Product Brief

Forcing the SAN to Grow Up

Date: October 2009 **Author:** Steve O'Donnell, Senior Analyst Data Center Practice

Abstract: *Last century, physical servers and physical storage were physically connected over the SAN; the resulting traffic patterns, while largely unmeasured, were relatively static and somewhat predictable. A conservative (albeit expensive) design approach to accommodating those patterns—throw hardware at the problem—was normally taken, resulting in a SAN that usually worked well; except, when it didn't, often with catastrophic consequences! Virtual Instruments, a private equity-backed spin-off from Finisar (a leading vendor in Fiber Optic Network Test Systems), has the solution: providing SAN managers much-needed LAN-style monitoring capabilities.*

Overview: The SAN

The Storage Area Network (SAN), a key component of IT Infrastructure, is surprisingly difficult to manage operationally and is disappointingly prone to in-life problems. It has evolved from simple beginnings to a complex, heterogeneous system, often of massive scale and resultant complexity, storing the vast majority of an enterprise's mission-critical data. Unfortunately, monitoring and management capabilities have not evolved at the same rate; while present in abundance in the LAN environment, these facilities are largely missing from the modern-day SAN.

SAN traffic between the application and storage array is significantly more sensitive to delay or loss than LAN traffic. Accordingly, the SAN is typically designed to avoid overloading of any component and the resultant delays that can cause application and, in turn, customer-affecting problems. One common solution is to architect and size the SAN so that it can accommodate every server port sending and receiving at full utilization. Unfortunately, SAN connections are expensive and large SANs are cost constrained, making it difficult or impossible to actually implement such designs. SAN architects must then balance cost and risk, recognizing that a mistake can result in performance problems and outright outages.

The primary SAN protocols, Fiber Channel (FC), Fiber Channel over Ethernet (FCoE), and Internet SCSI (iSCSI), are all mappings of the peer-to-peer SCSI (Small Computer Systems Interface) developed in the 1980s onto network communications protocols. Neither SCSI itself nor its mappings to these protocols includes a transport layer protocol for error correction and retransmission, so as a group they are particularly sensitive to data loss. Additionally, SCSI and its derivatives assume a direct peer-to-peer connection between the initiator and the target, so end devices, and the applications that use them, do not tolerate latency and delays well.

SANs are carefully designed to avoid frame loss, latency, and delay by integrating non-blocking switch fabrics and a judicious fan-out ratio between the hierarchies[1] of the connected host, edge switch ports, the inter-switch links (ISL) to the core switches and the connections to the storage array. Typically, SAN designs assume that all hosts may be communicating at full line rate simultaneously, implying significantly over-provisioned SAN capacity. Physical layer over-provisioning is further compounded when implementing complex, multi-tier, edge-core-edge switch designs in very large scale environments. Intended to provide physical reconfiguration flexibility to accommodate growth, these designs actually introduce substantial additional cost, complexity, and resultant risk.

SANs are typically designed around a fan-out ratio[2] between host connections and storage array ports of between 6:1 and 12:1, depending on the estimated intensity of the storage activity. Low-utilization hosts can be theoretically

[1] Typically, SANs are designed like LANs with hosts connected to edge switches that are, in turn, connected to core switches, that then connect to the storage arrays.

[2] The fan-out ratio refers to the ratio of the number of Host Bus Adapters (HBAs) connected to a SAN and the number of storage controller connections to the SAN. SAN designers determine this ratio based on rules of thumb provided by SAN vendors that presume (guess) the bandwidth rate that applications need to access storage, inferred latency, and performance requirements.

supported at the upper fan-out ratios, whilst highly utilized hosts[3] need much lower fan-out ratios. Good SAN design involves balancing high intensity hosts and low intensity hosts on the same edge switch to limit and balance switch utilization. Choosing the correct fan-out ratio is a hard enough decision at the initial implementation stage, but it becomes extremely difficult to maintain in a mature and growing SAN—and nearly impossible to manage with a virtualized workload that will be automatically and transparently moving in real time between physical hosts and their associated SAN connections.

Paradoxically, throwing hardware at the problem increases cost, complexity, and the statistical probability of failures—and, therefore, risk—as well as increasing the time taken to diagnose and address problems when they do occur. The over-provisioning approach becomes even less viable as the SAN and server virtualization increasingly converge (almost all VMware deployments have a SAN). Server virtualization drives much more dynamic SAN traffic patterns that can't be effectively predicted in support of over-provisioning and the very idea of over-provisioning the SAN stands in direct opposition to the fundamental server virtualization objective of increasing utilization and reducing costs.

