device’s Resource Access Document (RAD). If you are creating more than one similar Network Resource, create a separate RAD to specify that configuration imports are file-based. Then, import the configuration.
4. After the configuration is imported, you define the configuration as “golden”.
5. You then use the Netcool Configuration Manager - Compliance UI to create a Compliance Definition that uses the golden configuration.

6. (Optional) You can use the Test feature in the definition editor to check the golden configuration against a non-golden configuration. Only evaluations that contain a regular expression are tested.
7. If you are satisfied with the definition, you save it and create a Rule and Policy in the normal process in Netcool Configuration Manager Compliance.

4.3.1 Scenario details

In this section, creating and importing a file-based device are described.

The following artifacts are required to import a file-based device:

- Network device configuration file
- Netcool Configuration Manager File Transfer Resource
- Netcool Configuration Manager Authentication Resource
- Netcool Configuration Manager Resource Access Document (RAD)

File-based device

Create the device configuration file. As shown in Example 4-1, we use a portion of the entire device's configuration as the base for our compliance policy that we are creating.

Example 4-1 Creating the device configuration file

```
# START

snmp-server community @@@.*@@@ RO
snmp-server community @@@.*@@@ RW 1301
snmp-server location @@@.*@@@
snmp-server contact @@@.*@@@
snmp-server enable traps entity
snmp-server enable traps envmon fan shutdown supply temperature status
snmp-server enable traps config-copy
snmp-server enable traps config
snmp-server host @@@.*@@@ itecsrvr

# END
```

In Example 4-1, the `# START` and `# END` lines are used to inform Netcool Configuration Manager where the beginning and end of the configuration file are found.

The sections of text with the `@@@` symbols enclose a user-supplied regular expression. In each of the five examples in Example 4-1, the goal is to enforce a non-empty configuration without specifying the exact contents.

This configuration file should be placed on a server, which is referred to as the *file server*, that is reachable from the Netcool Configuration Manager Worker Server. The file format `date.name.extension` is used, such as `12May2016.snmpOne.txt`.

Important: The directory structure on the file server must match the directory structure on the Netcool Configuration Manager worker. For example, if the Netcool Configuration Manager worker FTP directory that is specified during the installation is `/home/icosftp`, the file server must have the same directory.

Complete the following steps:

1. On Netcool Configuration Manager, in an area of your choosing, create a File Transfer resource that contains the file server details. When the FTP resource is created, the VTMOS entries can be wild cards as all file-based devices use the same file server (see Figure 4-1).

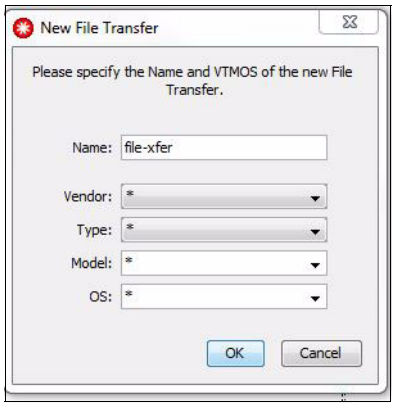

Figure 4-1 FTP resource VTMOS details

Figure 4-2 shows the FTP resource details.

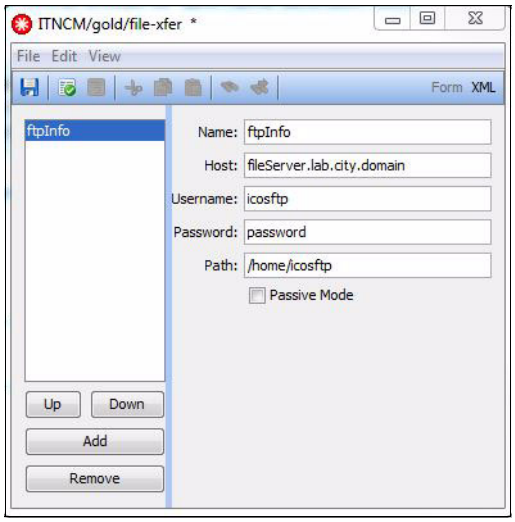

Figure 4-2 FTP resource details

2. Create an Authentication resource that specifies the username and password for the file server. Again, the VTMOS details all can be wild cards, as shown in Figure 4-3.

Figure 4-3 Authentication resource VTMOS details

The authentication resource that shows username and password is shown in Figure 4-4.

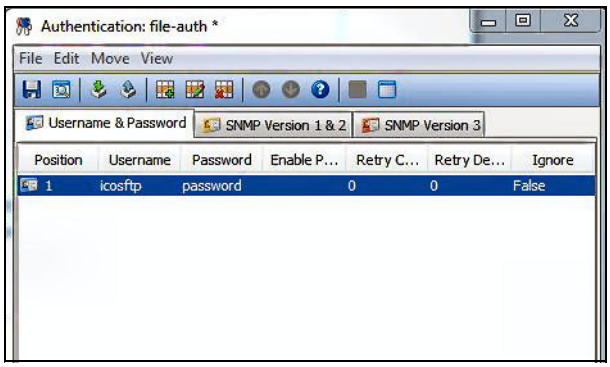

Figure 4-4 Authentication resource showing username and password

The final artifact is the Resource Access Document, which specifies the file-based access method.

Note: The file-based feature is available only in Netcool Configuration Manager Drivers 20 and beyond.

3. Create a RAD with the VTMOS of the target device. The VTMOS that was used in this case was `Cisco Switch 3750* *15.*`, as shown in Figure 4-5.

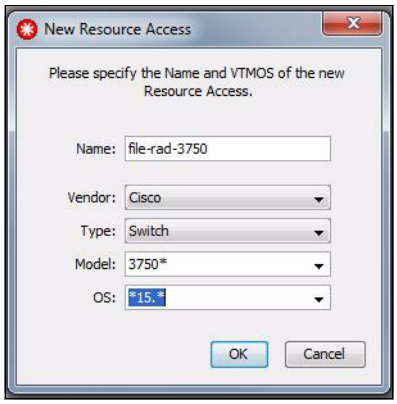

Figure 4-5 Resource Access Document VTMOS details

The contents of this RAD is replaced with the text that is shown Example 4-2.

Example 4-2 Resource Access Document

```
<?xml version="1.0" encoding="UTF-8"?>
<resource-access-data>
    <access-order>
        <name>file</name>
    </access-order>
    <rollback-options>
        <option name="NO_ROLLBACK">
            <description>No rollback</description>
            <required>false</required>
            <recommended>false</recommended>
        </option>
        <option name="REBOOT_ROLLBACK">
            <description>Reload the configuration and reboot the device</description>
            <required>false</required>
            <recommended>false</recommended>
        </option>
        <option name="MODELLED_ROLLBACK">
            <description>Use Modeled Rollback</description>
            <required>false</required>
            <recommended>false</recommended>
        </option>
    </rollback-options>
    <access-types>
        <access-type name="file">
            <hostname/>
            <port/>
            <sourceAddress/>
            <socketConnectTimeout/>
            <import-char-streaming-flag>false</import-char-streaming-flag>
            <import-char-streaming-time-interval/>
            <streaming-put-flag>true</streaming-put-flag>
            <streaming-get-flag>true</streaming-get-flag>
            <ssh-type>ssh2</ssh-type>
            <ssh1>
                <cipher>des3</cipher>
            </ssh1>
            <ssh2>
```

```
                <cipher>aes128</cipher>
            </ssh2>
            <user/>
            <password/>
            <enable-password/>
            <prev-user/>
            <alt-user/>
            <prev-password/>
            <alt-password/>
            <prev-enable-password/>
            <alt-enable-password/>
            <script-id>deviceScript-file-default</script-id>
            <prompt/>
            <enablePrompt/>
            <configEditPrompt/>
            <native-compare-flag>true</native-compare-flag>
            <import-prewrite-flag>false</import-prewrite-flag>
            <sync-prewrite-flag>false</sync-prewrite-flag>
            <import-report-diffs-flag>false</import-report-diffs-flag>
            <reboot-on-config-load>false</reboot-on-config-load>

<update-Resource-Info-On-Config-Change>false</update-Resource-Info-On-Config-Change>

<update-Resource-Config-On-Config-Change>true</update-Resource-Config-On-Config-Change>
            <lbl-mode-flag-forncs>true</lbl-mode-flag-forncs>
            <configDataType>CLI</configDataType>
            <retry-errors/>
            <additional-errors/>
        </access-type>
    </access-types>
    <scripts>
        <script name="deviceScript-file-default"><![CDATA[#turn on flags to get info from
File server.
getConfigFileServer=true
getModelFileServer=true
getDiagFileServer=true
getVersionFileServer=true
getBinaryFileServer=true
putConfigFileServer=true

### Defaults for sending commands. Errors must be separated by , not spaces
default.prompt=$
log-in.prompt=$
default.error=Error,Invalid
default.errorResponse=Error sending command

### Connection Global
# if connect.* is not present use connect.all.properties
connect.errorResponse=Unable to connect to router
connect.09.wait=$

# check for Running config and stored config values multipleConfigs or SingleConfig
config.check.end=singleConfig

# Signals start of config
config.start=# Generated
# Signals end of config
config.end=# Finished

# Identifies error retreiving config.  Must be separated by ,
```

```
config.fail=Error,Invalid

# Info
# these commands are used to gain some information on the hardware installed in the device
diag.01.setReturnBuff=test#
diag.end=$

# Model
## using the command below, or similar, select the text that will allow determination of a
model
model.01.setReturnBuff=3750
model.end=$
model.FIND-BEGIN=
model.FIND-END=

# Version
## using the command below, or similar, select the text that will allow determination of an
OS version
config.version.setReturnBuff=##15.1##%%IPBase%%
config.version.end=$
config.version.FIND-BEGIN=##
config.version.FIND-END=##
config.version.FEATURE-FIND-BEGIN=%%
config.version.FEATURE-FIND-END=%%

# Running config
###
## The setFtpFileName parameter can be used to specify a single file if needed.
###
#config.running.01.setFtpFileName=<filename>
###
## The addressName parameter will use the name of the resource.  If you modify the
directory structure
## on the source file server, then there is one edit to the device script required.  The
"f" parameter in
## the "cut" operation needs to be incremented to match the number of directory levels. If
you add a level
## on the source file server, then "-f4" is changed to "-f5" and so on.  The other change
required is an
## addition of a corresponding directory to the /home/icosftp directory to match the source
directory.  This
## is required due to the way the file-based access method uses the ftp resource.
###
## The end parameter, either first or last, specifies the first or last instance of the
file based on
## a file listing of "ls -r" reverse time.  Changing this to "ls -l" would require the
parameter to be
## changed to last to obtain the latest version of the file.  The reverse time listing was
used to reduce
## the number of results that would be examined.
###
config.running.01.getFtpFileName=ls -t $ftpPath$/*$addressName$* | cut -d"/" -f4,last
config.running.02.getSCP=$ftp_username$:$ftp_password$@$ftp_hostname$:$ftpPath$/$ftp_filena
me$
config.running.end=$
config.running.FIND-BEGIN=# START
config.running.FIND-END=# END

# Stored config
config.stored.send=show startup-config\r
```

```
config.stored.end=\nend\r
config.stored.FIND-BEGIN=version
config.stored.FIND-END=\nend\r

### Ftp file
ftp.01.send=putSCP://$ftp_username$:$ftp_password$@$ftp_hostname$/$ftp_filename$
$ftp_filename$\r
ftp.09.wait=$
]]></script>
    </scripts>
</resource-access-data>
```

4. The RADs are unique to a VTMOS and you must edit the line model.01 to conform to the device or devices that you are importing. Because we are not interacting with a live device, the device script must provide the appropriate responses to mimic the interaction with a device with the use of the **setReturnBuff** command, as shown in Figure 4-6.

```
# Model
## using the command below, or similar,
model.01.setReturnBuff=3750
model.end=$
model.FIND-BEGIN=
model.FIND-END=
```

Figure 4-6 RAD model details

Figure 4-7 shows the RAD version details.

```
# Version
## using the command below, or similar, select the text
config.version.setReturnBuff=##15.1###%IPBase%%
config.version.end=$
config.version.FIND-BEGIN=##
config.version.FIND-END=##
config.version.FEATURE-FIND-BEGIN=%%
config.version.FEATURE-FIND-END=%%
```

Figure 4-7 RAD version details

5. The RAD get running configuration mechanism uses the default `/home/icosftp` or two-deep directory structure. If `/home/icosftp/gold` was used as a directory structure on the file server and the Netcool Configuration Manager worker server, the `config.running` section must be edited to increase the field parameter in the **cut** command by the number of directory levels from the default value of “-f4” (in this case, with one extra directory that is “-f5”).

   You must specify the beginning and ending tags for the configuration file that is on the file server. In this case from Example 4-1 on page 85, we used “# START” and “# END”, as shown in Figure 4-8.

```
config.running.01.getFtpFileName=ls -t $ftpPath$/*$addressName$* | cut -d"/" -f4,last
config.running.02.getSCP=$ftp_username$:$ftp_password$@$ftp_hostname$:$ftpPath$/$ftp_filename$
config.running.end=$
config.running.FIND-BEGIN=# START
config.running.FIND-END=# END
```

Figure 4-8 RAD Running Configuration details

6. When finished editing, save the RAD by clicking **File → Save**, as shown in Figure 4-9.

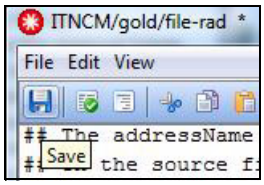

Figure 4-9 Saving the RAD

Importing the device into Netcool Configuration Manager

Complete the following steps to import the “device” into Netcool Configuration Manager.

1. In the same realm that you created the RAD, FTP resource, and Authentication resource, create a network resource with the same name as shown in Figure 4-10.

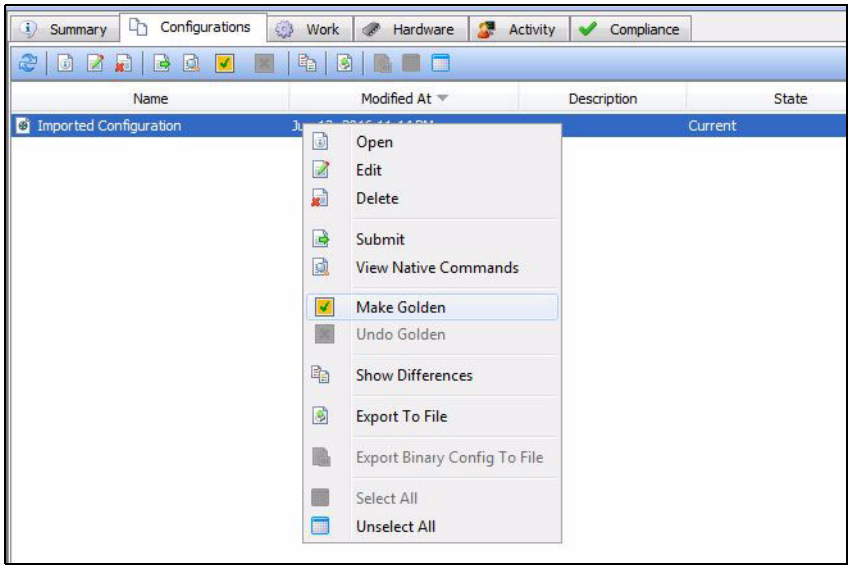

Figure 4-10 Setting a Golden configuration

2. Import the device.
3. After the device is imported, you can view the configuration in the Netcool Configuration Manager Configuration editor or by using View Native Commands. The configuration commands are shown with the `@@@` symbols.
4. Select the network device and click the **Configurations** tab.
5. Select the current configuration and set it as a “golden” configuration by right-clicking the configuration and selecting **Make Golden** from the pop-up menu, as shown in Figure 4-10.

   Alternatively, you can click the **Make Golden** icon in the middle menu bar, as shown in Figure 4-11.

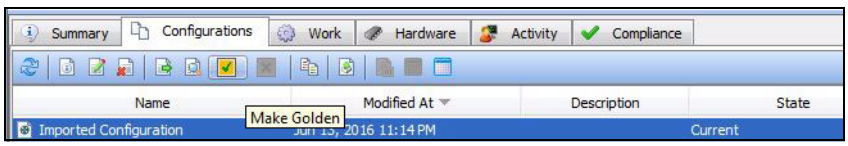

Figure 4-11 Make Golden menu icon

The golden configuration work in Netcool Configuration Manager Configuration is now complete. Next, we create a compliance policy that uses this configuration.

Creating a golden configuration policy

Complete the following steps to create a compliance policy that uses the golden configuration from “Importing the device into Netcool Configuration Manager” on page 92:

1. In the Netcool Configuration Manager Compliance application, browse to the Definitions pane and select **Create** → **Definition** from the menu bar, as shown in Figure 4-12.

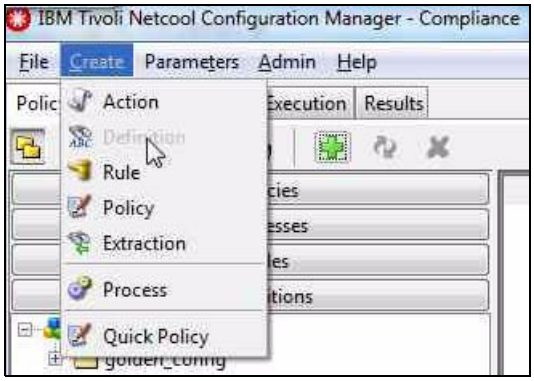

Figure 4-12 Creating a Definition

2. Select the **Create Compliance Definition using a Golden Configuration** option in the next pane, as shown in Figure 4-13.

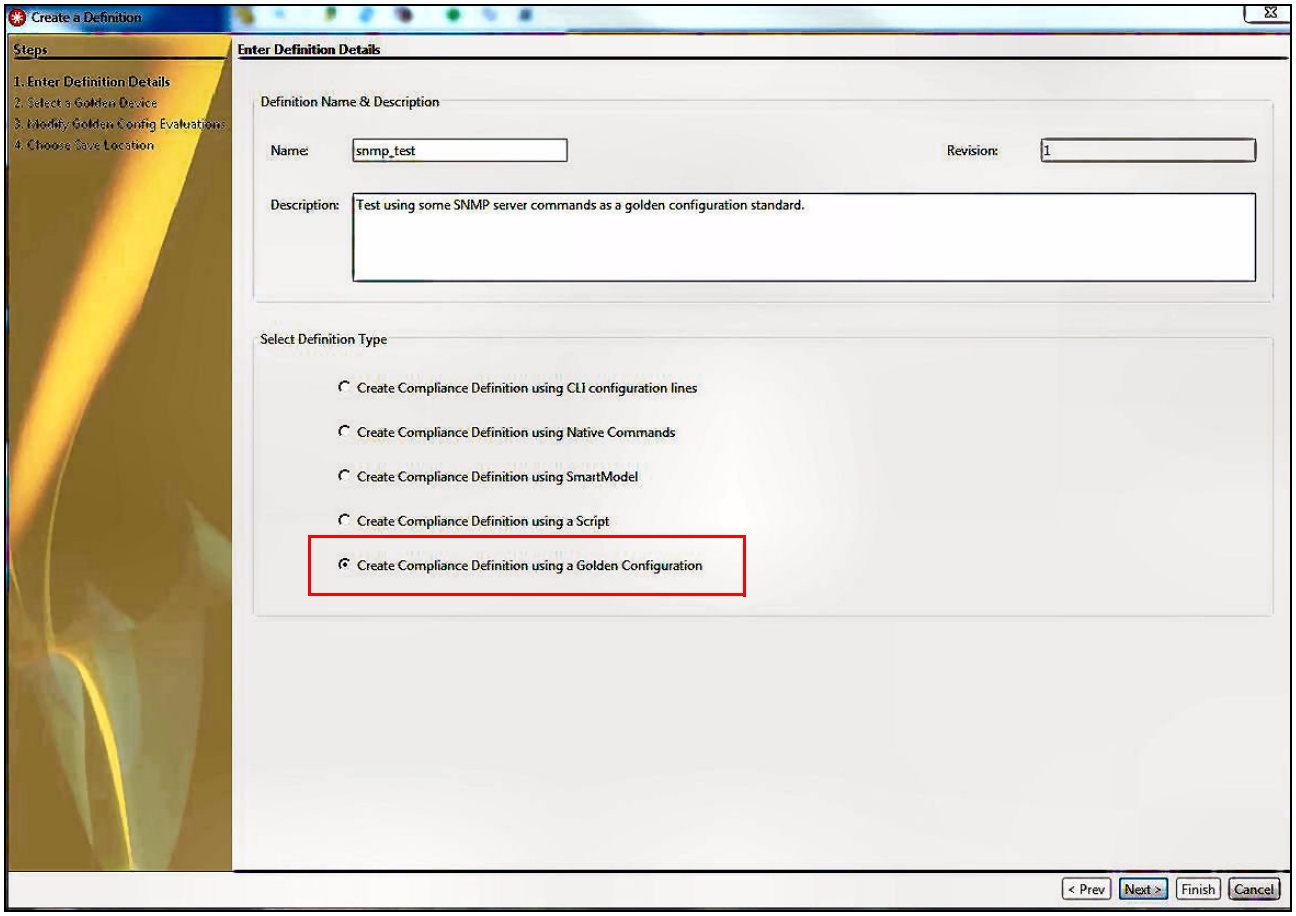

Figure 4-13 Golden configuration definition type

3. Click **Next**.

4. In the next pane, select the designated golden configuration, `snmpOne`, to use in this definition by using the By Realm or By VTMOS selection panes. Then, click **Next**, as shown in Figure 4-14.

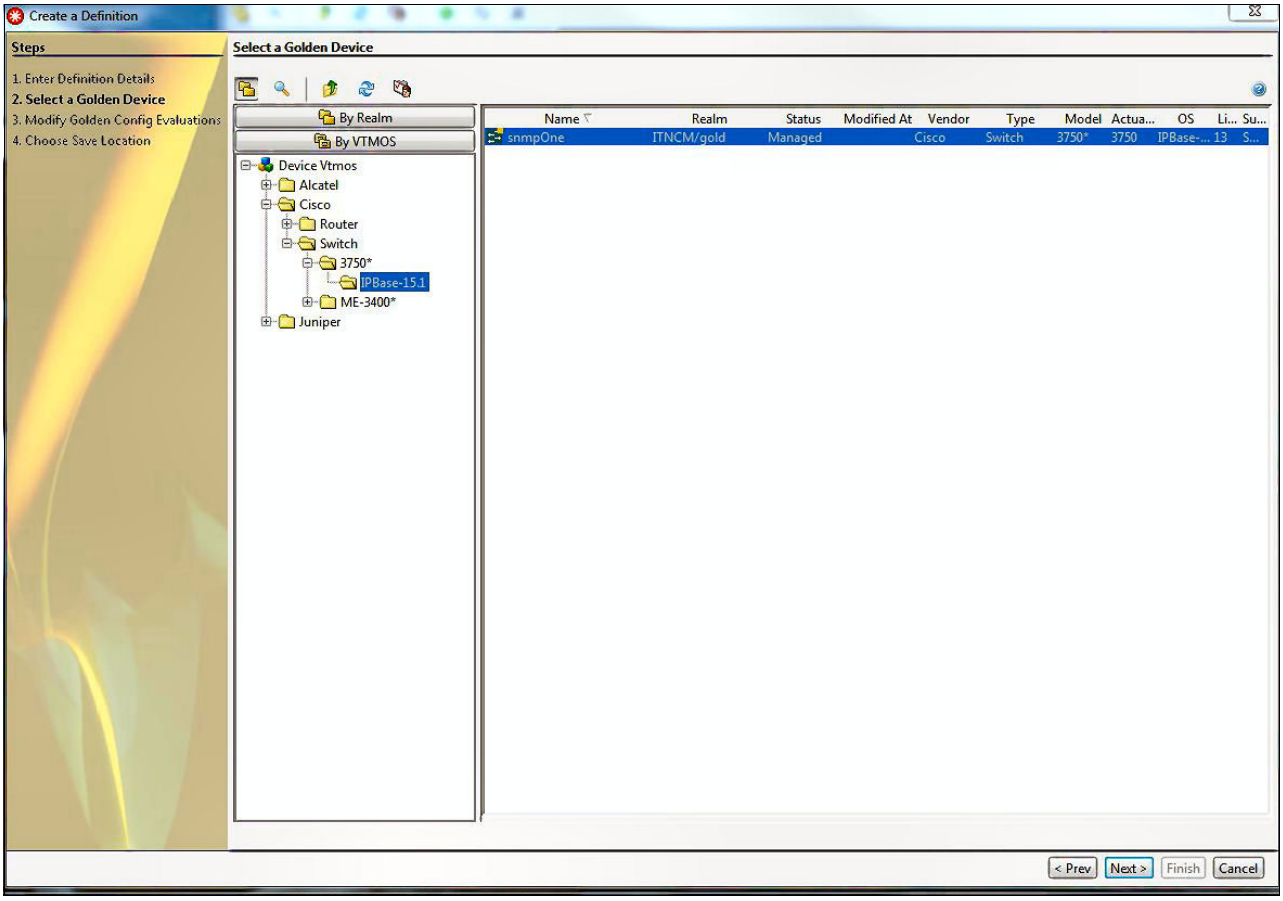

Figure 4-14 Golden Configuration selection

Note: Only those configurations that are set as golden are displayed.

5. In the next pane, you see the evaluations for only those configuration commands that include the `@@@` regular expression designator. If needed, you can edit the evaluations. By clicking the **Test** button, you can test this Definition against an imported device; however, only those evaluations that are displayed are tested.

   The remainder of the evaluations is generated dynamically at run time; therefore, testing at this point does not yield the same results as running a Compliance Policy.

As shown in Figure 4-15, the `@@@` annotation is converted into an appropriate compliance evaluation with regular expression matching syntax.

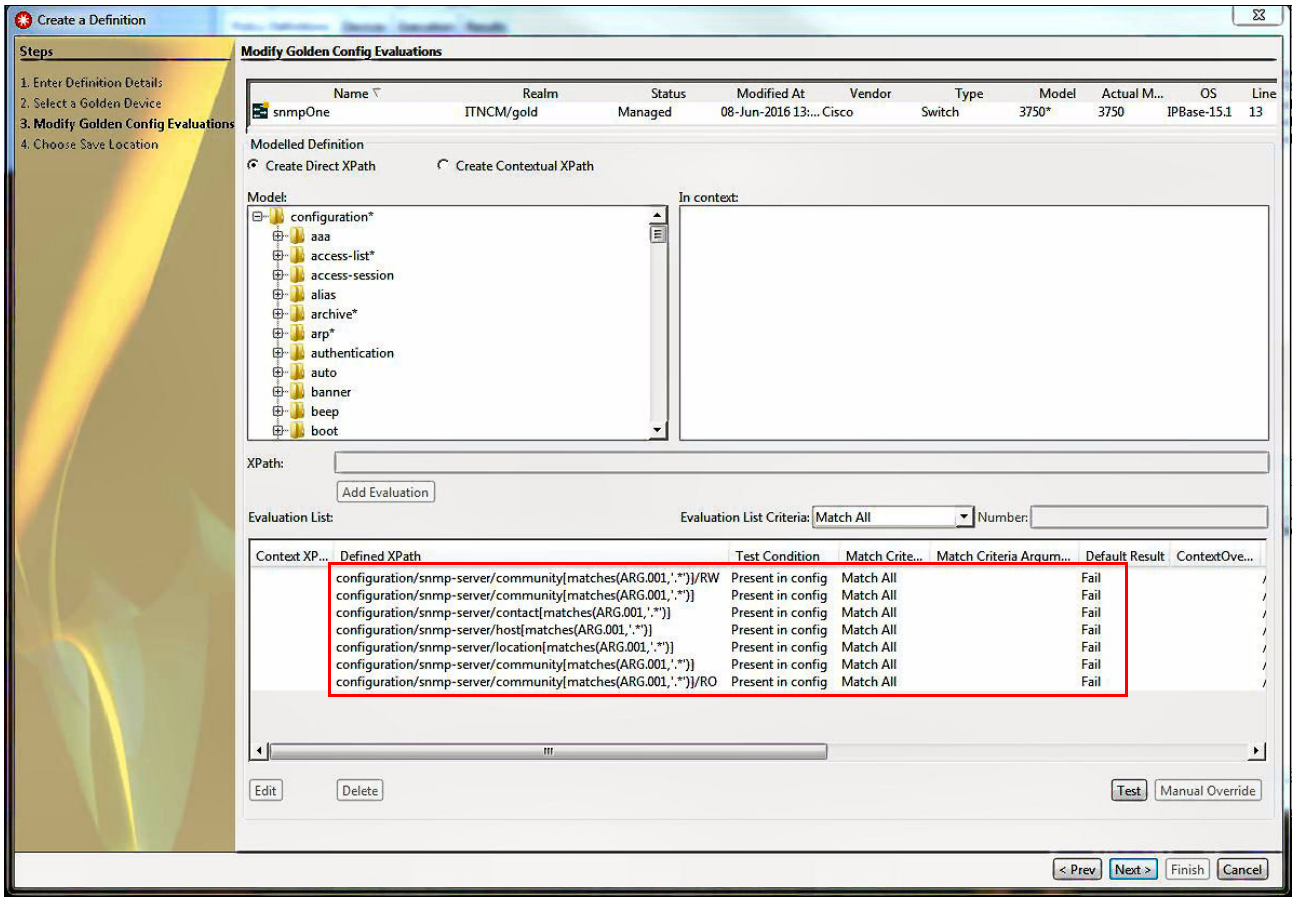

Figure 4-15 Evaluations with regular expressions

6. If no edits are required, click **Next** and save the Definition by clicking **Finish**.

   The steps to create a Rule, Policy, and optional Process are not described here because they do not change with the introduction of the golden configuration option.

Running the golden configuration policy and getting results

Complete the following steps for running the golden configuration policy execution and getting the results:

1. Click the **Execution tab** at the top and select the SNMP-related policy that was created. Click the **Execute** button. You can enter the details for the Name and Description on the next pane and click **Next** when finished, as shown in Figure 4-16.

