WHITE PAPER

BROADCOM

Interpreting Events with Intelligence to Find Root Cause

DX NetOps Spectrum Event Correlation and Root Cause Analysis

Table of Contents

Executive Summary

Challenge

Today's complex IT infrastructures are dynamic, multi vendor engines made of frequently changing components and technologies. The complexity of the infrastructure and the continual changes caused by business demands often lead to faults within the infrastructure. A fault in a single device can have a ripple effect that causes performance and availability problems for many users. The ripple effect also makes it difficult to pinpoint the root cause of the fault.

To remain relevant and competitive in the marketplace, companies must minimize outages and performance degradations and this requires effective performance and availability management. Unfortunately, many management tools are not adaptive and have not kept pace with the dynamics of real-time, on-demand IT. Other tools often used are niche tools that manage only a portion of the infrastructure, making them largely ineffective in large, interconnected environments.

Opportunity

DX NetOps Spectrum is an infrastructure management solution that provides integrated service, fault and configuration functionality for modeling, monitoring and reporting across multiple network device types and technologies. Using a "trust but verify" methodology, DX NetOps Spectrum provides an automated and intelligent management approach for your particular environment—whether you are as a service provider or an enterprise.

DX NetOps Spectrum leverages three types of problem solving to comprehensively manage your infrastructure:

- Model-based Inductive Modeling Technology (IMT)
- Rules-based Event Management System (EMS)
- Policy-based Condition Correlation Technology (CCT)

Benefits

Broadcom's model-, rules- and policy-based analytics understand relationships between IT infrastructure assets and the users and services they are designed to support. This insight has enabled DX NetOps Spectrum to deliver real benefit to customers. A large service provider realized a 70% reduction in downtime while resolving 90% of availability or performance problems from a central location. Patented root cause analysis has been able to reduce the number of alarms by several orders of magnitude while significantly reducing Mean-Time-to-Repair (MTTR).

The Broadcom integrated approach to fault and performance management has enabled enterprise, government and service provider organizations around the world to achieve reliability, efficiency and effectiveness in managing IT infrastructures as a business service.

Challenge: A Complex Problem in Need of a Solution

IT infrastructure management is an intensive undertaking with significant resource requirements. Most employees in an organization expect the infrastructure to work, not thinking of it as a dynamic, multi-vendor engine made up of frequently changing components and technologies. In fact, the complexity and dynamics of today's real-time, on-demand IT architectures present many opportunities for inconsistencies and failures. Invariably, the infrastructure will slow down or fail and when it does, tools are required to quickly pinpoint the root cause, suppress all symptomatic faults, prioritize based on business impact and accelerate service restoration.

The Infrastructure is the Business

The IT infrastructure is a collection of interdependent components including computing systems, networking devices, databases and applications. Within the set of infrastructure components are multiple versions of many vendors' products connected over a variety of networking technologies. In addition, each business environment is different from the next—there is no configuration or standard set of components that makes up an infrastructure.

There also is constant change in devices, firmware versions, operating systems, networking technologies, development technologies and tools. But this dynamic and complex infrastructure serves an important purpose; the infrastructure IS the business. Either the infrastructure works and evolves or the organization is out of business. Companies must also evolve their people, processes and management tools for greater efficiency or fall competitively behind.

To ensure the performance and availability of the infrastructure, most companies employ a dual-approach method consisting of:

- Highly available, fault-tolerant, load balancing designs for infrastructure devices and communication paths.
- A network management solution to ensure reliable operation.

High-availability environments further complicate the job of management solutions. The management solution must understand the load balancing capacity, be able to track primary and fault-tolerant backup paths, and understand when redundant systems are active.

The investment in the management solution is as important as the investment in the infrastructure itself. The solution must be broad, deep, integrated and intelligent to perform its intended function. The infrastructure is not static and the solution will need to embrace change while delivering an end-to-end integrated view of performance and availability across infrastructure silos. Unfortunately, many management tools are not adaptive and do not keep pace with the dynamics of IT reality.

The Need to Be Proactive

To remain relevant and competitive in the marketplace, companies must minimize outages and performance degradations. In order to do this, the individual or groups responsible for the care of the infrastructure (e.g., IT, Engineering, Operations) must be proactively notified of problems. There are many tools that monitor the availability and performance of infrastructure components and the business applications that rely on them.

Many of these tools simply identify that a problem exists and notify technicians of a problem after it has happened. They often give visibility into only a small slice of the technologies under management and have no ability to understand how the various components relate to each other. This is not enough. It is important that the management solution act as an early warning system to help avoid downtime and service level agreement (SLA) violations. After a problem has occurred is too late—users are dissatisfied and SLA penalties have been levied.

Before the true task of troubleshooting can begin, the troubleshooter has to isolate the problem. Simply knowing there is a problem and collecting all the problems on one screen is not enough. Troubleshooters need to know where the problem is (and where the problem is not) to effectively triage the issue. If multiple problems are happening simultaneously, issues should be automatically prioritized based on the criticality of the impacted service.

It is far too costly to rely on human intervention to determine the root cause of problems and it is also far too costly to sift through an unending storm of symptomatic problems. Knowing the root cause allows an organization to efficiently get problems fixed without wasting time pursuing symptomatic problems.

The Importance of Understanding Business Impact

The best management solutions will not only be able to identify problems, but also identify impacted services, assets and users. For the business, understanding impact is as important as understanding the root cause. When outages or performance degradations occur, people cannot do their jobs effectively, resulting in lower productivity and reduced efficiency.

Sometimes the products or services provided by the company to their customers are affected, resulting in lost revenue, SLA penalties, lost customers and even damaged brand reputation that can take years to repair. Knowing impact allows an organization to prioritize response efforts to fix what matters most, first.

Opportunity: Event Correlation and Root Cause Analysis—A Three-Pronged Approach for DX NetOps Spectrum Network Fault Management

Root Cause Analysis (RCA) can be defined as the act of interpreting a set of symptoms and events and pinpointing the source of the problem. A single problem often results in multiple events across the infrastructure. Events are typically local to a source, and without proper context do not always help with RCA because they are only symptoms of the problem. Many components provide events and events come in many forms: SNMP traps, syslog messages, application log file entries, TL1 events, ASCII streams, etc.

Many sophisticated management systems, including DX NetOps Spectrum Network Fault Management, even generate events based on proactive polling of component status to indicate parameter based threshold violations, response time measurement threshold violations, etc. Often, correlation of events is required to determine if an actionable condition or problem exists but correlation is almost always required to isolate problems, identify impacted assets and services and suppress symptomatic events.

Management software applications efficiently performing RCA should raise an alarm for the root condition and should minimize other events resulting from the same root condition to generate an alarm.

One service provider experienced a situation where they were receiving 500,000 daily problem notifications from their management tool. Clearly, no person or team of people could keep up with that many events. DX NetOps Spectrum root cause analysis technology helped this service provider reduce the number of daily problem notifications to 200 actual alarm conditions, while also automatically prioritizing issues based on impact. In this environment, 500,000 symptoms had only 200 causes. Average time to find and fix a problem was reduced from over four hours to less than five minutes.

Effective RCA must:

- Understand the relationship between information within the infrastructure and the services, assets and users that depend on that information
- Be proactive in its monitoring and not just rely on event streams
- Distinguish between a plethora of events and meaningful alarms
- Scale and adapt to the requirements of growing and dynamic infrastructures
- Work across multiple-vendor and multiple-technology environments
- Allow for extensions and customization

DX NetOps Spectrum employs multiple techniques working cooperatively to deliver its event correlation and root cause analysis capabilities. These include Inductive Modeling Technology (IMT), Event Management System (EMS) and Condition Correlation Technology (CCT). Each of these techniques is employed to diagnose a diverse and often unpredictable set of problems.

