



Hitachi IT Operations Analyzer

Root Cause Analysis for Supporting Fault Identification

By Yutaka Kudo and Saurabh (Manu) Batra

April 2010

Table of Contents

Executive Summary	3
Challenges with Root Cause Identification	3
Overview of Root Cause Analysis	4
How It Works	4
Ideal for Administrative Efficiency	5
Availability Related Root Cause Analysis Scenarios	6
Scenario: Identifying the Root Cause of a Hard Disk I/O Error Caused by Fibre Channel Switch Failure	6
Scenario: Identifying the Root Cause of Hard Disk I/O Error Caused by Cable Disconnection	7
Scenario: Identifying the Root Cause of Virtual Server Down	8
Performance Related Root Cause Analysis Scenarios	9
Scenario: Identifying the Root Cause of Disk I/O Slow Down Caused by Storage Device	9
Scenario: Identifying the Root Cause of Disk I/O Slow Down Caused by Fibre Channel Switch	11
Scenario: Identifying the Root Cause of Virtual Server Slowdown	11
Root Cause Analysis Topology View	12
Conclusion	13

Executive Summary

No data center can afford downtime with the manic pace of today's always available business operations. But for many midsized businesses, ensuring that mission critical systems perform efficiently and reliably amid the complexity of multivendor, multiplatform infrastructures can become an insurmountable task without the right tools. IT administrators need all the help they can get to quickly find and diagnose system disruptions or failures. The faster the problem is found, the faster it can be solved.

But getting to the root cause of the failure event is often a problem in itself. To remove the fault from a device in the data center, the administrator must be able to identify where the root cause of the fault has occurred. In most cases, this is a difficult task when that device is connected to other devices and has caused a chain reaction of subsequent alert messages. Also, the administrator may not have sufficient familiarity with the system configuration or related device dependencies. As a result, fault identification becomes time-consuming and the duration of disruption lengthens, which can affect production or data availability.

Most monitoring software alerts users on individual events and does not provide any correlation between events. Individual alerts for multiple devices may notify the administrator of a problem; however, they do not divulge why or where the fault has actually occurred. While some enterprise-class software uses intelligent algorithms and rules-based inference technology to guide administrators to the fault node, midmarket products usually offer separate tools for checking disruption events and correlating dependencies.

Hitachi Data Systems introduces Hitachi IT Operations Analyzer software with unique Root Cause Analysis functionality. The IT Operations Analyzer is a powerful, proven approach to simplifying data center monitoring, with comprehensive performance and availability monitoring of up to 750 server, network and storage nodes. IT Operations Analyzer software was designed from the ground up to meet the needs of medium business IT managers who are challenged to manage complex IT systems with the resources at hand, without requiring comprehensive knowledge of servers, networks or storage systems, or using multiple software products.

Root Cause Analysis provides administrators with a quick path to fault identification. Ideal for administrators wishing to reduce mean time to repair, Root Cause Analysis uses patent pending technology to usher the user directly to the node causing the fault, without requiring in-depth knowledge of separate environments or time-consuming investigation.

This application brief focuses on the outstanding benefits and business value of the Root Cause Analysis features of IT Operations Analyzer.

Challenges with Root Cause Identification

As businesses demand more from their IT infrastructures, administrators must meet the challenges of fine tuning performance, availability and recovery processes across the data center, usually with limited resources. Complexity within the data center can grow quickly and organizations may not have the deeper technical knowledge or the staff resources required to efficiently manage the multitude of devices, applications and environments. Administrators in midsized companies, both

IT generalists and specialists, are seeking ways to ease the burden of monitoring and managing complicated IT systems.

When a disruption occurs somewhere within the data center, it is imperative that the administrator understand where and why the fault has occurred in order to fix it. But getting to the root of a problem can be difficult, especially if the fault resides on a system with numerous servers, storage and switch dependencies or connections. The system will send out an alert, notifying the administrator that there is some sort of disruption or performance degradation. Alerts are initiated for various types of failures and also highlight symptoms that may put the system at risk down the road.

When a fault occurs on a system, the related devices also send out error or failure notifications, presenting the administrator with a labyrinth of alert messages to consider as possible suspects in the probable root cause of the event. The longer it takes for the administrator to identify the fault node, the greater the affect on business operations, including slow or no data access, or possibly data loss.