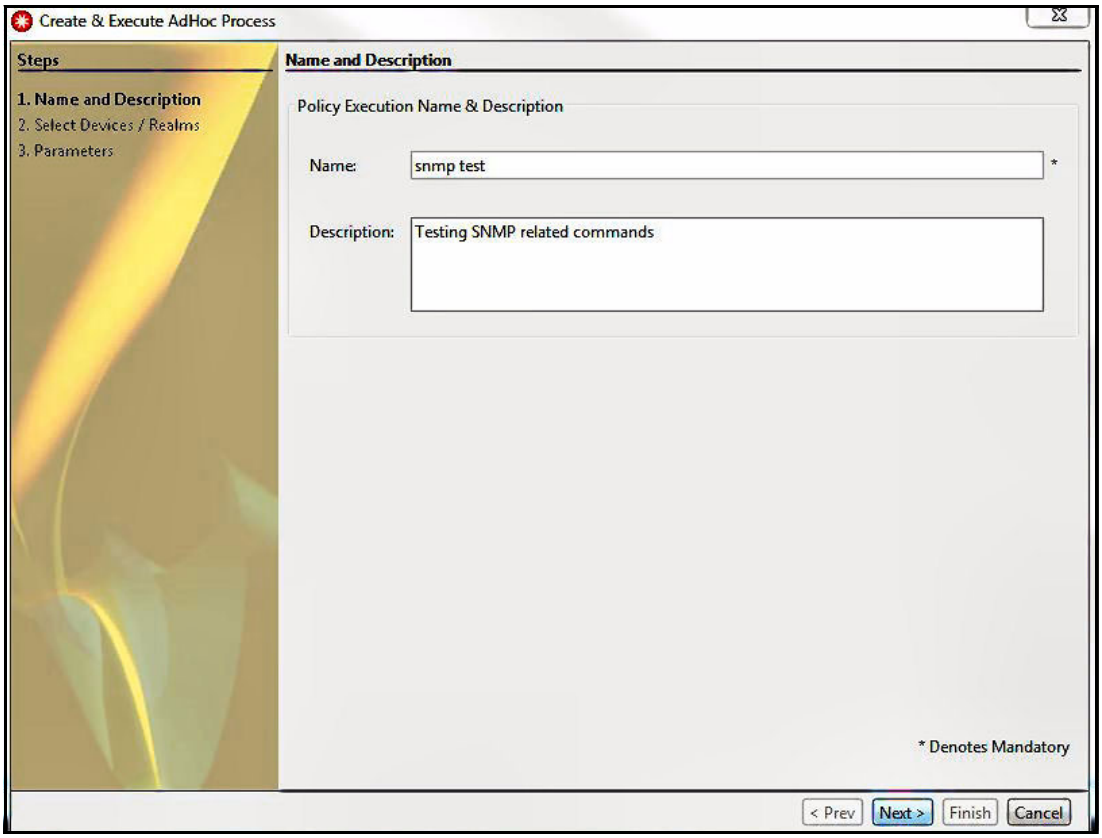

Figure 4-16 Execution details

2. Select the Realm or Devices that are to be used for testing and click **Next**, as shown in Figure 4-17.

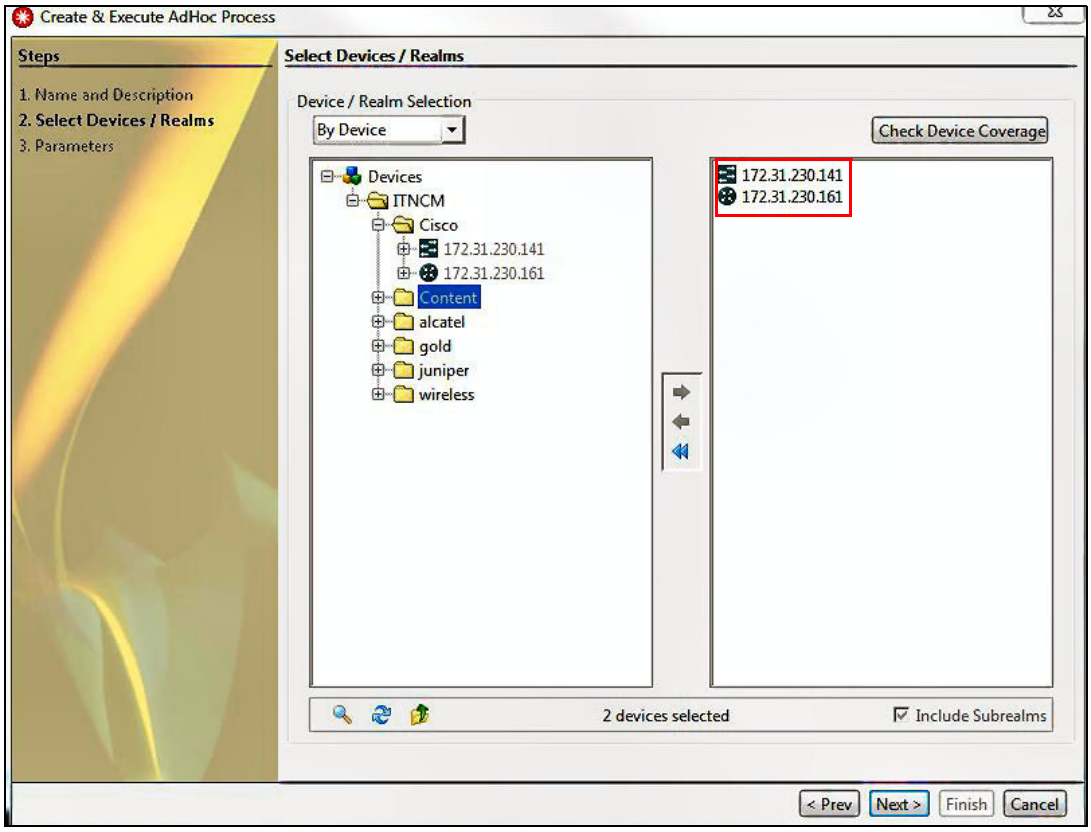

Figure 4-17 Device Selection

3. Answer **No** to the View Associated Parameters question and click **Finish**, as shown in Figure 4-18.

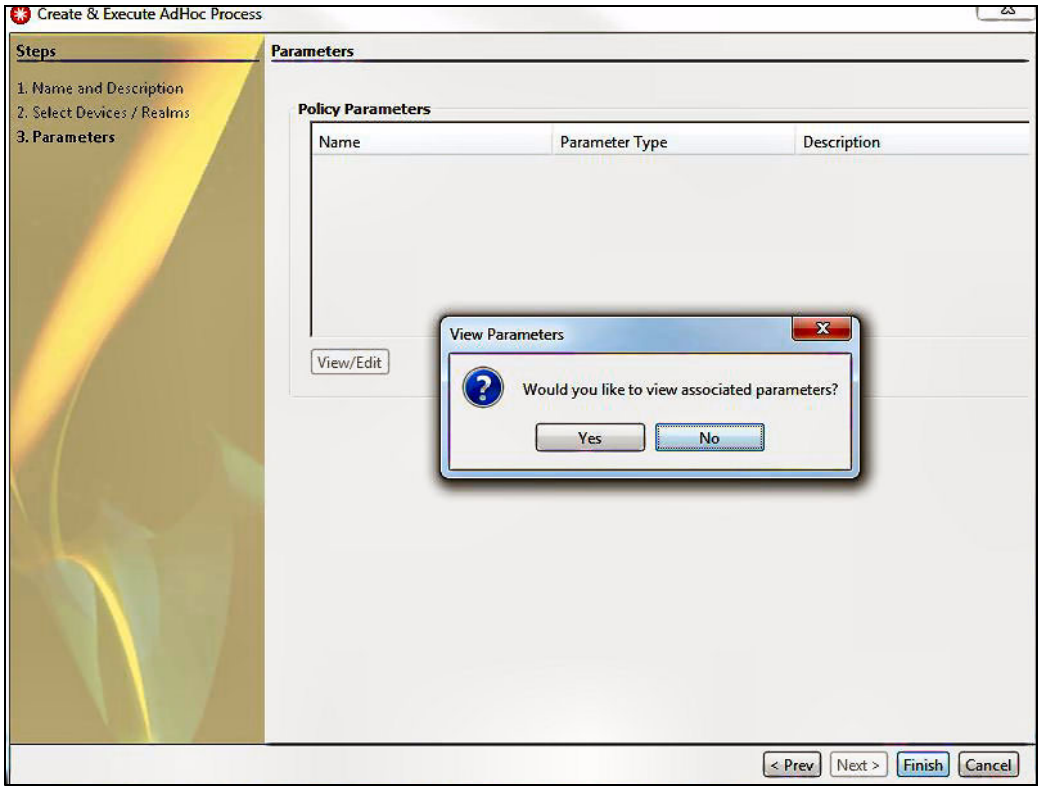

Figure 4-18 Answering No to View Parameters prompt

4. After the policy completes, click it in the Process Execution Summary pane and then, double-click the policy name in the Policy Validation Summary pane at the bottom of the window as shown in Figure 4-19.

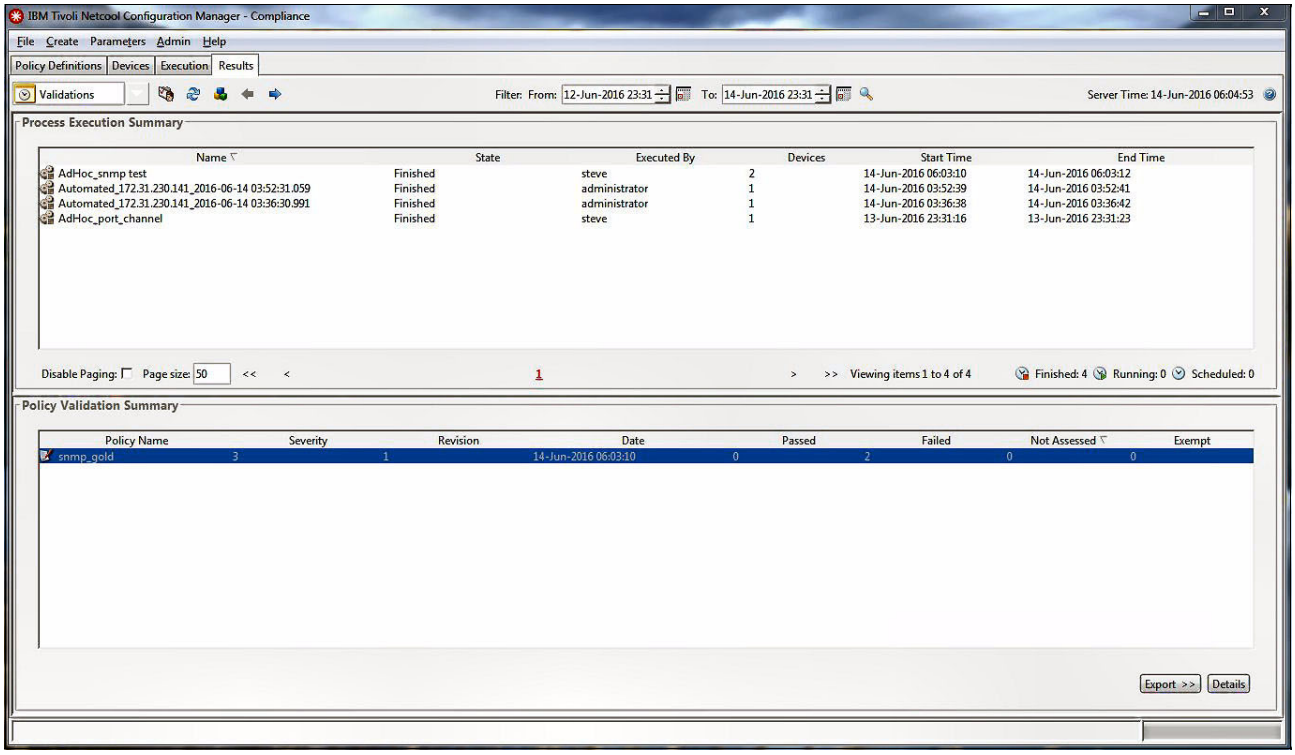

Figure 4-19 Selecting Policy Validation Summary

5. In the next window, select a device and view the results by double-clicking the device name. In the lower portion of the window, you can see a summary of the evaluation results by clicking the **Definition** in the right pane, as shown in Figure 4-20.

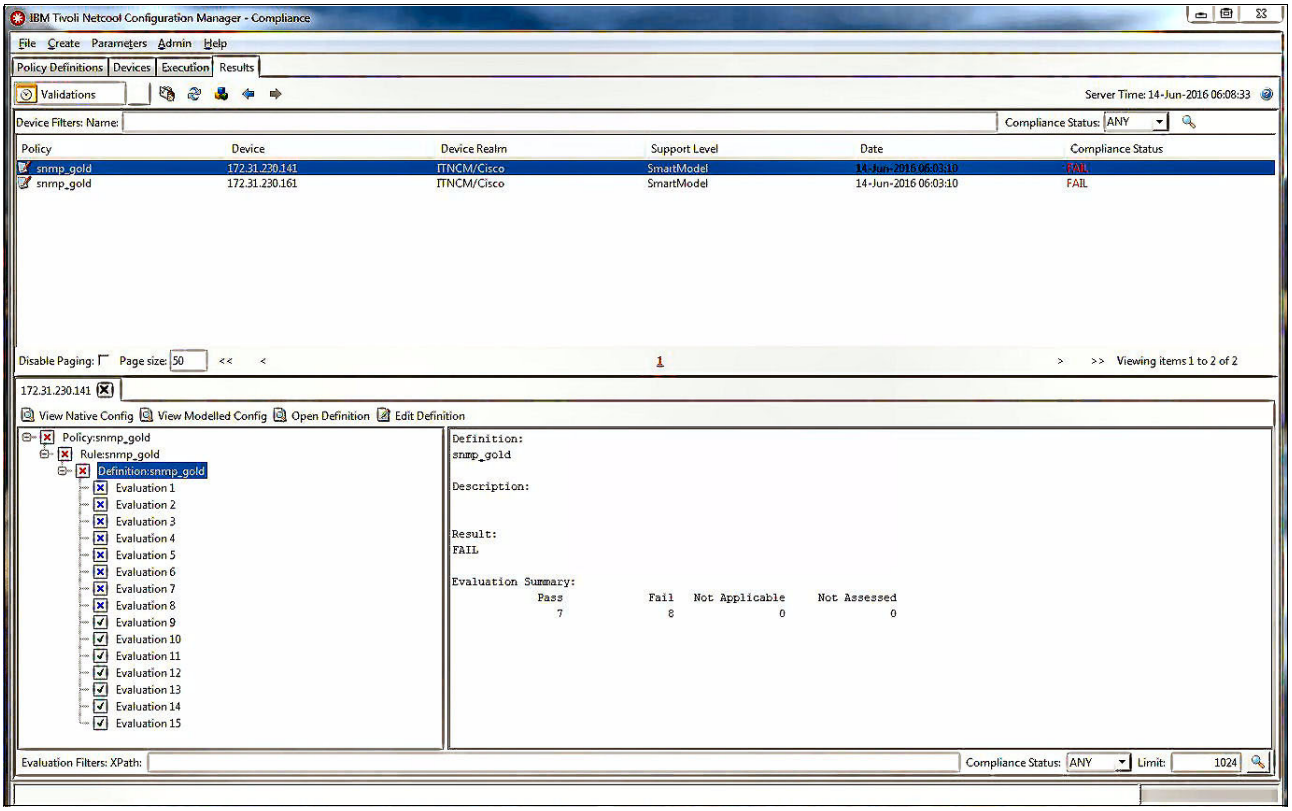

Figure 4-20 Execution Results Summary

There were a total of 15 evaluations: seven from evaluations with regular expressions as shown in Figure 4-15 on page 96 and eight evaluations that were dynamically created at run time from the remainder of the golden configuration.

6. Select an evaluation from the pane on the left side to show the details of the pass or fail results for that evaluation, as shown in Figure 4-21.

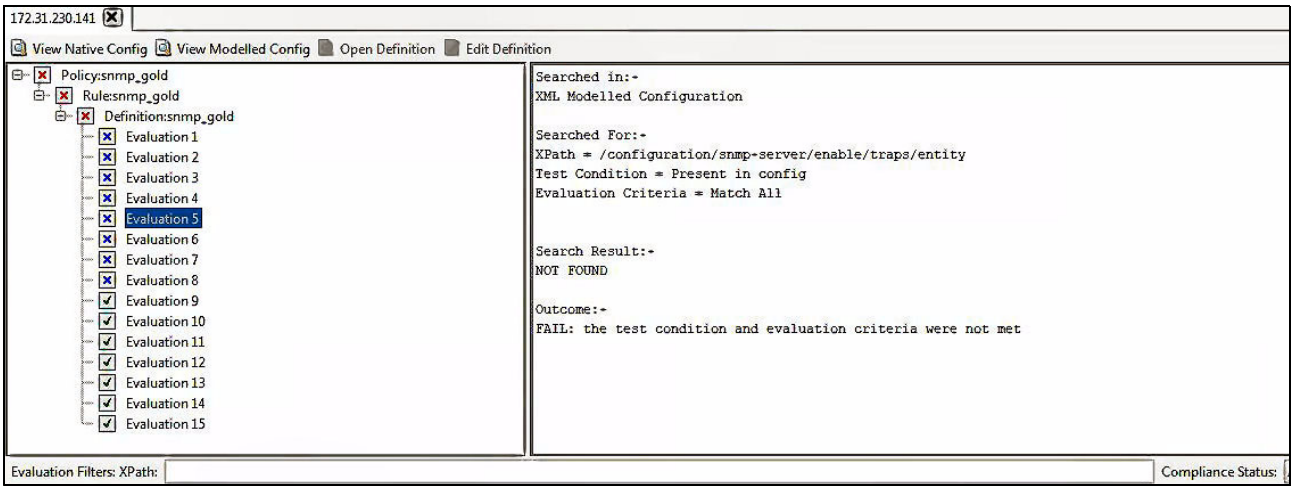

Figure 4-21 Evaluation failure

7. You also can use the Evaluation Filters: XPath: search feature at the bottom of the pane (see Figure 4-21) to search for particular items. This feature is useful in large configurations when you are trying to find certain evaluation results. You can also filter on Status or increase or decrease the evaluation row count.

Tip: Setting the limit to a large number negatively affects UI performance and memory usage.

Figure 4-22 shows the results after the Filter function is applied.

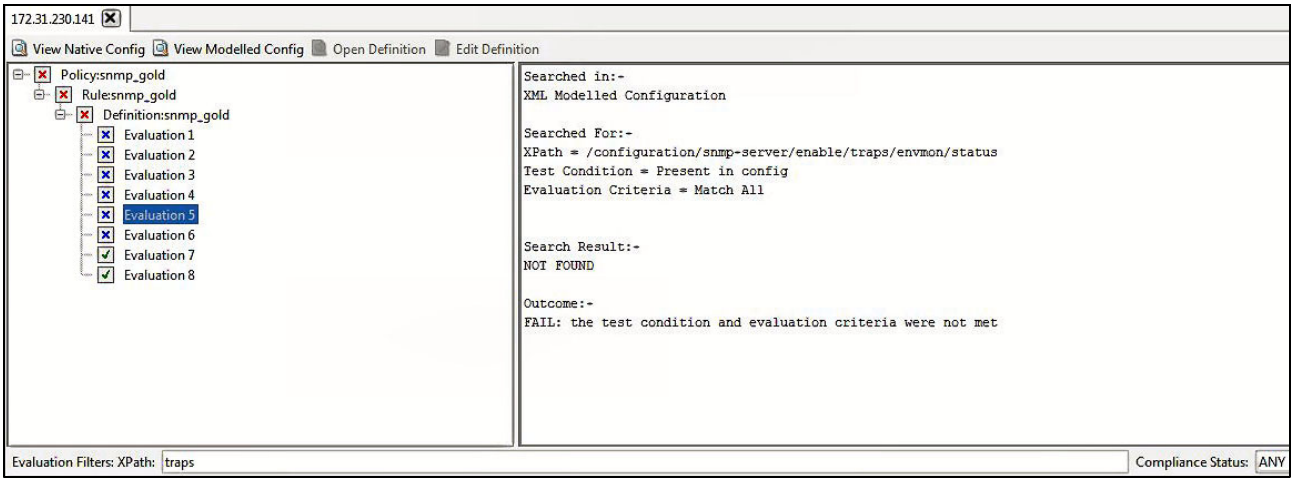

Figure 4-22 Using the Filter function to limit results displayed

Creating a contextual golden configuration and policy

The previous example used only direct xpath evaluations. By using a different regular expression designator, `@@@P@`, users can evaluate “in context” golden configuration commands. A common example is to determine whether an interface configuration contains the proper commands that are based upon a key designator within the interface configuration set of commands. In Example 4-3, we use the keyword “lag” in the interface description to set the scope of the definition. The keyword can be a VLAN ID or QoS policy name.

Example 4-3 Using the keyword lag in the interface description

```
## START

interface Port-channel11
 description @@@P@.*lag@@@
 no ip address
 logging event trunk-status
 mls qos trust cos
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 1499
 switchport trunk allowed vlan @@@.*@@@
 switchport mode trunk
 spanning-tree bpdufilter enable
 spanning-tree guard loop
 spanning-tree mst 0-7 cost 2000000

## END
```

The steps that were used to create the file-based device and then create the compliance definition that uses this golden configuration are the same. After the definition is created, the evaluations that are displayed show contextual xpaths versus direct xpaths, as shown in Figure 4-23.

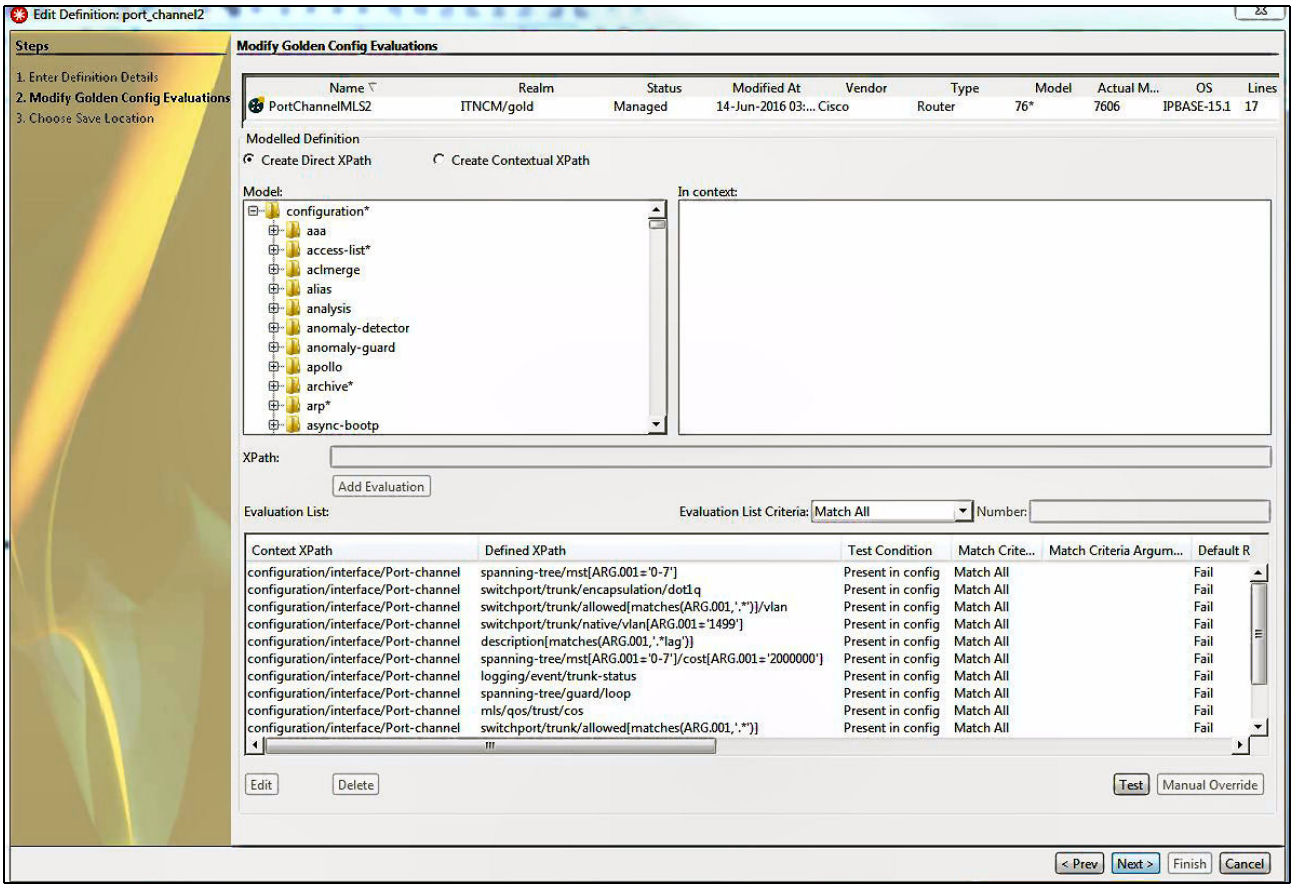

Figure 4-23 Contextual xpaths

Upon executing the policy against a device or set of devices as before, the results for each interface that contains the “lag” keyword are shown.

One evaluation was searching for a logging command under each Port-Channel interface that contained the “lag” keyword. On the particular device, there were 17 Port-channel interfaces that the **`show ip interfaces brief | include Port-channel`** command displays, as shown in Example 4-4.

Example 4-4 Output of the show ip interfaces brief | include Port-channel command

```
Port-channel1              unassigned      YES NVRAM  down                    down
Port-channel2              unassigned      YES NVRAM  down                    down
Port-channel3              unassigned      YES NVRAM  down                    down
Port-channel4              unassigned      YES NVRAM  down                    down
Port-channel5              unassigned      YES NVRAM  down                    down
Port-channel6              unassigned      YES NVRAM  down                    down
Port-channel7              unassigned      YES NVRAM  down                    down
Port-channel8              unassigned      YES NVRAM  down                    down
Port-channel9              unassigned      YES NVRAM  down                    down
Port-channel11             unassigned      YES NVRAM  down                    down
```

```
Port-channel13             unassigned      YES NVRAM  down                    down
Port-channel15             unassigned      YES NVRAM  administratively down down
Port-channel19             unassigned      YES NVRAM  down                    down
Port-channel20             unassigned      YES NVRAM  down                    down
Port-channel21             unassigned      YES NVRAM  down                    down
Port-channel121            unassigned      YES NVRAM  administratively down down
Port-channel256            unassigned      YES NVRAM  down                    down
```

Examine the results for the **`show run | include lag`** command to determine how many times the keyword “lag” is contained in the configuration, as shown in Example 4-5.

Example 4-5 Lag keyword occurrences in the configuration

```
description Port 11 lag_13
description Port 13 lag_13
description Port 20 lag_13
```

Finally, an examination of the evaluation results shows that the configuration was examined for the logging command only under the interfaces with “lag” in the description, as shown in Figure 4-24.

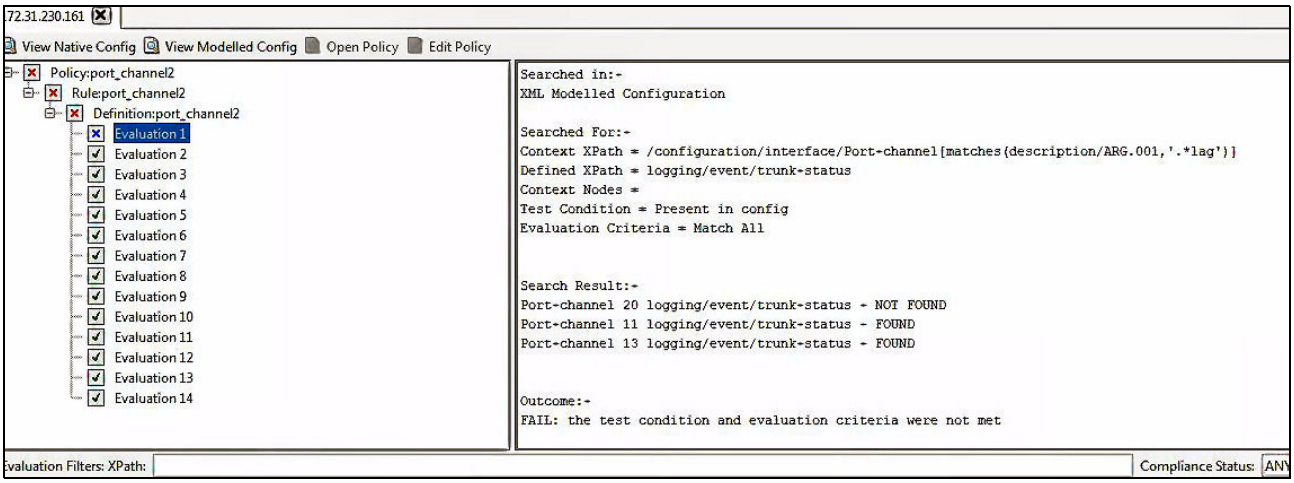

Figure 4-24 Evaluation results for contextual xpaths

4.4 Summary

Use of the golden configuration to create compliance policies greatly reduces the time to create the policies, especially for large configurations. Traceability from engineering or security best practices also is ensured while allowing compliance operations to be performed in the network with relative ease.

Part 3

Network event and cognitive-related scenarios

In this part, network event and cognitive-related scenarios are described.

5

Known slow traffic between two points in a network

The demonstration that is presented in this chapter guides you through the procedures of performing a search, selecting data from your search results, creating a graph that is based on that data, and starting a dashboard. The tasks are based on the sample data.

This scenario describes a user experience that manages problems with slow traffic between two points in a network. It helps to understand how Network Manager Insight Pack can be used effectively to search for OMNIbus events and how wanted information can be displayed in the IBM Operations Analytics - Log Analysis web console.

This chapter includes the following topics:

5.1 Scenario description

Company A is a large company with many facilities through out the Europe. Tom is one of operators of the consolidated Operations Center who is responsible for managing, evaluating, and resolving events throughout the enterprise. Tom is tasked with monitoring, resolving, and improving the efficiencies of the Operations Center.

This scenario shows how Tom uses Event Search Analytics for Operational Agility and Efficiency. In his daily work, Tom is informed about critical out-of-memory errors on intermediate devices between the end points. From Network Manager topology view, he selects the two end points and uses the Network Manager Insight Pack to search for OMNIbus events between the two nodes. Critical out-of-memory errors are found on intermediate devices between the end points.

5.1.1 Business value

Machine-driven, analytics-based event grouping assists Tom in several ways.

When you integrate IBM Operations Analytics - Log Analysis with IBM Tivoli Netcool/OMNIbus, you can use the text analytics features to find patterns and trends in event data. With the integration of these two products, you can view and search historical and real-time event data from IBM Tivoli Netcool/OMNIbus in the IBM Operations Analytics - Log Analysis user interface.

IBM Operations Analytics - Log Analysis parses event data into a format that is suitable for searching and indexing. The event data is transferred from IBM Tivoli Netcool/OMNIbus to IBM Operations Analytics - Log Analysis by the IBM Tivoli Netcool/OMNIbus Message Bus Gateway.

When a new event arrives or an event is reinserted in the ObjectServer, event data is sent to the Gateway for Message Bus with an Insert, Delete, Update, or Control (IDUC) signal or an accelerated event notification (AEN) channel. The gateway then sends the event through an HTTP interface to the IBM Operations Analytics - Log Analysis server.

This scenario is most suitable for customers who do not need to maintain their environment or want to start with a fresh approach. This option can also be a good option for a customer who no longer has (or never had) the skills in the Network Manager Insight Pack.

5.2 Scenario topology

For more information about system components and default settings in the test environment, see 1.4, “Our environment for the scenarios” on page 18. The solution that is used in this scenario includes system components that are installed on the following systems:

- `itnmrh61.test.ibm.com` server contains the following software:
  - IBM DB2 v10.5
  - Network Manager 4.2
  - Network Manager Health Dashboard 4.2
- `itnmrh62.test.ibm.com` server contains the following software**:**
  - Netcool/OMNIbus v8.1 FP 7
  - Netcool Web GUI v8.1 FP 5
  - Netcool/OMNIbus Gateway for Message Bus

  - Netcool/OMNIbus Syslog Probe
  - MTTrapd (SNMP) probe
- `itnmlogs.test.ibm.com` server features Operations Analytics - Log Analysis v1.3.2 software

The following dedicated settings are used on the `itnmrh61.test.ibm.com` machine**:**

- IBM DB2 v10.5 Enterprise Server Edition that is run in this scenario is hosting the NCIM Topology database and includes the following dedicated settings:

```
tnm.database.host=itnmrh61.test.ibm.com
tnm.database.dbname=ITNM
tnm.database.username=ncim
```

- ObjectServer includes the following dedicated settings:

```
server name = NM_LA
server host = itnmrh62.test.ibm.com
server port = 4100
```

Netcool simulated event generator

Netcool/OMNIbus can create event records by using information from many sources. For testing purposes, we use the event generator to send simulated events to the IBM Tivoli Netcool/OMNIbus ObjectServer that is generated with Netcool Event Generator. For more information about downloading the event generator for your environment, see the following resources (your IBM ID is required):

- NETCOOL Event Generator for Solaris:

  https://ibm.biz/BdrDSs

- NETCOOL Event Generator 1.0.1 for Linux

  https://ibm.biz/BdrDSi

- NETCOOL Event Generator for Windows

  https://ibm.biz/BdrDSZ

Note: For more information about Netcool Event Generator for Windows, see the following IBM developerWorks® website:

https://ibm.biz/BdrDS2

This software is delivered “as-is” and is not supported.

5.3 Scenario steps

This section describes the process that is used to recreate and solve the issue of a known slow traffic between two points in a network.

Netcool/OMNIbus can create event records by using information from many sources. The Netcool/OMNIbus SNMP probe receives the trap and events that are created in the ObjectServer. After the events show in the ObjectServer, the Message Bus Gateway retrieves them and sends them to Operations Analytics - Log Analysis.