Inductive Modeling Technology

The core of the DX NetOps Spectrum RCA solution is its patented Inductive Modeling Technology (IMT). IMT uses an object-oriented modeling paradigm with model-based reasoning analytics. DX NetOps Spectrum most often uses IMT for physical and logical topology analysis, as the software can automatically map topological relationships through its efficient Auto Discovery engine. The models created are software representations of a real-world physical or logical device. These software models are in direct communication with their real-world counterparts, enabling DX NetOps Spectrum root cause analysis to not only listen, but proactively query for health status or additional diagnostic information. Models are described by their attributes, behaviors, relationship to other models and algorithmic intelligence.

Intelligent analysis is enabled through the collaboration of models in a system. Collaboration includes the ability to exchange information and initiate processing between models within the modeling system. A model making a request to another model may in turn trigger that model to make requests on other models, and so on.

Relationships between models provide a context for collaboration. Collaboration between models enables:

- Correlation of the symptoms
- Suppression of unnecessary/symptomatic alarms
- Impact analysis

With DX NetOps Spectrum, a model is the software representation of a real-world managed device, or a component of that managed element. This representation allows DX NetOps Spectrum to not only investigate and query an individual element within the network, but also provides the means to establish relationships between elements in order to recognize them as part of a larger system.

By understanding the relationship between elements and the conditions of related managed elements, root cause analysis is simplified and problems are identified more quickly. Root cause and impact are determined through IMT's ability to both listen and talk to the infrastructure.

IMT is a very powerful analytical system and can be applied to many problem domains. A more in-depth discussion of IMT in action will be covered later in the paper.

A simple example of IMT in action can be demonstrated by a network router port transition from UP to DOWN. If a port model receives a LINK DOWN trap, it has intelligence to react by performing a status query to determine if the port is actually down. If it is in fact DOWN, then the system of models will be consulted to determine if the port has lower layer sub-interfaces. If any of the lower layer sub-interfaces are also DOWN, only the condition of the original port DOWN will be raised as an alarm. An application of this example can be described by several Frame Relay Data Link Control Identifiers (DLCIs) transitioning to INACTIVE. If the Frame Relay port is DOWN, IMT will suppress the symptomatic DLCI INACTIVE conditions and raise an alarm on the Frame Relay port model. Additionally, when the port transitions to DOWN, IMT will query the status of the connected Network Elements (NEs) and if those are also DOWN, those conditions will be considered symptomatic of the port DOWN, will be suppressed, and will be identified as impacted by the port DOWN alarm.

Event Management System

There are times when the only source of management information is through event streams local to a specific source. There may be no way to talk to the managed element, but only a way to listen to it. Any one event may or may not be a significant occurrence but, in the context of other events or information, may be an actionable condition.

Event Rules in DX NetOps Spectrum's Event Management System (EMS) provide a comprehensive decision-making system to indicate how events should be processed. Event Rules can be applied to look for a series of events to occur on a model in a certain pattern or within a specific time frame or within certain data value ranges. Event Rules can be used to generate other events or even alarms. If events occur such that the preconditions of a rule are met, another event may be generated allowing cascading events, or the event can be logged for later reporting/troubleshooting purposes, or it can be promoted into an actionable alarm.

DX NetOps Spectrum provides a series of customizable meta Event Rule types that form the basis of the EMS. These rule types are building blocks that can be used individually or cooperatively to effect an alarm on the most simple or sophisticated event-oriented scenarios. The rules engine allows for the correlation of event frequency/duration, event sequence and event persistence/ coincidence. More than 80% of rule conditions fall into one of the following three event types—frequency/duration, sequence or coincidence. Keep in mind that Event Rules can be run against the aforementioned models to avoid the need to constantly re-write rules to reflect infrastructure move/add/change activity. The Event Rule types are highlighted below, followed by usage examples later in the paper:

- **Event Pair (Event Coincidence)**. This rule is used when you expect certain events to happen in pairs. If the second event in a series does not occur, this may indicate a problem. Event rules created using the Event Pair rule type generate a new event when an event occurs without its paired event. It is possible for other events to happen between the specified event pair without affecting this event rule.
- **Event Rate Counter (Event Frequency)**. This rule type is used to generate a new event based on events at a specified rate in a specified time span. A few events of a certain type may be tolerated, but once the number of these events reaches a certain threshold within a specified time period, notification is

required. No additional events will be generated as long as the rate stays at or above the threshold. If the rate drops below the threshold and then subsequently rises above the threshold, another event will be generated.

- **Event Rate Window (Event Frequency)**. This rule type is used to generate a new event when a number of the same events are generated in a specified time period. This rule type is similar to the Event Rate Counter. The Event Rate Counter type is best suited for detecting a long, sustained burst of events. The Event Rate Window type is best suited for accurately detecting shorter bursts of events. It monitors an event that is not significant if it happens occasionally, but is significant if it happens frequently. If the event occurs above a certain rate, then another event is generated. No additional events will be generated as long as the rate stays at or above the threshold. If the rate drops below the threshold and then subsequently rises above the threshold, another event will be generated.
- **Event Sequence (Event Sequence)**. This rule type is used to identify a particular order of sequence in events that might be significant in your IT infrastructure. This sequence can include any number and any type of event. When the sequence is detected in the given period of time, a new event is generated.
- **Event Combo (Event Coincidence)**. This rule type is used to identify a certain combination of events happening in any order. The combination can include any number and type of event. When the combination is detected within a given time period, a new event is generated.
- **Event Condition (Event Coincidence)**. This rule type is used to generate an event based on a conditional expression. In keeping with the DX NetOps Spectrum "trust but verify" methodology, a series of conditional expressions can be listed within the event rule and the first expression that is found to be true will generate the event. Rules can be constructed to provide correlation through a combination of evaluating event data with IMT model data. For example, if a trap is received notifying the management system of memory buffer overload, to validate that an alarm condition has occurred, an Event Condition rule can initiate a request to the device to check actual memory utilization.

DX NetOps Spectrum implements a number of Event Rules out-of-the-box by applying one or more rule types to event streams. Users can create or customize Event Rules using any of the rule types. Further implementation of Event Rules using the Event Management System will be discussed later in this paper.

Condition Correlation

In order to perform more complex user-defined or user-controlled correlations, a broader set of capabilities is required. DX NetOps Spectrum policy-based Condition Correlation Technology enables:

- Creation of correlation policies
- Creation of correlation domains
- Correlation of seemingly disparate event streams or conditions
- Correlation across sets of managed elements
- Correlation within managed domains
- Correlation across sets of managed domains
- Correlation of component conditions as they map to higher order concepts such as business services or customer access

In order to understand these capabilities, the terminology is described in this context:

- **Conditions**. A condition is similar to state. Condition can be set by an event and cleared by an event. It is also possible to have an event set a condition but require a user-based action to clear the condition. The condition exists from the time it is set until the time it is cleared. A very simple example of a condition is "port down" condition. The port down condition will exist for a particular interface from the time the LINK DOWN trap or set event (such as a failed status poll) is received until the time the LINK UP trap or clear event (such as a successful status poll) is received. A number of conditions that may be of use for establishing domain level correlations are defined out-of-the-box and more can be added by the user.
- **Seemingly Disparate Conditions**. Many devices in an IT infrastructure provide a specific function. The device level function is often without context as it relates to the functions of other devices. Most managed devices can emit event streams but those event streams are local to each component. A simple example is when a response time test identifies a result exceeding a threshold. At the same time, an event may identify a condition of a router port exceeding a transmit bandwidth threshold. These conditions are seemingly disparate, as they are created independently and without context or knowledge of each other. In reality the two are quite related.
- **Rule Patterns**. Rule Patterns are used to associate conditions when specific criteria are met. A simple example is a "port down" condition caused by a "board pulled" condition—but only if the port's slot number is equal to the slot number of the board that's been pulled. Figure A illustrates this rule pattern. The result of applying a rule pattern can be the creation of an actionable alarm or the suppression of symptomatic alarms.
- **Correlation Policy**. Multiple rule patterns can be bundled or grouped into a Correlation Policy. Correlation Policies can then be applied to a Correlation Domain. For example, a set of rule patterns applicable to OSPF correlation can be created and labeled the OSPF Correlation Policy. This policy can be applied to each Correlation Domain as defined by each autonomous OSPF region and the supporting routers in that region.