In such situations, reducing the mean time to diagnose the fault is critical. Unlike most monitoring software for medium businesses, Hitachi IT Operations Analyzer unifies system monitoring and topology mapping for server, network (LAN and SAN) and storage to allow events and node dependencies to be checked as part of a comprehensive root cause analysis capability.

The Root Cause Analysis feature of Hitachi IT Operations Analyzer offers both generalists and specialists alike a fast and simple way to identify root cause nodes within complex IT system failures. By improving the fault identification process, administrators can find and fix problems faster to minimize system downtime and increase data center efficiencies.

Overview of Root Cause Analysis

In general, root cause analysis is a method of problem solving that attempts to identify the initiating factors that caused the issue and separate them from symptoms. The root cause often initiates a causal chain of outcomes. In the data center, that translates to a rapid fire of ensuing alert messages from devices related to the system where the root cause or fault occurred. To be effective at solving the originating problem, root cause analysis must be performed systematically, with conclusions and causes substantiated by documented evidence.

The Hitachi IT Operations Analyzer Root Cause Analysis feature amplifies this premise with a cohesive set of principles or rules designed to quickly identify the root cause node. Using a patent pending technology that includes sophisticated algorithms and rules-based inference, the Root Cause Analysis provides the process for examining failure events and deriving conclusions about the probable causes for those events. More importantly, the tool is able to quickly guide the administrator directly to the node at fault without a need for expert knowledge of system configurations or dependencies.

How It Works

IT Operations Analyzer software provides availability and performance monitoring of up to 750 server, network and storage nodes in the data center. Once the administrator defines the polling cycle frequency for performance, availability and configuration data collection, IT Operations Analyzer starts monitoring the environment for device failures and performance threshold violations.

The software treats three main types of events: device configuration change events, performance threshold violation events and SNMP Traps for root cause analysis. All these events are translated into internal events and matched against a set of rules and device dependency mappings to identify root cause. Events related to a single root cause are gathered together for easy identification and acknowledgement. IT Operations Analyzer has two main types of root cause analysis: configuration root cause, which is related to issues caused by configuration changes like device down, and performance root cause analysis (also known as bottleneck analysis), which indicates to performance related issues.

The system provides probability-based suggestion on root cause. The configuration root cause analysis works right out of the box, without user intervention or tuning. It's based on pre-defined set of heuristics. For performance root cause analysis the user needs to set performance thresholds via templates. Once the thresholds are set, IT Operations Analyzer correlates the performance threshold violation events with dependency maps to identify root cause.

As with expert system software that attempts to reproduce the analytical skill set of human experts in a particular field, the IT Operations Analyzer software uses event correlation rules for various situations and conditions of the IT systems. Root Cause Analysis calculates the certainty for every related rule to determine the most probable fault as the root cause of the event. Over time, as more information is accumulated and analyzed, the Root Cause Analysis can be requested to re-analyze, increasing the probability confidence factor of pinpointing the correct fault node.

Ideal for Administrative Efficiency

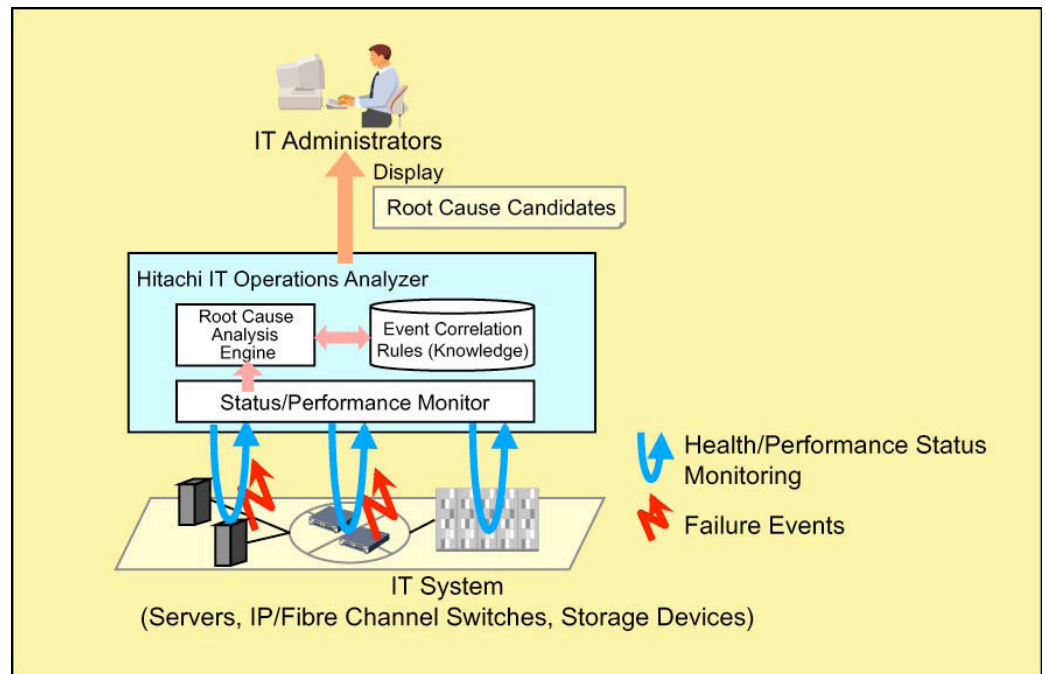
Because IT Operations Analyzer is designed to simplify root cause analysis, event correlation rules and other important identification rules are predefined based on the expert knowledge base within the software, alleviating the need for administrators to create rules or metrics themselves.