Tom performs the following steps:

1. He logs in to the following Integrated Portal console, as shown in Figure 5-1:

   https://itnmrh61.test.ibm.com:16311/ibm/console/logon.jsp

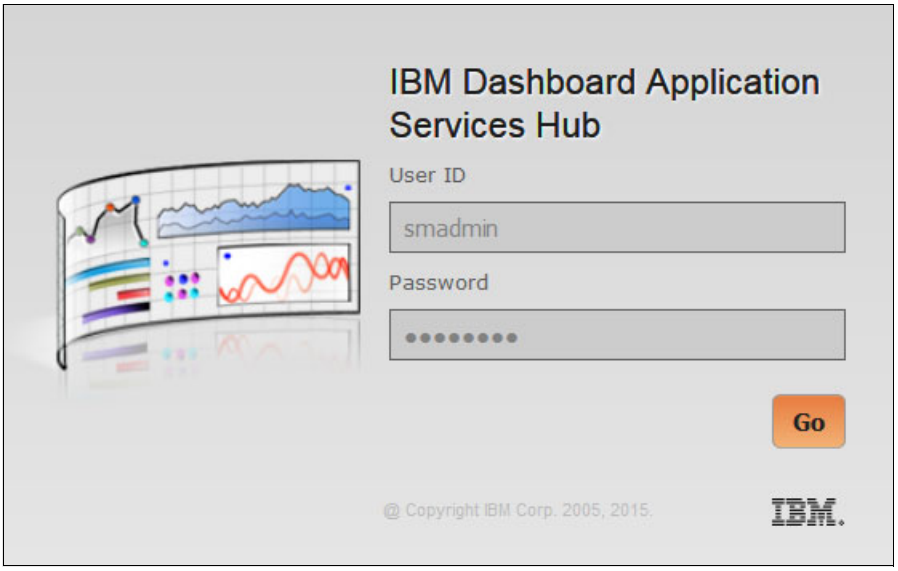

Figure 5-1 IBM Dashboard Application login

2. He clicks **Availability** → **Network Availability** → **Network Views** and selects the **Libraries** tab.
3. Tom opens the **Cisco Devices** view. He can see the topology view, as shown in Figure 5-2 on page 111. He might need to zoom in the view by using the top menu buttons or right-clicking any area in the window.

> **Tip:** It is easier to zoom in by using the keyboard and mouse. Tom can use mouse wheel for zooming. He can also use mouse wheel, in combination with keyboard Shift and Ctrl, to move in required position (up, down, right, and left).

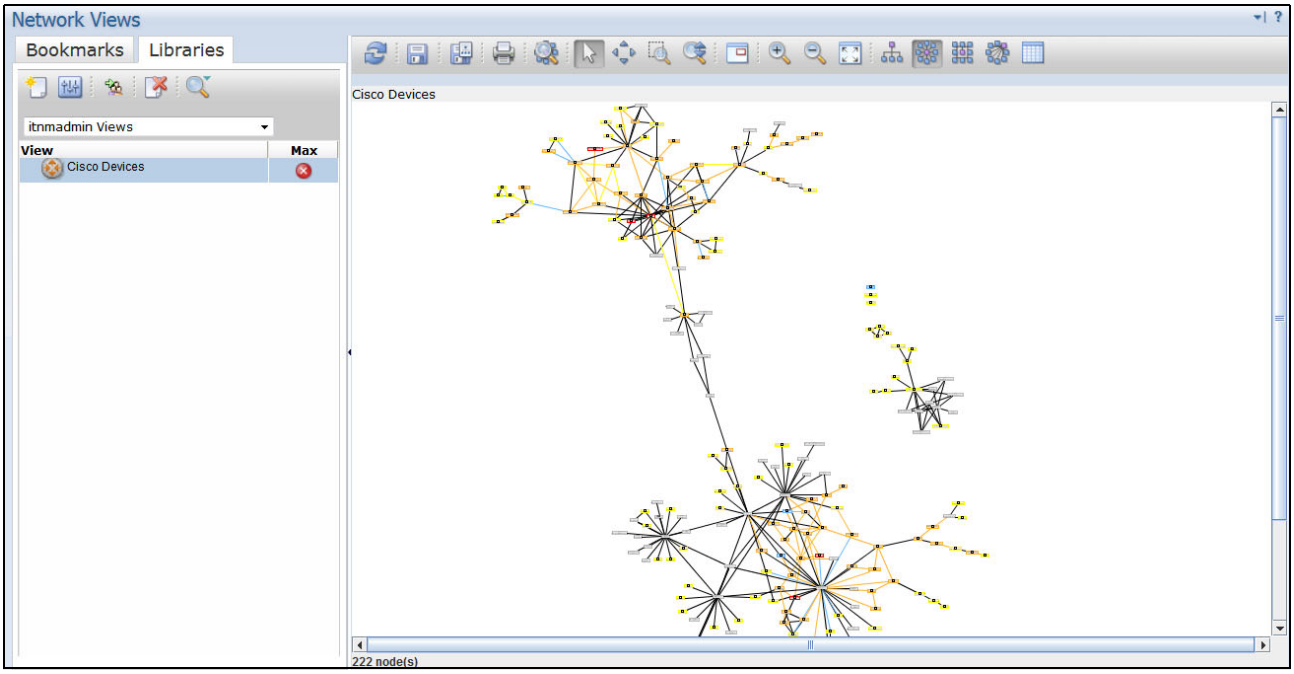

Figure 5-2 ITNM console view

4. To determine what is occurring between chosen devices, Tom presses and holds the Ctrl key and then, press the left mouse button to select the following devices:

   `london-asbr-cr72.uk.eu.test.lab`
   `paris-asbr-cr36.fra.eu.test.lab`

5. For the purposes of this scenario, we use NETCOOL Event Generator for Windows with a loaded `.xml` file that contains the event Table for the End-To-End Search Demonstration. Events are set to match the “Cisco Devices” Network View.

   As shown in Figure 5-3 on page 112, Tom selects the `paris-asbr-cr36.fra.eu.test.lab` device.

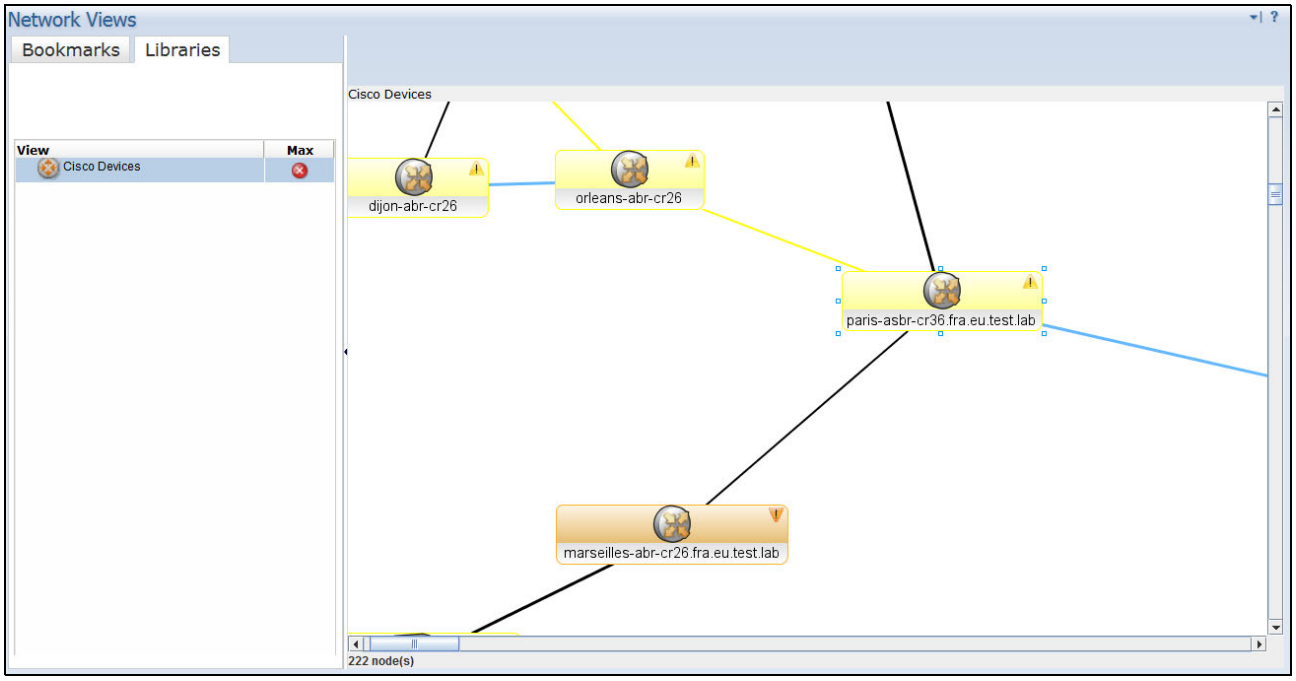

Figure 5-3 Selected Cisco device

6. Tom right-clicks one of selected devices and chooses **Event Search** → **Find events between two nodes** → **Layer 2 topology** → **Last 15 minutes**, as shown in Figure 5-4.

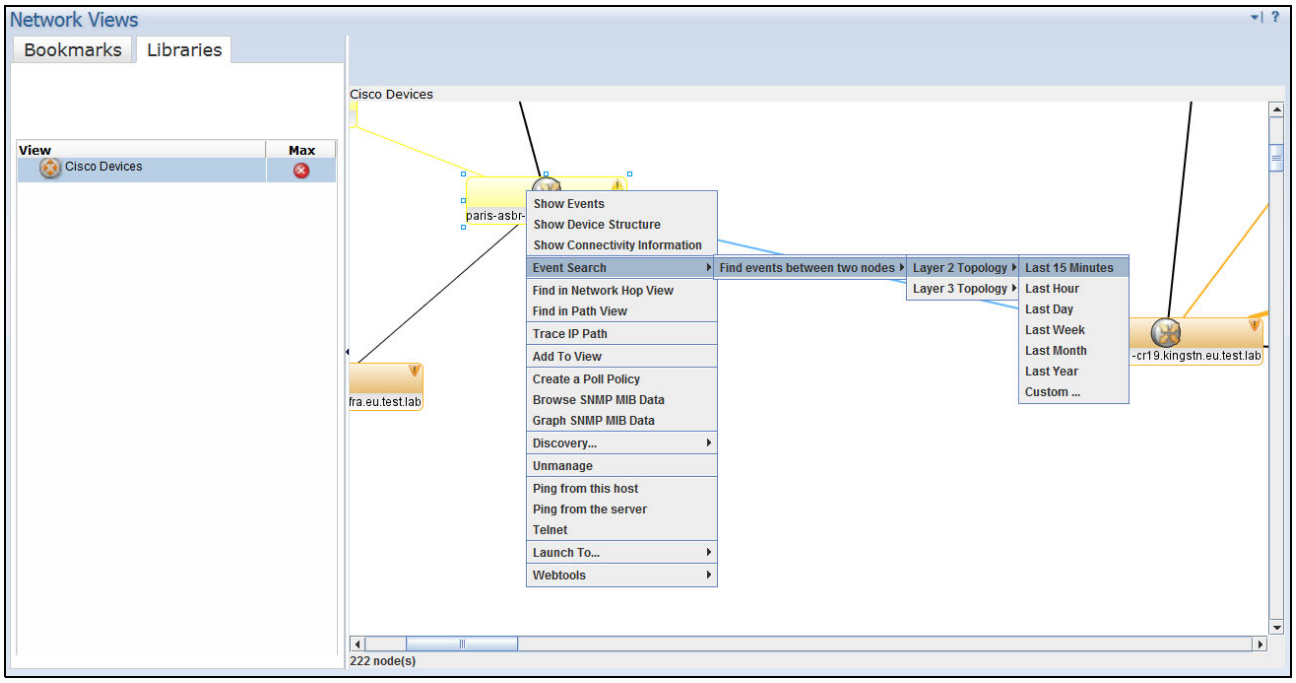

Figure 5-4 Find events between two nodes

The WebAnalysis site opens. (Tom might need to log in to access the site). The first search he performs after the IBM Operations Analytics - Log Analysis new processes were restarted might take longer to complete than subsequent searches.

7. The IBM Operations Analytics - Log Analysis console is shown. Tom can search the log files for keywords. Search results are displayed in a list or table format. Search results also are displayed in a distribution graph because he searched for events between two selected devices, as shown in Figure 5-5.

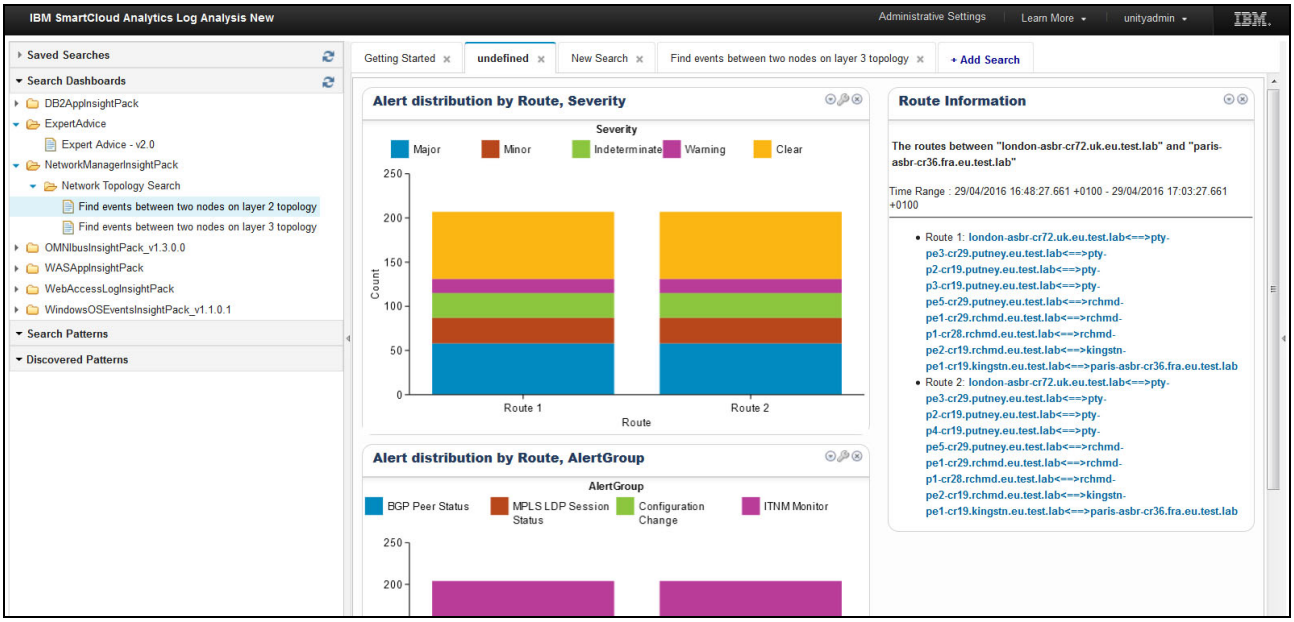

Figure 5-5 IBM SmartCloud Analytics - Log Analytics User Assistance

8. While clicking one of the routes, Tom can see a timeline and issues description. Because he investigates the memory issue now, the issue that interests him is low memory error between the devices, as shown in Figure 5-6.

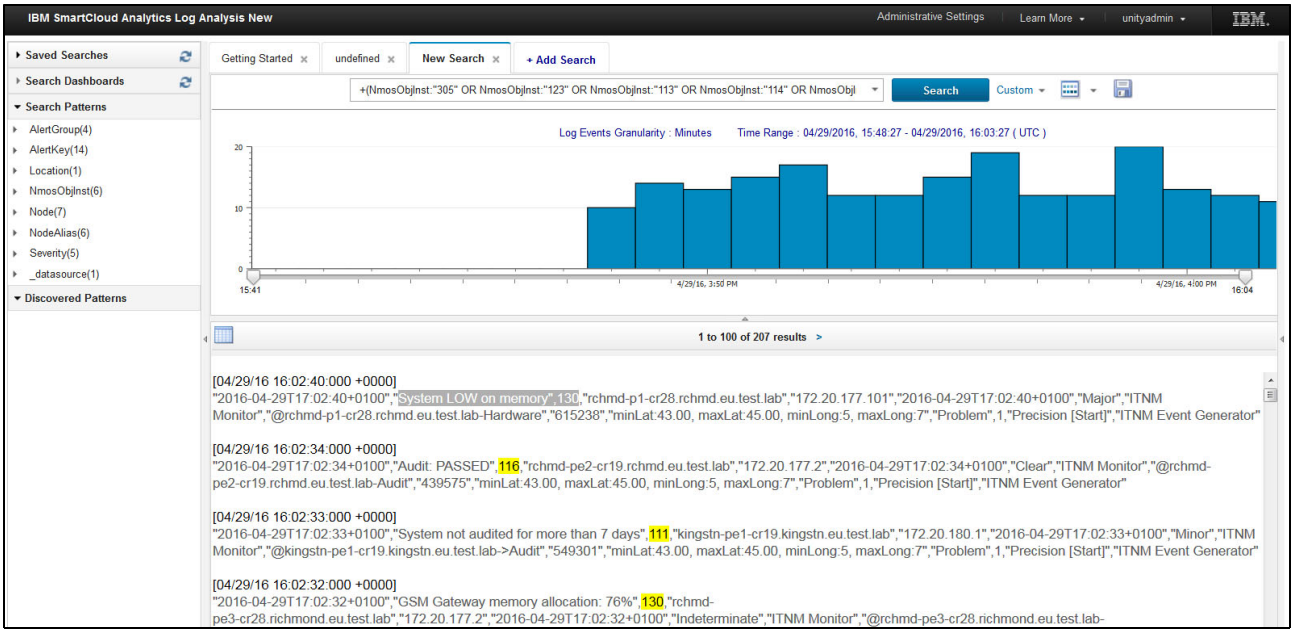

Figure 5-6 Dashboard view

9. After he displays the events on one of the routes, Tom clicks the **Grid view** icon and identifies which device is generating the highest severity alerts, as shown in Figure 5-7.

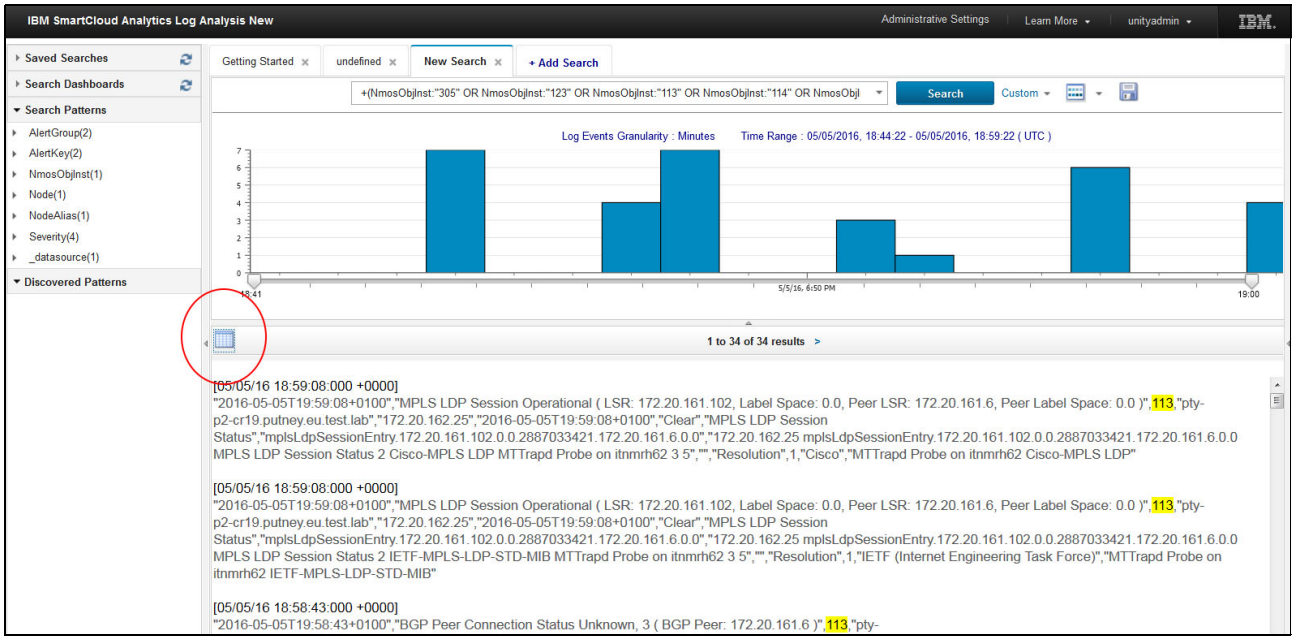

Figure 5-7 Identifying the problematic device on the Topology Map

He can see which device is problematic and note its entity ID, as shown in Figure 5-8.

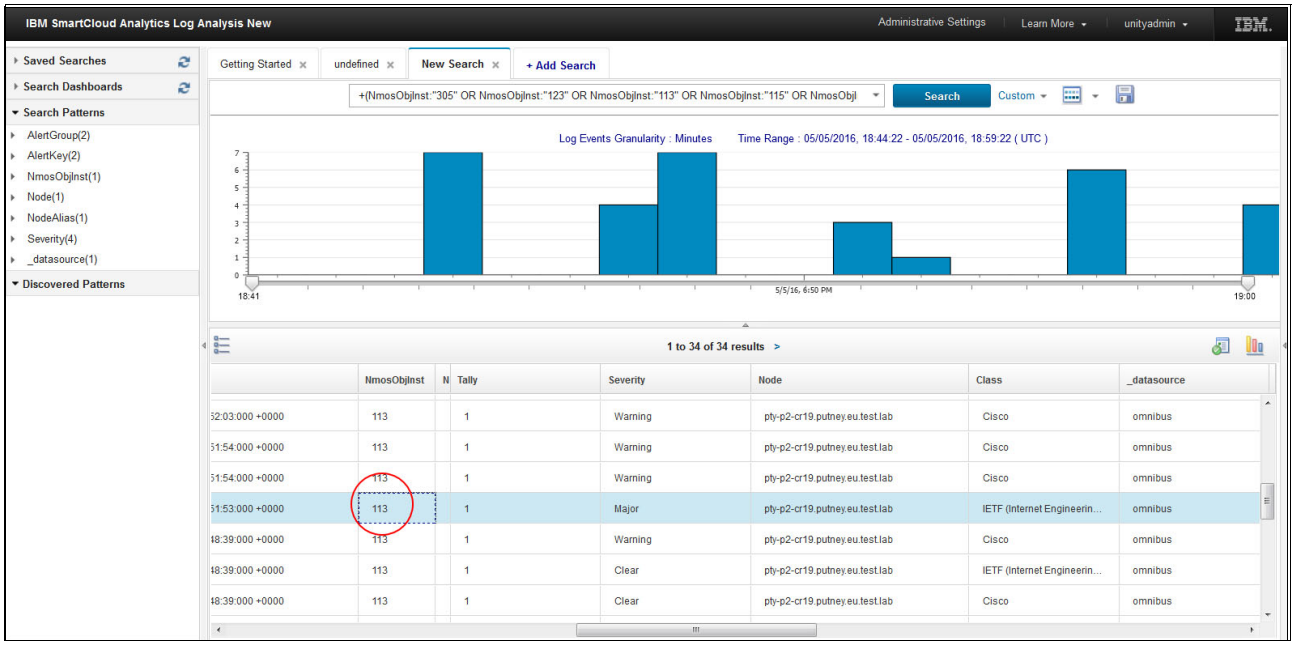

Figure 5-8 Finding entity ID of the problematic device

10. He returns to the Topology Map view and clicks the **Search** icon to find the problematic device, as shown in Figure 5-9.

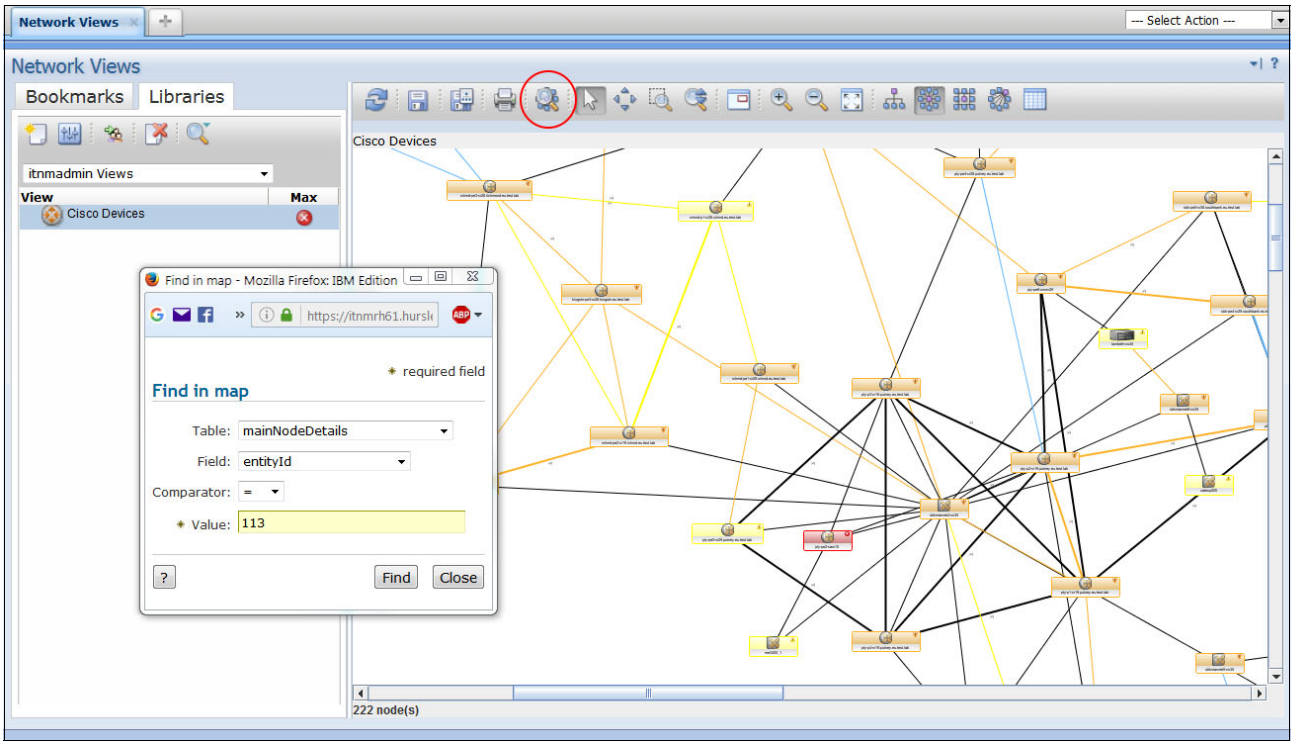

Figure 5-9 Looking for the problematic device from Topology Map

11. The device is found and highlighted, as shown in Figure 5-10.

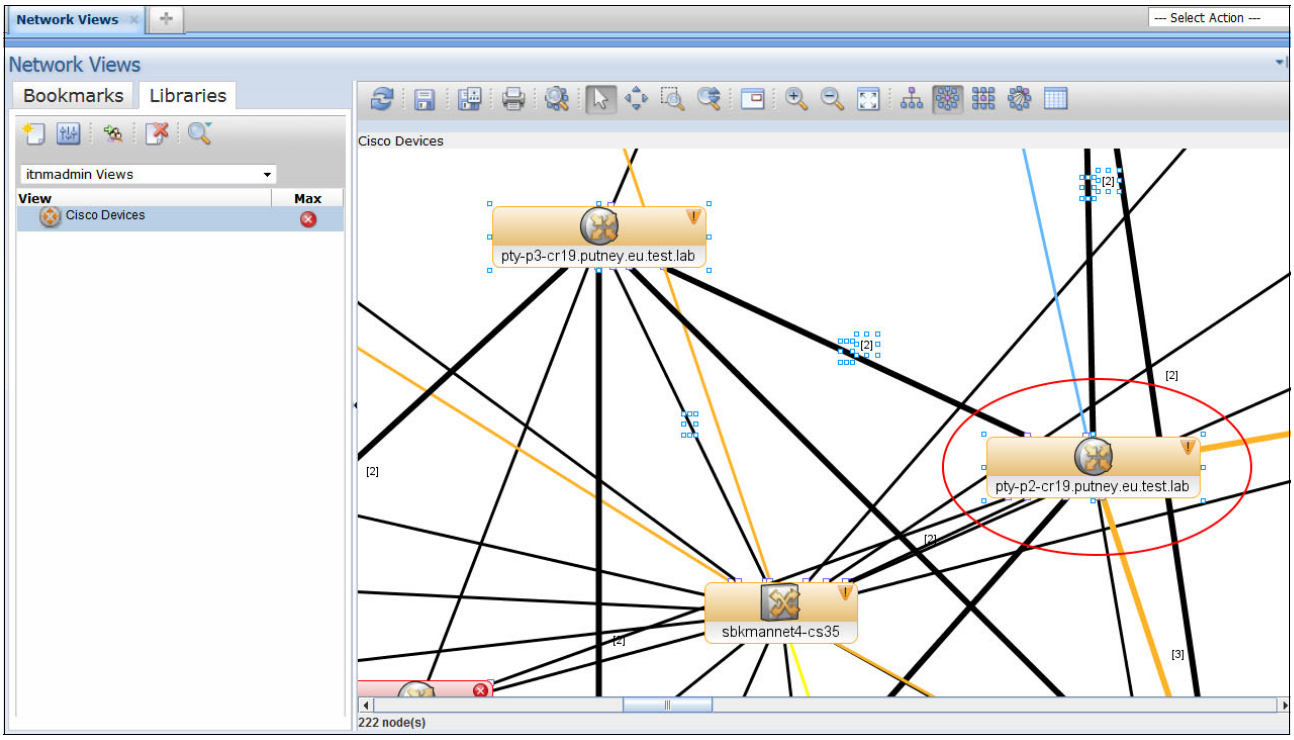

Figure 5-10 Problematic device found

12. Tom right-clicks the highlighted device and chooses **Show Events**, which shows the details of the state of this device (see Figure 5-11).

