WHITE PAPER

Creating Performance-based SAN SLAs Using Finisar's NetWisdom

HIGHLIGHTS

- Using SLAs, SAN managers can track SAN performance and manage user expectation
- NetWisdom allows SAN managers to improve configuration and resource management, ensuring enterprises get the most out of their SAN investment
- Tracking resource usage relative to capacity, SAN managers can spot usage trends and add capacity when and where needed

EXECUTIVE SUMMARY

Storage area networks (SANs) play an increasingly critical role in application performance and the function of information throughout an enterprise. More-and-more, organizations are looking to service level agreements (SLAs) as a means to quantify and enforce SAN effectiveness. Currently, many enterprises think of SAN SLAs in terms of uptime alone. However, a SAN SLA must encompass performance and resource usage metrics as well. Otherwise, enterprises can find themselves with a SAN that's up but performing so poorly as to make applications unstable or unusable.

Like any SLA, a SAN SLA will have enterprise and site-specific characteristics and include details about escalation procedures and potential penalties in the event of a breach. Likewise, a SAN SLA must clearly spell out the components encompassed by the agreement, identifying the boundary between the storage network and other aspects, such as application servers and LANs.

The key to establishing any SLA is having the tools necessary to accurately measure and confirm compliance. SAN managers need tools to help them monitor and report on performance metrics, resource usage and troubleshooting practices, as well as uptime to deliver the expected service levels. Finisar's NetWisdom monitoring tool and Xgig analyzer provide the functionality SAN managers need to establish and maintain SLAs for Fibre Channel SANS.

NetWisdom's comprehensive end-to-end monitoring of the entire SAN fabric, including per-application conversation tracking, enables enterprises to view a wide range of performance, resource utilization, and problem solving information. NetWisdom facilitates all aspects of the SLA process, from allowing managers to baseline SAN performance in order to set realistic performance objectives, to monitoring the SAN for ongoing SLA compliance for users. In addition, NetWisdom's rich error tracking and resource trending capabilities and Xgig's troubleshooting features enable SAN managers to quickly identify and resolve problems, ensuring maximum uptime, performance, and resource utilization.

SLAs ensure users are getting the application performance they require and enterprises are getting the most out of their SAN investment.

Finisar

THE NEED FOR SAN SLAS

Applications – everything from Enterprise Resource Management and Web portals to Customer Relationship Management and email – are the lifeblood of today's organizations. And as many enterprises have discovered the hard way, application performance must be up to snuff or they risk decreased employee productivity, low customer satisfaction, and lost business opportunities.

Consequently, IT managers are looking for ways to guarantee that users have access to the critical applications and data they need when they need it – what analysts at Enterprise Management Associates (EMA) call "end-user quality of experience" or QoE. Establishing a service level agreement (SLA) is one way to provide these guarantees. EMA analysts note that QoE-based SLAs are possible because management tools are available that can monitor a wide range of application-related performance metrics.

Clearly no one SLA can encompass all the infrastructure components along an application's path. After all, application performance can be impacted by variables ranging from poorly-written code and overburdened servers to slow storage systems and sluggish network links. However, IT managers can define SLAs for key infrastructure components, including local and wide area networks and storage area networks (SANs), ensuring that each infrastructure component is on-line and operating within the expected levels of performance.

In this paper, we examine the key business and technical issues that a SAN SLA should encompass. Given its critical role in data delivery, a SAN can have a significant impact on application performance. Too often, though, lack of insight into the SAN's actual operation and performance make the SAN an easy scapegoat when application performance issues arise. With Finisar's NetWisdom monitoring tool and Xgig analyzer for Fiber Channel SANs, storage managers can readily track a SAN's performance and determine what impact the SAN is – and isn't – having on application performance.

NetWisdom is unique in that it monitors performance end-to-end across the entire heterogeneous SAN fabric. In addition to providing detailed performance data for all SAN components in a consistent, easy to interpret fashion, NetWisdom gives SAN managers a system-level view. It can track a variety of metrics in real time, including data rates, latency through the SAN, link utilization, and Fibre Channel and SCSI errors.