These dependencies are automatically updated the moment a new node comes online or the system topology changes. IT Operations Analyzer software is perpetually conducting discovery throughout the data center systems to ensure that monitoring and analysis tasks are based on the latest IT systems configuration. Each time a change is detected, the software updates its own database and notifies the rules to expand to incorporate the new device or configuration. In this way, Root Cause Analysis is always up to date when evaluating alerts.

Once the Root Cause Analysis identifies the most probable cause or causes of the fault occurrence, IT Operations Analyzer sends a notification to the administrator. Results of the root cause analysis are displayed in an easy-to-understand view along with a probability confidence rating and helpful explanation window with more details. (See Figure 1.)

The Root Cause Analysis can significantly reduce mean time to diagnose the fault. Because IT Operations Analyzer employs an intuitive, unified interface, users can quickly recognize the root cause nodes and the impacts to applications or business units. By providing swift and accurate fault identification, the administrator is better able to fix the problem node and simultaneously arrest the chain of successive alerts, thereby improving administrative efficiencies and reducing downtime or performance degradation.

Figure 1. IT Operations Analyzer Root Cause Analysis Process



Availability Related Root Cause Analysis Scenarios

Hitachi IT Operations Analyzer monitors availability and performance across the data center and its Root Cause Analysis feature performs critical analytic tasks to help keep everything running smoothly. This section highlights three example scenarios of root cause analysis related to availability of IT systems; it explains how an administrator might typically address the situation and how Root Cause Analysis would manage the same scenario. In each case, Root Cause Analysis spans servers, network switches and storage devices and both physical and virtual nodes.

Scenario: Identifying the Root Cause of a Hard Disk I/O Error Caused by Fibre Channel Switch Failure


This scenario assumes the IT system comprises servers, Fibre Channel switches and storage devices. The storage machines are providing logical volumes to servers through a Fibre Channel SAN.

A hard disk drive (HDD) I/O error has occurred on some servers in the system. In a typical scenario, the administrator would have to trace the path of the volumes through server, Fibre Channel network switches and storage to identify the point of failure. This may cause significant delay in problem identification. But in the case employing the Root Cause Analysis feature, the software will correlate events from the environment to identify the root cause; therefore, there is no need for the administrator to manually trace the path to identify root cause. Root Cause Analysis is able to

correlate the events that occurred in a certain time period based on the predefined correlation rules. This identification process is automatically initiated. The administrator only needs to view the Root Cause Analysis window on IT Operations Analyzer screen to see the root cause candidate.

Root Cause Analysis Result: In this case, the root cause would be identified as Fibre Channel Switch Down because no events occurred on any storage device, and dependencies were found between the servers with the HDD I/O error and the Fibre Channel switch. (See Figure 2.)