The SAN Meets Virtualization

In a virtualized world, applications are untethered from the underlying physical hardware and SAN connections, enabling them to be moved dynamically across the infrastructure. For example, VMware VMotion (deployed in production at 60-70% of VMware customers, according to VMware) enables the migration of an entire running virtual machine, application and all, from one physical host and physical storage system to another—seamlessly and without user interruption.

Virtualization offers many benefits to the Enterprise: obscuring the physical complexity of servers, storage and networks from applications, enabling rapid deployment, enhancing availability through outages, and increasing asset utilization. Virtualization does not, however, offer a free lunch; inadequacies in the provisioning and management of the SAN infrastructure will be exposed very quickly as more and more workload is virtualized and automated tools dynamically migrate critical, IO-intensive business applications between nodes on ever-larger clusters of servers and storage.

Virtualization obscures the physical complexity of the underlying hardware from applications, but in doing so; it also adds several additional layers of abstraction that can significantly obscure the causes of performance problems and outages from IT Operations, increasing the time to resolve and reducing right first time diagnostics. It is not unusual for a fairly simple fault, one that can be repaired in a few minutes, to take many hours to first diagnose, with much technical hand-off between the various support teams along the way.

The choice is stark: over-engineer, pay the price, and hope for the best or adopt a more scientific approach and measure, manage, and optimize the SAN proactively.

The Virtual Instruments Approach

Virtual Instruments has developed a set of SAN instrumentation and monitoring solutions that enable organizations to truly optimize their SANs. The VirtualWisdom solution consists of hardware, software, and services designed to simultaneously reduce both risk and cost. It delivers a unique solution, providing real-time operational, utilization, and performance monitoring and root-cause diagnostics for the SAN—capabilities directly comparable to those available for the LAN. Leveraging Virtual Instruments' optical networks test legacy, VirtualWisdom passively mirrors traffic flowing across the SAN to examine and decode an out-of-band copy of the traffic, combining the results with monitoring data gathered from both VMware ESX servers and the SAN switches to present a comprehensive set of granular, accurate monitoring data.

VirtualWisdom can be delivered either temporarily as a troubleshooting service when organizations have pressing SAN issues that they need to fix or, more commonly, on a permanent basis, enabling a cycle of continuous improvement across the levels of a well-defined SAN Maturity Model (SMM). Virtual Instruments' SMM has been derived from ESG's Data Center Efficiency Maturity Model.

[3] Such as those running large database applications, VMware, or VDI.

Figure 1 shows the major components of the VirtualWisdom solution and the interactions between them.