NetWisdom can also monitor specific application "conversations" by tracking statistics for individual LUNs. This capability enables SAN managers to understand and document SAN utilization on a per-application basis and define application-specific SLAs as needed. With NetWisdom, SAN managers can establish realistic performance expectations based on actual application usage and validate compliance with a performancebased SLA. For its part, Xgig provides the troubleshooting and problem resolution capabilities managers need to keep the SAN operating within expected service bounds.

THE INS AND OUTS OF A SAN SLA

A SAN SLA can provide value on several levels. First and foremost, SAN SLAs enable SAN operators, whether they're in-house or outside service providers, to manage user expectations of the SAN's performance. Service levels are agreed upon up front, and the SAN operator ensures that those levels are met. Such agreements may be limited to critical applications that have stringent performance requirements, or written broadly to encompass all traffic in and out of the SAN.

Second, SLAs ensure that business managers get full value for their IT dollar. For organizations that outsource their SAN, an SLA provides written assurance that storage services will be provided as specified. Likewise, internal SAN managers can use SLAs to assure business managers that they are getting the SAN services they paid for.

Third, an SLA requires that SAN performance be monitored. Consequently, when application performance isn't meeting users' expectations, the SAN manager is in a position to quickly pinpoint whether the SAN is the cause of performance degradation, expediting resolution of performance problems. In addition, regular monitoring enables the SAN manager to more accurately track resource usage for capacity planning and change management purposes.

Whether an organization operates a SAN in-house or outsources SAN services, a SAN SLA should encompass both operational and technical issues. The key to creating meaningful SLAs is determining the salient service factors to include and having a reliable means of measuring service levels. Beyond spelling out the specific availability and performance metrics to be supported, an SLA must also list

the actions to be taken to return the SAN to normal operation in the event of an abnormal condition. Details such as who should be notified as well as a predefined set of actions for problem identification, diagnostics, troubleshooting, and problem resolution should be included.

As with other types of SLAs, SAN SLAs will have enterprise- and site-specific characteristics, such as the particular performance metrics to be monitored and the acceptable range of values for those metrics. For the most mission-critical applications, SAN managers may find it necessary to define performance metrics on an application-by-application basis based on individual LUN statistics.

It's also important to delineate the boundaries of the SAN so that all interested parties (i.e., IT management, business management, etc.) know exactly which components the SAN SLA encompasses. Typically, a SAN SLA will cover all the components of the storage network that facilitate data transfer, from the host bus adapters (HBA) in servers to the interfaces to storage devices, including Fibre Channel switches, cables, and any gigabit interface converters (GBICs).

Lastly, a SAN SLA must encompass uptime and performance metrics as well as resource availability. While uptime is a key factor in an overall service evaluation, it's not sufficient. A SAN may be up but performing so poorly as to make an application unstable or unusable, for example. Performance metrics should be based on actual SAN usage and application requirements, not arbitrary numbers.

SAN managers can establish the appropriate performance ranges for each SAN component by monitoring the SAN and establishing a performance baseline. NetWisdom, for example, provides comprehensive trending information that SAN managers can use to determine the high and low ranges of performance and resource utilization for each SAN component as well as aggregate end-to-end data, providing an accurate picture of application performance across the storage network.

SAN managers can use this baseline data to spell out expected service levels within an SLA and as the standard against which to evaluate compliance. As noted above, an SLA should include details about escalation procedures in the event an SLA is breached. Likewise, it may spell out penalties in the event of an outage or extended performance degradation.

UPTIME ISSUES

Given the highly redundant nature of most SAN designs, it may seem quite reasonable to ask for 99.999 percent uptime in a SAN SLA. Five 9s of uptime allows for literally minutes of downtime outside of scheduled maintenance and other planned outages. Though SANs rarely fail completely, catastrophic failures do occasionally occur.