Figure 2. Fibre Channel Switch Failure on Fibre Channel SAN Topology



Scenario: Identifying the Root Cause of Hard Disk I/O Error Caused by Cable Disconnection

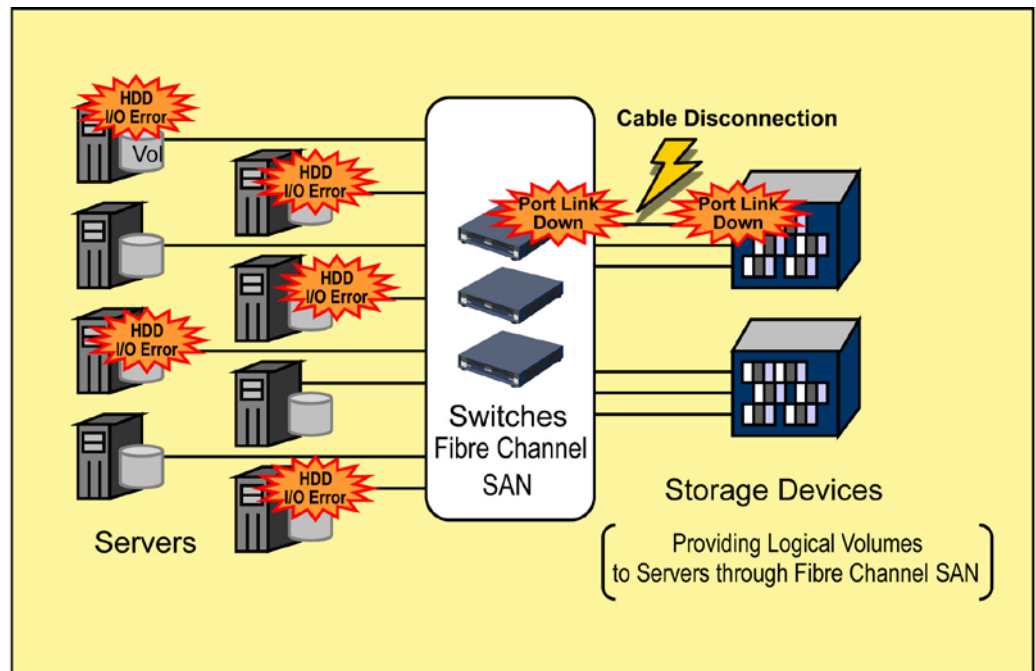
This scenario assumes the same IT environment as that of the previous scenario. Storage devices are providing logical volumes to servers through the Fibre Channel SAN as well.

An HDD I/O error has occurred on some servers on the system. The administrator observed Port Link Down events, so the administrator would need to refer to the network connection map and manually check all suspect connections to identify and verify if port disconnect caused the disk I/O failure. This may cause significant delay in problem identification. But in the case of the Root Cause Analysis, the software will correlate events from the environment to identify the root cause; therefore, there is no need for the administrator to manually trace the path to identify root cause.

Root Cause Analysis Result: In this case, the root cause would be identified as cable disconnection between the Fibre Channel switch and the storage device. The cable disconnection would signal a Port Link Down alert without any other port error being detected. (See Figure 3.)

Once the Root Cause Analysis result is accepted by the administrator, all the associated error messages and events would also be flagged as acknowledged and eliminated from subsequent error reporting.

Figure 3. Cable Disconnection between Fibre Channel Switch and Storage Device



Scenario: Identifying the Root Cause of Virtual Server Down

This scenario assumes the IT system comprises virtual servers, Fibre Channel switches and storage devices. These storage devices are providing logical volumes to the host (with a hypervisor) through a Fibre Channel SAN.