Figure 1. Virtual Instruments Connectivity Diagram

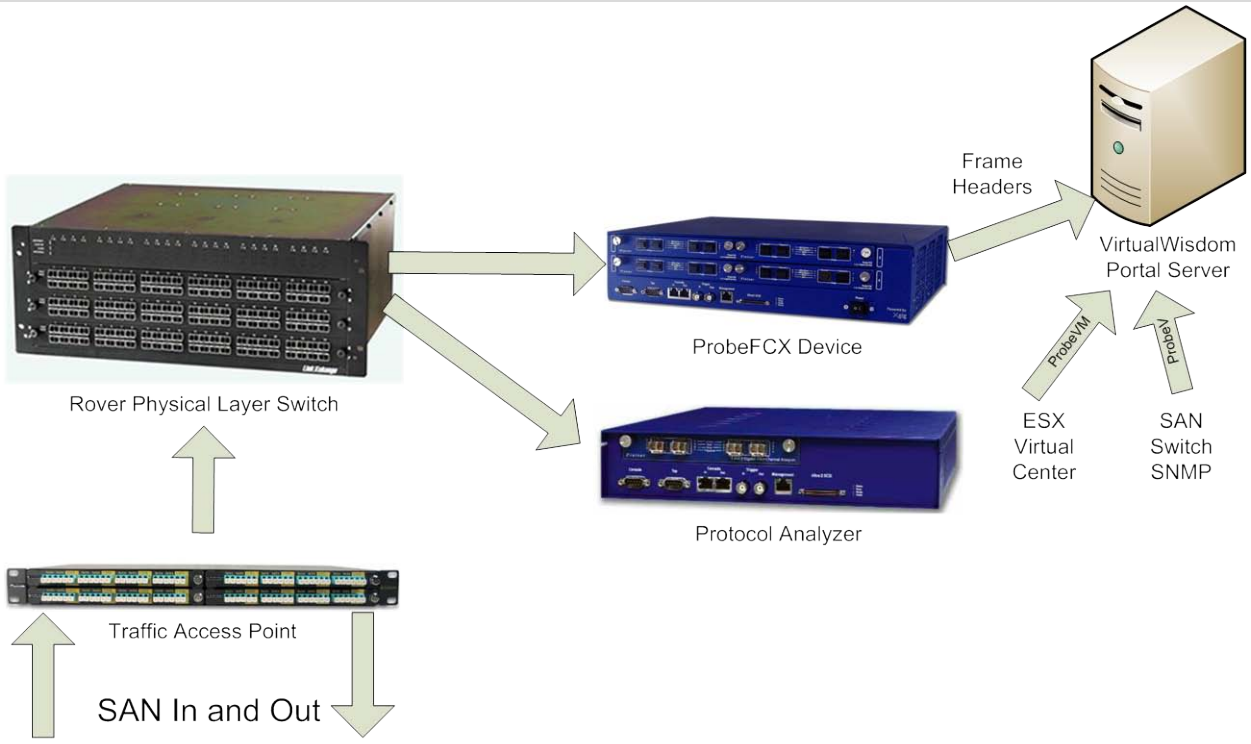

Source: Enterprise Strategy Group, 2009.

The most comprehensive and accurate monitoring and diagnosis of SAN traffic is performed directly against a copy of that traffic, provided by a passive (non-powered) Traffic Access Point (TAP) connected inline within the SAN link to create an out-of-band mirror copy of the optical signal. The signal copy can be fed directly into a ProbeFCX monitoring device or Protocol Analyzer, or routed into a multi-port Rover (not shown) to enable dynamic routing to a shared ProbeFCX or Protocol Analyzer. Because all monitoring activity is performed on this traffic copy downstream of the TAP, monitoring has no impact on the operation or performance of the live links. TAP installation is the only disruptive step of implementing VirtualWisdom, as the live links must be interrupted whilst the TAPs are inserted. TAPs are therefore typically installed at the same time as patch panels or during normal maintenance cycles.[4]

The ProbeFCX is an appliance that performs real-time monitoring of the traffic moving across multiple SAN links, tracking more than 200 metrics every second that measure errors at all layers of the FC stack: link, port, and device utilization and end-to-end SCSI performance and latency. The ProbeFCX generates these metrics by operating only against the Fibre Channel (FC) frame headers and primitives and stores only the metrics themselves; all data in the frame payloads is discarded at a very low level in the hardware, obviating any potential security issues.

By contrast with the ProbeFCX, the Protocol Analyzer is used to capture the traffic itself from the signal copy; the resulting trace files can be viewed, decoded, and analyzed using the industry's leading software tools. Tight integration allows VirtualWisdom to automatically connect the Protocol Analyzer to the appropriate link (if using a Rover), trigger the capture, and upload analysis of the resulting output into the Portal Server, where it becomes a permanent part of the monitoring record.

In addition to the ProbeFCX appliance, VirtualWisdom includes two software probes: the ProbeV collects monitoring data from the SAN switches via SNMP and the ProbeVM collects data from VMware's Virtual Center Server. Both

[4] Note: ESG strongly recommends that TAPs are installed on all SAN links as a standard best practice, even if you don't have near-term plans to implement VirtualWisdom. Having TAPs already in place greatly simplifies connecting a Protocol Analyzer or ProbeFCX in the event of a future problem.

software probes also feed their information into the Portal Server to provide a rich, holistic, and correlated view across the hosts and fabric.

Realizing the Benefits

Virtual Instruments recognizes that monitoring and instrumentation are only part of the story in getting to a fully optimized SAN environment, so the firm also delivers a SAN Maturity Model framework that gives enterprises a template of Maturity Levels, Key Process Areas (KPAs), Goals (KPIs), and Leading Practices. Fully adopting the recommendations in the SAN Maturity Model, as well as implementing the technology solution to monitor and manage the SAN, will lead down an assured path towards a lower cost and lower risk future.