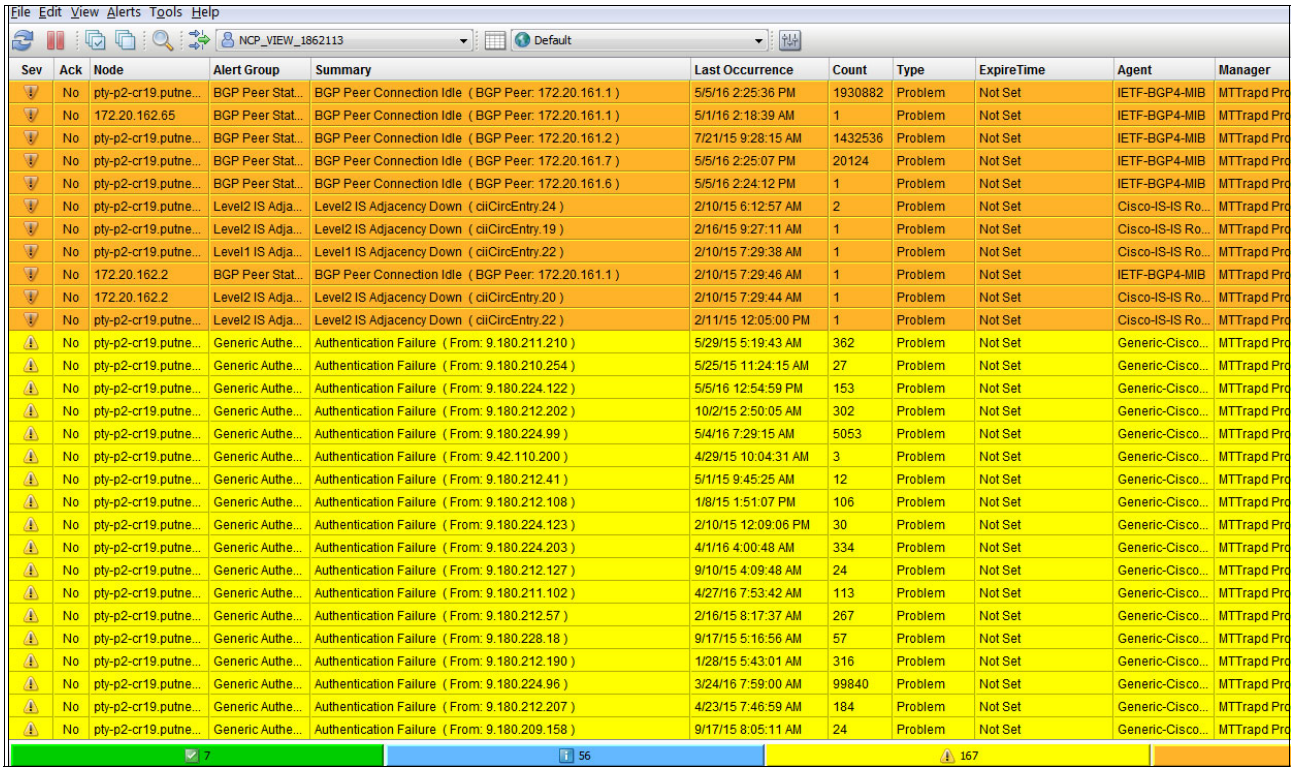

Sev	Ack	Node	Alert Group	Summary	Last Occurrence	Count	Type	ExpireTime	Agent	Manager
	No	pty-p2-cr19.putne...	BGP Peer Stat...	BGP Peer Connection Idle (BGP Peer: 172.20.161.1)	5/5/16 2:25:36 PM	1930882	Problem	Not Set	IETF-BGP4-MIB	MTTrapd Pro
	No	172.20.162.65	BGP Peer Stat...	BGP Peer Connection Idle (BGP Peer: 172.20.161.1)	5/1/16 2:18:39 AM	1	Problem	Not Set	IETF-BGP4-MIB	MTTrapd Pro
	No	pty-p2-cr19.putne...	BGP Peer Stat...	BGP Peer Connection Idle (BGP Peer: 172.20.161.2)	7/21/15 9:28:15 AM	1432536	Problem	Not Set	IETF-BGP4-MIB	MTTrapd Pro
	No	pty-p2-cr19.putne...	BGP Peer Stat...	BGP Peer Connection Idle (BGP Peer: 172.20.161.7)	5/5/16 2:25:07 PM	20124	Problem	Not Set	IETF-BGP4-MIB	MTTrapd Pro
	No	pty-p2-cr19.putne...	BGP Peer Stat...	BGP Peer Connection Idle (BGP Peer: 172.20.161.6)	5/5/16 2:24:12 PM	1	Problem	Not Set	IETF-BGP4-MIB	MTTrapd Pro
	No	pty-p2-cr19.putne...	Level2 IS Adja...	Level2 IS Adjacency Down (ciiCircEntry.24)	2/10/15 6:12:57 AM	2	Problem	Not Set	Cisco-IS-IS Ro...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Level2 IS Adja...	Level2 IS Adjacency Down (ciiCircEntry.19)	2/16/15 9:27:11 AM	1	Problem	Not Set	Cisco-IS-IS Ro...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Level1 IS Adja...	Level1 IS Adjacency Down (ciiCircEntry.22)	2/10/15 7:29:38 AM	1	Problem	Not Set	Cisco-IS-IS Ro...	MTTrapd Pro
	No	172.20.162.2	BGP Peer Stat...	BGP Peer Connection Idle (BGP Peer: 172.20.161.1)	2/10/15 7:29:46 AM	1	Problem	Not Set	IETF-BGP4-MIB	MTTrapd Pro
	No	172.20.162.2	Level2 IS Adja...	Level2 IS Adjacency Down (ciiCircEntry.20)	2/10/15 7:29:44 AM	1	Problem	Not Set	Cisco-IS-IS Ro...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Level2 IS Adja...	Level2 IS Adjacency Down (ciiCircEntry.22)	2/11/15 12:05:00 PM	1	Problem	Not Set	Cisco-IS-IS Ro...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Generic Authe...	Authentication Failure (From: 9.180.211.210)	5/29/15 5:19:43 AM	362	Problem	Not Set	Generic-Cisco...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Generic Authe...	Authentication Failure (From: 9.180.210.254)	5/25/15 11:24:15 AM	27	Problem	Not Set	Generic-Cisco...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Generic Authe...	Authentication Failure (From: 9.180.224.122)	5/5/16 12:54:59 PM	153	Problem	Not Set	Generic-Cisco...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Generic Authe...	Authentication Failure (From: 9.180.212.202)	10/2/15 2:50:05 AM	302	Problem	Not Set	Generic-Cisco...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Generic Authe...	Authentication Failure (From: 9.180.224.99)	5/4/16 7:29:15 AM	5053	Problem	Not Set	Generic-Cisco...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Generic Authe...	Authentication Failure (From: 9.42.110.200)	4/29/15 10:04:31 AM	3	Problem	Not Set	Generic-Cisco...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Generic Authe...	Authentication Failure (From: 9.180.212.41)	5/1/15 9:45:25 AM	12	Problem	Not Set	Generic-Cisco...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Generic Authe...	Authentication Failure (From: 9.180.212.108)	1/8/15 1:51:07 PM	106	Problem	Not Set	Generic-Cisco...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Generic Authe...	Authentication Failure (From: 9.180.224.123)	2/10/15 12:09:06 PM	30	Problem	Not Set	Generic-Cisco...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Generic Authe...	Authentication Failure (From: 9.180.224.203)	4/1/16 4:00:48 AM	334	Problem	Not Set	Generic-Cisco...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Generic Authe...	Authentication Failure (From: 9.180.212.127)	9/10/15 4:09:48 AM	24	Problem	Not Set	Generic-Cisco...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Generic Authe...	Authentication Failure (From: 9.180.211.102)	4/27/16 7:53:42 AM	113	Problem	Not Set	Generic-Cisco...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Generic Authe...	Authentication Failure (From: 9.180.212.57)	2/16/15 8:17:37 AM	267	Problem	Not Set	Generic-Cisco...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Generic Authe...	Authentication Failure (From: 9.180.228.18)	9/17/15 5:16:56 AM	57	Problem	Not Set	Generic-Cisco...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Generic Authe...	Authentication Failure (From: 9.180.212.190)	1/28/15 5:43:01 AM	316	Problem	Not Set	Generic-Cisco...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Generic Authe...	Authentication Failure (From: 9.180.224.96)	3/24/16 7:59:00 AM	99840	Problem	Not Set	Generic-Cisco...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Generic Authe...	Authentication Failure (From: 9.180.212.207)	4/23/15 7:46:59 AM	184	Problem	Not Set	Generic-Cisco...	MTTrapd Pro
	No	pty-p2-cr19.putne...	Generic Authe...	Authentication Failure (From: 9.180.209.158)	9/17/15 8:05:11 AM	24	Problem	Not Set	Generic-Cisco...	MTTrapd Pro

7 56 167

Figure 5-11 Problematic device event view

13. Tom can limit your search to one or more data sources. To further limit his search, he can create a time filter.

14. To include data from warning and error messages, Tom adds a logical operator value of OR to the search box. He clicks in the search box and adds the string OR to the end of the search box value. He ensures that a space is added before and after the OR string.

5.4 Summary

As described in this scenario, Operations Analytics - Log Analysis is a significant and beneficial feature of Networks for Operations Insight. By using it, an IT practitioner, such as Tom, can get a consolidated view of network devices and identify and troubleshoot network outages fast and resolve them quickly.

Event search applies the search and analysis capabilities of Operations Analytics - Log Analysis to events that are monitored and managed by Tivoli Netcool/OMNIbus. Events are transferred from the ObjectServer through the Gateway for Message Bus to Operations Analytics - Log Analysis, where they are imported into a data source and indexed for searching. After the events are indexed, you can search every occurrence of real-time and historical events.

The Tivoli Netcool/OMNIbus Insight Pack is installed into Operations Analytics - Log Analysis and provides custom modules that search the events based on various criteria. By using keyword searches and dynamic drill-down functions, you can more closely review event data for more information.

Tooling can be installed into the Web GUI that starts the modules from the right-click menus of the Event Viewer and the Active Event List. An *event reduction wizard* is also supplied that includes information and applications that can help you analyze and reduce volumes of events and minimize the “noise” in your monitored environment.

Note: For more information about Netcool Operations Insight 1.4.0.1 - Event search, see the following IBM Knowledge Center site:

https://ibm.biz/BdrDvk

6

Analytics-based event grouping and seasonality

The scenario that is presented in this chapter describes how you can use analytics-based event grouping and seasonality reporting to reduce overall event volume that is displayed onto operational dashboards.

This chapter includes the following topics:

6.1 Introduction

In this section, the concepts of analytics-based event grouping and seasonality are introduced.

6.1.1 Analytics-based event grouping

The analytics-based event grouping, also known as *related event grouping*, functionally analyzes the historic event archive (REPORTER schema that is populated by events from OMNIbus) and looks for groups of events that always occur together.

Identifying incidences of groups of events always occurring together, particularly when this issue occurs several times, provides strong evidence that the events are in some way related to each other and potentially the same fault. Even if causation cannot be implied from this relationship, correlation can be inferred. Knowing that a group of events always occurs together can provide valuable insights to the underlying infrastructure and any faults that occur.

After the analytics engine performs an analysis of the historic event archive and looks for groups of events that always occur together, a Subject Matter Expert (SME) can examine the resulting groupings that are found. The results dashboard provides a convenient portal through which to inspect the discovered groupings, the times the groupings occurred previously, and the individual events that were present in each case.

The value of analytics-based event groupings is that it builds relationships among events that were previously unknown. By identifying these relationships, remediation or preventative actions can be put in place similar to “well-known” event groupings and patterns.

6.1.2 Seasonality

The seasonality function works by analyzing the historic event archive (REPORTER schema that is populated by events from OMNIbus) and looking for individual events that occur with any sort of degree of regularity. For example, this occurrence might be at the same minute of the hour, or the same hour of the day, or the same day of the week, or day of the month or a combination of these instances.

The value of this analysis is that it helps identify chronic issues in our environment; that is, issues whose temporal characteristics often are not noticed by operators in the NOC, such as recurring critical disk space alarms, network congestion, or degraded application performance. In many cases, the characteristics of the seasonality are clues to the cause of the underlying problem.

When the seasonality results are reviewed, a SME might ask: “What happens every Monday at 4:00 AM?” This SME with some institutional knowledge who is reviewing the seasonality results knows that Monday at 4:00 AM is when the backups run. The SME also might provide the starting point for a discussion amongst other SMEs who might have that knowledge. A simple assessment of this issue might result in, for example, more disk being allocated to the backup job, network reconfiguration, or improved event management that is based on date or time evaluation.

By resolving chronic issues such as this reoccurring issue, a user can relatively easily calculate the monetary savings to the business by no longer submitting problem tickets.

6.2 Scenario description

Enterprise ABC is a large company with many facilities throughout the US. Helen is the manager of the consolidated Operations Center and is responsible for managing, evaluating, and resolving events throughout the enterprise. Helen is tasked with improving the efficiencies of the Operations Center and reducing the mean time to resolution on service-effecting events.

She identified the following list of challenges that are facing the Operations Center:

- There are too many events for operators to manage.
- The types and number of events are not easily categorized.
- It is difficult to prioritize the resolution activities.
- Multiple tickets often are opened from the many events.
- It is costly to the business to close all of the duplicate tickets.

The result of the challenges is a higher mean time to resolution rate than is wanted and inefficiencies in prioritizing and handling events. Helen believes that she can improve her mean time to resolution rates by providing technology domain SMEs with categorized analysis reports and event groupings.

Helen enlisted the help of Marco (a SME from the facilities engineering team) to review a set of seasonal and related event groups.

6.2.1 Business value

Machine-driven, analytics-based event grouping assists Enterprise ABC in the following ways:

- Large volumes of historical data can be processed much faster than manual reporting.
- Repeating patterns of events are easily displayed and delivered to the appropriate SMEs.
- Previously unknown relationships can be discovered.
- Event suppression or automated remediation can be applied before human interaction is required, which reduces the overall visible volume of events.
- Exception analysis can be applied; therefore, an alarm can be raised if expected events do not occur.

Through the application of machine-driven event analytics, Helen believes that she can reduce the overall “noise” in the event stream by providing only actionable events to the operators. By doing so, efficiencies will be improved and mean time to resolution reduced.

6.3 Scenario topology

For this scenario, we used the environment that is described in 1.4, “Our environment for the scenarios” on page 18. IBM DB2 v10.5 Enterprise Server Edition that is run in this scenario is hosting the Historical Event Archive (REPORTER) database, which is used in event analysis.

6.4 Scenario steps

This section describes the steps that were used to define actionable-related events by using analytics-based grouping.

The following assumptions were made in producing this scenario:

- All components of the topology are installed and running.
- SNMP Events were received from Infrastructure devices and are enriched by the OMNIbus SNMP probe.
- Events were populated into the Historical Event Archive.

6.4.1 Creating an event analytics configuration

The first step in applying analytics to the historical event records is to create an analytics configuration.

Because Helen observed many infrastructure-related events, she chose to limit her first analysis to run against SNMP events that were related to infrastructure. The OMNIbus SNMP probe rules assigns a “Manager” value of “SNMP SCADA” to events that are captured matching the appropriate Enterprise Object Identifier (OID). By using this attribute, Helen can constrain her first event analytics configuration rather than trying to build relationships across every event and type.

Configuring the Analytics portlet

Helen completed the following steps in the Configure Analytics portlet:

1. She selects **Insights** → **Configure Analytics** to access the Analytics Configuration portlet, as shown in Figure 6-1.

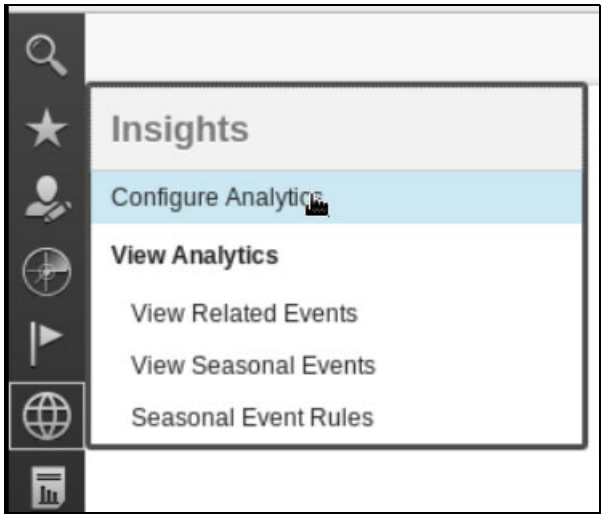

Figure 6-1 Selecting the Configure Analytics portlet

2. After the Configure Analytics portlet opens, Helen clicks the **Create New Configuration** icon, as shown in Figure 6-2.

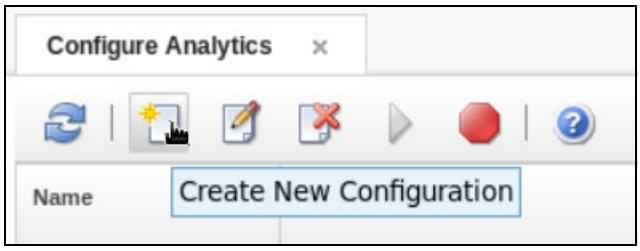

Figure 6-2 Create New Configuration icon

3. The New Configuration window opens. The appropriate parameters are entered and the configuration is saved and run. Because Helen chooses to limit the analysis to SNMP events from SCADA sources, a filter of `Manager = 'SNMP SCADA'` is entered for this configuration, as shown in Figure 6-3.

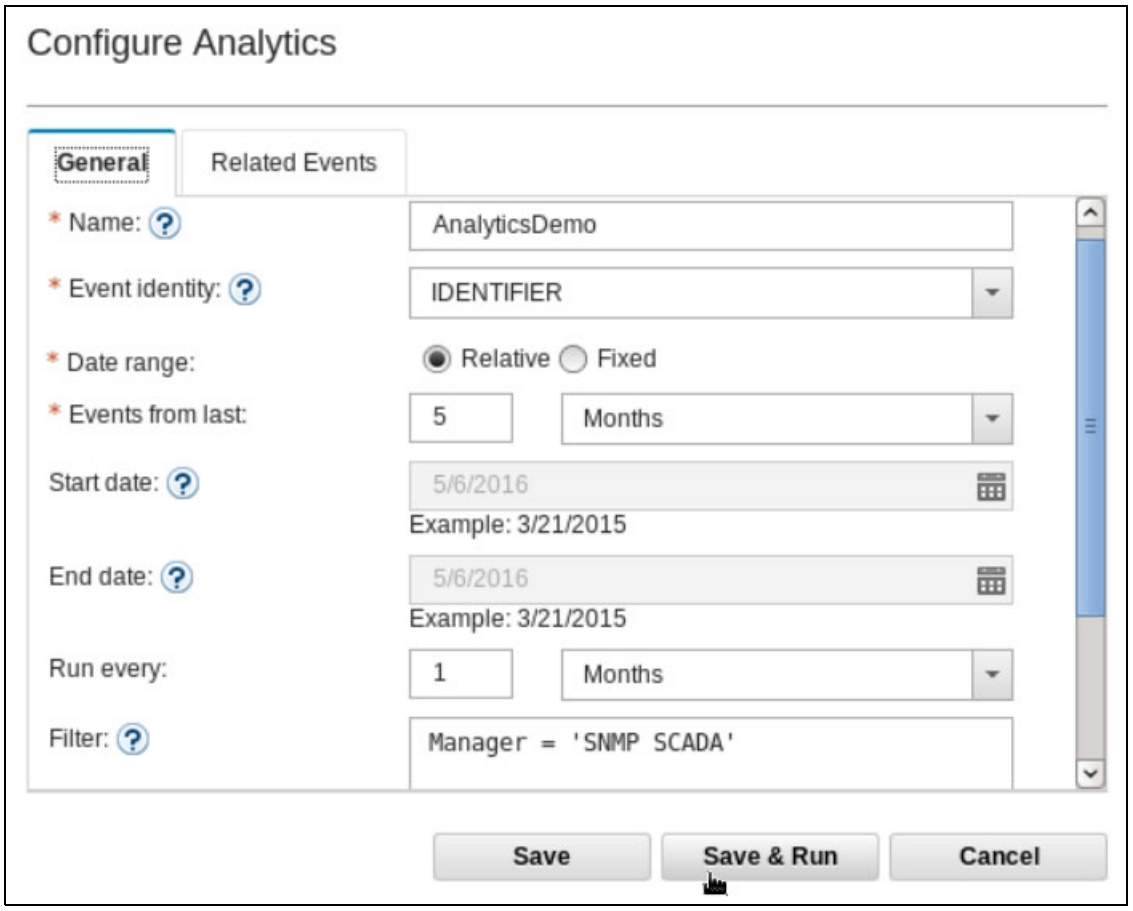

Figure 6-3 Analytics Configuration window

6.4.2 Actionable seasonal event reports

Now that the analytics configuration is complete and the configuration was run, Helen can share the results with Marco, the SME from the Infrastructure Engineering team.

The first set of results that are reviewed are the *seasonal events*. If there is a set of events that are showing a high rate of seasonality, they are good candidates for review and remediation.

Marco completed the following steps to review the season events:

1. He selects **Insights** → **View Seasonal Events** from the navigation menu to display the Seasonal Event page, as shown in Figure 6-4.

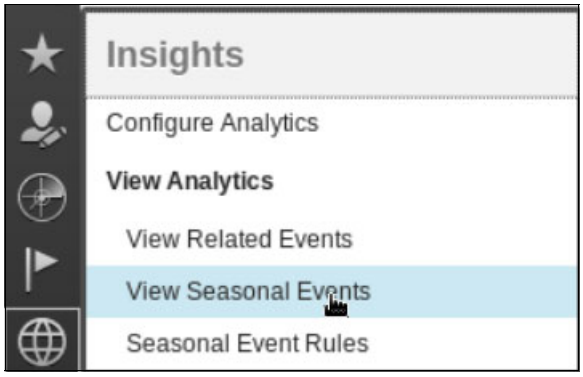

Figure 6-4 View Seasonal Events

The View Season Events page is displayed and the seasonal events that are discovered by the Analytics configurations are shown (see Figure 6-5).

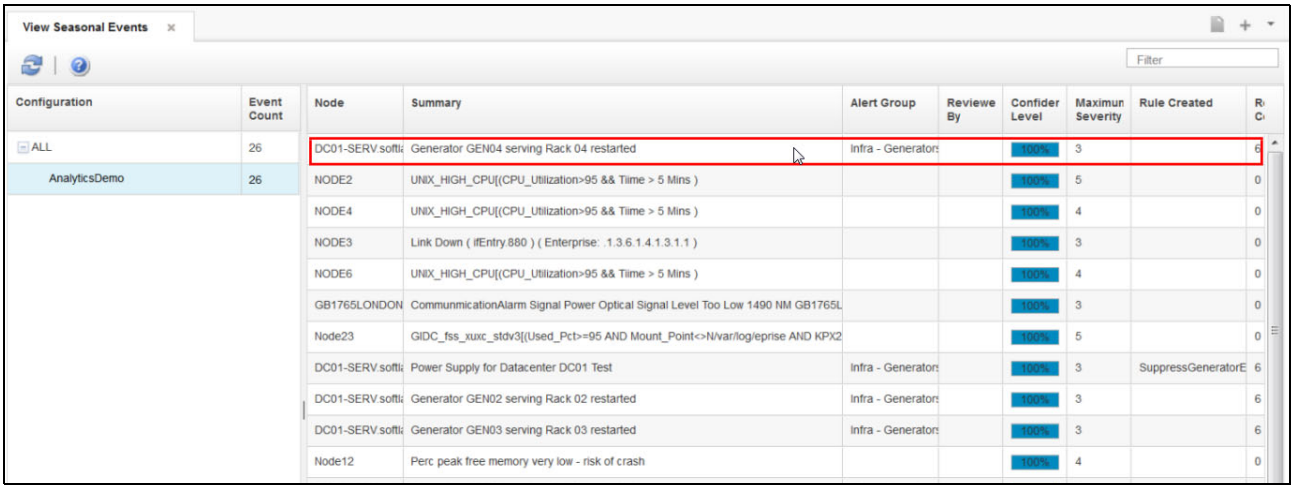

Figure 6-5 View Seasonal Events Page

2. Marco notes that there are several generator events that show a high degree of seasonality. To review the seasonality information in more detail, he right-clicks the first of these events and chooses **Show Seasonal Event Graphs**, as shown in Figure 6-6 on page 125.

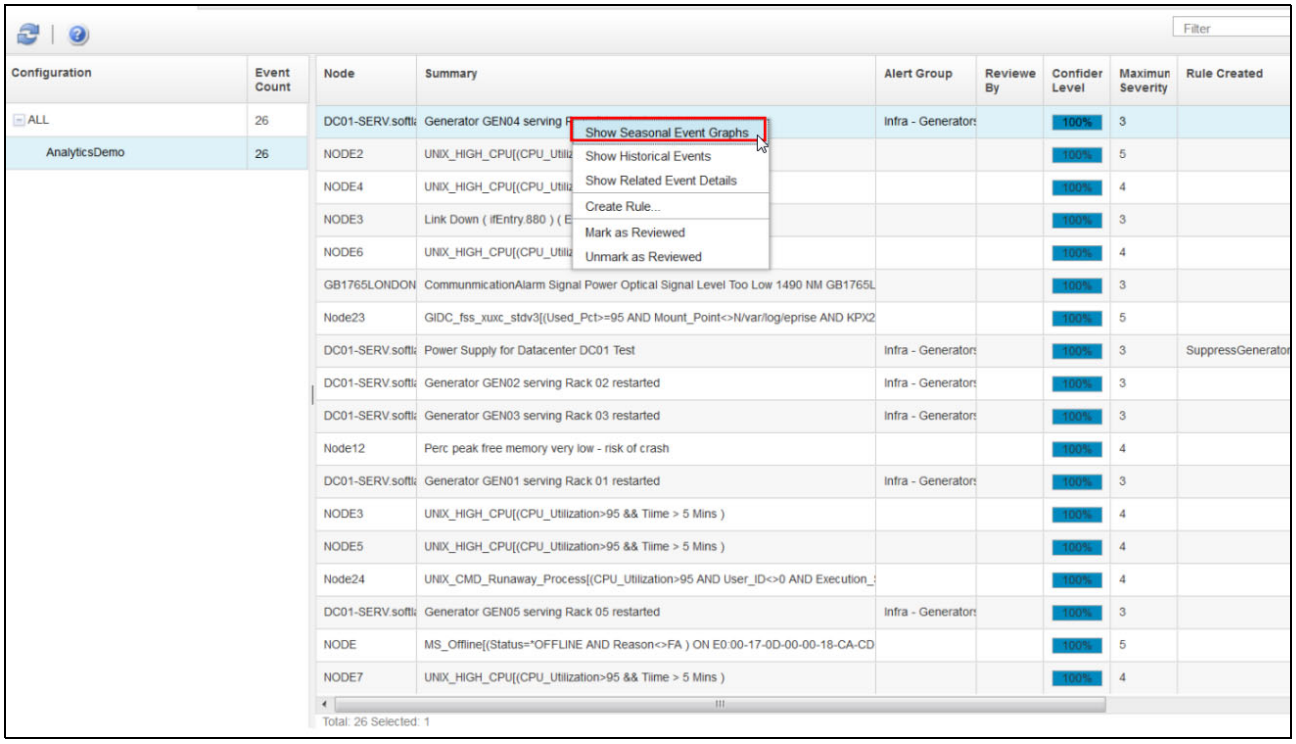

Figure 6-6 Right-click Show Seasonal Event Graphs

3. The seasonal event graph page opens and Marco can review hourly, day of week, and day of month graphs for this particular event, as shown in Figure 6-7.

Figure 6-7 Seasonal event graph page

4. The graphs help Marco understand that these events are occurring regularly every Monday at 9:00 AM.

   To drill down further into the information, Marco wants to view all of the events for the specific graphs. Marco performs the following steps to show the historical events behind the graph:

   a. He selects the bar in the Hour of the Day graph (9:00).

   b. In the Action section in the top right corner, he selects **Show Historical Events for Selected Bars** from the drop-down menu.

   The Historical Events page opens with the appropriate data, as shown in Figure 6-8.

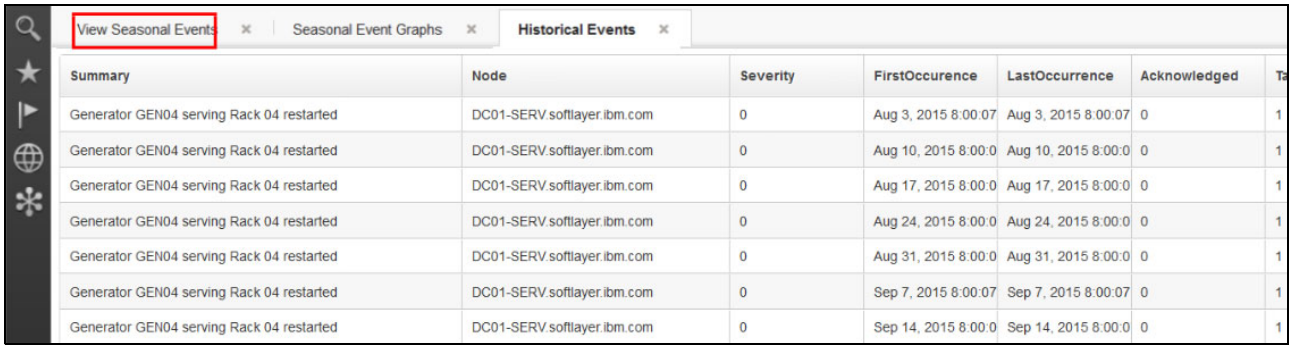

Figure 6-8 Historical Event details

6.4.3 Related event details analysis

Marco now knows that generator restart events are being created every Monday, but he wants to see whether Analytics determined whether other regularly occurring events might be related to the generator restart events. He can get this information by viewing related event details for an event that is listed in the View Season Events page.

Marco performed the following steps to show related event details for the Analytics configuration:

1. He returns to View Seasonal Events page.
2. He selects the appropriate generator restart event in the list.
3. Marco right-clicks the event to display the selection menu and chooses **Show Related Event Details**, as shown in Figure 6-9.

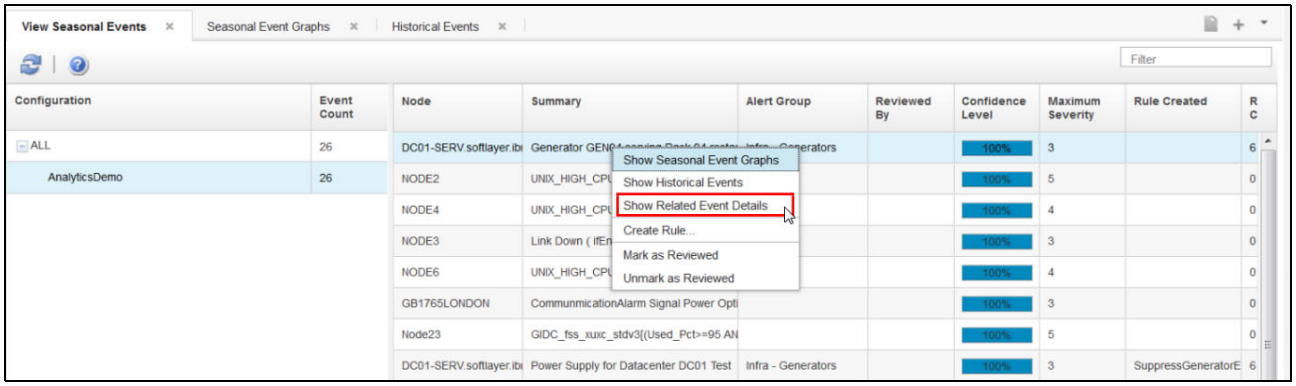

Figure 6-9 Related Event details

The Related Event Details page opens (see Figure 6-10) and a list of related events for a specific occurrence of the pivot event (the event that is selected in the Seasonal Event page) is shown. If there is more than one occurrence of the pivot event, the other occurrences are listed in the left frame.

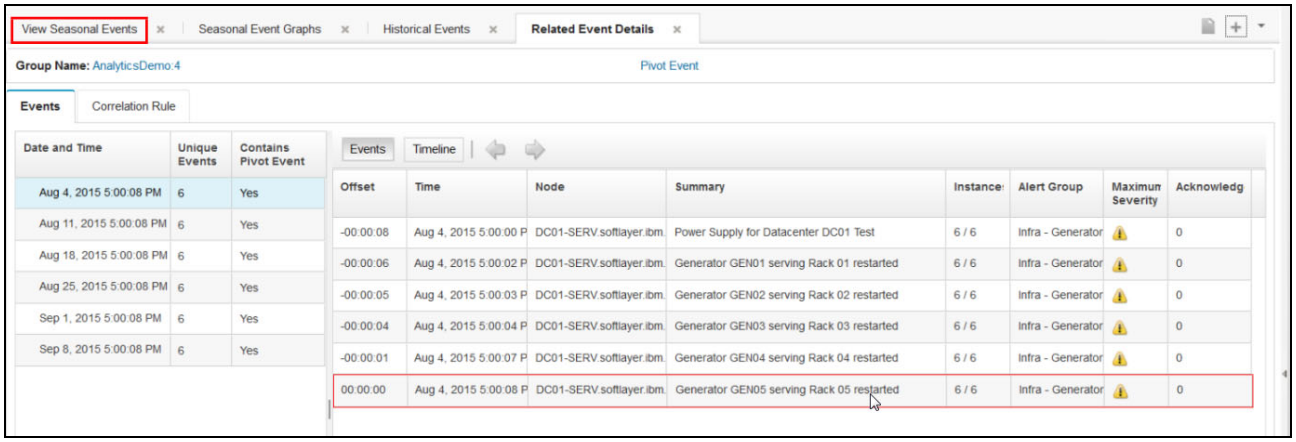

Figure 6-10 Related Event Details

The details list shows Marco a series of five other related events, which are listed in chronological order. A quick observation shows that a Power Supply test is originating the series of generator restart messages for that datacenter.