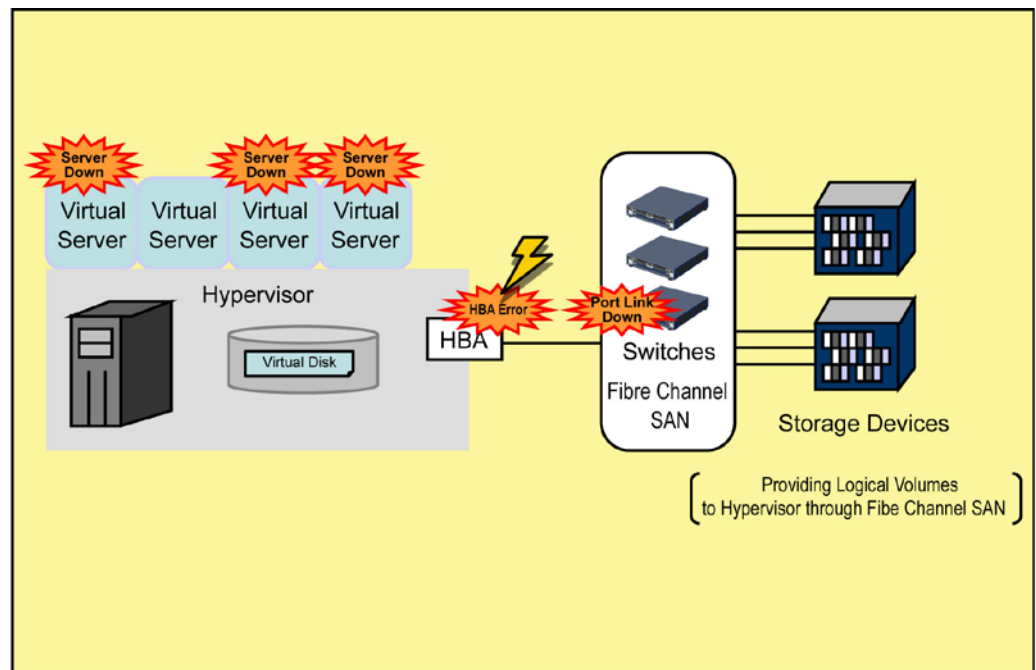
Some of the virtual servers have become inaccessible. In order to check the dependencies between virtual servers and the hypervisor, the administrator would need to first consult the management console used for the virtual server environment. If the virtual server management software does not include a topology map of the connections among the hypervisor, Fibre Channel switch and storage devices, the administrator may need to refer to the software tool that manages the physical systems environment.

If required to toggle between multiple management consoles to understand the dependencies for virtual and physical nodes, the administrator risks extending the time to find and diagnose the problem and the chance for error.

The Root Cause Analysis feature helps to identify and analyze the root cause in this complicated situation using the always updated topology mapping information that pinpoints both physical and virtual nodes.

Root Cause Analysis Result: A host bus adapter (HBA) provides connectivity between a storage device and a server, in this case, the connectivity is through the hypervisor to the unreachable virtual servers. For this scenario, both the HBA on the hypervisor and the Fibre Channel switch would indicate Port Link Down, and the root cause would be identified as HBA Err (HBA Error) on the hypervisor (See Figure 4.)

Figure 4. Virtual Server Error caused by HBA Error on Host



Performance Related Root Cause Analysis Scenarios

This section examines three example scenarios of root cause analysis related to performance problems, how the administrator might manage the situations and how the Root Cause Analysis feature of Hitachi IT Operations Analyzer software would resolve them. In each case, Root Cause Analysis spans servers, network switches and storage devices and both physical and virtual nodes.

Scenario: Identifying the Root Cause of Disk I/O Slow Down Caused by Storage Device

This scenario assumes the IT environment comprises servers, Fibre Channel switches and storage devices. Storage devices are providing logical volumes to servers through the Fibre Channel SAN.

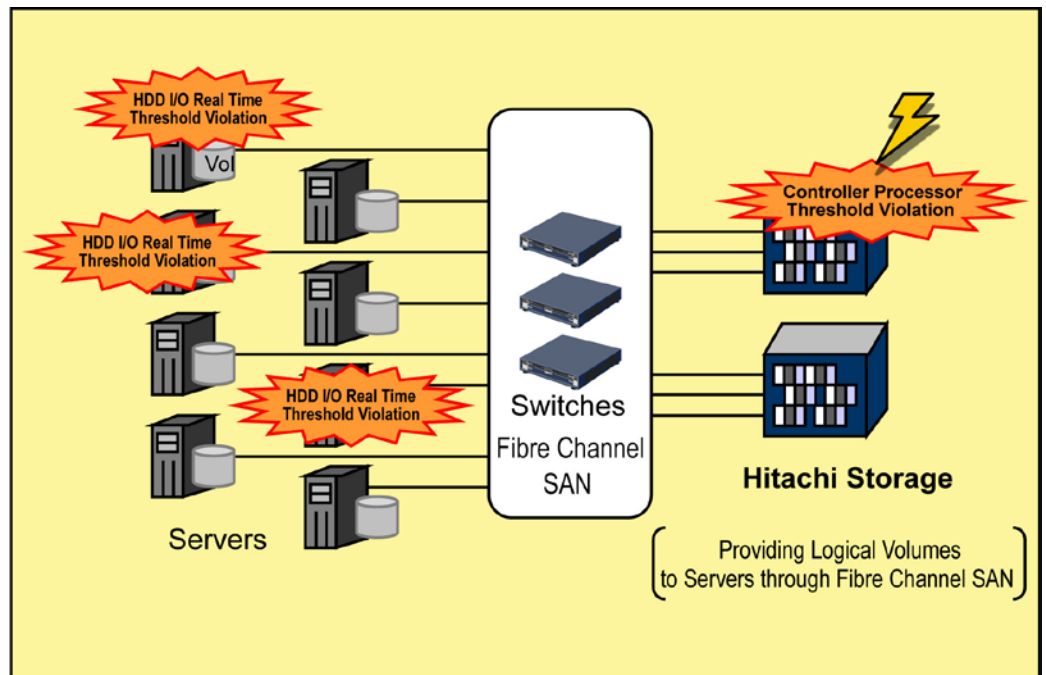
HDD I/O performance degradation has occurred on some servers. In this case, the administrator would need to check the performance status and dependencies of those servers, as well as the storage nodes and Fibre Channel switches, to determine where the bottlenecks reside.

