Destroying the Myths surrounding Fibre Channel SAN

Introduction

There are common myths amongst the storage network community that Fibre Channel Storage Area Networks (FC SAN) are costly, complex and hard to manage and troubleshoot. Taking advantage of these myths are heavily marketed agendas to move customers away from FC SAN to new and *allegedly* more cost effective solutions.

This white paper sets out to address the reasons behind these myths and demonstrates how FC SAN is not only a high performance connectivity infrastructure but is also highly reliable, secure and scalable.It argues that many of the "problems" associated with FC SAN are, in reality, a result of poor resource utilization, lack of insight and "business as usual" working practices.

Ensuring that FC SAN performance is correctly monitored for optimum resource usage and maximizing all the available I/O resource to the entire FC SAN, utilizing I/O Virtualization (IOV) will not only improve the performance and management of existing FC SAN deployments but also allow for scalability to meet accelerating demands from organizations without the need to invest heavily in new storage infrastructures.

In summary this white paper will demonstrate that there is indeed a way to…

- Reduce the amount of physical adapters, FC cables, SAN ports and Storage ports while concurrently improving application response time, availability and SLAs.
- Simplify server FC I/O provisioning, enabling a more agile, scalable and dynamic deployment model.
- Gain the insight that, on average, most FC SAN and Storage ports are greatly underutilized by 90-95%.
- Drive higher levels of resource utilization and performance out of an already deployed solution.
- Know your FC SAN Storage related problems before they occur and shorten the time to troubleshoot your FC SAN Storage environment from days to minutes.
- Reduce the OPEX and CAPEX of your existent FC SAN Storage infrastructure to the extent that it's the most reliable, secure and cost effective solution in your environment.

Myth #1 – Fibre Channel bandwidth capacity is not sufficient

'The pipe just isn't big enough for my growing demands'

A common complaint amongst users of Fibre Channel (FC) is that the bandwidth capacity is not large enough to handle data transfer at the expected levels. One of the reasons for this is that when designing, for example, a core edge SAN fabric, a SAN architect has to apply assumptions when determining adequate fan-in and fan-out ratios. In addition there is the tendency to be overly cautious which inevitably results in overprovisioned ISLs, storage and switch ports.

The metrics displayed in Figure 1 are the inventory of an actual Enterprise environment and are typical to many Enterprise end users. What this shows is that:

- More than 50% of SAN ports are not utilized at all
- Over 40% of the HBAs have little or no traffic
- Almost 40% of the storage ports have little or no traffic.

This equates to almost 40% of the environment's FC resources being completely unused. Many end users are wholly unaware of these performance statistics in relation to their FC SAN environment.