RULE PATTERN

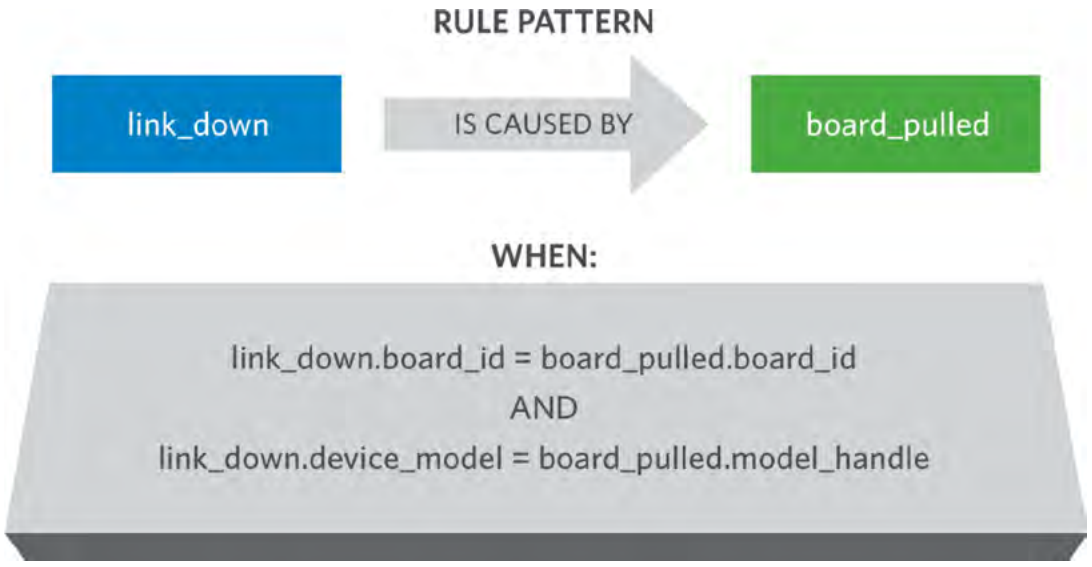

Figure A: Rule patterns determine the sequence of investigation that will result in either the creation of an alarm or the suppression of symptomatic alarms.

- **Correlation Domain**. A Correlation Domain is used to both define and limit the scope of one or more Correlation Policies. A Correlation Domain can be applied to a specific service. For example, in the cable broadband environment, a return path monitoring system may detect a return path failure in a certain geographic area. This "return path failure" condition is causing subscribers' high speed cable modems to become unreachable and causing Video on Demand (VoD) pay-per-view streams to fail. The knowledge that the return path failure, the modem problems and the failed video streams are all in the same correlation domain is essential to correlating the events and ultimately identifying the root cause. However, it is also important to have the ability to distinguish that a "return path failure" condition occurring in one correlation domain (Philadelphia, PA) is not correlated with VoD stream failure conditions occurring in a different correlation domain (Portsmouth, NH).

Condition-based correlations are very powerful and provide a mechanism to develop Correlation Policies and apply them to Correlation Domains. When applied to Business Service Management, Correlation Policies can be likened to metrics of an SLA and Correlation Domains can be likened to service, user or geographical groupings.

There are times when the only way to infer a causal relationship between two or more seemingly disparate conditions is when those conditions occur in a common Correlation Domain. These mechanisms are necessary when causal relationships cannot be discovered through interrogations or receipt of events to/from the infrastructure components.

Use Case Scenarios

Out-of-the-box, DX NetOps Spectrum addresses a wide range of different scenarios where it can perform root cause analysis. This section provides specific scenarios where the techniques described in the previous section are employed to determine root cause and impact analysis. The detail will be limited to the basic processing for the sake of simplicity and brevity. Also for the purpose of the discussion and figures, the following table is provided showing the color of alarms that are associated with the icon status of models at any given time.

MODEL STATUS COLORS

STATUS OF MODEL	COLOR OF ALARM
Normal Operation	GREEN
Critical Fault	RED
Major Fault	ORANGE
Minor Fault	YELLOW
Unknown or Suppressed	GRAY
Down for Maintenance	BROWN
Intial Condititon	BLUE

Figure B: Alarms are color coded reflecting the model status.

Inference and Inductive Modeling Technology

Communication outages are often described as "black-outs" or "hard faults." With these types of faults, one or more communication paths are degraded to the point that communication is no longer possible. The cause could be broken copper/fiber cables/connections, improperly configured routers or switches, hardware failures, severe performance problems and security attacks. Often the difficulty with these hard communication failures is that there is limited information available to the management system, as it is unable to exchange information with one or more managed elements.

The DX NetOps Spectrum system of sophisticated models, relationships and behaviors available through IMT allows it to infer the fault and impact. IMT inference algorithms are also called inference handlers and a set of inference handlers designed for a purpose is labeled as an intelligence circuit or simply "intelligence". This section will outline how intelligence is applied to isolate communication outages.

Building the Model

The accurate representation of the infrastructure through the modeling system is the key to being able to determine the fault and the impact of the fault. DX NetOps Spectrum has specific solutions for discovering multi-path networks over a variety of technologies supporting different architectures. It offers support for meshed and redundant, physical and logical topologies based on ATM, Ethernet, Frame Relay, HSRP, ISDN, ISL, MPLS, Multicast, PPP, VoIP, VPN, VLAN and 802.11 wireless environments—even legacy technologies such as FDDI and Token Ring. Its modeling is extremely extensible and can be used to model OSI Layers 1-7 in a communication infrastructure.

DX NetOps Spectrum provides four different methods for building the physical and logical topology model and interdependent connectivity for any given infrastructure:

- The AutoDiscovery functionality can be used to automatically interrogate the managed infrastructure about its physical and logical relationships. AutoDiscovery works in two distinct phases (although there are many different stages within each phase that are not covered here) and dynamically.

  When initiated, AutoDiscovery automatically discovers the devices that exist in the infrastructure. This provides an inventory of devices that could be managed.

  The second phase is Modeling. AutoDiscovery uses management and discovery protocols to query the devices it has found to gain information that will be used to determine the Layer 2 and Layer 3 connectivity between managed devices. For example, AutoDiscovery uses SNMP to examine route tables, bridge/switch tables and interface tables, but also uses traffic analysis and vendor proprietary discovery protocols such as Cisco Discovery Protocol (CDP).

  AutoDiscovery is a very thorough, accurate and automated mechanism to build the infrastructure model.

- Alternately, the DX NetOps Spectrum Modeling Gateway can be used to import a description of the entire infrastructure's components, as well as physical and logical connectivity information from external sources, such as provisioning systems, network topology databases or configuration management databases (CMDBs)

- The Command line interface or programmatic APIs can also be used to build a custom integration or application to import information from external sources.

- DX NetOps Spectrum's OneClick Network Console can be used to quickly point and click to manually build the model.

DX NetOps Spectrum allows a single managed element to be logically broken up into any number of sub-models. This collection of models and the relationships between them is often referred to as the semantic data model for that device. Thus, a typical semantic data model for a device may include a chassis model with board models related to the chassis. Associated to the board models would be physical interface models. Each physical interface model may have a set of sub-interface models associated below them.

DX NetOps Spectrum has a set of well-defined associations that define how different semantic data model sets act with one another. When the software determines the connectivity between two devices, a relationship is established between the two ports that form the link between them, as well as the relationships that form between device models and to the corresponding interface and port models of other devices. This is depicted in Figure C.