Root Cause Analysis can help the administrator promptly identify the HDD I/O bottleneck, by scrutinizing performance thresholds. To obtain performance status for storage devices, Root Cause Analysis will initially evaluate information collected by IT Operations Analyzer using SMI-S (Storage Networking Industry Association standard) on those devices. Root Cause Analysis will also examine the performance status of other types of monitored nodes within the data center, such as servers and network switches. In this way, Root Cause Analysis is able to eliminate from suspicion those nodes performing normally and draw a more accurate conclusion about the performance degradation.

For instance, Root Cause Analysis may determine that storage performance has violated thresholds but that server CPU performance and port performance of switches are normal. Note: Root Cause Analysis is highly optimized to provide greater detail and probability when analyzing Hitachi modular storage devices because IT Operations Analyzer is able to collect more detailed performance status information about them.

Root Cause Analysis Result: In this case, the root cause would be identified as Controller Processor of Storage Device because the performance of the Fibre Channel Switch is normal and the processor utilization of storage controller violates its threshold. (See Figure 5.)

Figure 5. HDD Slow Down Caused by Storage Device



Scenario: Identifying the Root Cause of Disk I/O Slow Down Caused by Fibre Channel Switch

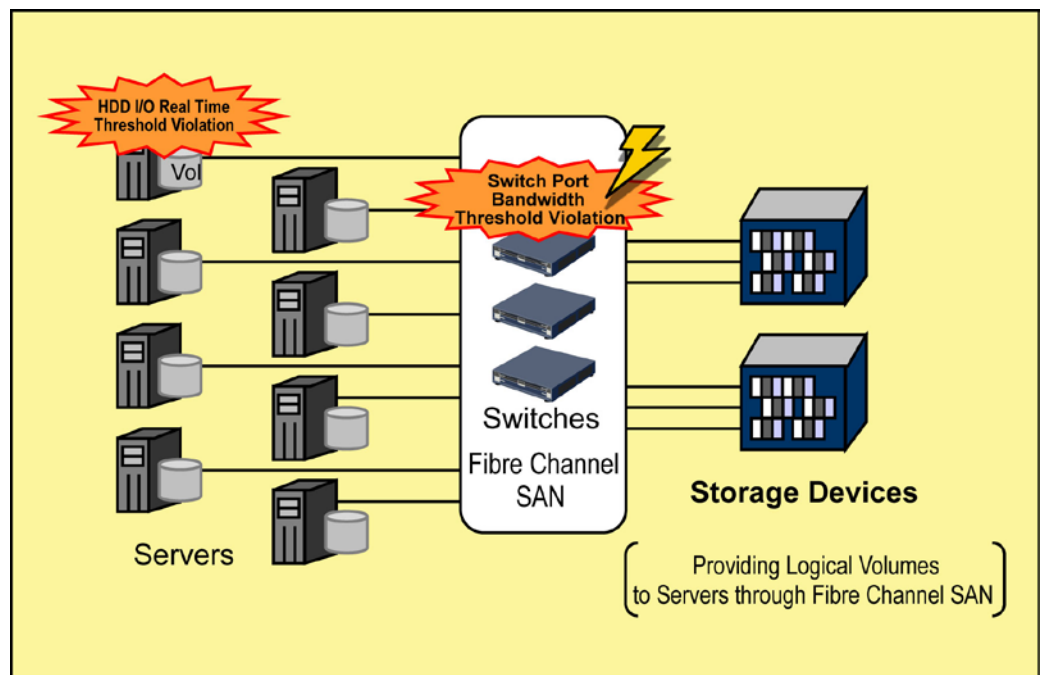
This example assumes the same IT system as that of the previous scenario. Storage devices are providing logical volumes to servers through a Fibre Channel SAN.