Now that Marco can see that the Power Supply test is the cause of six events that are occurring every Monday at 9:00 AM, he and Helen's team can take steps to reduce the noise and produce only actionable events. Marco suggests that they suppress the generator restart events; however, because a power test continues to occur, he wants to ensure that all generators are restarted after the test.

Helen's team can accomplish both goals by creating a seasonal event rule.

6.4.4 Creating and deploying seasonal event rules

Netcool Operations Insights allows administrators to deploy validated event rules without writing code in other applications or interfaces. By selecting a Seasonal event within the View Seasonal Event page, a Netcool Operations Insights administrator can start creating the rule from a right-click menu option.

Macro completed the following steps to create and deploy a seasonal event rule for the Power Supply Test event and related events:

1. He returns to the View Seasonal Events page.
2. He searches for or selected an event.

Figure 6-11 shows starting the Seasonal Rule Creation panel.

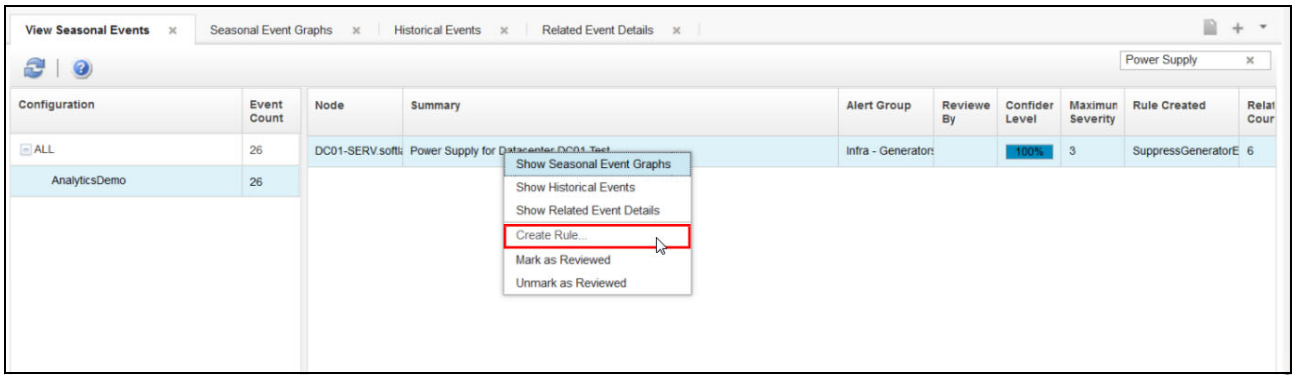

Figure 6-11 Launch Seasonal Rule Creation

3. Marco right-clicks the event and chose **Create Rule...** from the menu. The Season Rule Creation window opened, as shown in Figure 6-12.

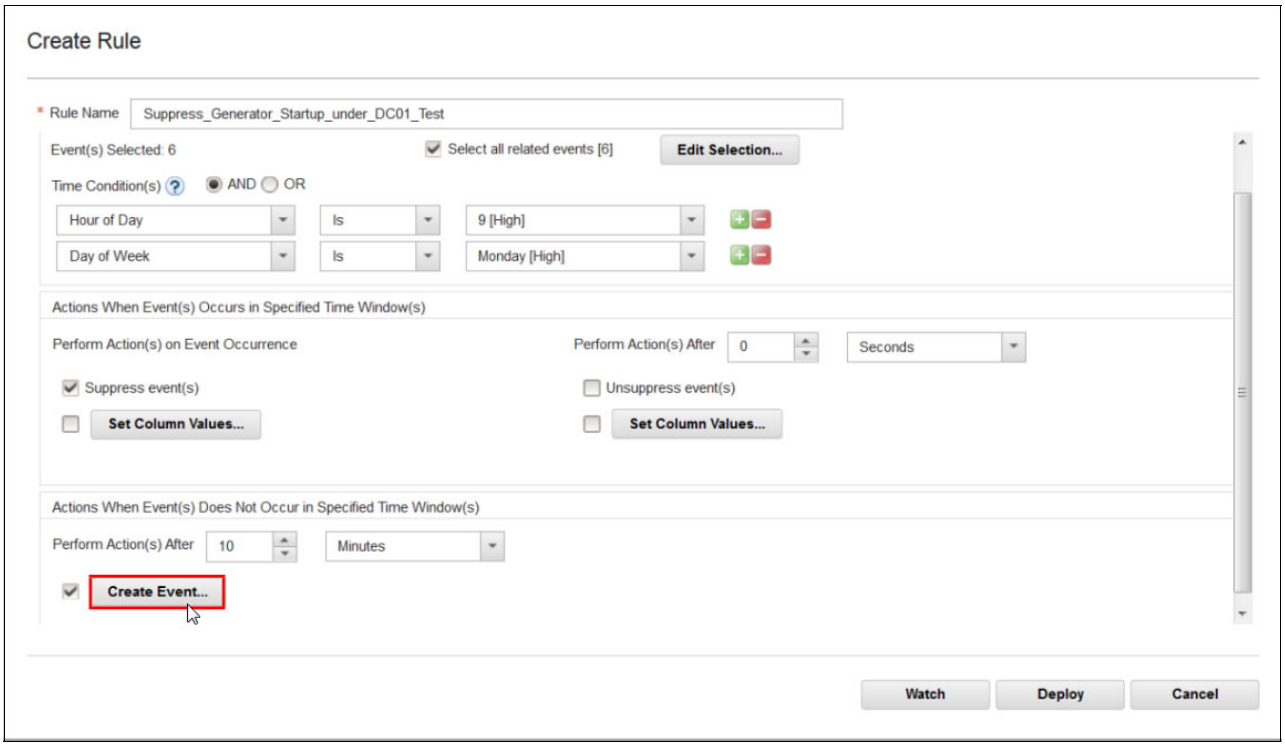

Figure 6-12 Create Seasonal Rule dialog window

4. Marco names the rule.
5. He selects all related events.
6. He chooses the appropriate Time Conditions. This step ensures that he is acting only on the known Seasonal Events at the expected day and time. For this rule, the Hour of Day was set to `9` and the Day of Week to `Monday` because both parameters have a high degree of seasonality for these events.
7. Marco selects **Event suppression** because he no longer wants these events to appear.
8. The Perform Action(s) option was set to `0` seconds because Marco wants suppression to occur immediately.

9. If the expected Power Supply test and Generator restart events do not occur, Marco wants to be notified. To set up this notification, he creates a Non-Occurrence event by creating an Event in the Actions When Event Does Not Occur... field (see Figure 6-13).

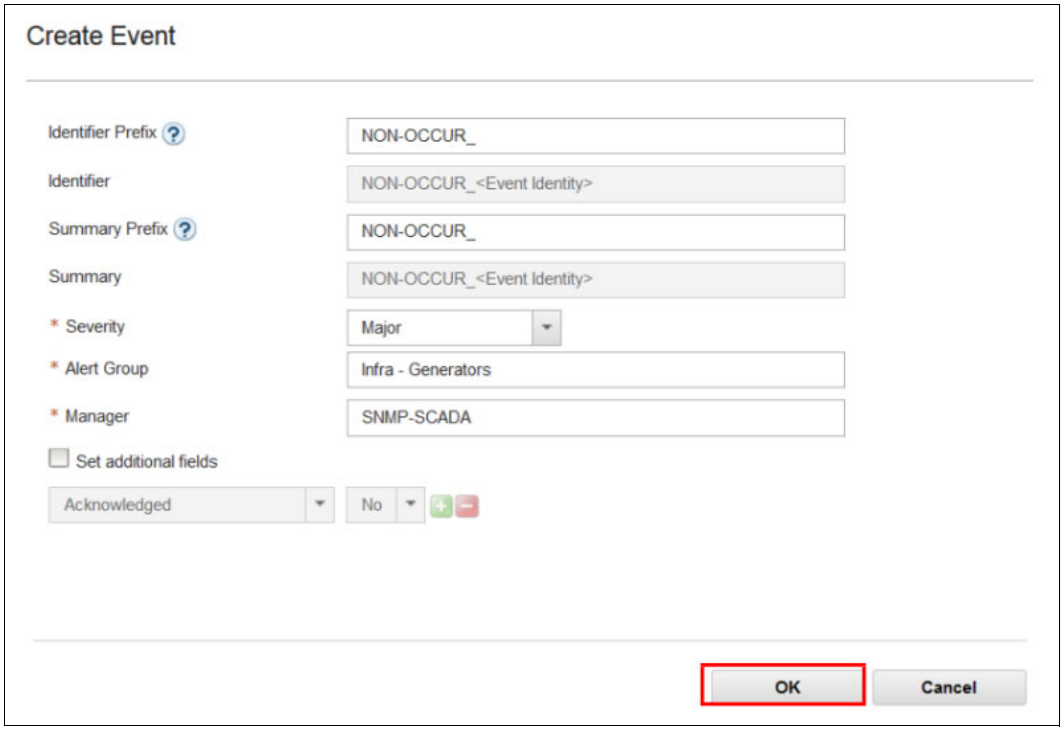

Figure 6-13 Non-Occurrence Event dialog

10. To ensure accurate event routing, the Alert Group and Manager are set to values that go to Marco's team.
11. Marco clicks **OK** in the Event Creation window to returns to the Seasonal Rule Creation window.
12. He clicks **Deploy** to immediately activate the new Seasonal Event Rule for processing against new incoming events.

By using only two windows with provided parameters, the Marco created a rule set that suppresses routine events that do not require action for resolution. He also created a synthetic event if the expected events do not occur.

Without Netcool Operations Insights, creating these rules requires complex policy coding in a separate interface.

6.5 Summary

In this chapter, seasonality and related event grouping in relation to Netcool Operations Insights was described. How to use it within your Netcool Operations Insights deployment to achieve greater value for your organization also was described.

Seasonality and Related Event analytics provides event management teams another tool set to help improve efficiencies in handling their event loads. Analytics can uncover previously unknown time and date patterns and relationships between events, which allows for refined events management.

The seamless integration of the analytics configuration and reporting pages into DASH reduces the time that analysts previously spent bouncing between applications. The ability to create, watch, and deploy seasonal event rules within the Netcool Operations Insights analytics pages speeds the pace of development and removes the need for custom coding to achieve the same event rules.

Part 4

Network event-related scenarios

In this part, network event-related scenarios are described.

7

Flood event detection

This scenario demonstrates how automations in IBM Netcool Operations Insight can detect and manage event floods. It shows how the event visualization tools can graphically depict the flood by using gauge, maps, and monitor boxes on dashboards.

Probes can be coded to discard events or send them to a backup ObjectServer if a flood occurs. You can see how network operators can drill into the events that are causing the flood by using the Active Event List, from which tools can be started to instruct network engineers to resolve the underlying problems in the infrastructure.

This chapter includes the following topics:

7.1 Scenario description

This scenario describes how the event flood automations, which are included with the ObjectServer, can be used to automatically detect an event flood and prevent operators in the Operations Center and the system from being overwhelmed with the flood of events. It also describes how customized dashboards can warn the Operations Center manager that a flood is occurring. The event search integration with Log Analytics can be used to retrospectively analyze an event flood and isolate the cause.

You can see how network operators can drill into the events that are causing the flood by using the Active Event List, from which tools can be started to instruct network engineers to resolve the underlying problems in the infrastructure. You can also see how an administrator can retrospectively review the events via the Log Analysis tool and, by categorizing the event data, determine the root cause and effect of the events.

In this scenario, we introduce a fictional organization Company C that uses IBM Netcool Operations Insight to manage a globally distributed network of data centers. A change that is made to the configuration of one of the devices causes it to malfunction. As a result, many monitored systems and services go offline. Jane, one of the network operators in Company C, is assigned to work on the problem.

7.1.1 Business value

A mis-configured network device causes a flood of events. This flood can cause a high load on the monitoring infrastructure and flood operators with events, which affects the Operations Center's ability to monitor and deal with other events occurring in the infrastructure.

The centralized flood detection and management functionality of Netcool/OMNIbus can reduce the effect of the flood on operators who are monitoring the event management system. It also can provide out-of-band accelerated notification of key events during the flood and provide the tools that enable operators to take decisive action to fix the problem.

7.2 Scenario topology

System components and default settings in the test environment are described in Chapter 1, "IBM Netcool Operations Insight overview" on page 3. The solution that is used in this scenario includes the following system components that are installed in the IBM test environment:

- Web GUI:

  `https://testserv1.ibm.com:16311/ibm/console/logon.jsp`

- LogAnalysis server

  `https://testserv2.ibm.com:9987/Unity/login.jsp`

- WebGUI/ObjectServer/Gateway/Probe

  `https://testserv1.ibm.com`

The following steps were used to achieve the scenario assumptions:

1. Set up a simnet probe to issue a background stream of an event at four events per second.

2. Set up the syslog probe to receive events from the test cell. The test cell includes misconfigured devices and generates a high number of events.
3. Switched on the syslog probe for a short period.
4. Used the Log Analysis user interface to find the spike in events and report on the categories of events that were observed.

7.3 Scenario steps

This section describes the process that is used to manage the event flood issue.

The Tivoli Netcool/OMNIbus includes a set of resources that you can use to extend the product to include event flood detection. It also warns of flood events and informs users of actions that can be taken, as required, to prevent abnormal behavior from affecting the entire Netcool Operations Insight. The customization is added to the probe rules file and target Object Servers.

The following process is used:

1. A device's configuration was changed, which causes the device to malfunction. As a result, many monitored systems and services suddenly go offline. This issue results in approximately 2,000 critical outage events being sent to the operator.
2. The event management infrastructure is subjected to an event flood. The Netcool Operations Insight event flood automations detect the event flood and then centrally manage it by using bidirectional probe communication to implement flood management policies in the probe.
3. As the flood starts, the event rate starts to climb. The flood is detected. In the meantime, Object Server informs the probe to reroute the lower severity events, which reduces the number of events that are seen by operators by rerouting low severity events.
4. On the dashboard, Jane can see that event gauge goes critical. As more events are received, the event rate gauge starts to climb and it goes close to the red (critical section), as shown in Figure 7-1.

Figure 7-1 NOI KPI dashboard

5. Jane reviews the Active Event List (AEL) dashboard. By using this dashboard, she can view event details and run context sensitive tools on events. She double-clicks the alert record to display details, as shown in Figure 7-2 on page 136.

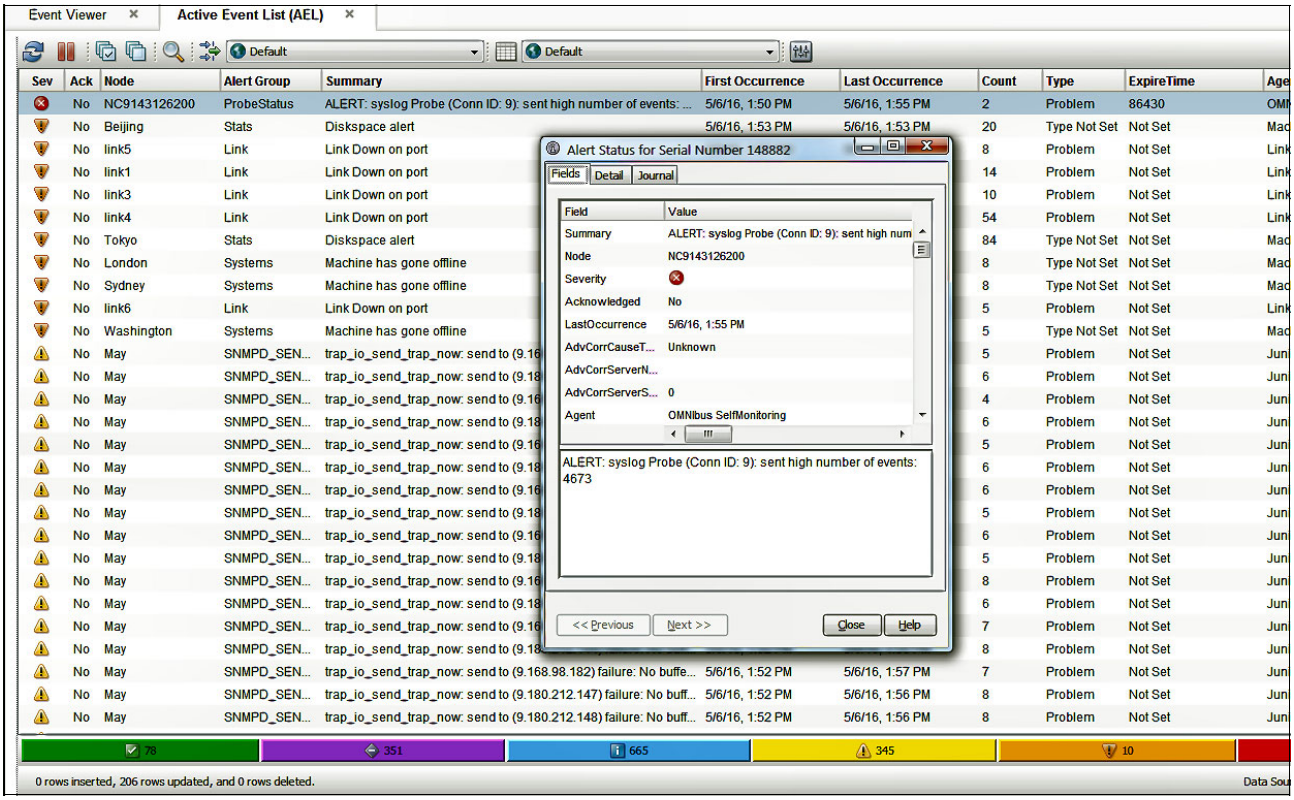

Figure 7-2 Active Event List

6. MTTraped Probe informs Jane that flood control ends and events are no longer rerouted.

   After approximately 10 minutes, the event rate gauge settles down at around 560 events per minute, as shown in Figure 7-3.

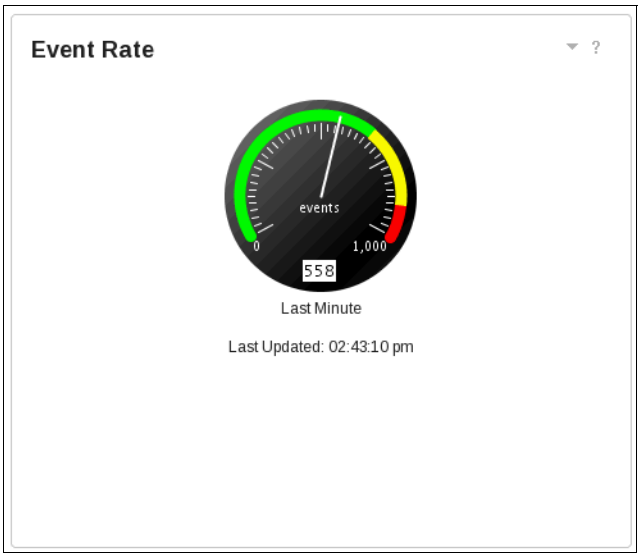

Figure 7-3 Last Minute event rate drops down

7. Jane switches to Event Dashboard. The network monitor window contains critical event alerts, as shown in Figure 7-4.

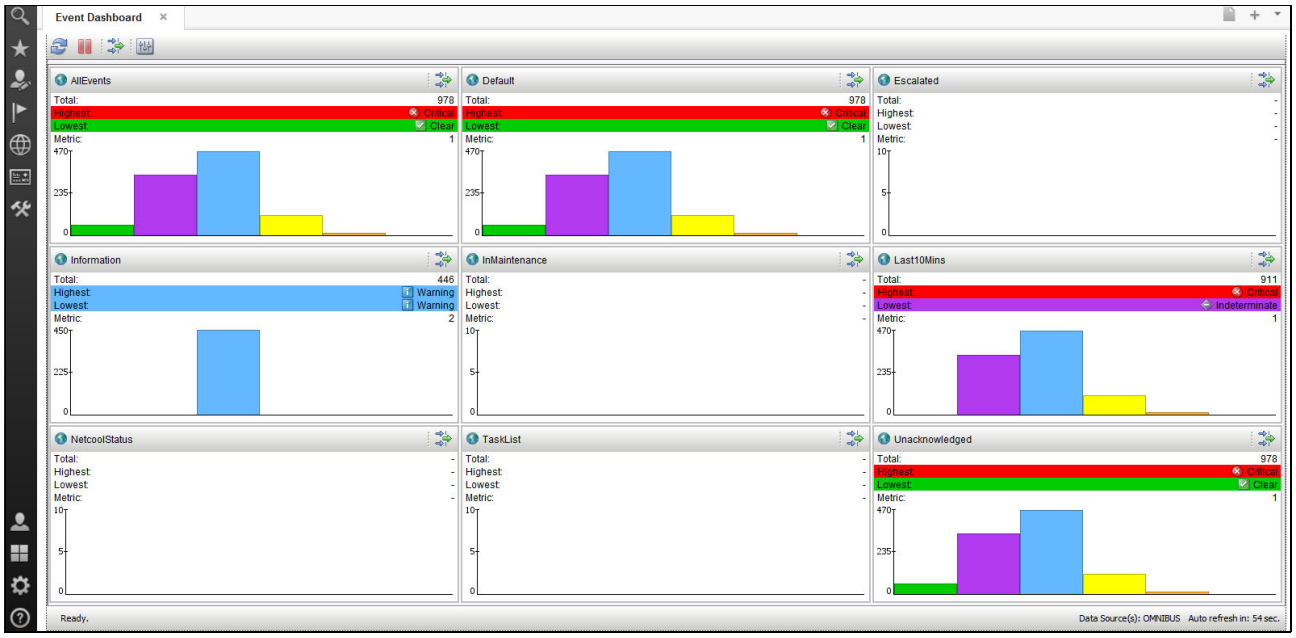

Figure 7-4 Event Dashboard

When IBM Operations Analytics - Log Analysis is integrated with IBM Tivoli Netcool/OMNIbus, it can be used the text analytics features to find patterns and trends in event data. With the integration of these two products, historical and real-time event data from IBM Tivoli Netcool/OMNIbus in the IBM Operations Analytics - Log Analysis user interface can be viewed searched.

IBM Operations Analytics - Log Analysis parses event data into a format suitable for searching and indexing. The event data is transferred from IBM Tivoli Netcool/OMNIbus to IBM Operations Analytics - Log Analysis by the IBM Tivoli Netcool/OMNIbus Message Bus Gateway.

8. Jane opens IBM Operations Analytics - Log Analysis console to see the event rate chart. She starts the Event Trend By Severity view by clicking **Search Dashboards** → **OMNIbusInsightPack** → **Event Analysis and Reduction** → **Event Trend By Severity** from the left top menu, as shown in Figure 7-5.

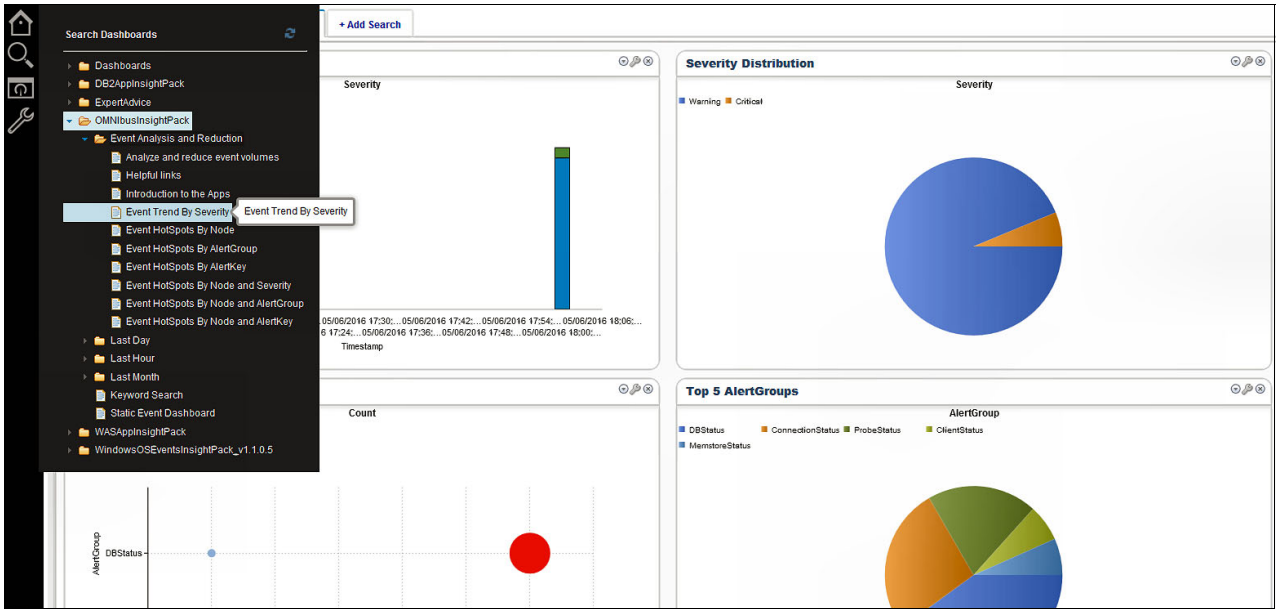

Figure 7-5 Event Trend by Severity menu

She sees the graphs that are similar to the graphs that are shown in Figure 7-6.

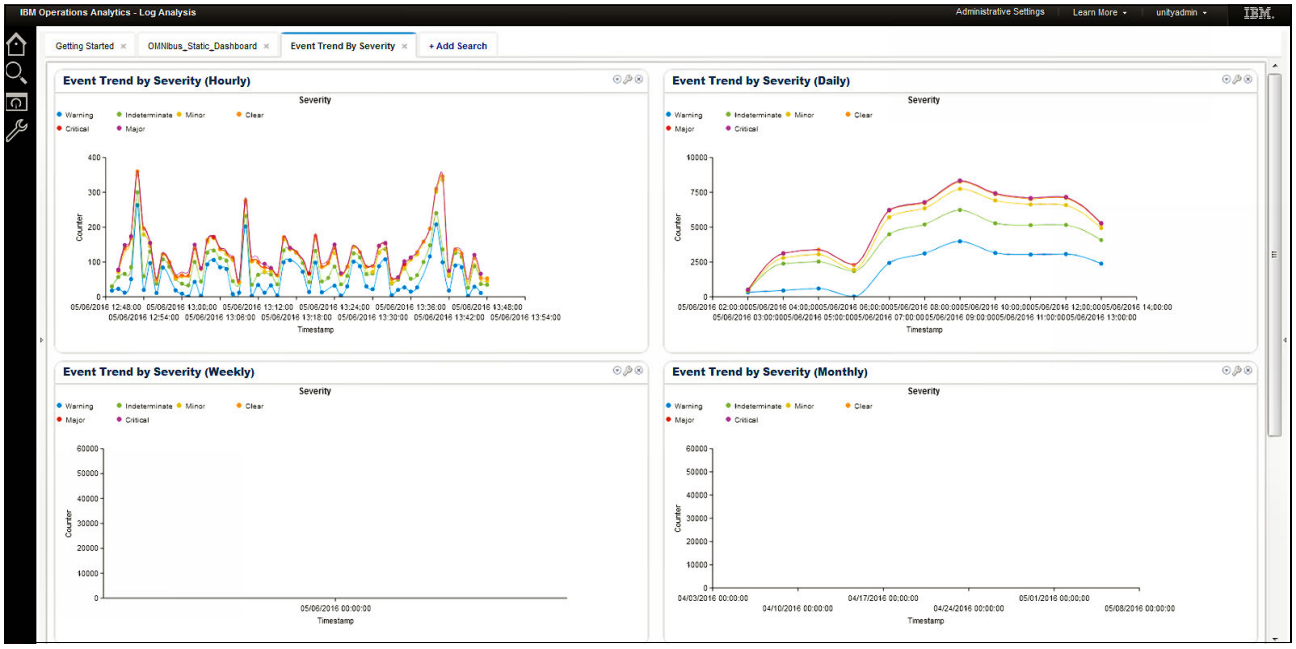

Figure 7-6 Event trends

9. To get into event details, Jane double-clicks a peak on the chart. The view changes. A “Discovered Patterns” menu is shown in the lower left part of the window.

10. Jane selects the node host name and drills down to see the events that contain that host name. By carrying out this search further, she sees the events that show where the problem is occurring. She clicks the **NOTPubType** tab and sees the devices or locations that are responsible for the flood, as shown in Figure 7-7.

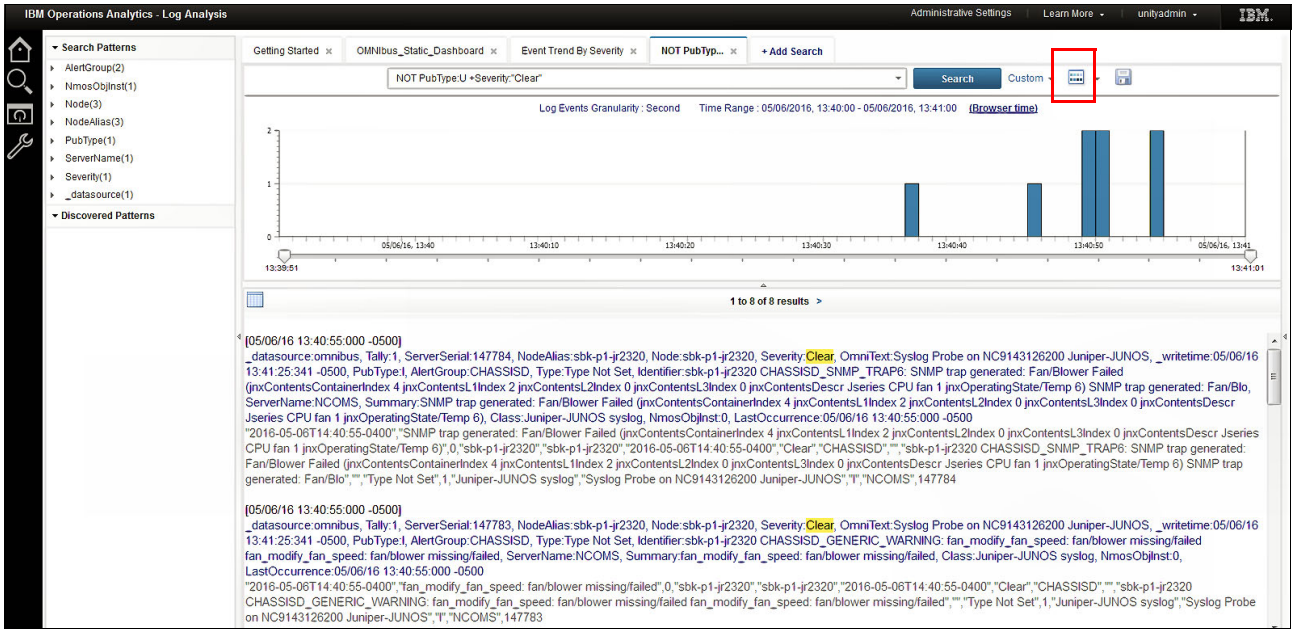

Figure 7-7 Event details

11. Jane clicks the small **grid** icon (as shown in the red box in Figure 7-7).
12. The window changes to a grid view. Jane clicks the **Node** column. She can see that the “May” node is problematic, as shown in Figure 7-8.
13. Jane clicks the **Chart** icon that can be see in the right upper corner (see the red box in Figure 7-8).