More commonly, however, individual component failures or rising traffic levels over time trigger overload conditions that can result in one or more applications being starved of SAN resources. Whether such resource deprivation will result in a brief dip in application performance or an outright failure requiring human intervention will depend on the sensitivities of the application relative to the resource deprivation. For example, an email application may fail if the client cannot access email for an extended period because a large MS SQL transaction has filled the pipe that the SQL application shares with the email server.

Many Fibre Channel errors, particularly those that are transient, are never reported to the host operating system. However, a pattern of errors, such as aborts, link resets, and fabric log-outs and log-ins, can signal problems with particular SAN components that could lead to component failures over time. As part of their best practices regimen for meeting SLA uptimes, SAN managers should routinely track trends for such errors and initiate diagnostics in order to identify the root cause of errors.

NetWisdom can monitor 250 different metrics for each conversation, giving managers a fine-grained view of SAN operations and error conditions. When errors are detected, SAN managers can run a trace and perform root cause analysis using Finisar's Xgig analyzer. With these tools, SAN managers can quickly identify and resolve problems and pro-actively schedule maintenance to replace or reconfigure a component, if needed, thus preventing a failure from impacting users and ensuring uptime compliance.

For economic reasons, multiple applications often share a SAN infrastructure. Because Fibre Channel has no quality of service capabilities, applications use SAN resources on a first-come-first-served basis. As we noted in the example above, this can result in a situation where one application hogs all the bandwidth from the SAN switch to a disk array, starving other applications dependent on the same array.

Consider various scenarios where this type of "application starvation" can occur. In one case, SAN managers have installed two 2Gb host bus adapters (HBAs) in a host and set up load balancing software. Initially, the host only required 100 MB/s of throughput, which could be handled by a single HBA and easily handled by two HBAs sharing the load. Over time the load gradually increases until the host needs an average of 220 MB/s of throughput, which is provided by the two HBAs. Then one HBA fails. The load balancing software sees this and takes the failed HBA off-line. Now a single HBA must accommodate 220 MB/s of consistent throughput, pushing it to its limits. Coincidentally, on the same day, accounting begins a major end-of-quarter financial roll-up and hogs the full link of the functioning HBA, "freezing out" users of other applications. These users experience either a work slowdown or a total halt in data exchange.

A similar kind of "freeze out" can occur through traffic growth of existing applications over time (for example, more users being added to the finance department), new applications being deployed on the same SAN infrastructure, or simply a spike in one application's usage of the SAN.

HITTING THE FIVE 9S

SAN managers can mitigate and potentially eliminate application degradation and downtime using NetWisdom to monitor application usage of SAN resources relative to their capacity. For example, Enterprise A has configured its MS SQL and Exchange servers to use the same storage array. Each host is connected to two SAN switches, with each switch in turn connected via a 2Gb/sec link to a storage array.

A baseline analysis shows that Exchange averages about 70 MB/sec load on the links between the switch and the array, while MS SQL typically runs about 80 MB/sec, for an average of 150 MB/sec of traffic between the switches and array.

The two links could easily handle 320 MB/sec of traffic at a comfortable 80 percent utilization. To prevent a large MS SQL job or a spike in email usage from consuming more than its share of the bandwidth, SAN managers use NetWisdom to establish a threshold of 50 percent utilization for both applications and set alarms. SAN managers could also establish a 75 percent maximum threshold for the links between the hosts and switches for added protection.

Should either application exceed its bandwidth threshold, SAN managers will be notified and can take appropriate action, such as taking a server out of rotation until the load drops. Should link loads remain high over an extended time (days, weeks), the SAN manager may decide to move the MS SQL server to its own link to avoid potential starvation of Exchange, or vice versa.

Using NetWisdom to monitor error rates, resource utilization, and other metrics, SAN managers can prevent transient problems from snowballing into major outages. Likewise, by tracking actual resource usage relative to capacity, SAN managers can easily spot usage trends and add capacity when and where needed, heading off downtime triggered by bandwidth and other resource constraints.