MODELING DEVICE AND INTERFACE LEVEL CONNECTIVITY

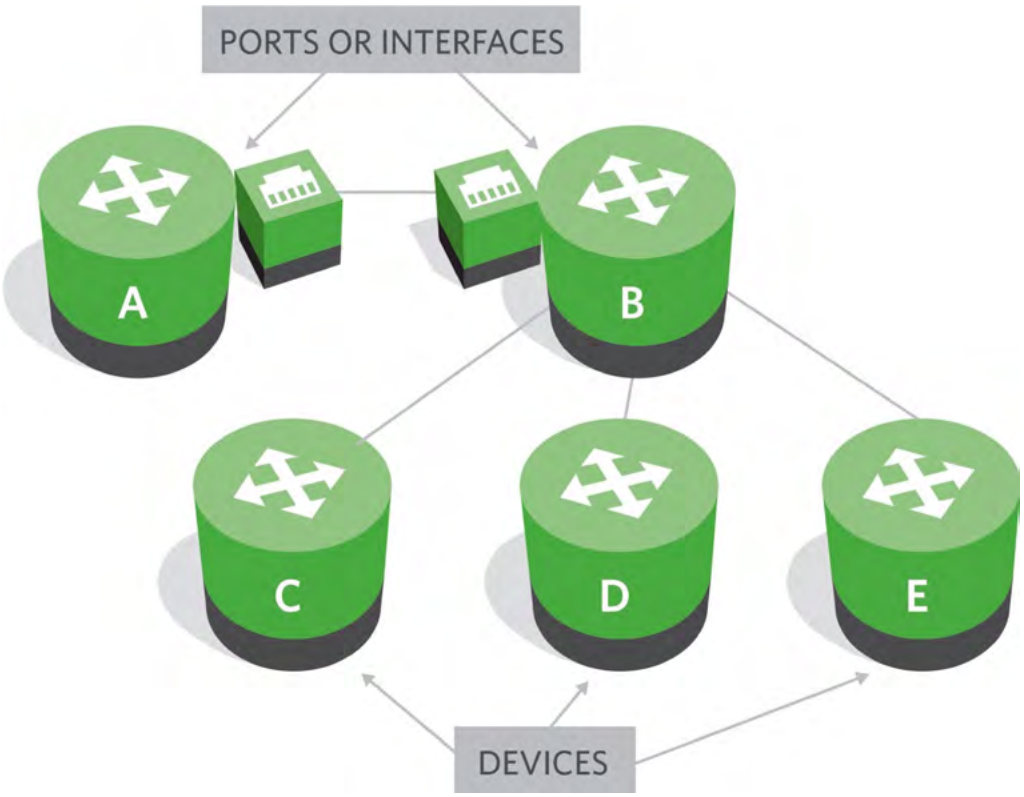

Figure C: Modeled relationships between devices are reflected in a relationship between the ports and interface that connect them.

When Does the Analysis Begin?

DX NetOps Spectrum can begin to solve a problem proactively upon receipt of a single event. Many problems share the same set of symptoms and only through further analysis can the root cause be determined. For communication outages, the analysis is triggered when a model recognizes a communication failure. Failed polling, traps, events, performance threshold violations or lack of response can all lead to this recognition. DX NetOps Spectrum validates the communication failures through retries, alternative protocols and alternative path checking as part of its "trust but verify" methodology.

DX NetOps Spectrum will refer to the model that triggered the intelligence as the initiator model, although more than one model can trigger the intelligent validation procedures. The initiator model intelligence requests a list of other models that are directly connected to it. These connected models are referred to as the initiator model's neighbors.

THE INITIATOR MODEL AND NEIGHBORS

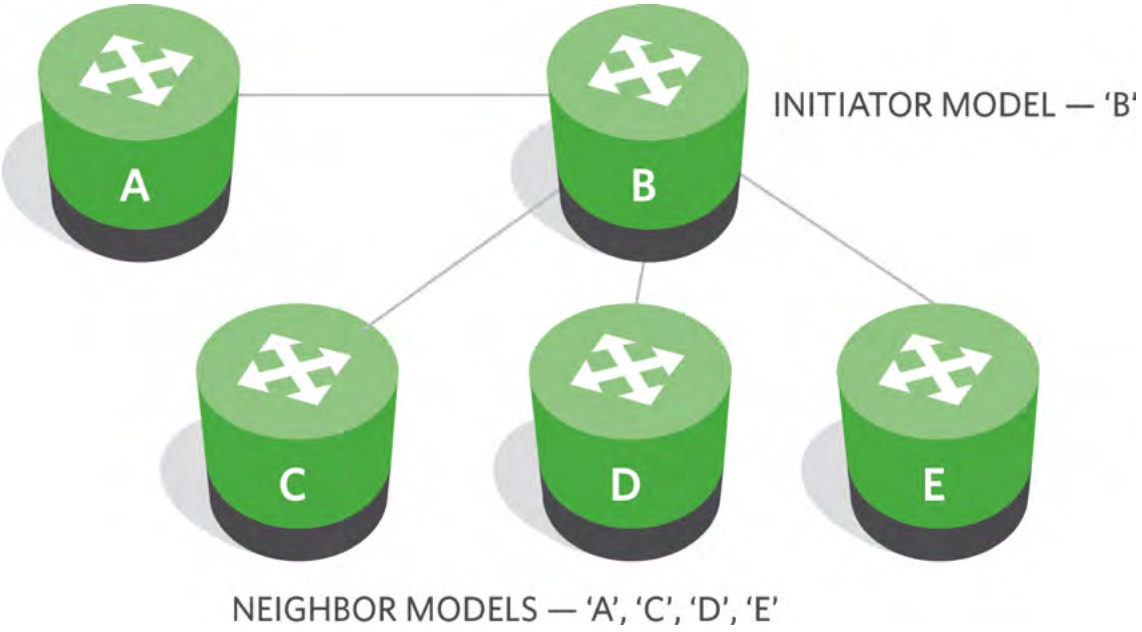

Figure D: The initiator model triggers the intelligent validation procedure. It requests a list of models that are directly connected and these are called the initiator model's neighbors.

With a list of neighbors determined, DX NetOps Spectrum directs each neighbor model to check its current status. This check is referred to as the "Are you OK?" check. "Are you OK" is a relative term, and a unique set of attributes related to performance and availability will vary from model to model based on the real-world capabilities of the device that the model is representing.

When a model is asked "Are you OK?", the model can initiate a variety of tests to verify its current operational status. For example, with most SNMP managed devices the check is typically a combination of SNMP requests but could be more involved by interrogating an Element Management System or as simple as an ICMP ping. A comprehensive check could include threshold performance calculations or execution of response time tests.

Each neighbor model returns an answer to "Are you OK?" and DX NetOps Spectrum then begins its analysis of the answers.

DETERMINING THE HEALTH OF NEIGHBORS

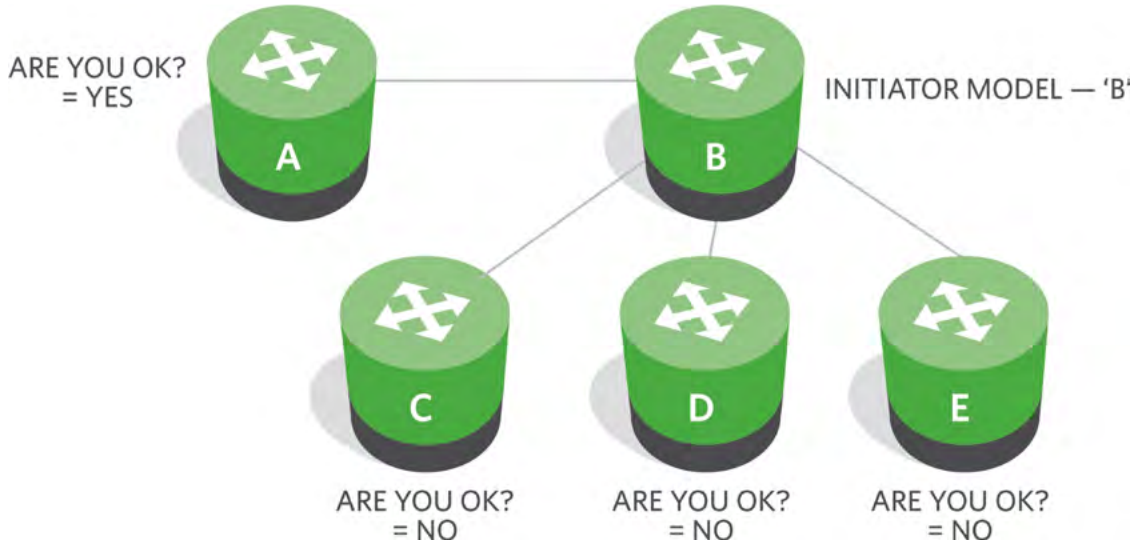

Figure E: Once the list of neighbors is established, the status of each neighbor is checked with the "Are you OK?" message.