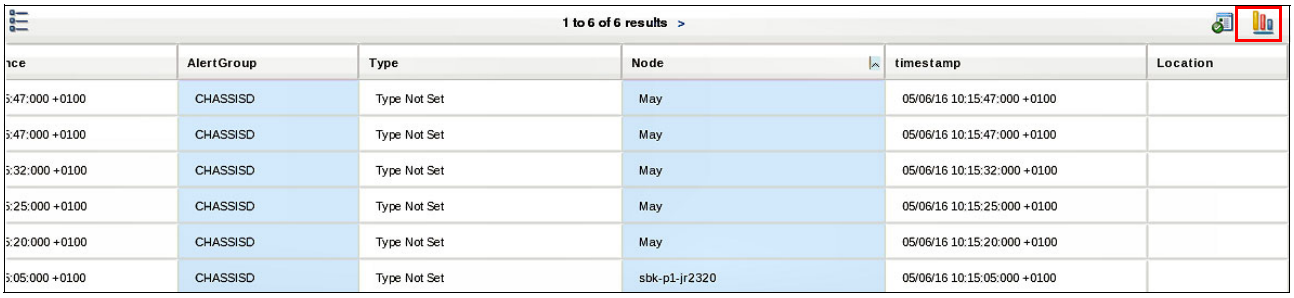

nce	AlertGroup	Type	Node	timestamp	Location
5:47:000 +0100	CHASSISD	Type Not Set	May	05/06/16 10:15:47:000 +0100	
5:47:000 +0100	CHASSISD	Type Not Set	May	05/06/16 10:15:47:000 +0100	
5:32:000 +0100	CHASSISD	Type Not Set	May	05/06/16 10:15:32:000 +0100	
5:25:000 +0100	CHASSISD	Type Not Set	May	05/06/16 10:15:25:000 +0100	
5:20:000 +0100	CHASSISD	Type Not Set	May	05/06/16 10:15:20:000 +0100	
5:05:000 +0100	CHASSISD	Type Not Set	sbk-p1-jr2320	05/06/16 10:15:05:000 +0100	

Figure 7-8 Alert group report

14. She makes sure that Generate Count option is selected and then chooses the **Plot Chart (All Data)** button, as shown in Figure 7-9 on page 140.

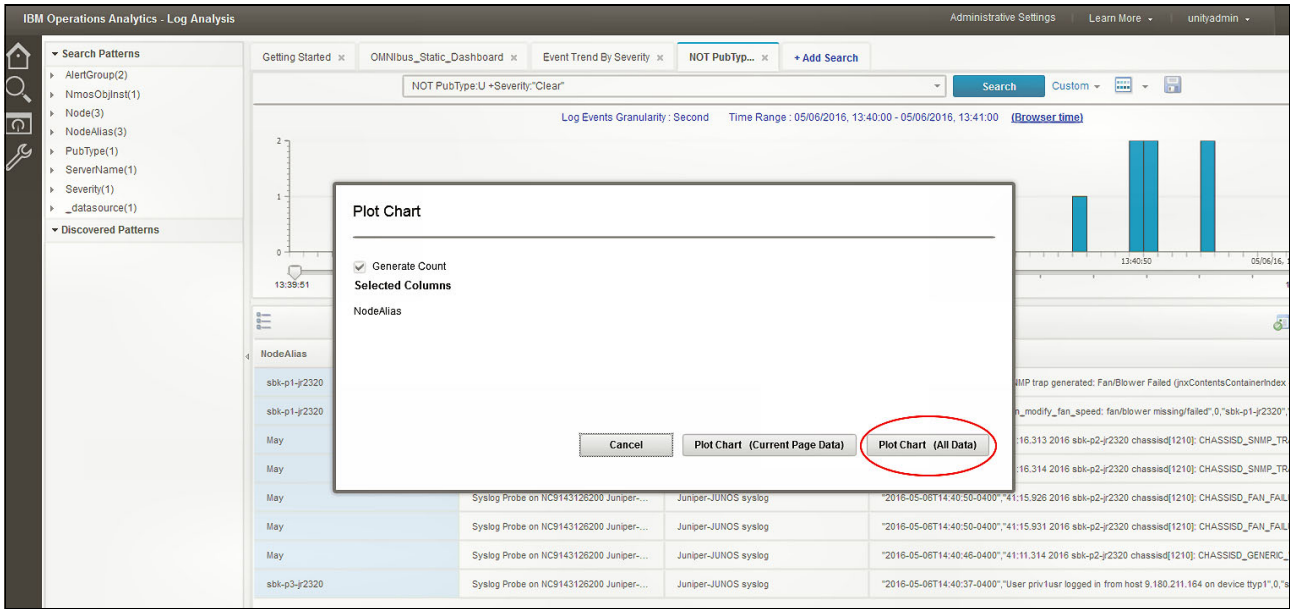

Figure 7-9 Generating plot with all data chart

15. She selects the **spanner** icon on the right side of the chart. She then clicks **Chart type** and sees a drop-down menu of the charts to create. Now, she can experiment with chart types to get meaningful data; for example, the bubble chart that is shown in Figure 7-10.

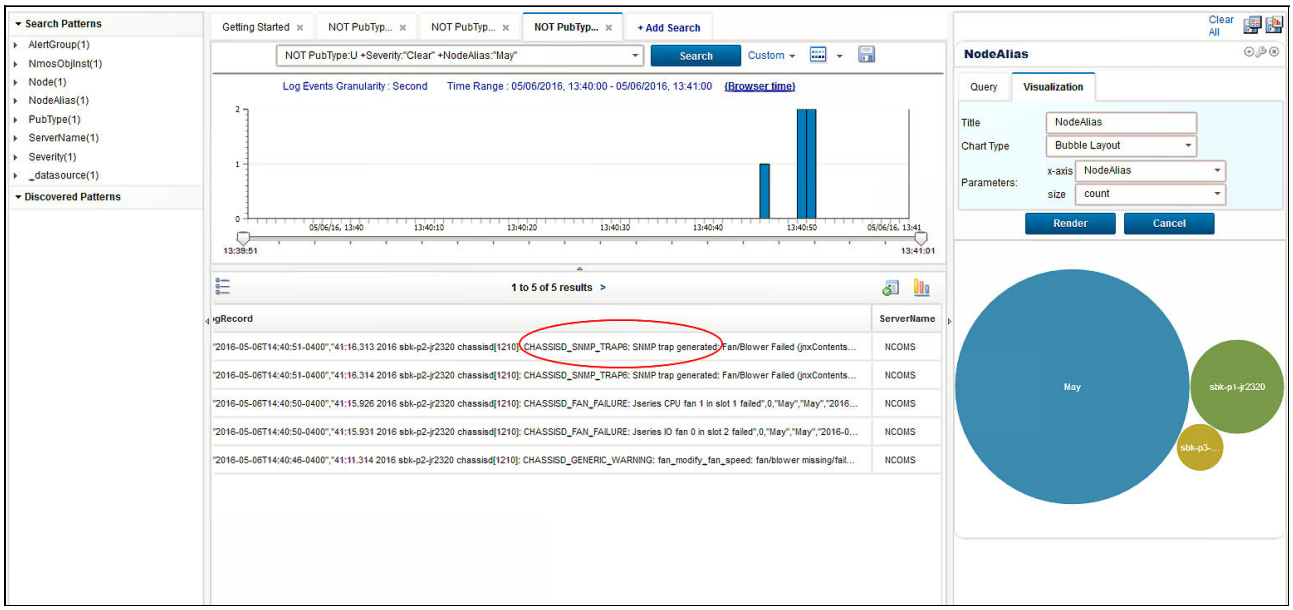

Figure 7-10 Bubble chart visualization

This chart shows that the event flood was caused by a CHASSIS warning on a node that is named “May”.

16. The new chart shows the relative quantity of each event severity for this node. Jane hovers over a data series in the chart. A tool-tip shows the actual event count for the corresponding severity.

17. It is clear to Jane that such a mis-configuration is a significant problem. She immediately dispatches an engineer to correct this problem. With the mis-configuration corrected, the monitored devices and services all auto-clear and the event flood is halted.

7.4 Using launch-in-context tools from Netcool Web GUI

The integration of IBM Tivoli Netcool/OMNIbus with IBM Operations Analytics - Log Analysis also provides right-click tools in Web GUI. These tools are for users to access from the Active Event List (AEL) or the Event Viewer within the Web GUI component of IBM Tivoli Netcool/OMNIbus.

The automated search that is implemented with the launch-in-context tools uses the FirstOccurrence time stamp in the event record as the basis of the search. FirstOccurrence is used because the tools are designed to find other events, not the event that is used as the basis for the search. The search criterion is designed to look for events with a time stamp that is less than the FirstOccurrence. This feature eliminates the possibility of finding the event that is used to start the search.

Complete the following steps to use the launch-in-context tools from the Netcool Web GUI:

1. Log in to Dashboard Application Services Hub as `ncouser`.
2. Click the flag icon and select **Event Viewer**, as shown in Figure 7-11.

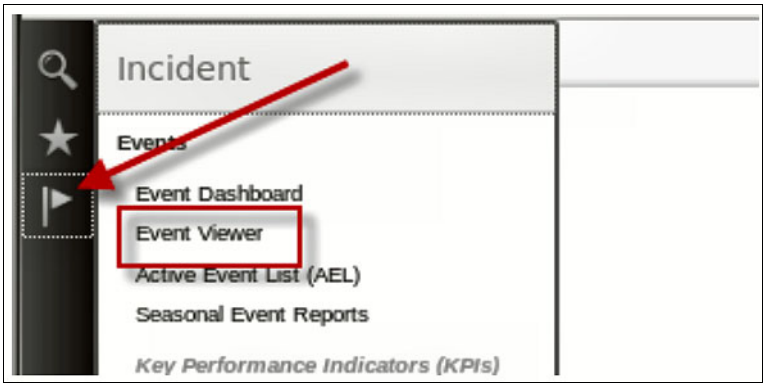

Figure 7-11 Selecting the Event Viewer option

When a problem with a device is investigated, one question that often comes to mind is whether other devices are experiencing the same or similar issues. The Search for similar events right-click tool is designed for this scenario.

3. Use the tool to find all devices with Critical Problems that are affecting the application. Locate a Critical event with APP1 in the AlertGroup field. Click the event to select it. Right-click and select **Event Search** → **Search for similar events** → 1 **day before event**, as shown in Figure 7-12 on page 142.

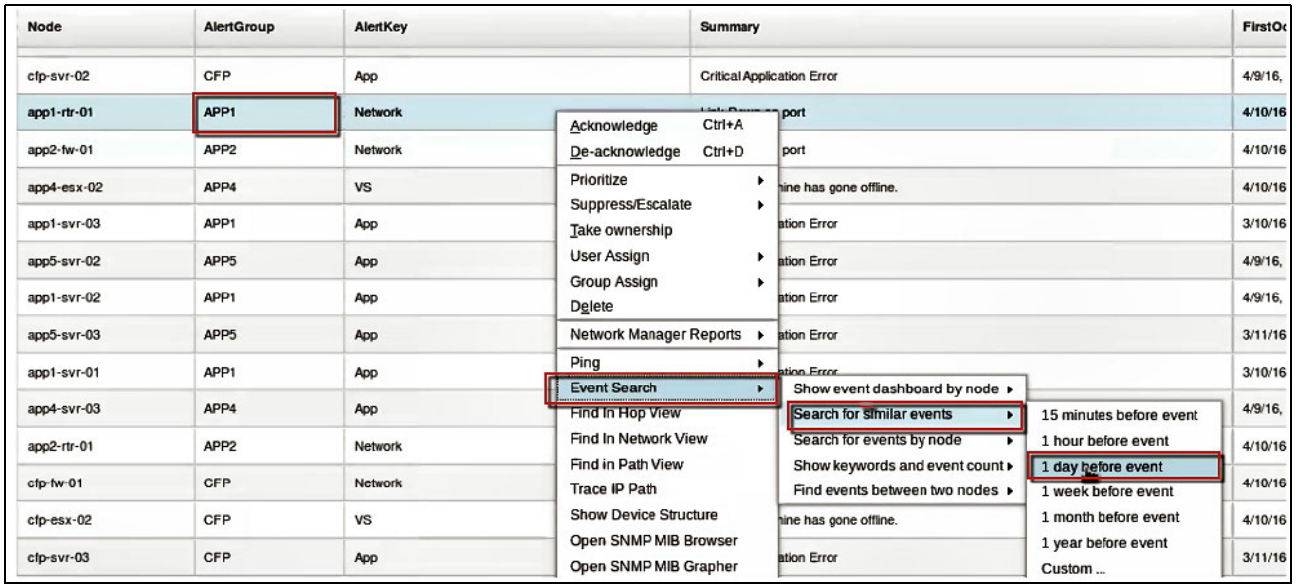

Figure 7-12 Search for similar events function

A new browser tab opens. You are logged in to Operations Analytics - Log Analysis and a search starts. After a short time, the results open in the window, as shown in Figure 7-13.

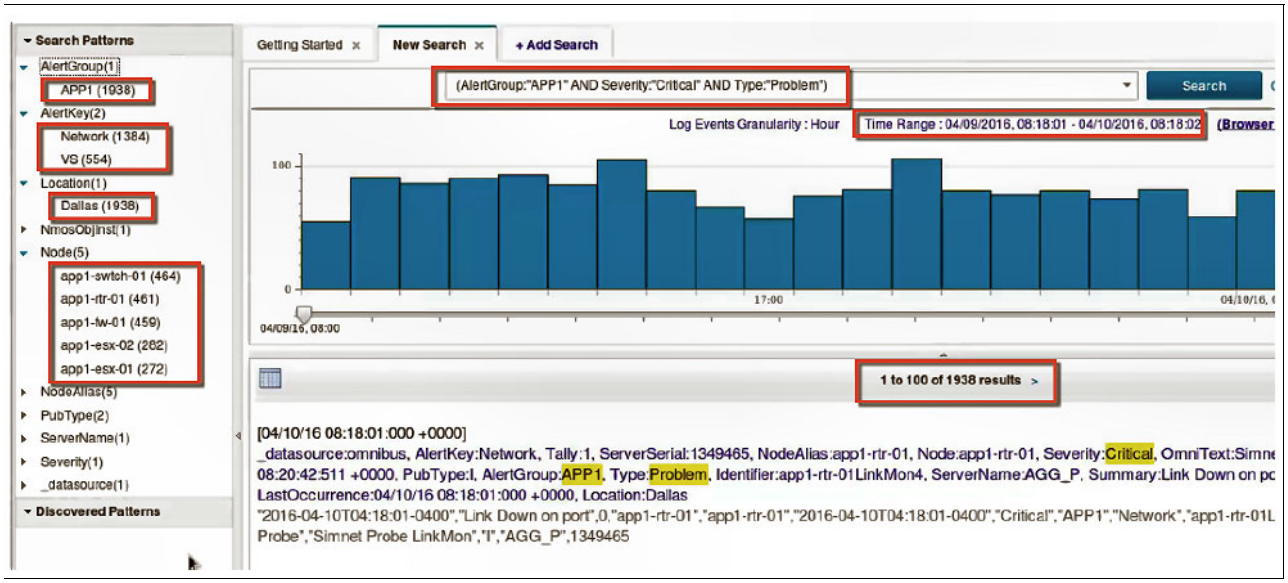

Figure 7-13 Results window

The following important points are shown in Figure 7-13:

- The search text is configured based on values for the AlertGroup, Type, and Severity event columns. The values are extracted from the event record.
- The time range is defined based on the value of FirstOccurence that is extracted from the event record.
- From the text analytics, we can see that application experienced Network and Virtual Server problems on five nodes.

7.5 Summary

There are instances when a failure can lead to a flood of events. An example of this issue can be an air conditioner failing in a specific data center, which can lead to many events from systems in this data center that can fail because of overheating.

The centralized flood detection and management functionality of Tivoli Netcool/OMNIbus can reduce the effect of the flood on operators who are monitoring the event management system. This example demonstrates how an operator, who often might take hours to manually review many hundreds of events, quickly found the root cause of the major outage within a few mouse clicks.

The ability to perform in-context keyword searches of the entire event history within a specified time window by using the Log Analysis tool allows an operator to distill down, summarize, and make sense of vast quantities of event data.

Notes: The Tivoli Netcool/OMNIbus includes a set of resources that you can use to extend the product to include event flood detection. For more information, see the following articles in the IBM Knowledge Center:

- Detecting event floods and anomalous event rates:

  https://ibm.biz/BdrEFm
- Protecting the ObjectServer against event floods:

  https://ibm.biz/BdrEFG
- Extending the functionality of Tivoli Netcool/OMNIbus:

  https://ibm.biz/BdrEFe

8

Using the WebGUI event search feature

This demonstration describes the procedure that is used to search for the historical reoccurrences of an event, view the data in-line in the EventViewer, and then, start the Log Analysis user interface. The tasks in this procedure use the sample data.

This chapter includes the following topics:

8.1 Scenario description

This scenario shows how to use the WebGUI event search feature in IBM Netcool Operations Insight V1.4.0.1.

A monitoring team working in Company C receives a critical CPU warning alarm that occurs every Sunday when the backup jobs are running. This warning creates an auto-ticket, which is assigned to an L1 operator. By the time the operator receives the ticket and checks the CPU usage on the target server, the CPU level returns back to normal levels, so the operator closes the ticket.

Next week on Sunday, the same issue occurs again, but there is a different operator on the duty who receives the ticket. This cycle goes on for weeks because the operations staff is large, so a different person receives the ticket each week it happens. Because of the volume of tickets that the operations team deals with, no one identifies the pattern. This issue, in turn, incurs significant and wasteful cost to the business.

Joe, the Chief of IT Monitoring Team, wants to get more information about the event before raising a next ticket in case it is indicative of a systematic problem that needs to be addressed. He can use the integration with Log Analysis in the Event Viewer feature in WebGUI to check event occurrences over the previous month.

8.1.1 Business value

By resolving chronic issues, such as the one described in this scenario, the monetary savings to the business can be calculated by identifying and rectifying issues that were causing seasonal events, which decreases the number of costly trouble tickets.

Analytics-based event grouping assists Company C in the following ways:

- Repeating patterns of events are easily displayed and delivered to the appropriate Subject Matter Experts (SMEs).
- Previously unknown relationships can be discovered.
- Event suppression or automated remediation can be applied before human interaction is required, which reduces the overall visible volume of events.
- Exception analysis can be applied so if expected events do not occur, an alarm can be raised.

8.2 Scenario topology

System components and default settings in the test environment are described in Chapter 1, “IBM Netcool Operations Insight overview” on page 3. The solution that is used in this particular scenario includes the following system components that are installed on the IBM test environment:

- WebGUI:

  `https://testserv1.ibm.com:16311/ibm/console/logon.jsp`

- LogAnalysis server

  `https://testserv2.ibm.com:9987/Unity/login.jsp`

- WebGUI/ObjectServer/Gateway/Probe

  https://testserv1.ibm.com

8.3 Scenario steps

Joe performs the following steps to recreate and eventually solve the re-occurring CPU event issue:

1. He logs in to the Integrated Portal console by using the following link, as shown in Figure 8-1:

   https://testserv1.ibm.com:16311/ibm/console/logon.jsp

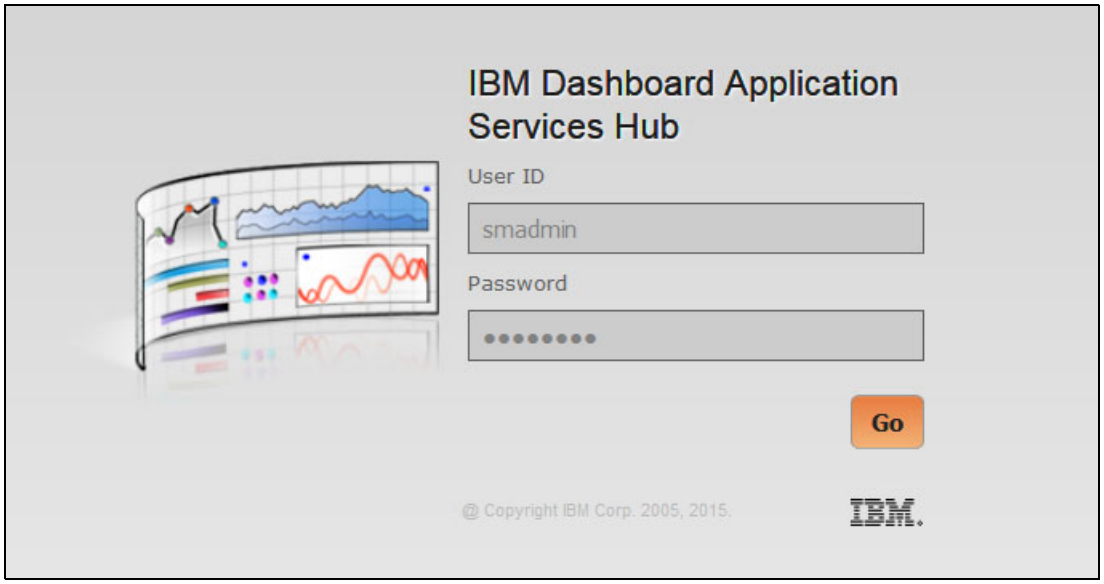

Figure 8-1 IBM Dashboard Application Services Hub

2. He selects the **Event Viewer** from the menu on the left side, as shown in Figure 8-2 on page 148.

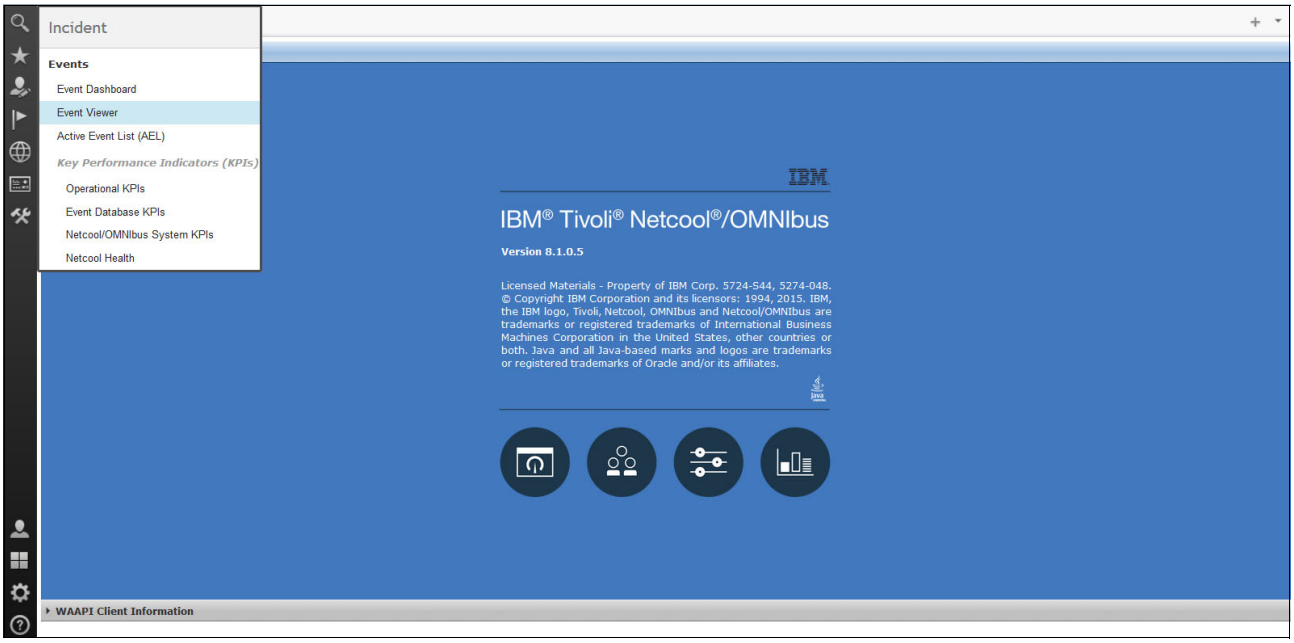

Figure 8-2 Opening the Event Viewer

Joe can see that CPU is used in 90%, as shown in Figure 8-3. He now wants to find more information about this event.

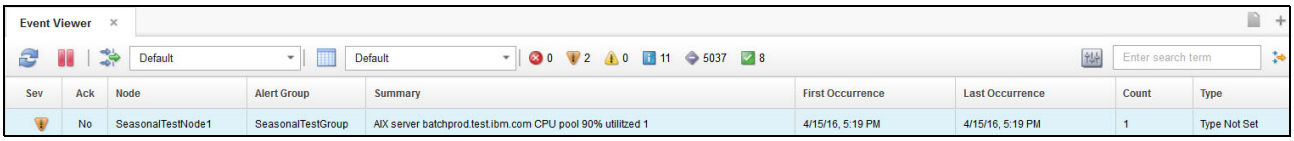

Figure 8-3 Event occurred in demonstration environment

3. To display more information, he double-clicks the alert. A new window opens and he clicks the **Event Search** tab to see whether this event occurred before, as shown in Figure 8-4. He then clicks **Search**.

> **Note:** Before reaching the generated report, the system might prompt for a login to IBM Operations Analytics - Log Analysis console.

Figure 8-4 Properties for selected event

Event search applies the search and analysis capabilities of Operations Analytics - Log Analysis to events that are monitored and managed by Tivoli Netcool/OMNIbus. Events are transferred from the ObjectServer through the Gateway for Message Bus to Operations Analytics - Log Analysis, where they are imported into a data source and indexed for searching. After the events are indexed, every occurrence of real-time and historical events can be searched.

The default search of last week shows that the event occurred six times on a Sunday, but no other day, as shown in Figure 8-5.

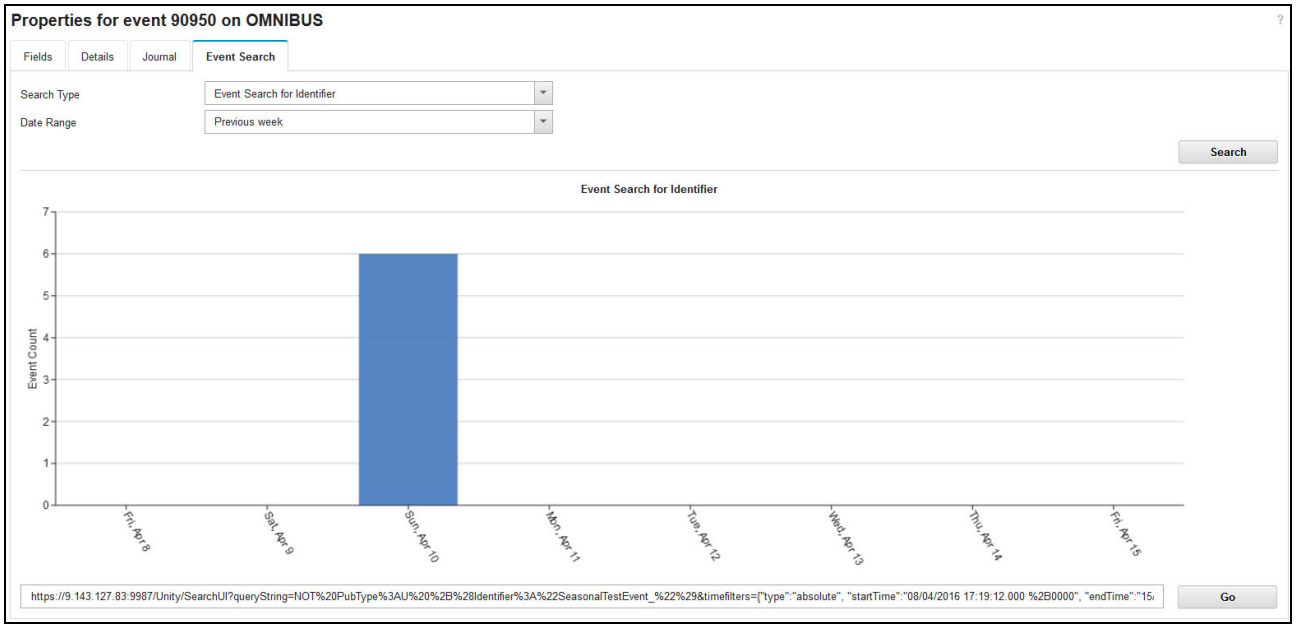

Figure 8-5 Event occurrence in previous week view

4. Because these results appear abnormal, Joe want to widen the search to the last month. He chooses the **Date range: Previous month** option and then, clicks **Search**.

   The previous month view result is shown in Figure 8-6.

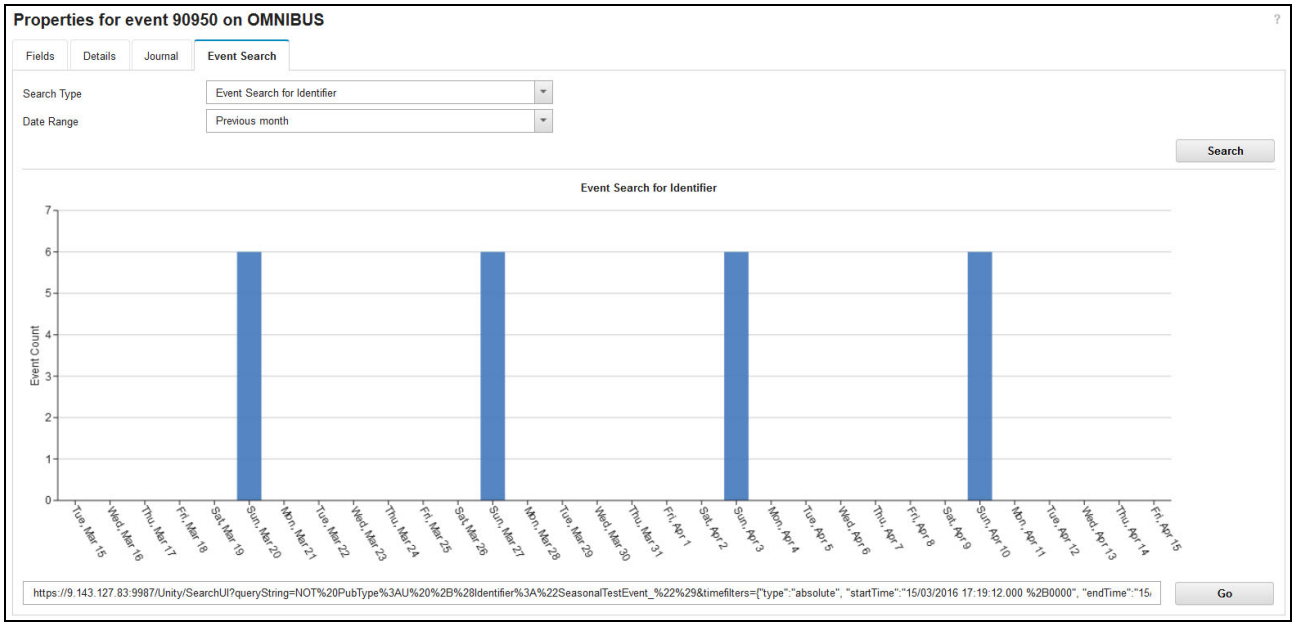

Figure 8-6 Event occurrence in previous week month

5. Current analysis shows that there is a pattern to the events, which occurs on a Sunday only. Joe wants to check whether some of these events were present last year. He chooses the **Date range: Previous year** option and clicks **Search**. The result is shown in Figure 8-7 on page 150.

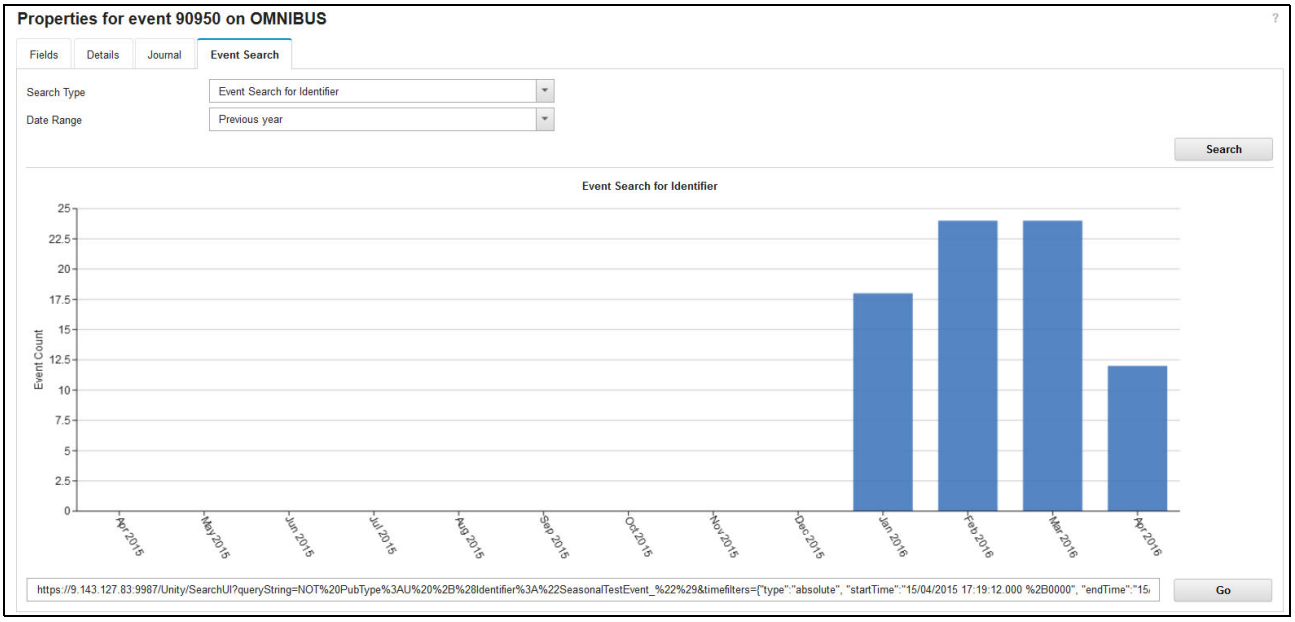

Figure 8-7 Previous year view

Joe determines that the event always occurs on a Sunday and was occurring approximately 15 times a month over the last couple of months.