HDD I/O performance has been diminished on a server. In this case, the administrator would need to check the performance status and dependencies of that server, as well as the storage nodes and Fibre Channel switches, to determine where the bottleneck resides.

Root Cause Analysis verifies performance for the server, Fibre Channel switches and storage devices to identify the reason for the change in HDD I/O response time.

Root Cause Analysis Result: In this case, the root cause would be identified as Fibre Channel Switch Port because the performance of the switch port connected to that server exceeds its preset or user defined threshold. (See Figure 6.)

Figure 6. HDD Slow Down Caused by Network Overflow on Fibre Channel Switch Port



Scenario: Identifying the Root Cause of Virtual Server Slowdown

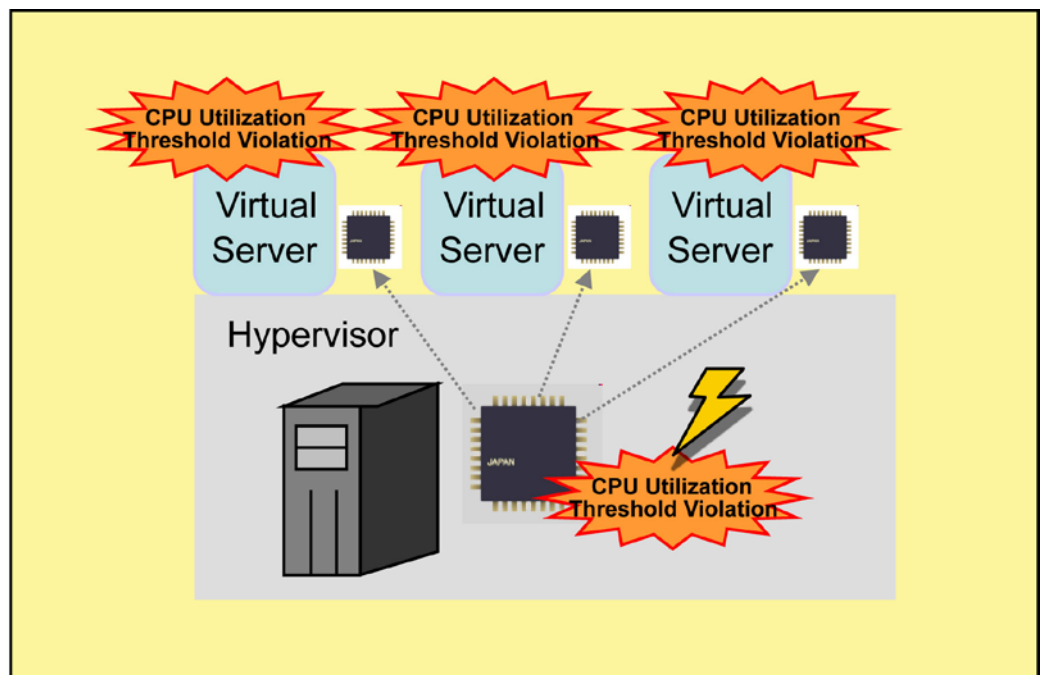
This scenario assumes the IT system comprises virtual servers, Fibre Channel switches and storage devices. These storage devices are providing logical volumes to the host (with a hypervisor) through a Fibre Channel SAN.

Some servers have slowed down. To begin diagnosing the problem, the administrator would first need to know whether these servers are virtual or physical machines. In the case of virtual servers, the administrators would then verify whether or not they are running on the same host (hypervisor). In order to check this type of dependency, the administrator must refer to any virtual server management software as well as the topology mapping information to reveal the connections among the hypervisor, Fibre Channel switches and storage devices.

Root Cause Analysis automates the manual processes of checking both topology mapping information and performance status that are updated in near real time. As a result, Root Cause Analysis can more accurately analyze the performance of both virtual servers and their host (hypervisor).