Figure 2. SAN Maturity Model Benefits

	1: Initial	2: Managed	3: Defined	4: Quantitatively Managed	5: Optimizing
Level Synopsis	SAN works well –except when it doesn't. Over-provisioned, but unable to prove it. Headcount can't keep up with the high growth/change rate.	The use of better tools and processes is keeping up operational workload; turning attention to proactive initiatives.	Smoothly managing both ongoing operation and new implementations using defined best practices and targets.	Delivering best-in-class capabilities using objective, transparent SLAs.	Proactively taking steps to further improve QoS & reduce costs; delivering best-in-class capabilities at benchmarked best-in-class costs.
Risk Levels					
Business Impacting Failure	**SEVERE**	**ELEVATED**	**GUARDED**	**LOW**	**LOW**
Outage Duration / Impact	**SEVERE**	**ELEVATED**	**LOW**	**LOW**	**LOW**
Overspend / Overprovision	**SEVERE**	**HIGH**	**ELEVATED**	**GUARDED**	**LOW**
Underprovision / Underperform	**SEVERE**	**HIGH**	**ELEVATED**	**GUARDED**	**LOW**
Representative Benefits					
OPEX		**$$**	**$$$**	**$$$$**	**$$$$$**
		Faster troubleshooting. Shorter outages. Improved SAN performance.	Reduced downtime. Higher SAN admin efficiency.	Fewer outages, reduced duration. Fewer trouble tickets. Improved app performance/ throughput	Maximum efficiency in admin resources. Lowest downtime. Application-to-spindle optimization
CAPEX			**$**	**$$$**	**$$$$$**
			Initial reduction in storage ports and related elements (switch ports, cabling) through better utilization of existing assets.	Further reduction in ports, related expenses (software, replication, etc.). Re-tiering data to lower cost storage.	Optimized use of existing and new capacity. Re-tiering across tiers and RAID configurations. Advance planning yields best prices.

Source: Enterprise Strategy Group, 2009.

Why This Matters

The SAN is at the center of all mission critical enterprise application architectures. The risk associated with being unable to measure and monitor the health of this major component is significant, real, and increasing with the widespread adoption of server virtualization and cloud computing. Ignoring the potential for SAN optimization because a serious outage has not yet occurred, or because the necessary tools to monitor and manage it effectively are lacking, is both foolhardy and short sighted.

In most cases the investment required to provide good SAN management tools, such as VirtualWisdom, can be recovered from reductions in the capital outlay required to accommodate organic SAN growth following the old architectural patterns and rules of thumb. It might be slightly more difficult to accurately assess the monetary value of avoiding even a single customer-affecting outage, but the significant impacts of unplanned outages can include:

- Lost orders and revenue
- Customer retention issues
- Brand and reputational damage
- Increased helpdesk demand
- Rework and lost time
- Service credits and damages claims
- Fines for compliance violations and more stringent operational requirements

Enterprises should not take the approach of pay and pray when deploying large scale SAN hardware overlaid with cloud and virtualized platforms. An investment made in SAN instrumentation will pay back in four key risk areas:

1. Reduction of business impacting failure risk
2. Reduction of Outage Duration risk
3. Reduction of Overprovision and Overspend risk
4. Reduction of Under-provision risk

Are There Alternative Approaches?

One alternative is to reduce the fan-out ratios used in the SAN design; this will have the effect of reducing the risk of under-provisioning the SAN and have an impact on the risk of business impacting failures. However, outage duration, over-provisioning and overspend risks become certainties. This is a particularly expensive solution and is not recommended.

There are a number of alternative SAN monitoring solutions; however none of them has the real-time instrumentation capabilities, the cross domain view or legacy of the many years' investment and experience in optical networking monitoring solutions that are part of the DNA of Virtual Instruments. This unique and high value additional feature adds significantly to the Virtual Instruments' VirtualWisdom solution making it best of breed in the market today.

The Bottom Line

By bringing together the critical sources of information about SAN health (link monitoring, switch monitoring and ESX Server monitoring), Virtual Instruments has been able to deliver a unique and compelling solution to in-life SAN management in the Virtualized Data Center. By building a fully integrated implementation plan in the form of a SAN Maturity Model, Virtual Instruments lays out a vision of SAN best practices that offers the promise of low operational risk combined with low cost of operation. For the first time, SAN and Virtualization Architects have the information they need to engineer service level targets at the lowest possible cost, and operations teams have the information they need to drive time to resolve towards zero and right first time diagnosis to 100%, dramatically improving application availability and performance. Enterprises that choose to follow the recommended path through the Virtual Instruments' SAN maturity model will reach an ideal level of optimization.