6. He adds a description to the Event Journal so that the information is available for the person who is resolving the ticket. The description is added by clicking the **AddToJournal** option from the right-click menu or pressing Shift+J, as shown in Figure 8-8.

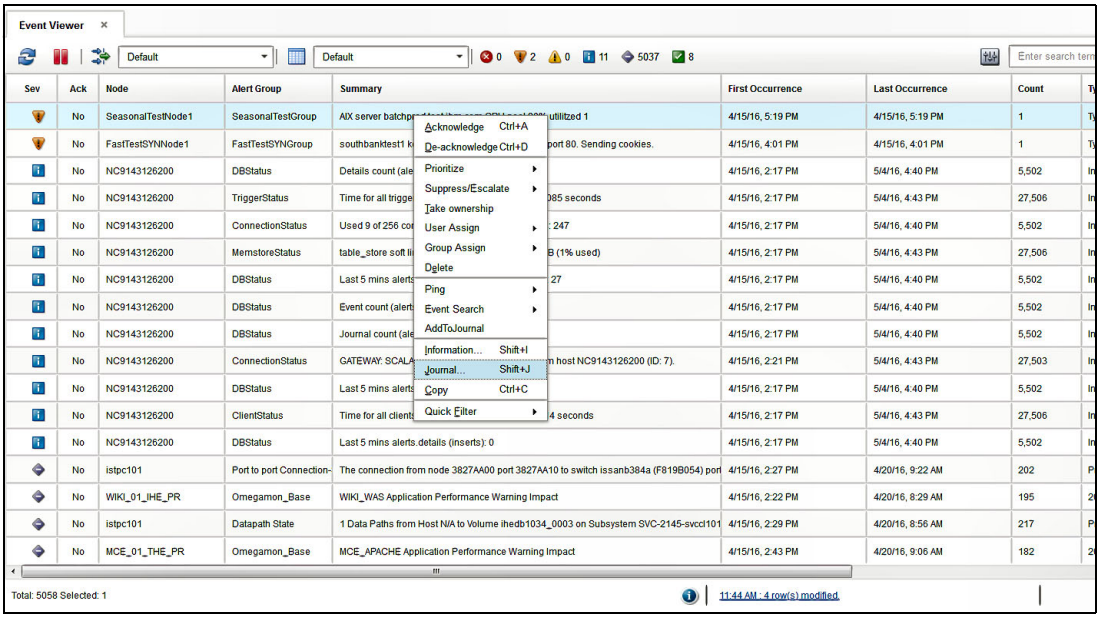

Figure 8-8 AddToJournal menu option

7. He enters a Journal Entry description manually and pastes the URL that is seen at the bottom of the search results (the URL next to the 'GO' button) that are displayed in the event properties, as shown in Figure 8-9 on page 151.

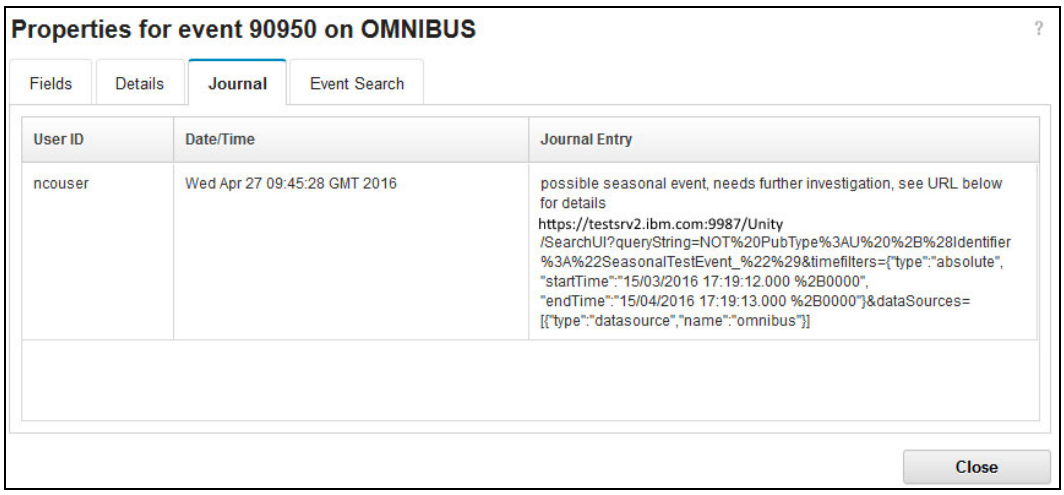

Figure 8-9 Adding entry to Journal

If information is added to the Event Journal, it is available for the person who is resolving the ticket. A linked URL allows the SME to reference the historical information and start the full Log Analysis user interface when they start to work on resolving the ticket.

8. Joe copies the generated URL by right-clicking the URL and then pressing Ctrl+C. He can paste the URL into web browser and start the Event Viewer directly in IBM Operations Analytics - Log Analysis console, as shown in Figure 8-10.

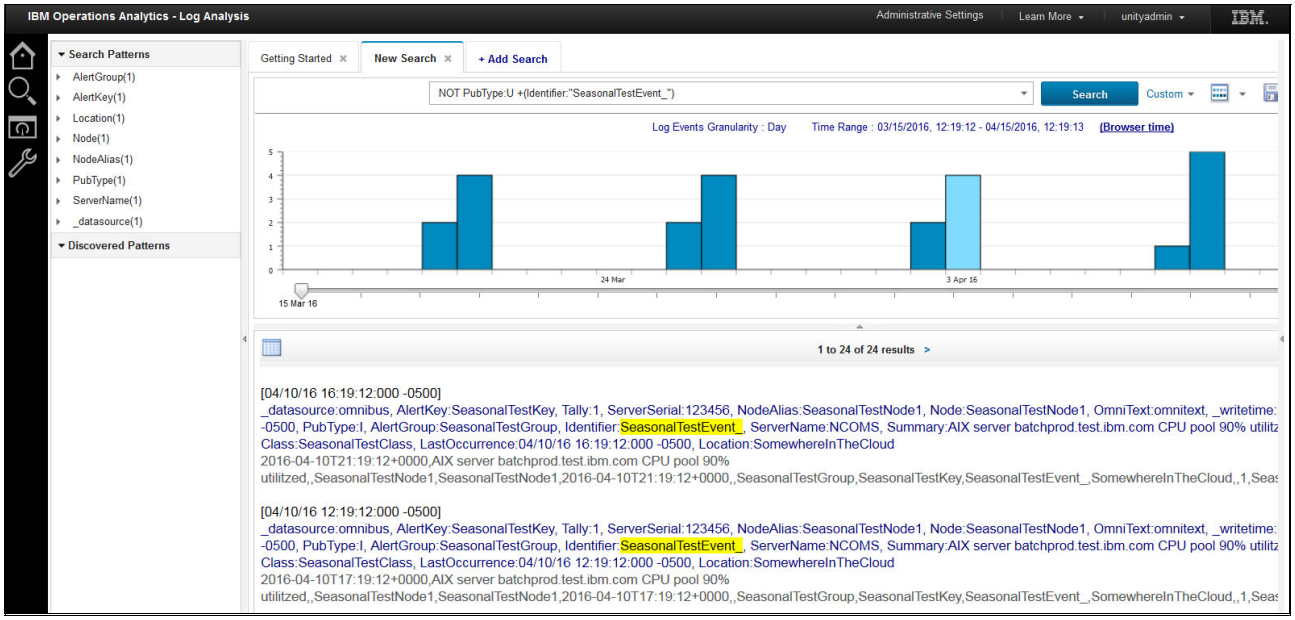

Figure 8-10 IBM Analytics console

In this scenario, it is determined that the event is caused by a batch process that uses a high amount of CPU. Further investigation shows that there was a script that runs some backup jobs on Sunday on too many machines. The solution was to change backup script and suppress the event.

8.4 Summary

IBM Netcool Operations Insight provides updated innovative features to provide deeper insight, event reduction, and automation capabilities to help drive further efficiencies for more agile and leaner operations.

It allows unique visibility of event relationships and recurring patterns by using machine learning, which offers the administrator a single-click option to implement a rule for automated correlation of future occurrences without coding.

In this scenario, the operator identified that the event occurred regularly and supplied the relevant information in the service ticket. This information helps the SME to deal with the systematic cause of the event.

More information: For more information, see the following resources:

- For more information about configuring event search, see this website:

  https://ibm.biz/BdrRnt

- For more information about configuring the Event Viewer to connect to IBM Operations Analytics - Log Analysis, see this website:

  https://ibm.biz/BdrRn6

9

Scope-based event grouping

This chapter describes a scenario for scope-based event grouping and includes the following topics:

9.1 Introduction

Scope-based event grouping is based on the premise that if you have a group of events that occur at the same place at the same time, it is likely that they are related to the same problem. In this context, scope is another way of referring to *same place*. In practice, this method proves to be effective for grouping events.

Grouping events includes the following goals:

- Bring order to the event list and logically grouping events by incident
- Provide a mechanism to create only one incident ticket per incident
- Keep all the related event information together to aid problem diagnosis
- Reduce mean time to repair (MTTR) and operations costs

Although the use of a scope that is based on geographic location is the natural choice for scope-based event grouping, the ScopeID field is a string and can be set to anything that makes sense in the context of the grouping scenario. Another way to think of the scope is the field or reach of influence.

The idea is that your scope is wide enough to include all the events that might be related to a problem without making it so large that the automation incorrectly groups too many events together. It is better to stray on the conservative side of not grouping too many events together rather than too many.

Scope-based event grouping seeks to group events that happen at the same place at the same time. The *same time* is defined in terms of a minimum period that needs to pass without further new events occurring before it is deemed that the incident finished. This period of "quiet time" is referred to as the QuietPeriod.

The term "new events" applies to the occurrence of new, unique events only. Recurrences of the same events (known as deduplication in Netcool terms) are not applicable in terms of resetting the QuietPeriod.

One example of a real-world scenario is a telecommunications company that defined the scope to be the cell site identification code, which is encoded into every event. Implementing this scope was simple and involved adding one line of code to the Probe rules to set the ScopeID to match the cell site identification code. Any events that came from that site at the same time were automatically grouped. This scope was convenient because the company has multiple different equipment vendors' events on each site and building management events.

The company reduced the number of events that was presented to operators by 77% and had an average of 12 events per grouping. In this scenario, a QuietPeriod of 10 minutes (600 seconds) was found to be optimal.

Another example of a real-world scenario is a large bank that defined the scope to be the line-of-business identifier. Within the bank, there are many lines-of-businesses that represent the "customers" of the bank's ITSM solution. Each line-of-business owns several servers that run business critical applications. IBM Tivoli Monitoring is used extensively on these servers and monitor everything from the applications to the hardware on which the applications are running.

This configuration generates a great deal of ITM events into Netcool, which makes it challenging for operations to manage. The bank discovered that, by setting the ScopeID to the line-of-business ID, they reduced the number of events that are presented to operators by more than 99% and had an average of 210 events per grouping. In this scenario, a QuietPeriod of 10 minutes (600 seconds) was found to be optimal.

The only settings that are required to start scope-based event grouping is the setting of an appropriate ScopeID and QuietPeriod if the default value is not appropriate. An optional extension to the grouping is to activate the automatic probable cause and affect determination via the weightings.

Note: The weighting function is optional and the grouping is not dependent on setting up this function.

9.2 Scenario description

Company A is a telecommunications company with an extensive wireless network. They use several different vendors' equipment across their widely distributed cell sites. The equipment they use creates a high volume of network traffic and Operations often must deal with many events. This challenge is exacerbated whenever there is a major outage.

Helen runs the tooling department for the Netcool based ITSM solution at Company A. She identified the following list of challenges that are facing Operations:

- There are too many events for operators to manage.
- There often are many events present that relate to each incident.
- Events for multiple incidents are all mixed, which makes it difficult to manage.
- Multiple tickets often are opened from the many events.
- The information from these events often is fragmented across tickets.
- It is costly to the business to close all of the duplicate tickets.

The net result of these challenges is that the MTTR is higher than it should be. The metaphorical "house" is messy and needs organizing. It is here that Helen believes that event grouping can help.

9.2.1 Business value

The scope-based event grouping feature helps Company A in the following ways:

- Automatically group events by incident, based on location and time
- Combine the selected event details into a single place for ticketing
- Allow the cutting of a single ticket per incident, which hopefully eliminates duplicates

By implementing these features, Helen predicts that she can considerably reduce the number of trouble tickets that are opened, and keep the related event information that pertains to an incident in one place. She believes that these changes can help operators pinpoint and resolve problems faster.

9.3 Scenario topology

For this scenario, we used the environment that is described in 1.4, "Our environment for the scenarios" on page 18 of this book.

9.4 Scenario steps

The following sections describe the step-by-step implementation of this scenario.

9.4.1 Analyzing the current event set

The operators review traditional filtered lists of events that include field-based sorting applied. Although they often work on higher severity events first, they are mindful of ensuring that older events are dealt within a timely fashion.

The operations team frequently struggle to stay on top of the high volumes of events that come through and invariably end up cutting multiple duplicate trouble tickets, particularly whenever there is an event storm. Figure 9-1 shows the current environment.

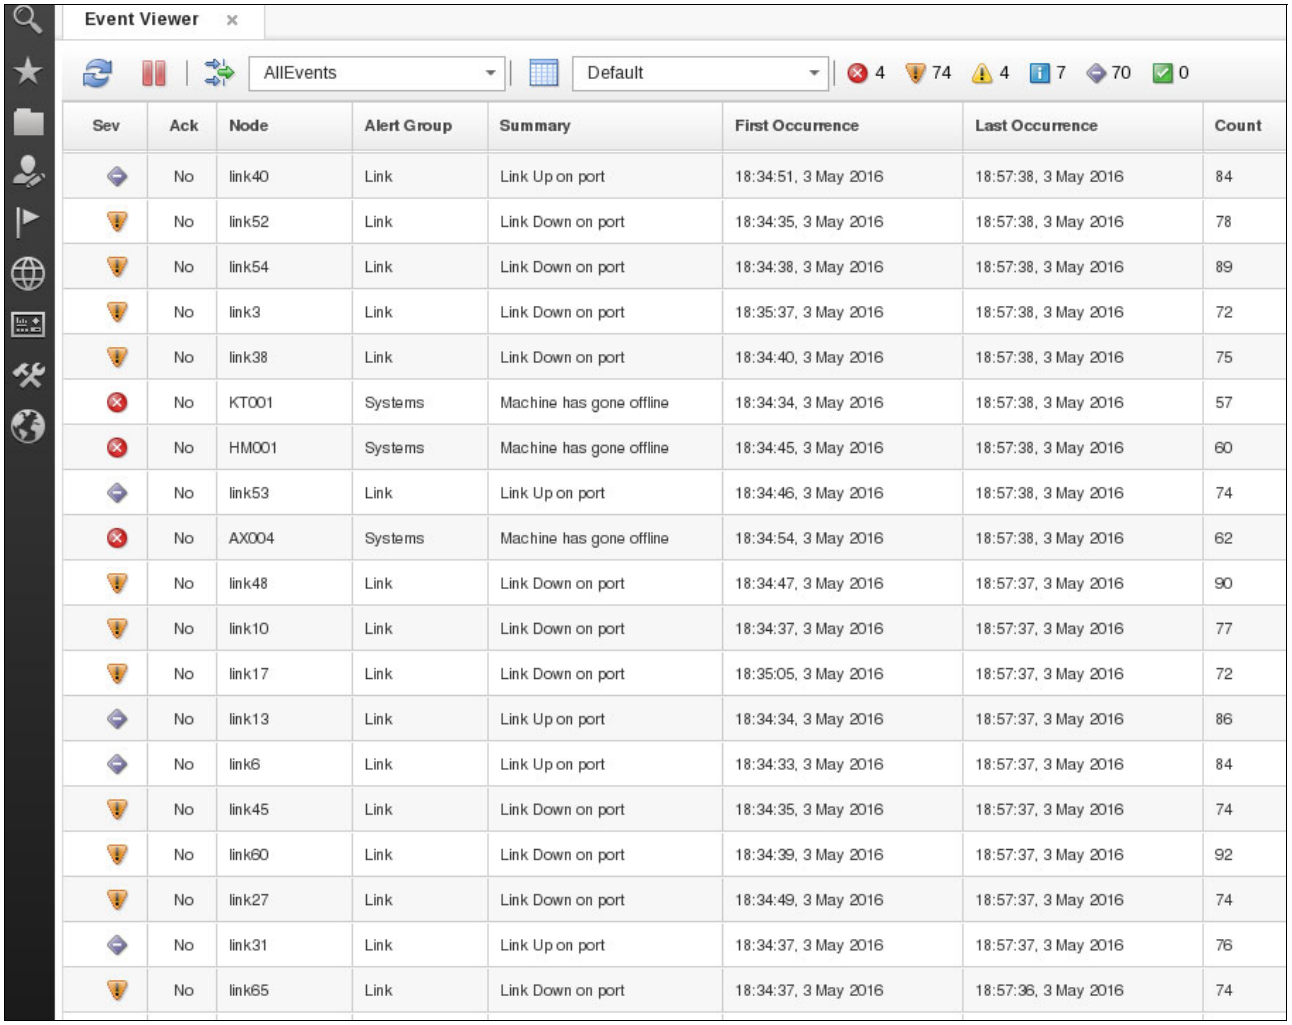

Sev	Ack	Node	Alert Group	Summary	First Occurrence	Last Occurrence	Count
	No	link40	Link	Link Up on port	18:34:51, 3 May 2016	18:57:38, 3 May 2016	84
	No	link52	Link	Link Down on port	18:34:35, 3 May 2016	18:57:38, 3 May 2016	78
	No	link54	Link	Link Down on port	18:34:38, 3 May 2016	18:57:38, 3 May 2016	89
	No	link3	Link	Link Down on port	18:35:37, 3 May 2016	18:57:38, 3 May 2016	72
	No	link38	Link	Link Down on port	18:34:40, 3 May 2016	18:57:38, 3 May 2016	75
	No	KT001	Systems	Machine has gone offline	18:34:34, 3 May 2016	18:57:38, 3 May 2016	57
	No	HM001	Systems	Machine has gone offline	18:34:45, 3 May 2016	18:57:38, 3 May 2016	60
	No	link53	Link	Link Up on port	18:34:46, 3 May 2016	18:57:38, 3 May 2016	74
	No	AX004	Systems	Machine has gone offline	18:34:54, 3 May 2016	18:57:38, 3 May 2016	62
	No	link48	Link	Link Down on port	18:34:47, 3 May 2016	18:57:37, 3 May 2016	90
	No	link10	Link	Link Down on port	18:34:37, 3 May 2016	18:57:37, 3 May 2016	77
	No	link17	Link	Link Down on port	18:35:05, 3 May 2016	18:57:37, 3 May 2016	72
	No	link13	Link	Link Up on port	18:34:34, 3 May 2016	18:57:37, 3 May 2016	86
	No	link6	Link	Link Up on port	18:34:33, 3 May 2016	18:57:37, 3 May 2016	84
	No	link45	Link	Link Down on port	18:34:35, 3 May 2016	18:57:37, 3 May 2016	74
	No	link60	Link	Link Down on port	18:34:39, 3 May 2016	18:57:37, 3 May 2016	92
	No	link27	Link	Link Down on port	18:34:49, 3 May 2016	18:57:37, 3 May 2016	74
	No	link31	Link	Link Up on port	18:34:37, 3 May 2016	18:57:37, 3 May 2016	76
	No	link65	Link	Link Down on port	18:34:37, 3 May 2016	18:57:36, 3 May 2016	74

Figure 9-1 Initial event listing

Scope-based event grouping works on the basis that events are grouped that occur at the same place at the same time. Helen identifies that the event stream includes location information that is encoded within the event stream that might be used to define the scope.

Tip: In many cases, the location information is not present in the event stream. If it is available in a database (for example, an asset database), the location or scope can be enriched into the event stream by using Netcool/Impact.

9.4.2 Configuring the system

Helen completed the following steps to configure the development system to test out how her plan to deploy scope-based event grouping might look:

1. Helen imports the scope-based event grouping automation into the Net-cool/OMNIbus ObjectServer. She opens a command line session to the ObjectServer and imports the automation functions, as shown in the following example:

   ```
   $OMNIHOME/bin/nco_sql -server AGG_P -user root -password abc123 \
   < $OMNIHOME/extensions/eventgrouping/objectserver/ \
   scope_event_grouping_aggregation.sql
   ```

   Helen repeats the process on the backup ObjectServer and makes some additions to the failover bidirectional Aggregation Gateway so that the automation control in-formation is replicated between the primary and backup Netcool/OMNIbus systems. For more information about this process, see this website:

   http://ibm.biz/seg_install

2. Helen modifies the Probe rules to set up the scope. Because the event stream contains the information Helen needs to use for the scope-based grouping, she makes a small addition to the Probes rules file to set up the ScopeID and, where possible, subgrouping.

   The primary location is stored in a data token that is known as *location* with some events that also contain secondary location information in the data token that is called *suburb*. Helen adds the following lines to the Probe rules file to assign the primary location into the ScopeID field and the suburb location into the SiteName field:

   ```
   @ScopeID = $location
   @SiteName = $suburb
   ```

   Tip: Subgrouping is done automatically if the `$suburb` token is populated with a non-null value. If it is null, subgrouping is not done. Subgrouping can be forced in any case, however, by setting a `SEGNoSiteNameParentIfSiteNameB-lank` property to 0. With this setting, a subgrouping is created to subgroup events that do not have a subgrouping value defined.

3. Helen edits the Event Viewer view so that it includes IBM Related Events in the view. This change allows the relationships to be rendered in the Event Viewer.

   If Netcool Operations Insight is not installed, IBM Related Events relationship might not be installed. In this case, you can create a relationship, as shown in Figure 9-2 on page 158.

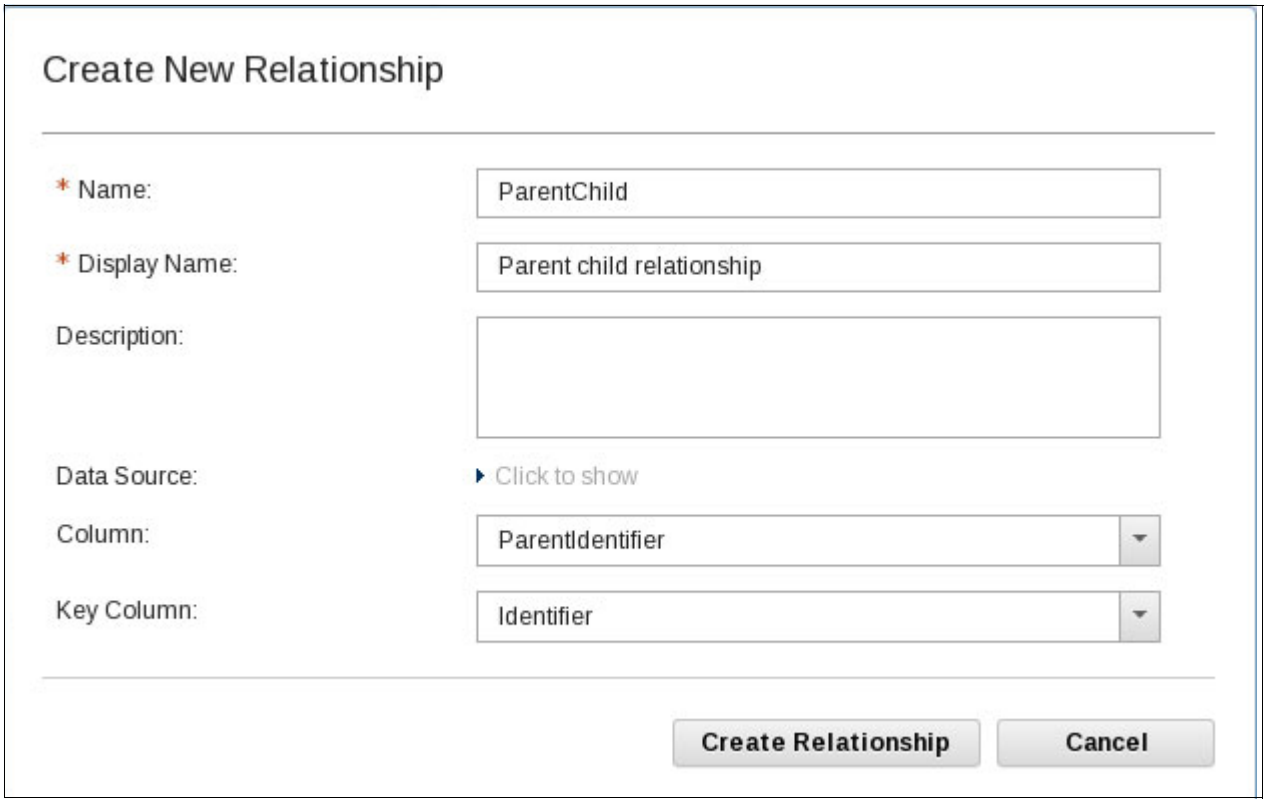

Figure 9-2 Create New Relationship window

9.4.3 Viewing the grouping

Helen adds the ScopeID and SiteName fields to her Event List view and renames the SiteName field to something more appropriate, such as Suburb. She then replays a sample of the organization's event data through the new Probe rules file (the results are shown in Figure 9-3).

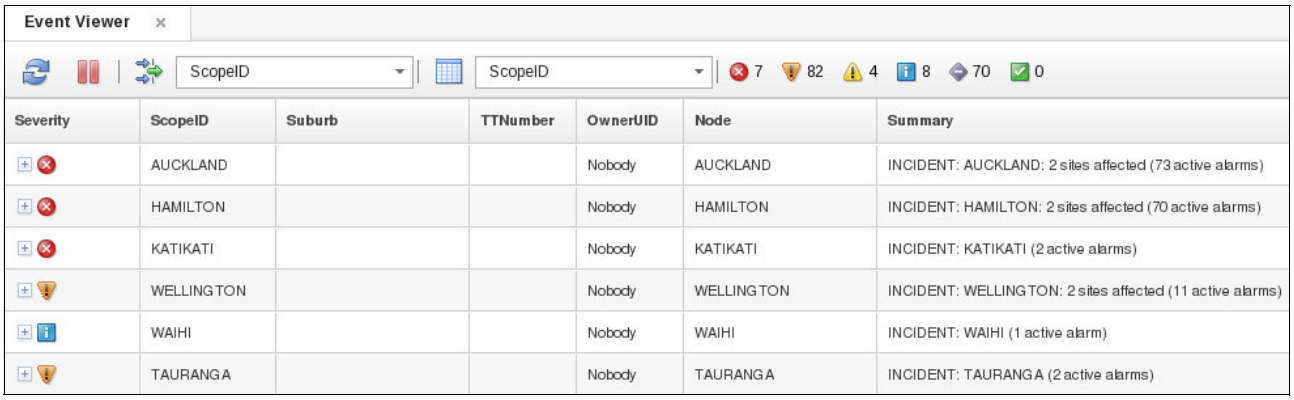

Severity	ScopeID	Suburb	TTNumber	OwnerUID	Node	Summary
	AUCKLAND			Nobody	AUCKLAND	INCIDENT: AUCKLAND: 2 sites affected (73 active alarms)
	HAMILTON			Nobody	HAMILTON	INCIDENT: HAMILTON: 2 sites affected (70 active alarms)
	KATIKATI			Nobody	KATIKATI	INCIDENT: KATIKATI (2 active alarms)
	WELLINGTON			Nobody	WELLINGTON	INCIDENT: WELLINGTON: 2 sites affected (11 active alarms)
	WAIHI			Nobody	WAIHI	INCIDENT: WAIHI (1 active alarm)
	TAURANGA			Nobody	TAURANGA	INCIDENT: TAURANGA (2 active alarms)

Figure 9-3 List of groupings

The sample contains 159 unique events. By applying scope-based event grouping, the large number of events collapses down to six rows or groups, which is comparable to the number of incidents to which the event set relates. For the specific sample set, the number of rows that are being presented to operations is reduced by 96%.

9.4.4 Modifying the properties

Helen performs the following steps to modify the properties:

1. She opens the AUCKLAND grouping and inspects the appearance of the groupings, as shown in Figure 9-4.

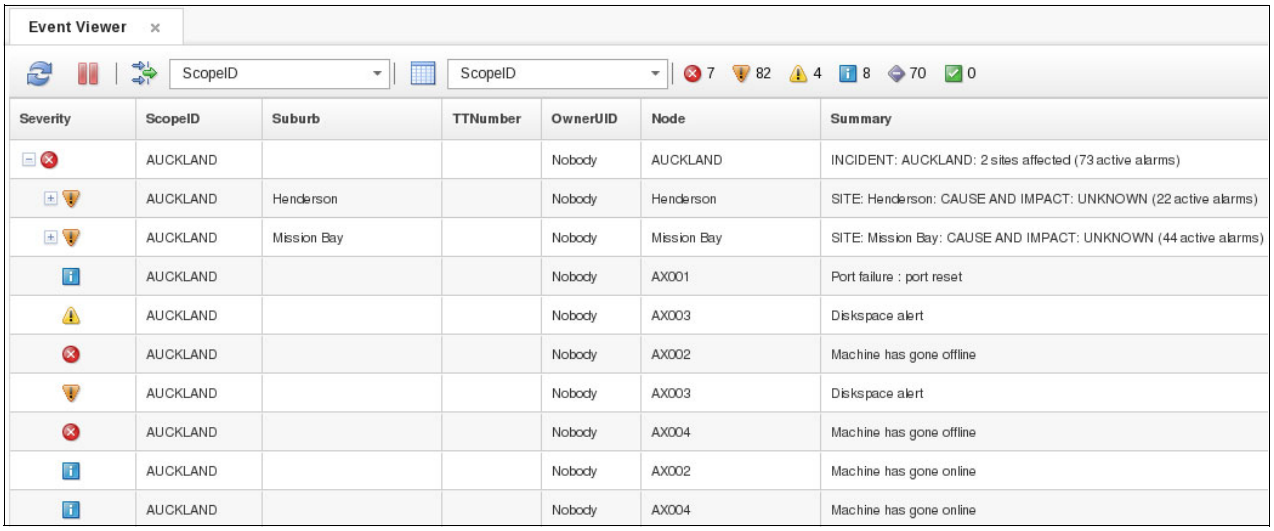

Severity	ScopeID	Suburb	TTNumber	OwnerUID	Node	Summary
	AUCKLAND			Nobody	AUCKLAND	INCIDENT: AUCKLAND: 2 sites affected (73 active alarms)
	AUCKLAND	Henderson		Nobody	Henderson	SITE: Henderson: CAUSE AND IMPACT: UNKNOWN (22 active alarms)
	AUCKLAND	Mission Bay		Nobody	Mission Bay	SITE: Mission Bay: CAUSE AND IMPACT: UNKNOWN (44 active alarms)
	AUCKLAND			Nobody	AX001	Port failure : port reset
	AUCKLAND			Nobody	AX003	Diskspace alert
	AUCKLAND			Nobody	AX002	Machine has gone offline
	AUCKLAND			Nobody	AX003	Diskspace alert
	AUCKLAND			Nobody	AX004	Machine has gone offline
	AUCKLAND			Nobody	AX002	Machine has gone online
	AUCKLAND			Nobody	AX004	Machine has gone online

Figure 9-4 AUCKLAND grouping

Although Helen is pleased with the grouping that results, she wants to modify the appearance of the event groupings to provide more useful information to the operators.