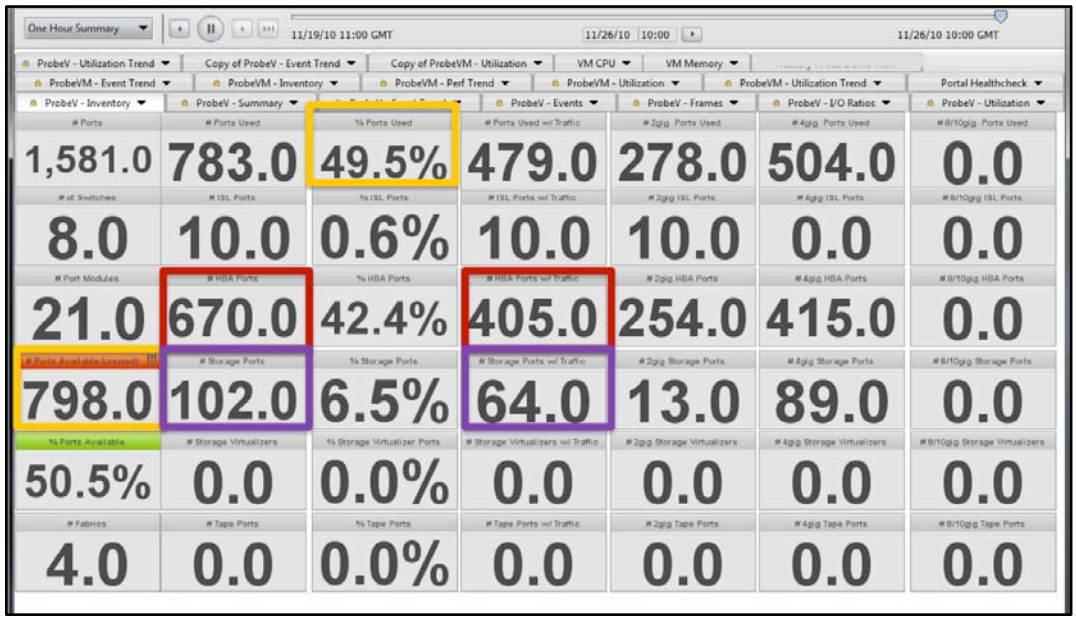

Figure 1

Another typical example can be seen in Figure 2,where we see an end user's HBA and storage port ratio utilization. We can see that only 4.8% of the HBAs are responsible for more than 82% of the reads while less than 2% of the HBAs are responsible for the writes. In a dual redundant 200 server environment this would equate to only 8 HBAs out of the 400 doing most of the writes! On the storage side we see only 31% of the storage ports responsible for nearly 90% of the reads and 12% responsible for more than

half of the writes. Such disparities typify the over cautious fan-in fan-out ratios seen in the majority of FC environments and equate to a huge waste of valuable resources and a gross under utilization of Fibre Channel's capabilities.

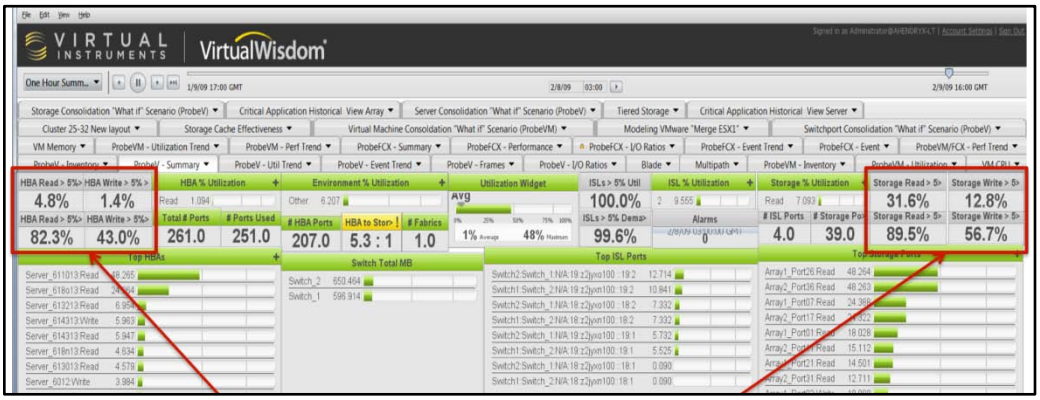

Figure 2

'The problem lies in severely imbalanced FC SAN resource utilization'

The "business as usual" approach of end users buying multiple storage ports, front end processors and cache in the belief that these are necessary to meet their applications needs will make, in fact, little difference as only a few links ever go beyond 90% utilization or exceed thresholds while the majority of storage ports rarely have more than 10% utilization.

The reason for this false idea of a lack of bandwidth availability in FC environments lies in the severely imbalanced utilization of HBAs, ISLs and Storage ports. As can be seen in Figure 3, only a few ports have a significant amount of traffic at any given time. As a result, when these minority of ports hit thresholds at peak times it creates the false impression of FC not being able to cope with increased throughput demand whereas, in reality, the fault lies with a *poorly* balanced infrastructure that fails to effectively utilize the available resources.