Root Cause Analysis Result: In this case, the root cause would be identified as CPU Overload on the host because all of the CPU utilization values of virtual servers on that host exceed their thresholds. (See Figure 7.)

Figure 7. Virtual Server Slow Down Caused by CPU Overload on the Host



Root Cause Analysis Topology View

Hitachi IT Operations Analyzer software includes the unique patent-pending Topological List View feature to provide administrators with a simplified and universal end-to-end display of all logical connections and dependencies through an intuitive, unified graphical user interface (GUI). One of the elements of the Topological List View is the Root Cause Analysis Topology View.

Figure 8 shows a view based on the Cable Disconnection scenario explained earlier in this application brief. The root cause is identified as a cable disconnection at the Fibre Channel storage port because IT Operations Analyzer has detected a Link Down alert for both the storage device and the Fibre Channel switch. The root cause node is displayed on the screen on a red background adjacent to the yellow thunderbolt symbol. These easy-to-understand icons demonstrate yet another way the IT Operations Analyzer is designed to quickly guide the administrator to the root cause of problems within even complex IT environments.

Figure 8. GUI for Cable Disconnection between Fibre Channel Switch and Storage Device

The screenshot displays the Hitachi IT Operations Analyzer interface. The main window shows a network diagram with nodes for Backbone, LAN, Computer, SAN, and Storage. A red lightning bolt icon indicates a detected fault at the Storage node. The right-hand pane shows 'Detected Faults: #1 (AVL-d) Root Cause' with a red background and a yellow lightning bolt icon. Below this, 'Events Detected' lists several events related to the storage device and its connections.

State	Descr.	Date/Time	Category	So...	Gr...	Device
Not Ack	The...	02/23/2009 16:54:52	Performance	W...	Acco...	Computer
Not Ack	The...	02/23/2009 16:50:52	Performance	W...	Acco...	Computer
Not Ack	The...	02/23/2009 16:49:23	Status	D...	Acco...	Computer
Not Ack	The...	02/23/2009 16:49:23	Status	D...	Acco...	Computer
Not Ack	The...	02/23/2009 16:48:56	Performance	W...	Acco...	Computer
Not Ack	The...	02/23/2009 16:48:56	Performance	W...	Acco...	Computer
Not Ack	The...	02/23/2009 16:48:56	Performance	W...	Acco...	Computer
Not Ack	The...	02/23/2009 16:48:55	Performance	W...	Acco...	Computer
Not Ack	The...	02/23/2009 16:48:55	Performance	W...	Acco...	Computer
Not Ack	The...	02/23/2009 16:48:55	Performance	W...	Acco...	Computer

Conclusion

Hitachi IT Operations Analyzer software uniquely delivers Root Cause Analysis functionality to automate and simplify fault identification for up to 750 server, network and storage nodes. Ideal for the administrator with limited resources to effectively monitor the data center, Root Cause Analysis helps reduce the steps and complexity of finding the root of the problem.

Using patent-pending technology to calculate the probable cause of an event, Root Cause Analysis provides the administrator with a quick path to fault identification, which can reduce the mean time to diagnose the problem. The faster the root cause can be identified and fixed, the sooner IT operations can resume, which helps to minimize system downtime and costs, and improve administrative efficiencies.

Meeting both availability and performance expectations across the data center is often vital to maintaining business commitments. Event alerts notify administrators of failures or negative changes, such as hardware faults, performance bottlenecks of networked storage, slowed server response and other infringements of normal IT activity. Finding the causal event can be time-consuming and difficult if the administrator does not have knowledge of the system configurations and dependencies.

Unlike other management software for midsized businesses, Root Cause Analysis unifies IT system configuration information and topology mapping details to correlate events and identify root cause nodes within complex IT system failures.

Hitachi Data Systems is a world leader in reliability, quality and innovation. With IT Operations Analyzer, midsized businesses are able to leverage Hitachi strengths in storage and systems management for smart, simple data center performance.

© Hitachi Data Systems Corporation

Corporate Headquarters

750 Central Expressway
Santa Clara, California 95050-2627 USA
www.hds.com

Regional Contact Information

Americas: +1 408 970 1000 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries.

All other trademarks, service marks and company names in this document or website are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems Corporation.

© Hitachi Data Systems Corporation 2010. All Rights Reserved. DS-332-B DG April 2010