Scope-based event grouping makes available the 43 properties that allow its appearance and behavior to be modified. The properties are stored in the `master.properties` table in the Netcool/OMNIbus ObjectServer and can be updated via nco_sql or the Netcool Administrator tool. For more information about the properties, see the following IBM Knowledge Center site:

`http://ibm.biz/seg_docs`

2. Helen wants to modify some of the properties to change the appearance of some elements of the synthetic parent events. She also wants to activate the journaling feature in scope-based event grouping so that a single ticket can be cut from each subgrouping.

   Helen modifies the following properties to the values that are shown:

   – `SEGSiteNamePrefix = SUBURB` (CharValue)
   – `SEGScopeIDSitesAffectedLabel = suburb` (CharValue)
   – `SEGJournalToScopeIDParent = 1` (IntValue)
   – `SEGJournalToSiteNameParent = 1` (IntValue)

   These changes update the Event Viewer, as shown in Figure 9-5.

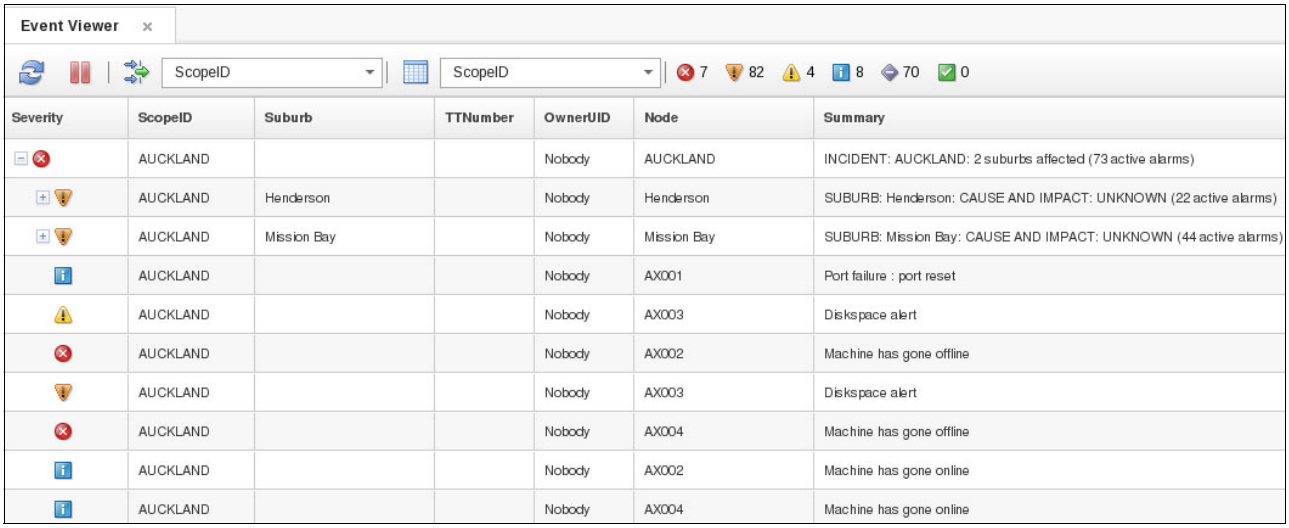

Severity	ScopeID	Suburb	TTNumber	OwnerUID	Node	Summary
	AUCKLAND			Nobody	AUCKLAND	INCIDENT: AUCKLAND: 2 suburbs affected (73 active alarms)
	AUCKLAND	Henderson		Nobody	Henderson	SUBURB: Henderson: CAUSE AND IMPACT: UNKNOWN (22 active alarms)
	AUCKLAND	Mission Bay		Nobody	Mission Bay	SUBURB: Mission Bay: CAUSE AND IMPACT: UNKNOWN (44 active alarms)
	AUCKLAND			Nobody	AX001	Port failure : port reset
	AUCKLAND			Nobody	AX003	Diskspace alert
	AUCKLAND			Nobody	AX002	Machine has gone offline
	AUCKLAND			Nobody	AX003	Diskspace alert
	AUCKLAND			Nobody	AX004	Machine has gone offline
	AUCKLAND			Nobody	AX002	Machine has gone online
	AUCKLAND			Nobody	AX004	Machine has gone online

Figure 9-5 Updated Event Viewer

3. When Helen double-clicks **SiteName event** (which is also known as a SUBURB event), she can see that the journals are now there and ready for ticketing, as shown in Figure 9-6).

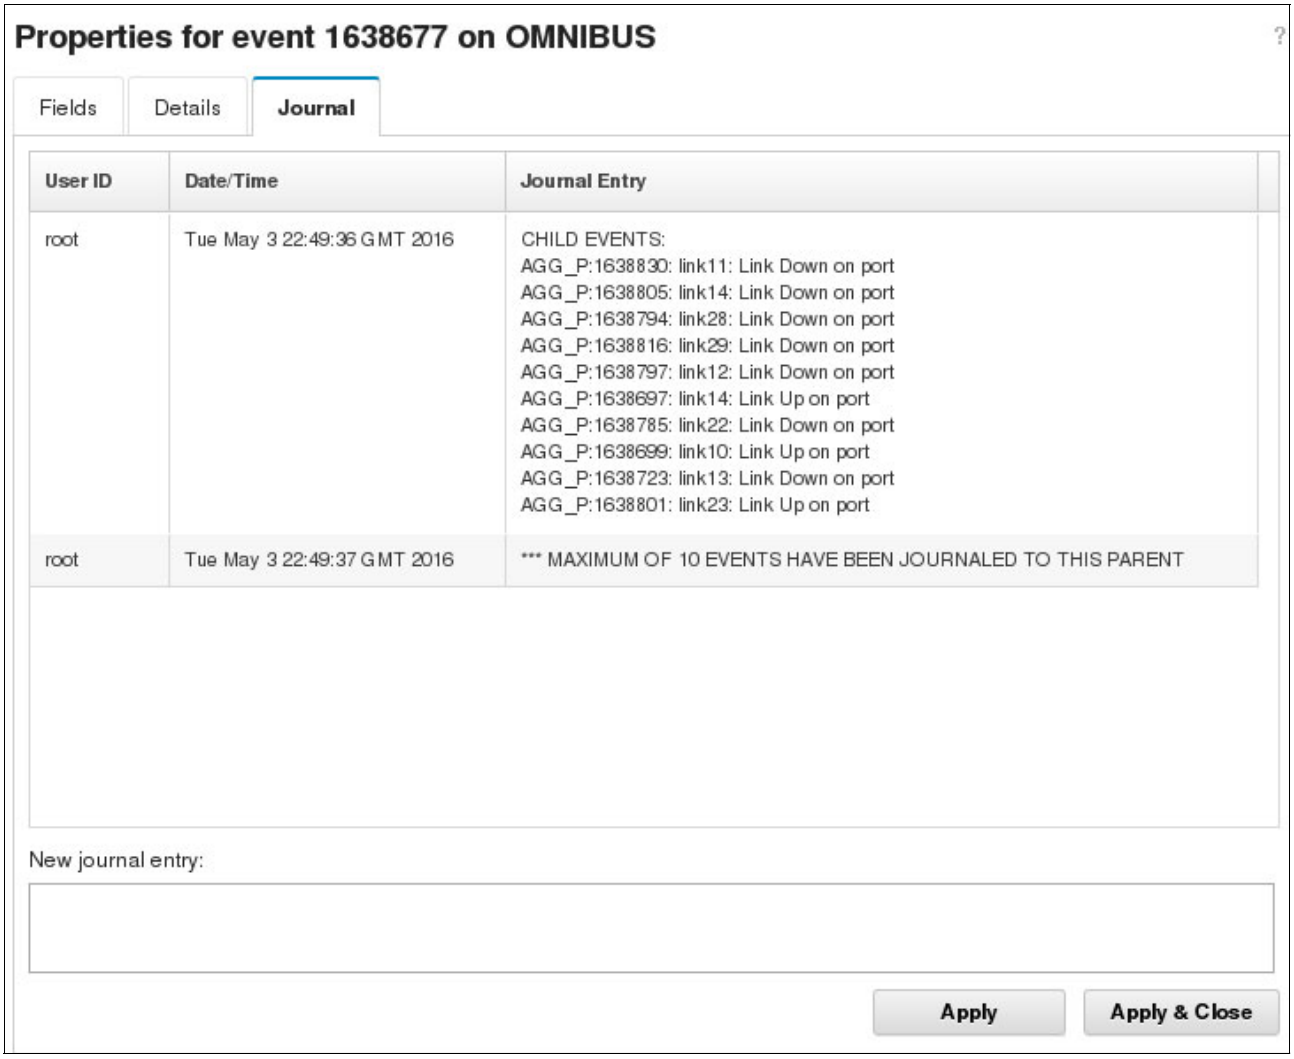

Figure 9-6 Journals ready for ticketing: SiteName event

Tip: The default setting for the maximum number of events to send to the journal of a SiteName event is 10. This setting can be modified via the `SEGMaxSiteNameJournals` property.

4. When Helen double-clicks a **ScopeID event**, she can see that the journals are now there and ready for ticketing, as shown in Figure 9-7.

Figure 9-7 Journals ready for ticketing: ScopeID event

Tip: The default setting for the maximum number of events to send to the journal of a ScopeID event is 50. This setting can be modified via the `SEGMaxScopeIDJournals` property.

5. Helen verifies that when the synthetic parent event is assigned to a user, it is assigned to a group or has a ticket assigned to it and that the OwnerUID, OwnerGID, and TTNumber all individually propagate to the child events, as shown in Figure 9-8.

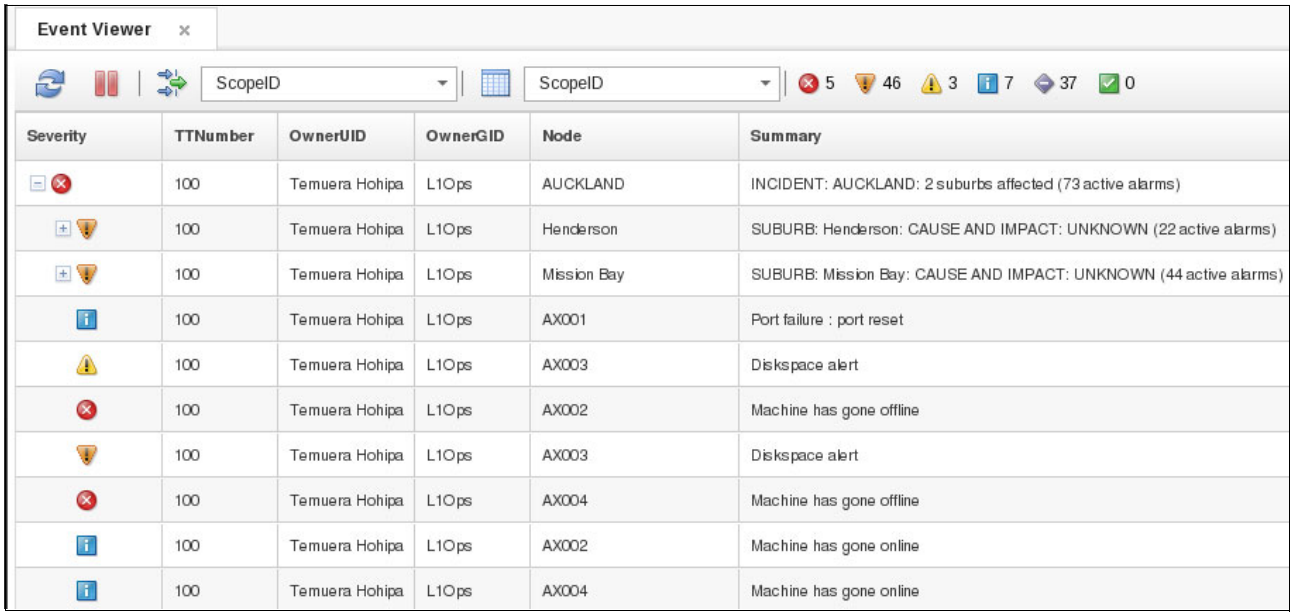

Severity	TTNumber	OwnerUID	OwnerGID	Node	Summary
	100	Temuera Hohipa	L1Ops	AUCKLAND	INCIDENT: AUCKLAND: 2 suburbs affected (73 active alarms)
	100	Temuera Hohipa	L1Ops	Henderson	SUBURB: Henderson: CAUSE AND IMPACT: UNKNOWN (22 active alarms)
	100	Temuera Hohipa	L1Ops	Mission Bay	SUBURB: Mission Bay: CAUSE AND IMPACT: UNKNOWN (44 active alarms)
	100	Temuera Hohipa	L1Ops	AX001	Port failure : port reset
	100	Temuera Hohipa	L1Ops	AX003	Diskspace alert
	100	Temuera Hohipa	L1Ops	AX002	Machine has gone offline
	100	Temuera Hohipa	L1Ops	AX003	Diskspace alert
	100	Temuera Hohipa	L1Ops	AX004	Machine has gone offline
	100	Temuera Hohipa	L1Ops	AX002	Machine has gone online
	100	Temuera Hohipa	L1Ops	AX004	Machine has gone online

Figure 9-8 OwnerUID, OwnerGID, and TTNumber propagate

9.4.5 Using ScopeAlias

Although scope-based event grouping is data-driven, there are occasions where it makes sense to merge two or more different scopes and reflect them as a single entity. This merge can be done by defining a ScopeAlias.

Helen used the following process:

1. After reviewing the resulting groupings, Helen identifies that three of the smaller towns (KATIKATI, WAIHI, and TAURANGA) overlap in terms of the underlying network infrastructure. Helen decides to combine the events from these three subscopes because together they make up the definition of same place, as shown in Figure 9-9.

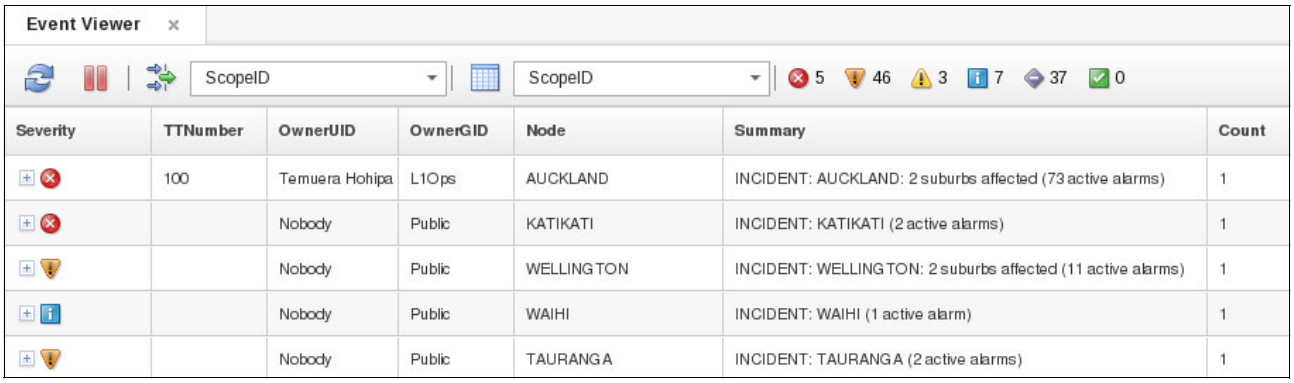

Severity	TTNumber	OwnerUID	OwnerGID	Node	Summary	Count
	100	Temuera Hohipa	L1Ops	AUCKLAND	INCIDENT: AUCKLAND: 2 suburbs affected (73 active alarms)	1
		Nobody	Public	KATIKATI	INCIDENT: KATIKATI (2 active alarms)	1
		Nobody	Public	WELLINGTON	INCIDENT: WELLINGTON: 2 suburbs affected (11 active alarms)	1
		Nobody	Public	WAIHI	INCIDENT: WAIHI (1 active alarm)	1
		Nobody	Public	TAURANGA	INCIDENT: TAURANGA (2 active alarms)	1

Figure 9-9 Combining the events

2. Helen creates a ScopeAlias of BAY OF PLENTY for the three smaller towns; which is the name of the larger region in which the three smaller towns are located.
3. Helen adds one entry for each of the three towns via the Netcool Administrator. She also prepares an SQL file to check into the company version control system for future use, as shown in the following example:

```
insert into master.correlation_scopealias_members (ScopeAlias, ScopeID) values
('BAY OF PLENTY', 'TAURANGA');
go
insert into master.correlation_scopealias_members (ScopeAlias, ScopeID) values
('BAY OF PLENTY', 'WAIHI');
go
insert into master.correlation_scopealias_members (ScopeAlias, ScopeID) values
('BAY OF PLENTY', 'KATIKATI');
go
```

4. Upon replaying the test data through the system, Helen sees the results that are shown in Figure 9-10. All three of KATIKATI, WAIHI, and TAURANGA retain their original ScopeID values in each's respective ScopeID fields; however, they are grouped under BAY OF PLENTY alias instead. The BAY OF PLENTY label is an alias to all three scopes, hence the term ScopeAlias.

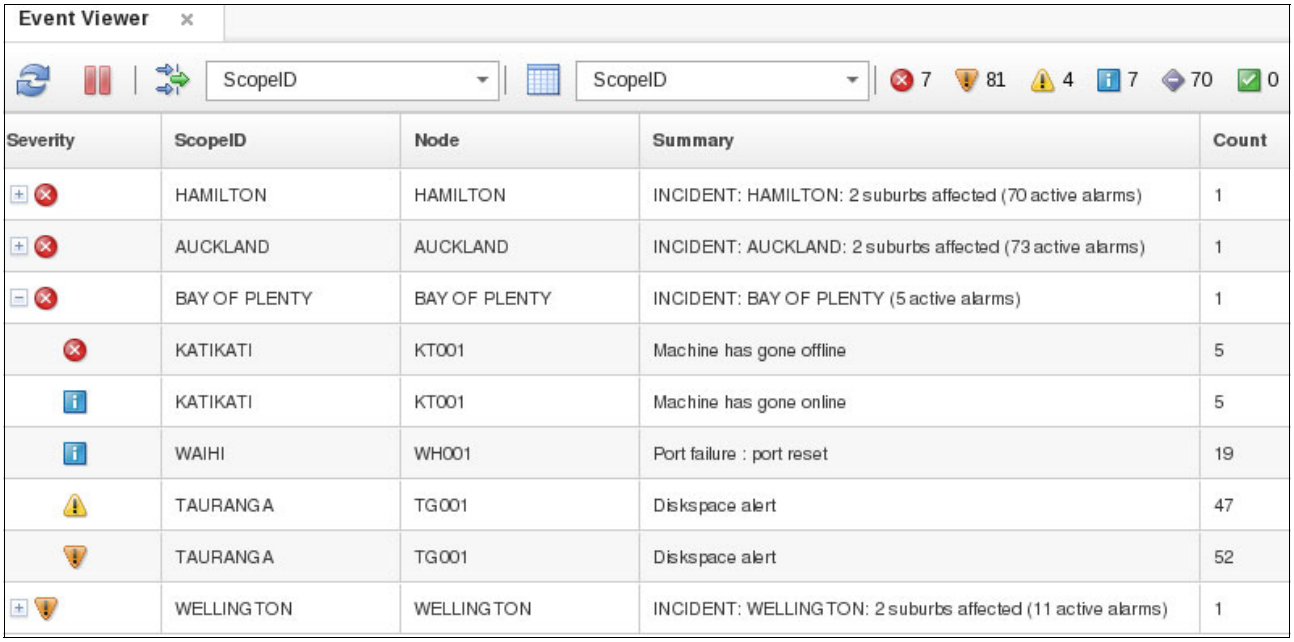

Severity	ScopeID	Node	Summary	Count
	HAMILTON	HAMILTON	INCIDENT: HAMILTON: 2 suburbs affected (70 active alarms)	1
	AUCKLAND	AUCKLAND	INCIDENT: AUCKLAND: 2 suburbs affected (73 active alarms)	1
	BAY OF PLENTY	BAY OF PLENTY	INCIDENT: BAY OF PLENTY (5 active alarms)	1
	KATIKATI	KT001	Machine has gone offline	5
	KATIKATI	KT001	Machine has gone online	5
	WAIHI	WH001	Port failure : port reset	19
	TAURANGA	TG001	Diskspace alert	47
	TAURANGA	TG001	Diskspace alert	52
	WELLINGTON	WELLINGTON	INCIDENT: WELLINGTON: 2 suburbs affected (11 active alarms)	1

Figure 9-10 Events that are grouped under BAY OF PLENTY alias

Using CauseWeight and ImpactWeight

Helen wants to enhance the resulting groupings with weightings so that the events can be prioritized for the operators in terms of cause and affect. This change helps the operators to more easily pinpoint the events in a group that represent the probable causes of each incident. Until now, the synthetic subgrouping parent events showed CAUSE AND IMPACT: UNKNOWN in the Summary field because none of the child events included assigned weightings.

Tip: This cause and affect text can be switched on or off the scope and subgrouping parent events separately via the Properties menu.

Scope-based event grouping provides a standard method to weigh events in terms of the likelihood that they are a high-impacting event to businesses or services, and each one's likelihood that it is a contributing cause of an incident. This information can then be used to enrich the Summary field of the synthetic parent events to both guide operators, and provide more information to any ticket headlines.

Helen completed the following steps:

1. Helen creates a copy of the Probe rules template that is provided with scope-based event grouping into the main Probe directory to work with so that she can allocate the event categories.
2. Helen edits the top section of the template to standardize the setting:
   - ScopeID
   - SiteName (where available)
   - Event category (NormalisedAlarmCode)
   - OSI level of the event
3. In the first subsection, Helen sets the default values of the following fields per the template:

```
# SET / INITIALISE MANDATORY VARIABLES
@ScopeID = $location
@NormalisedAlarmCode = 0
$OSILevel = 9
# SET / INITIALISE OPTIONAL VARIABLES
@SiteName = $suburb
```

4. Helen edits the next section of the file that switches on $EventCode. In the case of Company A, $EventCode is not a valid token. Therefore, Helen modifies the switch statement to use other tokens to determine event categorization, and setting a sensible OSI level for the events:

```
switch ($MyField) {

        case "INFO": # EXAMPLE - Informational events

                @NormalisedAlarmCode = 10
                $OSILevel = 3

        case "A1400": # EXAMPLE - Workarounds in execution

                @NormalisedAlarmCode = 20
                $OSILevel = 3
        ...
```

Tip: For more information about the 16 event categories (from purely informational to controlled shutdown), see this website:

http://ibm.biz/seg_fields

The table at this site shows how categorizations and OSI levels combine to establish the weightings. These weightings are implemented in the second half of the Probe rules file template and are not edited. It is important to use the standard weighting method so that events from any source can be compared with any other events, in terms of their cause and affect, regardless of their source.

5. Helen modifies the following other properties to enable the display of cause and affect analysis text to automatically appear in the Summary fields of the synthetic containment events, wherever direct child events exist to each respective parent event:
   - `SEGUseScopeIDImpactCause = 1`
   - `SEGUseSiteNameImpactCause = 1`
6. Helen clears the ObjectServer and replays the event data through the Probe. She sees the events that are shown in Figure 9-11 with weightings preset and cause and affect diagnosis text that appears in the Summary fields of the ScopeIDParent events and the SiteNameParent events.

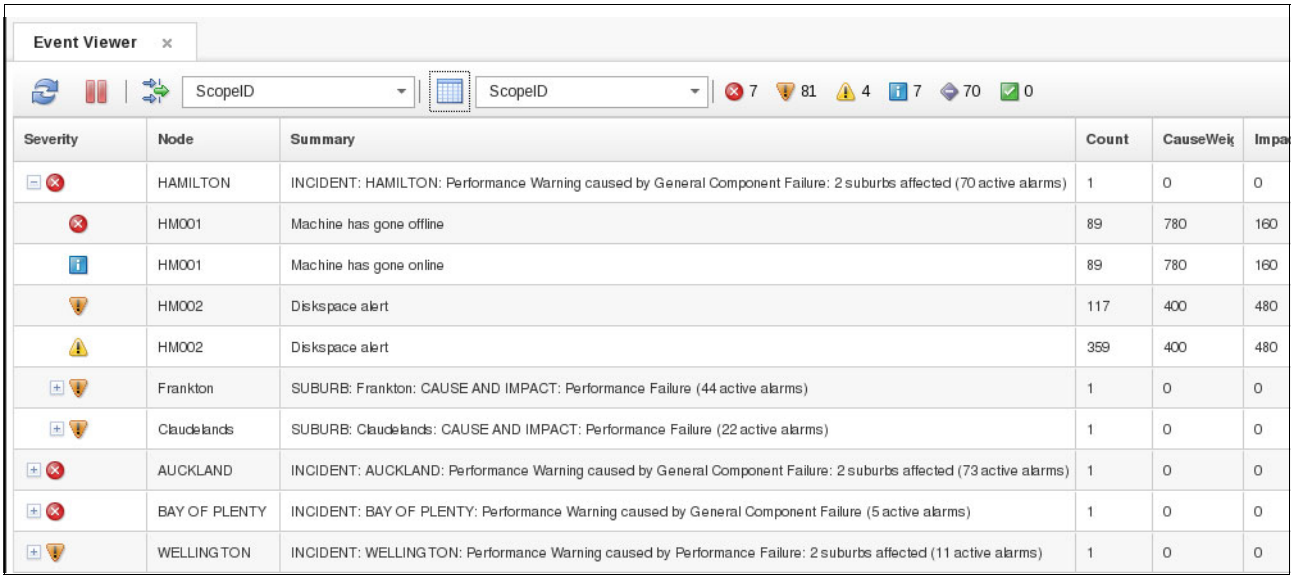

Severity	Node	Summary	Count	CauseWeig	Impa
	HAMILTON	INCIDENT: HAMILTON: Performance Warning caused by General Component Failure: 2 suburbs affected (70 active alarms)	1	0	0
	HM001	Machine has gone offline	89	780	160
	HM001	Machine has gone online	89	780	160
	HM002	Diskspace alert	117	400	480
	HM002	Diskspace alert	359	400	480
	Frankton	SUBURB: Frankton: CAUSE AND IMPACT: Performance Failure (44 active alarms)	1	0	0
	Claudelands	SUBURB: Claudelands: CAUSE AND IMPACT: Performance Failure (22 active alarms)	1	0	0
	AUCKLAND	INCIDENT: AUCKLAND: Performance Warning caused by General Component Failure: 2 suburbs affected (73 active alarms)	1	0	0
	BAY OF PLENTY	INCIDENT: BAY OF PLENTY: Performance Warning caused by General Component Failure (5 active alarms)	1	0	0
	WELLINGTON	INCIDENT: WELLINGTON: Performance Warning caused by Performance Failure: 2 suburbs affected (11 active alarms)	1	0	0

Figure 9-11 Cause and affect diagnosis text

9.4.6 Using data from the highest ranked child event

Finally, Helen wants to enhance the Summary field of the subgroup parent events with the Node value of the highest weighted child. Although the properties limit which default items can go into the Summary line of the synthetic parent events, it can be done via the CustomText field if there is any other text that needs to be included.

The process includes the following tasks:

- We can put any text that we want in each child event's CustomText field.
- We can auto-select the child with the highest weighted cause, the highest weighted impact, the first FirstOccurrence (for example, first event in the group), or the last LastOccurrence (most recent recurrence in the group). The CustomText from the auto-selected event is copied to its direct parent event.
- We can opt to display a synthetic parent event's CustomText in its Summary field.

The following properties perform the auto-select of the priority child event for a ScopeID parent event in the order of precedence listed:

- `SEGPropagateTextToScopeIDParentCause = 1`
- `SEGPropagateTextToScopeIDParentImpact = 1`
- `SEGPropagateTextToScopeIDParentFirst = 1`
- `SEGPropagateTextToScopeIDParentLast = 1`

Tip: If SEGPropagateTextToScopeIDParentCause is set to 1, the rest are ignored. Similarly, if SEGPropagateTextToScopeIDParentCause is set to 0 but SEGPropagateTextToScopeIDParentImpact is set to 1, the rest are ignored, and so on.

Similarly, the following properties perform the auto-select of the priority child event for a ScopeID parent event in the order of precedence listed:

- `SEGPropagateTextToSiteNameParentCause = 1`
- `SEGPropagateTextToSiteNameParentImpact = 1`
- `SEGPropagateTextToSiteNameParentFirst = 1`
- `SEGPropagateTextToSiteNameParentLast = 1`

The following properties enable showing the CustomText field of the synthetic parent in its own Summary field:

- `SEGUseScopeIDCustomText = 1`
- `SEGUseSiteNameCustomText = 1`

Helen completed the following steps:

1. Because Helen wants to propagate the highest cause weighted child to the subgrouping parent and display its Node in the event, she sets the following properties only:
   - `SEGPropagateTextToSiteNameParentCause = 1`
   - `SEGUseSiteNameCustomText = 1`
2. Helen enters the value of the CustomText fields for the children events in the Probe rules so that when the events are present in the ObjectServer, they are holding the text that they need to pass should they happen to be the highest weighted child, as shown in the following example:

   ```
   @CustomText = "high node: " + @Node
   ```
3. Helen clears the ObjectServer and replays the event data through the Probe and sees the groupings that are shown in Figure 9-12 with the top weighted child event Node displayed in the Summary line of each of the subgrouping synthetic parent events.

   Figure 9-12 also shows that the subgroupings of Frankton and Claudelands have top nodes of link42 and link65 because of the modifications Helen made.

Event Viewer

ScopeID | ScopeID | 7 | 81 | 4 | 7 | 70 | 0

Severity	Node	Summary	Count	CauseWei	ImpactWei
	HAMILTON	INCIDENT: HAMILTON: Performance Warning caused by General Component Failure: 2 suburbs affected (70 active alarms)	1	0	0
	HM001	Machine has gone offline	89	780	160
	HM001	Machine has gone online	89	780	160
	HM002	Diskspace alert	117	400	480
	HM002	Diskspace alert	359	400	480
	Frankton	SUBURB: Frankton: CAUSE AND IMPACT: Performance Failure: high node: link42 (44 active alarms)	1	0	0
	Claudelands	SUBURB: Claudelands: CAUSE AND IMPACT: Performance Failure: high node: link65 (22 active alarms)	1	0	0
	AUCKLAND	INCIDENT: AUCKLAND: Performance Warning caused by General Component Failure: 2 suburbs affected (73 active alarms)	1	0	0
	BAY OF PLENTY	INCIDENT: BAY OF PLENTY: Performance Warning caused by General Component Failure (5 active alarms)	1	0	0

Figure 9-12 Subgroupings of Frankton and Claudelands

9.5 Summary

Helen finished all her configuration modifications and is ready to begin user acceptance testing. She is confident the new groupings will help the operations team more easily make sense of the large volumes of events because of the logical grouping that is occurring and the cause and affect analysis that is done automatically by the system. She expects ticket counts to drop significantly and MTTR to improve dramatically. This improvement will help to save the business money and provide a better service to its customers.

Note: For more information about setting up scope-based event grouping, see the documentation that is available at the following IBM Knowledge Center site:

http://ibm.biz/seg_docs

Related publications

The publications that are listed in this section are considered particularly suitable for a more detailed discussion of the topics that are covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide more information about the topic in this document. Note that some publications that are referenced in this list might be available in softcopy only:

- *IBM Netcool Operations Insight Version 1.4 Deployment Guide*, SG24-8365
- *Delivering Consistency and Automation with Operational Runbooks,* REDP-5347

You can search for, view, download, or order these documents and other Redbooks, Redpapers, Web Docs, draft, and other materials, at the following website:

ibm.com/redbooks

Online resource

The IBM Netcool Operations Insight Version 1.4.0.1 Knowledge Center documentation website also is relevant as an information source:

https://ibm.biz/BdrFcE

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Redbooks

IBM Netcool Operations Insight: A Scenarios Guide

(0.2”spine)
0.17”<->0.473”
90<->249 pages

SG24-8352-00

ISBN 0738441856

Printed in U.S.A.