PINPOINTING PERFORMANCE METRICS

As SLA experts will tell you, performance has many aspects, making it more difficult to measure than uptime. Which performance metrics an enterprise includes in its SAN SLA will depend on the nature of their applications and the QoE they need to deliver to users. Technical and business experts need to collaborate to define these metrics. Here we present a few scenarios and discuss several metrics that SAN managers should consider when drafting the performance portion of an SLA.

Some organizations, such as entertainment firms and health-care providers, have very high performance environments requiring the movement of large amounts of time-sensitive data. For these organizations, the key performance metrics may be the volume and speed of data transferred over X amount of time.

For example, Enterprise B has a video editing group working with a high-definition video processing application. Monitoring and analysis show that the SAN needs to be able to transfer 1 TB of data per hour, or 277 MB/sec, and that I/O times cannot exceed 30 milliseconds (msec). These values would be used to establish the performance requirements in the SLA, and NetWisdom can monitor the data transfer and I/O rates to ensure compliance.

General business applications typically have less stringent performance requirements, and a more broadly applicable metric to track is Exchange Completion Time (ECT). ECT

measures the time from when a host requests data to when it receives that data. This transaction time measurement encompasses all portions of the SAN data transfer, from host to switch to disk reads and back up to the host. From the user's perspective, ECT is the time the application must wait between requesting data and receiving data and thus greatly impacts the time that the user waits for a file to open, a spreadsheet to be populated, etc. Roughly speaking, it's the SAN's response time.

In many cases, an ECT value may be a sufficient performance metric for an SLA. As we noted earlier, establishing a baseline for performance is key. Using NetWisdom, a SAN manager can determine the baseline ECT for each application and each disk array. Using this information, the SAN manager can define minimum and maximum ECT values based on the application's sensitivity. This information will also allow SAN managers to ensure that a given disk array is capable of delivering the required performance. For example, an older array may have an ECT of 1200 msec, making it a poor choice to house data for an application that requires an ECT of 200 msec.

Once the SAN manager has established baselines for both the applications and the disk arrays, he or she can set an alarm threshold to ensure they're notified in the event a threshold is exceeded. Likewise, the business owner could be sent a report verifying the SLA has been met, or showing why it has been violated.

OTHER PERFORMANCE METRICS

As we noted earlier, ECT reflects the total round trip time, so to speak, from when a host first requests data to when that data is delivered. It's an aggregate of the response times for all the components along the SAN data path, including the HBA to switch link, the switch itself, and the link between the switch and disk array. It's often helpful for SAN managers to have insight into the performance of these subcomponents as well as other performance metrics for tuning and planning purposes. Depending on the performance requirements of a particular application or user environment, one or more of these additional metrics may be used to characterize performance within an SLA.

For example, one subset of ECT is a metric called Command to First Data, which measures the time from when a host requests data to when the disk responds. By using NetWisdom to compare the Command to First Data and ECT times, a SAN manager can gain insight into cache usage and seek times. For instance, a SAN manager can determine if cache is being used inefficiently and thus impacting application performance. If the write cache is limited, an application may be unable to write all the data it needs to the cache and have to page out to disk. In this situation, a SAN manager can increase the cache size or re-balance the cache that's allocated to writes vs. reads.

Similarly, by tracking seek times, a SAN manager may determine that reconfiguring the data on an array could significantly reduce seek time, thus improving application performance. In one example, an enterprise was able to accelerate the response of a Web-based application 300 percent by increasing write cache and moving the application to a disk with faster seek times.

A third metric SAN managers should track is queue depth. If the queue depth is set too high or too low, it can impact both SAN performance and stability. Storage vendors typically set a maximum queue depth for a given model or code revision. SAN staff then set controls on the host and HBA to ensure the queue on the storage device isn't overrun.

If the queue depth is set too low (e.g., at 4), the HBA must wait until the queue is empty to send more data, creating a bottleneck for the HBA to get data to storage. For example, if an organization has only four LUNS and the port can handle 1024 requests, then theoretically a SAN manager can set the per-LUN queue-depth to 256. However if the storage port was configured with 256 LUNs, the queue-depth would be limited to a maximum of 4 and the HBA would only perform 50 percent of the I/O operations/sec that one would expect from that drive. Setting the queue depth higher, such as 28 or 32, would increase the I/O operations/sec to expected levels.

