

## **EMC Storage Insight for Availability**

April, 2006



In the typical large IT shop, SAN management has become massively more complex in the past 3 years. In addition to a full boat of heterogeneous host environments, we see a range of HBAs and FC switch types, increasingly complex multi-path schema, all linking into multiple types of storage arrays. Along with this SAN element complexity comes equally complex fault management. Because of all of the “moving parts” in a SAN fabric, when a failure event transpires today, the typical storage team will spend hours or days sifting through the myriad of component alerts, trouble tickets or even phone calls to interpolate the root cause, and its impact on their storage users. This is a situation begging for improvement. The old manual approach of writing individual rules for every SAN component simply does not fly. It is too brittle, incomplete and non-scalable. The alternative is to create a true *model-based approach* to root cause analysis, to literally automate how the vast numbers of components in a SAN deployment report about their respective statuses, interconnections and relationships to the services that they enable. This is precisely the transition that has taken place with automated fault management software. It promises to radically reduce the amount of time that teams spend recovering from fault events by removing literally hundreds or thousands of discrete human-intensive correlation activities and providing superior real-time insight into the overall SAN’s state of functioning. EMC Smarts Storage Insight for Availability (EMC SI-A) is a perfect example of where this technology is going. We think EMC SI-A points the way for even deeper cross-domain capabilities that link the FC SAN to the overall IP network, as well as up the stack to include business impact analysis. In this profile, we will provide an introduction to this important new class of storage software and take a deeper look at EMC SI-A.

### **The Root of the Problem**

Ask any SAN manager to explain their process for “Root Cause Analysis” (RCA) today, and you will probably witness a mix of angst, anger and frustration wash over their face in the course of just a few moments. This is because that SAN manager is likely reliving a ridiculously complex and manual process for identifying and resolving fault events in today’s typical SAN environment. While every storage team may have a different specific fault management process, they all have to deal with the infamous “sea of red”

alerts that can bombard their management console during a failure.

For any typical fault event, that swarm of alerts can easily encompass hundreds of components spanning applications, IP networks, HBAs, FC switches, FC connectivity, and storage devices. On top of all of these alerts, multiple human teams must communicate with each other -and their respective trouble tickets- to determine which IT domain bears responsibility for the potential fault event. This often requires

**P R O D U C T P R O F I L E**

complex co-ordination and planning between network administrators, application managers, database administrators, and storage teams, all of whom have a stake in meeting any Service Level Agreements (SLA) that the enterprise may have in place.

Because of this complexity, the act of finding the root cause and potential impact of a SAN-based fault event is undeniably one of the single most dreaded tasks facing IT teams today. Some of the key challenges that teams face can be summarized as follows:

**Challenge: Alert Overloading**

When a fault event transpires today in a FC switched fabric (e.g. an HBA failure in a database server), an entire wave of sympathetic alerts from a vast range of directly and indirectly connected SAN elements hits the administrative alert console. These alerts all lack intelligence. They are simply devices describing their failure state, with no context for remedy. This overload of alert data then forces the IT team to sort through what can easily amount to several hundred discrete messages in order to identify which are relevant and which are merely SAN elements responding to a domino of preceding failures messages. This non-contextual overload is one of the principle shortcomings bogging down traditional root cause analysis efforts.

**Challenge: Manual Correlation**

Once the IT team can sort through their alert feeds and use their expertise to identify obviously irrelevant alerts from the immediately relevant, they can then begin to manually correlate the alert messages and back into their root cause. Sometimes, this

may be a very obvious remedy, but other times it may be maddeningly difficult, especially in larger or high-growth SANs that require frequent topological changes. This manual process can regularly take several hours for even easily identifiable failure events, and much longer when the failure is not easily discernable.

**Challenge: Domain Gaps**

Another key downfall of current approaches is that they can lead to unintentional redundancy as multiple IT teams unwittingly pursue the same root cause analysis. It is not uncommon for the network administration team and the storage team to overlap each other as they both scramble to identify where a particular fault transpired. This silo-based separation of the IP networks and storage fault management can lead to significant inefficiencies and risk exposures.

**Challenge: Impact Blindness**

Even after an IT team has successfully identified a fault issue in the topology, it is not inherently clear from that fault identification what its relative *criticality* is to storage delivery, and for what users. As one storage manager we know describes it, “while each is painful, not all failures are created equal.” Especially in a complex or large fabric, unaided human “eyeballing” is not sufficient to determine how a particular fault will flow up to impact service delivery and SLAs. Traditional approaches to fault management simply fall short when it comes to linking technology failures with business impact. They are literally blind to the business.