Fault Isolation

If the initiator model has a neighbor that responds that it is "OK", (Figure F, Model A), then it can be inferred the problem lies between the unaffected neighbor and the affected initiator (Figure F, Model B). In this case, the initiator model that triggered the intelligence is a likely culprit for this particular infrastructure failure. The result? A critical alarm will be asserted on the initiator model and it is considered the "root cause" alarm.

FAULT ISOLATION IN PROGRESS

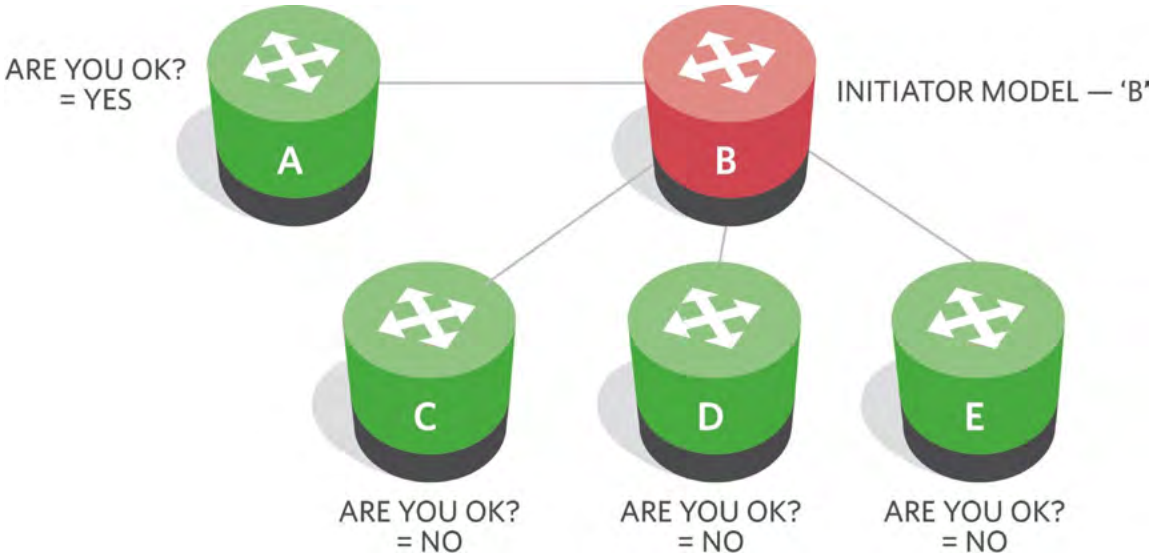

Figure F: The root cause alarm is established through a sequence of sharing status between models.

Alarm Suppression

As the analysis continues beyond isolating the device at fault (Figure G, Model B), the next step is to analyze the effects of the fault, the goal of which is intelligent alarm suppression. If a neighbor (Figure G, Models C, D or E) of the initiator model responds, "No, I am not OK", then this particular neighbor is considered to be affected by a failure that is occurring somewhere else. As a result, DX NetOps Spectrum will place these models into a suppressed condition (Grey Color) because any alarms from this device are symptomatic of a problem elsewhere.

FAULT ISOLATION COMPLETE

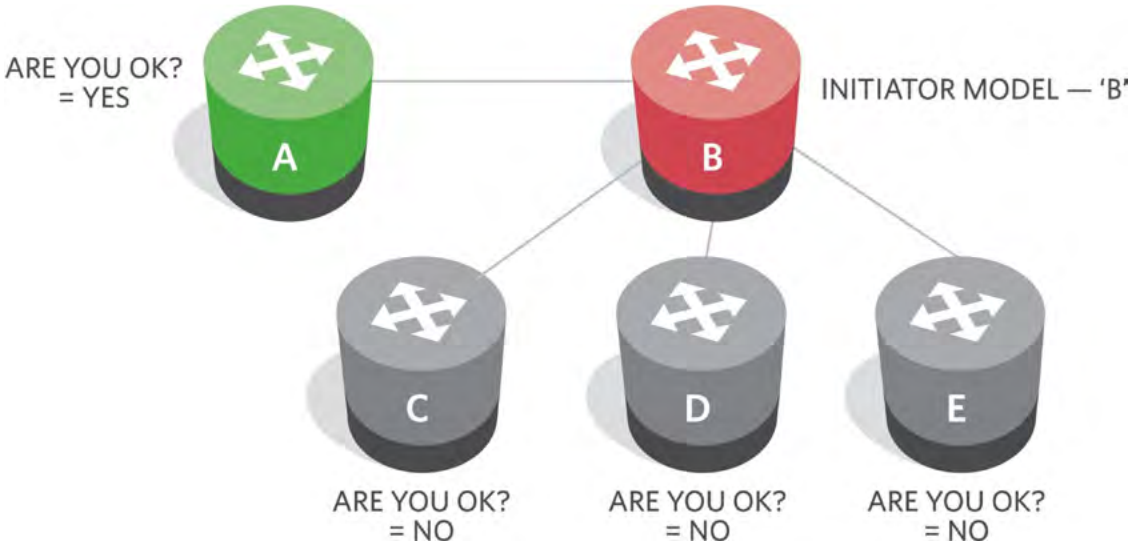

Figure G: Models that respond with a "No, I am not OK" status are put in a suppressed condition to suppress the alarms that are symptomatic of a problem elsewhere in the infrastructure.

Impact Analysis

DX NetOps Spectrum continues to analyze the total impact of the fault. It will analyze each Fault Domain, a Fault Domain being the collection of models with suppressed alarms that are affected by the same failure. These impacted models are linked to the root fault for presentation and analysis. The intelligence provides the impact measurement this fault is creating, by examining the models that are included within this Fault Domain and calculating an Impact Severity value. The ranking allows operators to quickly assess the relative impact of each fault and prioritize corrective actions.

Creatively Using the Event Management System

There are many applications of Event Rules that will allow higher order correlation of event streams. Event Rule processing is required for situations when the event stream is the only source of management information. For example, this situation can occur when DX NetOps Spectrum accepts event streams from devices and applications that it does not directly monitor, so that DX NetOps Spectrum can only listen—it cannot talk. DX NetOps Spectrum provides many out-of-the-box event rules, but also provides easy-to-use methods for creating new rules using one or more of the event rule types. This section highlights a couple of out-of-the-box event rules and also a few customer examples of event rule applications.

An Out-of-the-Box Event Pair Rule

DX NetOps Spectrum has the ability to interpret Cisco syslog messages as event streams. Each syslog message is generated on behalf of a managed switch or router and is directed to the model representing that managed element. One of the many Cisco syslog messages indicates a new configuration has been loaded into the router. The "Reload" message should always be followed by a "Restart" message, indicating the device has been restarted to adopt the newly loaded configuration. If not, a failure during reload is probable. An event rule based on the Event Pair rule type is used to raise an alarm with cause ERROR DURING ROUTER RELOAD if the restart message is not received within 15 minutes of the reload message. Figure H diagrams the events and timing.