Setting the queue depth too high (for many drives this occurs above 36 or 38) has no positive impact on I/O processes/sec and results in a condition called "excessive queuing." While not harmful, excessive queuing wastes resources.

Setting queue depth too high can also result in an HBA sending more requests than an array queue can handle, a condition referred to as "queue-full." This condition can trigger SCSI retries, aborts and, in some cases, resets. Sometimes the HBA will throttle itself back to a queue depth of 1 or 2 in an effort to prevent overloading the queue again. Some

operating systems will have to be rebooted to reset the queue, others with slowly creep back up from 1. Unfortunately, disk arrays are not required to handle this situation gracefully, and any requests may be dropped, leading to data corruption.

With NetWisdom, SAN managers can gain insight into queue behavior by monitoring pending exchanges either on a given link or on a per-LUN basis. With this trend information in hand, SAN managers can tune queue depths. Likewise, they can set an alarm to trigger if the queues reach 75 percent full, thus ensuring optimal performance while avoiding overload conditions.

NetWisdom gives SAN managers tremendous flexibility both in terms of the types of performance data and the granularity of the traffic flows they can track. For example, NetWisdom can indicate what host the data came from and which LUNs are being accessed, allowing SAN managers to set up filters so they can see the performance statistics for a single application. With this capability, SAN managers can establish application-specific SLAs.

For each initiator-target-LUN combination, NetWisdom can provide performance information including: average, minimum, and maximum ECT; average, minimum, and maximum Command to First Data; read and write throughput in MB/sec; issued vs. completed reads and writes; and average, minimum, and maximum pending exchanges.

PUTTING IT ALL TOGETHER

Once the SAN team determines the performance metrics to include in an SLA, the next step is to define the minimum and maximum (or the average) values required for those metrics to deliver the desired service level. In addition, SAN managers should set alarm thresholds so that they're immediately alerted when performance falls outside the specified bounds. Setting thresholds is straightforward with NetWisdom.

Figure 1 shows the maximum ECT for writes set to 35 ms; if this threshold is exceeded, a notification will be sent to the appropriate administrator.

Figure 2 illustrates that threshold minimums can also be set. In this case, if the data rate for writes falls below 130 MB/sec, an email alert will be sent.

Monitoring pending exchanges is the best way to get a handle on queue depths. NetWisdom gives SAN managers the flexibility to set the threshold for pending exchanges on a per-Target or per-LUN basis. Figure 3 shows the maximum pending exchange threshold set at 512 for a given link. While SAN managers may also want to set an alarm at the SCSI level, the link level alarm acts as a "catch all." With this alarm set, even if someone increases the queue-depth on a host, the SAN manager will be altered before a storage port experiences a queue-full situation.