With these general challenges in mind, it makes sense to take a look at how they can

**P R O D U C T P R O F I L E**

actually be remedied through a new approach that actually makes IT managers' lives easier. This approach is called automated fault management.

## **Automated Fault Management**

Resolving the challenges stated above (and beyond) is exactly why the entire category of automated fault management software has emerged. The key insights driving this new software category are the following:

### **Insight: Build With Models, Not Rules**

Automated fault management makes a very critical shift to a *model-based* as opposed to a *rules-based* methodology. In traditional fault management, the IT team is required to specify reporting rules for how each infrastructure element and its linkages are related to each other in the event of a failure. This approach is the source of the "alert overload" described above, as specific rules regarding a wide range of IT elements across multiple domains simply do not have a context for their overall relation to their neighbor elements. This rules-based approach also has the key shortcoming that any changes made to the IT topology will require a readjustment of any pre-existing rules, a task that many teams rely on custom scripting to solve today.

By contrast, a model-based approach gathers a totalizing view of the IT topology (in our specific case here, the FC SAN) via an automated discovery interface, and then reconstructs the relationships of the topology in the automated fault management software itself. Many of the key advantages of

automated fault management result from this shift to modeling rather than rule creation.

### **Insight: Federate Across Domains**

Automated fault management software also leverages another key insight stemming from the model-based approach: Enable the federation of multiple domains (e.g. the IP network, the application environment, and storage) into one architecture so that the "domain gaps" described above are eliminated. This prevents multiple teams from responding to multiple trouble tickets, or from wading through the alerts generated outside of their sphere of operations. This is a key improvement that removes one of the most obvious and time-consuming inefficiencies of traditional approaches.

### **Insight: Link to Business Impact**

Automated fault management software goes beyond identifying a root cause. It creates a linkage between the fault event and the impact a fault will have on specific IT services being delivered to users. This linkage between technology and business value has been one of the missing pieces in root cause analysis in the manual rules-based world. Automated fault management can address this by removing the necessity of human guesswork. This represents a significant generational leap over the "impact blind" RCA methods most shops have in place today.

### **Insight: Build With Modularity**

One of the unique advantages of a model-based approach to automated fault management is that it opens up the opportunity to extend functionality in new directions over time, while maintaining

**P R O D U C T P R O F I L E**

common data collection and correlation engines. This means that additional capabilities that might leverage a modeling based approach such as business impact, performance or cross-domain functionality can be easily included. In addition to making the solution more intrinsically valuable, it also helps customers to use their fault management tools as real ROI generators.

All of these collective insights, when marshaled together, enable an entirely new type of software platform that can transform how IT teams resolve *all* fault issues, getting to root cause immediately, and with a far deeper insight into the pervasive impact this will have across their operating environments. We will now take a look at how EMC is actually working with these same insights to shape its entry in the automated fault management category.

### **Meet EMC Storage Insight for Availability**

The EMC Smarts product family is designed to provide a full range of automated fault management software to enterprise IT environments. Note that the initial offering in the EMC Smarts family, IP Availability Manager, was designed specifically IP network availability. The most recent offering and focus of this profile is EMC Storage Insight for Availability (EMC Smarts SI-A), which extends the proprietary Smarts automated root cause and impact analysis technologies into the FC SAN world.

### **The EMC SI-A Deployment**

Of course, this offering deploys as 100% software on industry standard architectures.

EMC SI-A leverages EMC Control Center 5.2, for Fibre Channel SAN discovery and resides on a standard Windows Server 2003 device. The product also requires EMC Smarts Service Assurance Manager, which provides a point of integration and correlation of management information across SI-A. For cross-domain IP correlation capabilities, the offering requires the deployment of IP Availability Manager (IP-network discovery is then provided by SNMP traps for Smarts IP Availability Manager.) The entire deployment is designed to work as a modular addition to the existing EMC environment.

### **EMC SI-A: The Technology**

This offering leverages a patented model-based approach to fault management. Rather than rely on an amalgamation of user-driven rules that are scripted for every conceivable component and failure event, the EMC SI-A architecture recreates a usable model of the all of the elements and relationships in the customer's SAN. The scope of element coverage includes all port links and unit statuses for supported HBA cards, FC SAN switches, and Symmetrix or CLARiiON arrays. Any changes to the SAN environment can be quickly discovered and integrated into the model schema. It is hard to overstate the critical difference that this manner of perpetually adaptive modeling brings to fault management. It is literally a total automation of what has been historically a very labor-intensive and error-prone exercise.