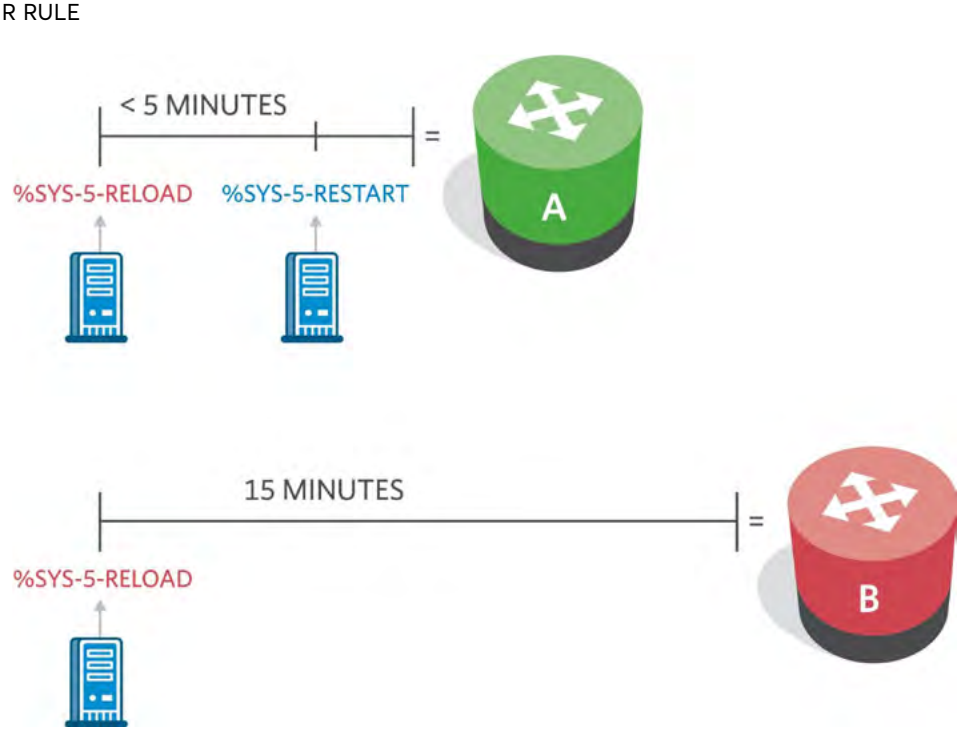

Figure H: This figure depicts an example of an event pair rule in operation. Reload messages indicating new router configurations should always be followed by restart messages to indicate the router has adopted the new configuration. An alarm is raised to indicate a probable failure if the restart message is not received within the expected time period.

Managing Security Events Using an Event Rate Counter Rule

DX NetOps Spectrum is often used to collect event feeds from many sources. Some customers send events from security devices such as intrusion detection systems and firewalls. These types of devices can generate millions of log file entries. One customer utilizes an Event Rate Counter rule to distinguish between sporadic client connection rejections and real security attacks. The rule was constructed to generate a critical alarm if 20 or more connection failures occurred in less than one minute. Figure I depicts this alarm scenario.

EVENT RATE COUNTER RULE

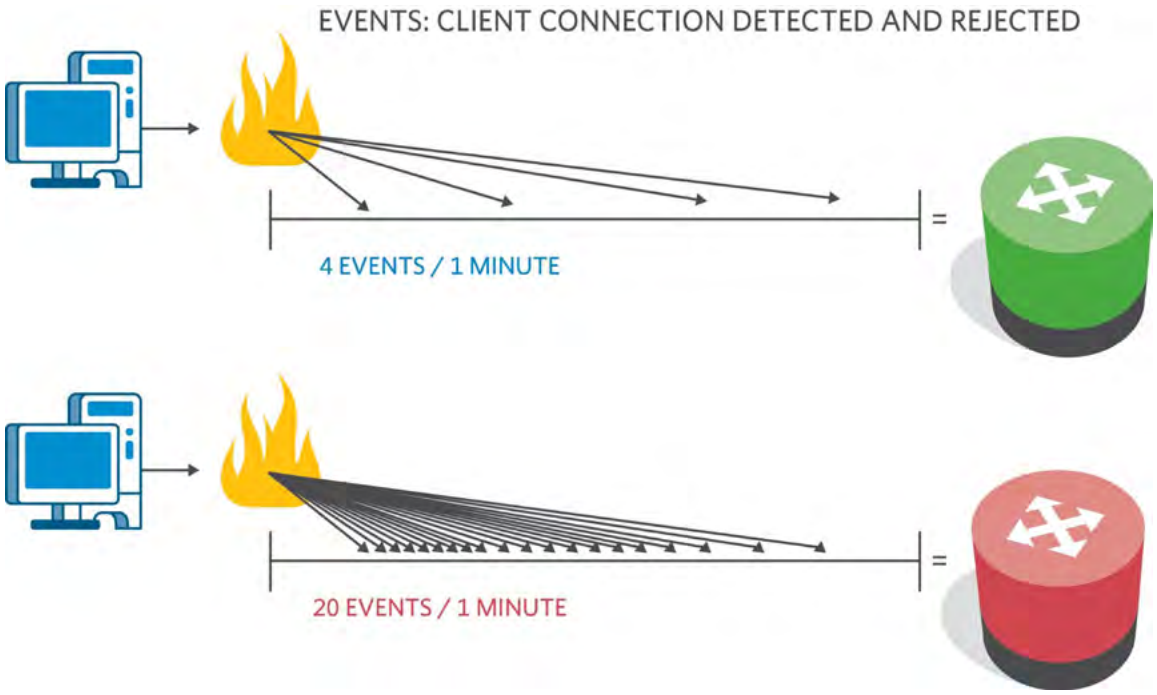

Figure I: This figure depicts an example of the Event Rate Counter rule. Events are often sent by other devices, such as intrusion detection systems, that can generate millions of events. To separate event "noise" from real events, this rule limits alarms to situations where more than 20 connection failures are logged within a minute.

Managing Server Memory Growth Using an Event Sequence Rule

A common problem with some applications is the inability to manage memory usage. There are applications that will take system memory and never give it back for other applications to reuse. When the application does not return the memory, and also no longer requires the memory, it is called a "memory leak." The result is that performance on the host machine will degrade and eventually cause the application to fail.

At one customer environment this problem regularly occurs on a Web Server application. The customer has a standard operating procedure to reboot the server once a week to compensate for the memory consumption. However, if the memory leak occurs too quickly, there is a deviation from normal behavior and the server needs to be rebooted before the scheduled maintenance window. The customer employs a combination of progressive thresholds with an Event Sequence rule to monitor for abnormal behavior. Monitoring was set to create events as the memory usage passed threshold points of 50%, 75% and 90%. If those threshold points are reached in a period of less than one week, an alarm is generated to provide notification to reboot the server prior to the scheduled maintenance window. Figure J depicts the fault scenario.