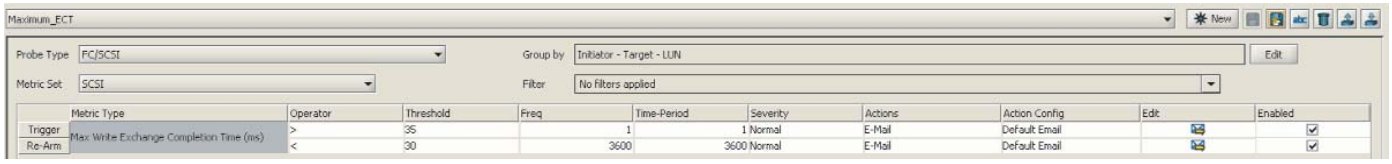

Figure 1: Maximum ECT for writes set to 35 ms

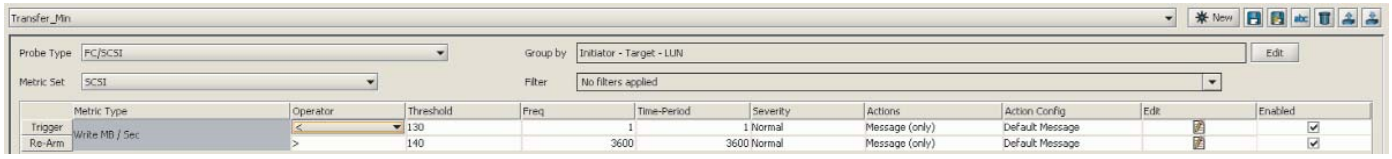

Figure 2: Threshold minimums set

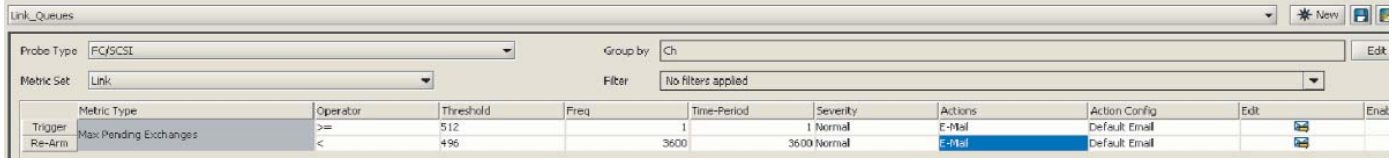

Figure 3: Maximum pending exchange threshold set at 512 for a given link

Finally, SAN managers need to measure and report on the metrics defined in the SLA to prove compliance or non-compliance. Daily, weekly, and/or monthly reporting may be specified as part of the SLA. Since user requirements, IT infrastructure, and other factors that impact SAN performance may change over time, SAN managers should periodically re-visit and confirm or modify the SLA parameters. Monitoring SAN operations is the best way to get a picture of current performance and to identify areas where tuning or reconfiguration could improve performance.

REAPING THE BENEFITS OF SAN SLAS

SANs are the first link in the application performance chain; their availability and performance directly impact a user's ability to read and write data. Given the critical role of applications in today's organizations and the growing dependence on SANs, organizations are increasingly looking to SLAs as a means to guarantee SAN service levels.

Each organization is unique; therefore SLAs are by nature site-specific. With its rich monitoring and trending capabilities, NetWisdom can help SAN managers pinpoint metrics relevant to their organization at every phase of the SLA process, from performance baselining and thresholding to ongoing compliance reporting.

NetWisdom enables SAN managers to define SLAs with uptime, performance, and resource availability requirements, and to establish performance expectations based on actual application usage. Its conversation tracking capability allows for application-specific SLAs, so SAN managers can guarantee that mission-critical applications have the resources they need to meet users' "quality of experience" expectations.

Likewise, NetWisdom allows SAN managers to improve configuration and resource management, ensuring enterprises get the most out of their SAN investment and add resources only when and where they're needed.

In addition to making more efficient use of SAN resources, NetWisdom lets SAN managers proactively identify and resolve performance and resource problems. Coupled with Xgig, NetWisdom can pinpoint errors and expedite problem resolution. As a result, managers can quickly return the SAN to normal operations, preventing business losses that can easily mount into the hundreds of thousands of dollars due to application performance degradation or outage.

Enterprises spend a significant amount of money on storage. With their comprehensive, end-to-end monitoring and troubleshooting capabilities, NetWisdom and Xgig help enterprises get the most out of that investment through SLAs that ensure SAN service levels are up to snuff. The bottom line for enterprises is guaranteed application performance and SAN uptime, which translates into greater user productivity, higher customer satisfaction, and increased business opportunities.

Finisar

1389 Moffett Park Drive
Sunnyvale, CA 94089
Phone (US Toll Free): 888.746.6484
Phone Intl: 408.400.1000
Email: networktools-sales@finisar.com
www.finisar.com

SNIA
Storage Networking Industry Association

FCIA
FIBRE CHANNEL INDUSTRY ASSOCIATION