The key to this unique modeling approach rests in two co-operating technologies at the heart of EMC SI-A which impact both data collection and data correlation. Specifically, the offering utilizes an automatically

## PRODUCT PROFILE

updating “Codebook” Correlation Technology and Common Information Model (CIM) interface to identify and interpret all SAN network elements and their respective relationships. The Codebook Correlation Technology is a set of software algorithms that hones in on and identifies the SAN environment elements that *must* be associated with a given fault. Next, the CIM interface takes these identified SAN elements and interprets their position and role relative to each other, then surfaces this information to EMC SI-A. This process enables a definitive determination of any fault issue’s root cause (Note that EMC SI-A’s UI presents a certainty value of its root cause determination as a percentage value. In the vast majority of common fault cases, it will register “100%” certainty.)

Literally, all of the hours of intensive manual correlation of the “red sea” of alerts are automatically linked based on the SI-A model to the responsible element. The difference to the IT administrator is profound: Going from hours of step-by-step hunting to nearly immediate determination of real-time root causes for fault events.

### **Beyond Root: Impact Analysis**

EMC SI-A’s modeling approach extends beyond just eliminating the manual process of root cause and looks to the absolute service *impact* of that fault event. With the add-on Business Impact Manager (BIM), EMC SI-A enables the IT administrator to associate each element (and therefore, its potential fault events) with all of the storage services in the enterprise. Most importantly, the administrator can then customize

business values for each of these elements in their service context. The net result of this functionality is a linkage of technology availability *and* business criticality. This would not be achievable in a rules-based fault management paradigm without significant levels of customized and ongoing investment that most teams would never pursue.

BIM is an important add-on that we suspect most users will opt to include. They certainly should. For senior IT executives or team members in other domains, this module constitutes an easily referenced “eyeball” on a given fault event. For example, a series of switch port fluctuations in the SAN fabric could either spell disaster for the enterprise, or merely a minor routine fault scenario. With BIM installed, anyone with permission can know with confidence the relative criticality of those fault events without hours of waiting and guesswork.

### **Cross-Domain: The Big Picture**

One observation that jumps out to anyone who plays with the EMC SI-A user interface is its modularity and flexibility. It is easy to add additional functionality like BIM and access these new modules from the same management console in an adjacent UI frame. Beyond expansion for BIM, this also applies to cross-domain functionality for IP network fault management, as well. EMC SI-A is based on the same Smarts software platform as the IP Availability Manager, and they are designed to inter-operate. In fact, utilizing these two offerings together, while not required, is far and away the best means of achieving a significant and rapid improvement in any EMC customer’s cross-

**P R O D U C T P R O F I L E**

domain fault management schema. This removes overlapping, redundancy and guesswork by framing and guiding administrators through multiple domains based on automatically updating models of both their IP and storage infrastructures. If *that* isn't an improvement over traditional approaches, then nothing is.

### **Taneja Group Opinion**

The pain of conducting manual root cause analysis in the SAN world has finally started to find some relief. We believe very strongly that the aggressive automation of fault management through model-based approaches is clearly a superior way to deal with one of the hardest challenges in the data center.

From an economic standpoint, the value that EMC's Smarts family, and EMC SI-A in particular, can bring to the data center is real and can be measured. We highly encourage EMC customers who are looking at deploying this technology to actually account for how much valuable team time they sacrifice to fault resolution on an annualized basis. Even more dramatically, the reduction in downtime that an automated fault management solution brings to the table should help customers to frame up their ROI analyses, as well.

This concept of "cross-domain analysis" also merits explicit mention here. A major trend that we are seeing across the board in infrastructure management software is this concept of interpolating data across more than one tier or domain of data in order to provide accurate and actionable information.

In the context of automated fault management and the EMC Smarts family, this is extremely powerful. We wager that once most IT teams really understand just how much easier a cross-domain fault schema makes their lives, they will never look back. We encourage customers to explore where their "domain gaps" are exposing them to unneeded labor and risk. The good news is that there are solutions now!

EMC SI-A is the first product in what will clearly be an entire suite of integrated offerings from the EMC Smarts Storage Insight family. Obvious and expected extensions of the Storage Insight platform include performance monitoring and change management modeling for "what if" scenarios. Additionally, as EMC extends their heterogeneous storage array coverage beyond the core Symmetrix and CLARiiON lines, they will seal the deal in winning over customers with the Storage Insight platform.

Storage Insight for Availability is a profound improvement over historical rules-based approaches, and can provide measurable, defensible gains to a wide range of SAN environments. If you are a current EMC customer still wrestling with a lackluster, time-consuming fault resolution schema, this offering will definitely get to the root of your problem! This is a technology that merits a deep dive.



## P R O D U C T P R O F I L E

---

***NOTICE:** The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the TANEJA GROUP and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. The TANEJA GROUP, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document.*