EVENT SEQUENCE RULE

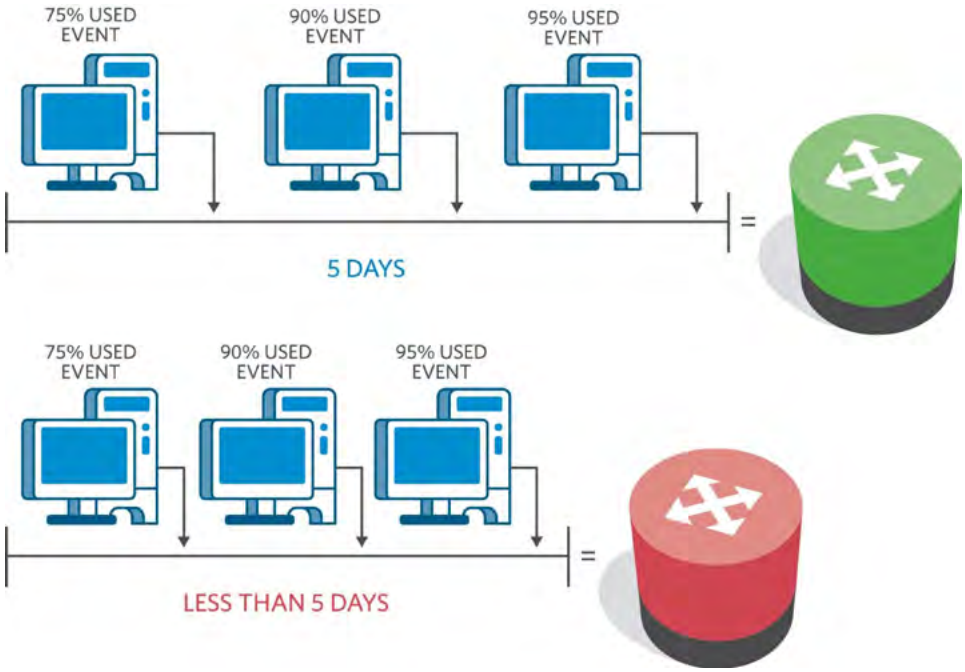

Figure J: Reaching threshold points is sometimes acceptable if the thresholds are reached over a period of time that will ensure they are reset by regular maintenance schedules before reaching critical levels. If they are reached sooner, however, they may cause outages. The Event Sequence rule will measure threshold attainment and generate an alarm only if required.

Condition Correlation Technology

There are many uses for policy-based Condition Correlation Technology (CCT). For example, consider the complexities of managing an IP network that provides VPN connectivity across an MPLS backbone with intra-area routing maintained by IS-IS and inter-area routing maintained by BGP. Any physical link or protocol failure could cause dozens of events from multiple devices. Without sophisticated correlation capability applied carefully, the network troubleshooters will spend most of their time chasing after symptoms, rather than fixing the root cause.

An Is-Is Routing Failure Example

A specific example experienced by one of our customers can be used to describe the power of Condition Correlation. The failure scenario and link outages are illustrated in Figure K. The situation occurs where a core router, labeled in the figure as R1, loses IS-IS adjacencies to all neighbors (labeled in the figure as R2, R3, R4). This also results in the BGP session with the route reflector (labeled in the figure as RR) being lost. This condition, if it persists, will result in routes aging out of R1 and adjacent edge routers R3 and R4. Eventually, the customer VPN sites serviced by these customer premise edge (CPE) routers will be unable to reach their peer sites (labeled in the figure as CPE1, CPE2, CPE3).

IMPACT OF AN IS-IS ROUTING FAILURE

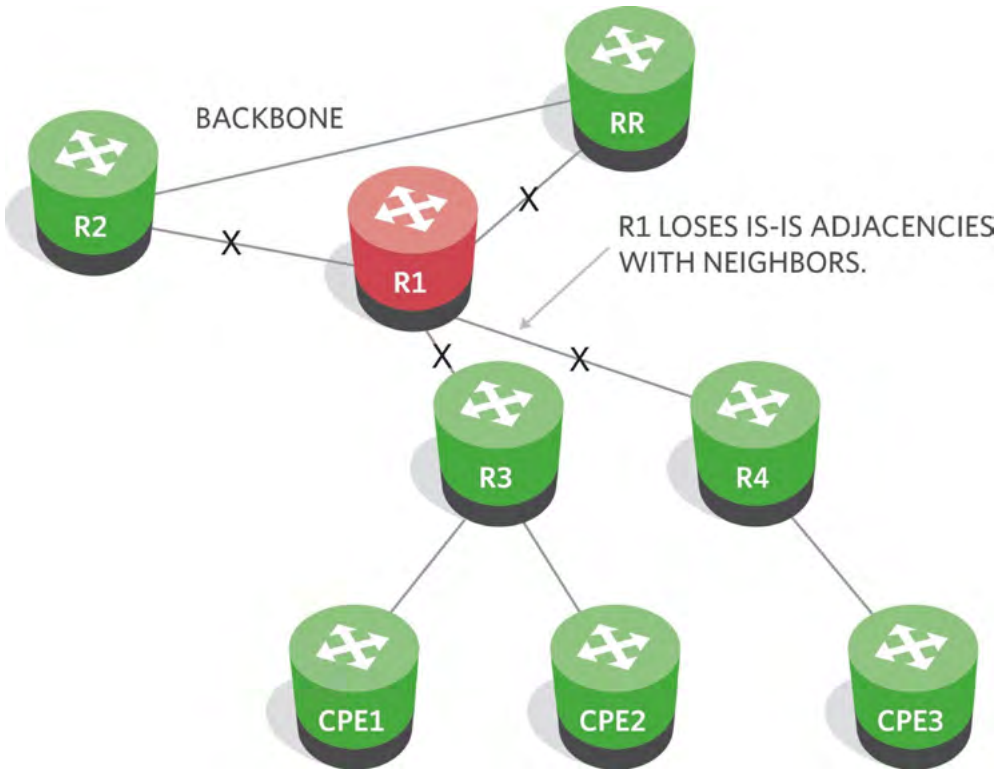

Figure K: This diagram shows how a failure in a core router can ripple through the network, causing numerous events and alarms, if not intelligently managed.

This failure causes a series of syslog error messages and traps to be generated by the routers. The messages and traps that would be received are outlined in Figure L.

SYSLOG ERROR MESSAGE AND TRAP SEQUENCE

Figure L: Error messages and traps cascade from a single core router failure.

SOURCE	TYPE	MESSAGE
R1	Syslog message	%CLNS050ADJCHANGE: ISIS: Adjancency to R2 (P0S5/0/0) Down, hold time expired
R1	Syslog message	%CLNS050ADJCHANGE: ISIS: Adjancency to R3 (P0S5/0/0) Down, hold time expired
R1	Syslog message	%CLNS050ADJCHANGE: ISIS: Adjancency to Rn (P0S5/0/0) Down, hold time expired
RR	Syslog message	%BGP-5-ADJCHANGE: neighbor R1 Down BGP Notification sent
RR	Syslog message	%BGP-3-NOTIFICATION: sent to neighbor R14/0 (hold time expired) 0 bytes
RR	Trap	BGP Backwards Transition trap, neighbor = R1
R2	Syslog message	%CLNS050ADJCHANGE: ISIS: Adjancency to R1 (P0S5/0/0) Down, hold time expired
R3	Syslog message	%CLNS050ADJCHANGE: ISIS: Adjancency to R1 (P0S5/0/0) Down, hold time expired
Rn	Syslog message	%CLNS050ADJCHANGE: ISIS: Adjancency to R1 (P0S5/0/0)

The root cause of all these messages is an IS-IS outage problem related to R1. For many management systems the operator would see each of these traps as seemingly disparate events on the alarm console. A trained operator or experienced troubleshooter may be able to deduce, after some careful thought, that an R1 routing problem is indicated. However, in a large environment these alarms will likely be interspersed with other alarms cluttering the console. Even if the operator were capable of making the correlation manually, there would be significant effort and time spent doing so. That time is directly related to costs, lower user satisfaction and lost revenue.

Using a combination of an Event Rule and Condition Correlation, a set of rule patterns can be applied to a Correlation Domain consisting of all core, label switch routers, enabling DX NetOps Spectrum to produce a single actionable alarm. This alarm will indicate that R1 has an IS-IS routing problem, and a network outage may result if this is not corrected. The seemingly disparate conditions that were correlated by the software, resulting in this alarm, will be displayed in the "symptoms" panel of the alarm console as follows:

- A local Event Rate Counter rule was used to define multiple 'IS-IS adjacency change' syslog messages reported by the same source as a routing problem for that source.
- A rule pattern was used to make an IS-IS adjacency lost event "caused by" an IS-IS routing problem when the neighbor of the adjacency lost event is equal to the source of the routing problem event.
- A rule pattern was used to make a BGP adjacency down event "caused by" an IS-IS routing problem when the neighbor of the adjacency down event is equal to the source of the routing problem event.
- A rule pattern was used to make a BGP backward transition trap event "caused by" an IS-IS routing problem when the neighbor of the backward transition event is equal to the source of the routing problem event.

Applying Condition Correlation to Service Correlation

It is common to have several services running over the same network. As an example, in the cable industry, telephone service (VoIP), Internet access (high speed data), video on demand (VoD) and digital cable are delivered over the same physical data network. Management of this network is quite a challenge. Inside the network (cable plant), the video transport equipment, video subscription services and the Cable Model Termination System (CMTS) all work together to put data on the cable network at the correct frequencies. Uncounted miles of cable along with thousands of amplifiers and power supplies must carry the signals to the homes of literally millions of subscribers.

With the flood of events and error messages that will be provided by the managed elements, the fact that there is a problem with the service will be obvious. The challenge is to translate all this data into root cause and service impact actionable information.

Service impact relevance goes beyond understanding what is impacted; it is also important to understand what is not impacted. It's possible for the video subscription service to fail to deliver VoD content to a single service area, and yet all other services to that area could be fine. Or, a return path problem in one area could cause Internet, VoIP and VoD services to fail and digital cable to degrade, yet analog cable could still function normally. In the case of a media cut in one area, the return path monitoring system and the head end controller would report return path and power problems in that area. The CMTS would provide the number of cable modems off-line for the node. The video transport system would generate errors for video subscriptions in that area. Lastly, any customer modems that are being managed will become lost to the management system.

DX NetOps Spectrum can make sense of the resulting deluge of events by using the service area of the seemingly disparate events as a factor in the Condition Correlation. If the service areas and services are appropriately modeled, Condition Correlation can be used to determine which services in which areas are affected and the root cause or causes.

SERVICE CORRELATION

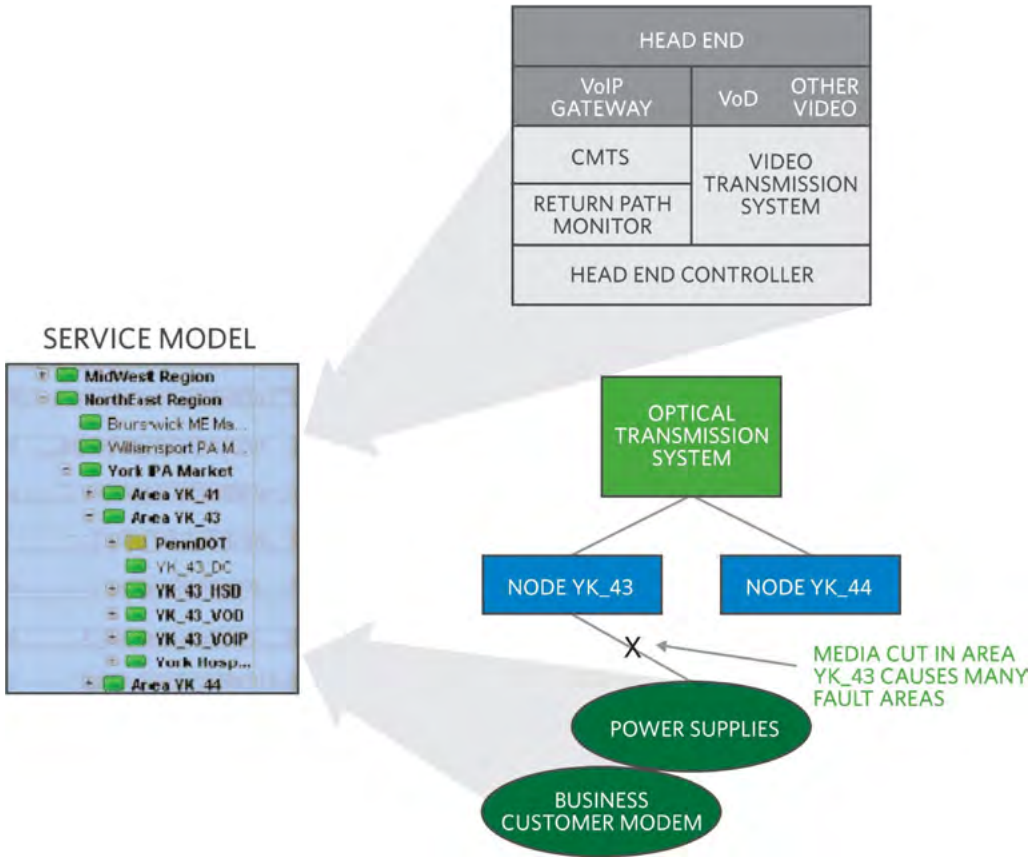

Figure M: Condition Correlation enables DX NetOps Spectrum to analyze a deluge of events to determine which service is impacted.

Benefits: DX NetOps Spectrum for High Performing Infrastructures

A high performing IT infrastructure is at the core of today's successful businesses. Whether your business is online retail, financial services or manufacturing, your infrastructure is essential. Keeping the infrastructure running, avoiding outages, quickly finding the causes of degradation and outages, and simplifying management and event data for your IT staff are all factors in maintaining high performance.

Patented Software Elevates DX NetOps Spectrum Capabilities

DX NetOps Spectrum provides intelligence, multiple methods and patented solutions to apply the best in event correlation and root cause analysis to your infrastructure. Event correlation is at the heart of root cause analysis. With large and complex infrastructures, events flood event logs and your IT staff can be overwhelmed by attempting to correlate events manually. The time wasted in this effort has direct effect on your business and its bottom line. DX NetOps Spectrum uses intelligence and event rules to separate true root causes from associated, symptomatic causes, thereby minimizing the amount of information and maximizing the quality of information your IT staff must address.

Benefits for Experienced Users and New Users

DX NetOps Spectrum provides out-of-the-box utilization, performance and response time thresholds that act as an early warning system when a problem is about to happen or when a service level guarantee is about to be violated. While these thresholds can be tuned for a specific customer environment, there is tremendous value in having these out-of-the-box thresholds. They enable DX NetOps Spectrum to deliver value on day one.

DX NetOps Spectrum makes it easy for experienced users to add their own thresholds and watches such that after a unique problem happens in the environment once, new watches can help predict or prevent it from happening again. Through Event Rules and combinations of Event Rules, even complex behaviors can be captured and managed.

Conclusion

Change is a constant, requiring any management system to be automated, adaptable, and extensible. The number of multi-vendor, multi-technology hardware and software elements in a typical IT environment exponentially increases the complexity of managing a real-time, on-demand IT infrastructure. DX NetOps Spectrum currently supports several thousand distinct means to automate root cause analysis across over hundreds of product families and device types from today's leading infrastructure vendors.

Knowing about a problem is no longer enough. Predicting and preventing problems, pinpointing their root cause, and prioritizing issues based on impact are requirements for today's management solutions. The number and variety of possible fault, performance and threshold problems means that no single approach to root cause analysis is suited for all scenarios. For this reason, DX NetOps Spectrum incorporates model-based IMT, rules-based EMS, and policy-based CCT to provide an integrated, intelligent approach to drive efficiency and effectiveness in managing IT infrastructure as a business service.

To learn more about DX NetOps, visit **www.broadcom.com/products/software/aiops/network-monitoring-software**

About Broadcom
Broadcom Inc. (NASDAQ: AVGO) is a global technology leader that designs, develops and supplies a broad range of semiconductor and infrastructure software solutions. Broadcom's category-leading product portfolio serves critical markets including data center, networking, enterprise software, broadband, wireless, storage and industrial. Our solutions include data center networking and storage, enterprise and mainframe software focused on automation, monitoring and security, smartphone components, telecoms and factory automation. For more information, go to www.broadcom.com.

0420