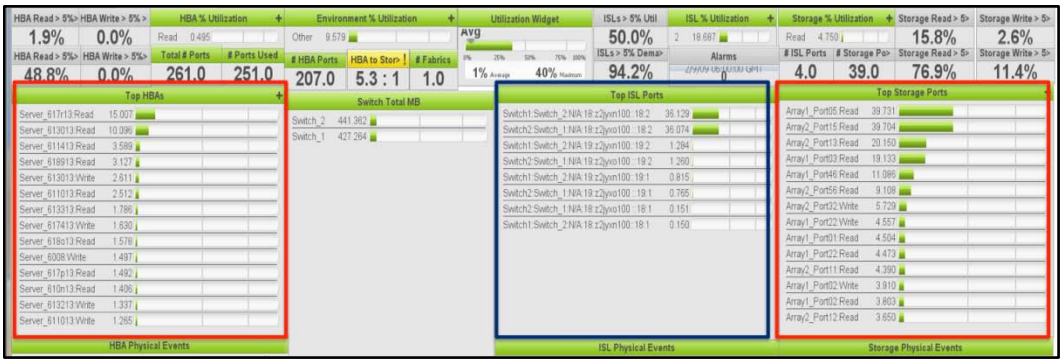

Figure 3

Myth #2 - It's difficult to deploy and maintain

'Over 70% of the server admin's time is absorbed by routine tasks'

When provisioning and managing FC, in the "typical" data center, it is commonplace to find greater than 50% of the server administrators time is absorbed by "routine" tasks – by this we mean physical cabling of servers, installing HBAs, switch port configuration and HBA firmware updates.

Figure 4 shows a traditional data center connectivity environment, including FC, where multiple I/O adapters, network and storage cables and network and storage switches are deployed. Supporting this environment with its multiple points of physical and logical management and maintenance is both resource intensive and strategically unproductive. In the mistaken belief that more bandwidth is required to support growing enterprise demands, more and more HBAs are installed in the server which further compounds this issue.

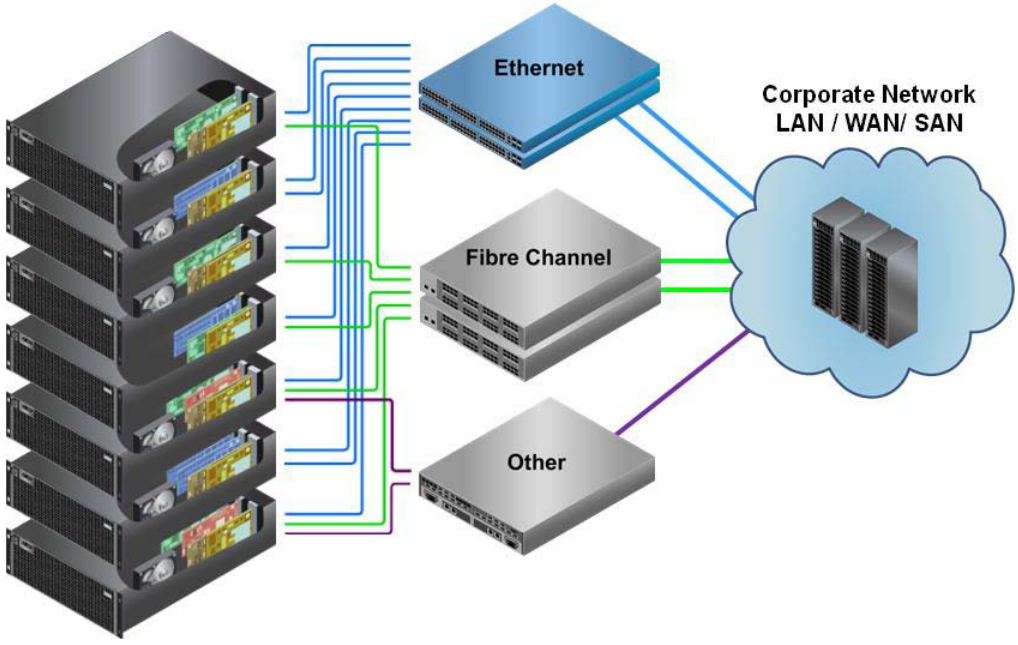

Figure 4

I/O Virtualization (IOV) removes the adapters from the servers, aggregates them into an IOV appliance as fewer high powered adapters and shares them back to the same server (and to as many as 16 server connections per appliance) as virtualized instances of those same standard I/O adapter resources. To the server's OS (operating system) they appear and behave exactly as if physical, installed local I/O adapters. Figure 5 shows the same data center environment but utilizing an IOV solution to provide connectivity. This elegant solution means that:

- Physical traditional I/O adapters are removed from the servers
- Fewer high powered interfaces are consolidated and virtualized within the IOV appliance

- PCI extension cables connect servers to the appliance, allowing “Wire Once” for all Ethernet and FC ports (provides a typical 4:1 cable consolidation and at least a 2:1 card reduction)

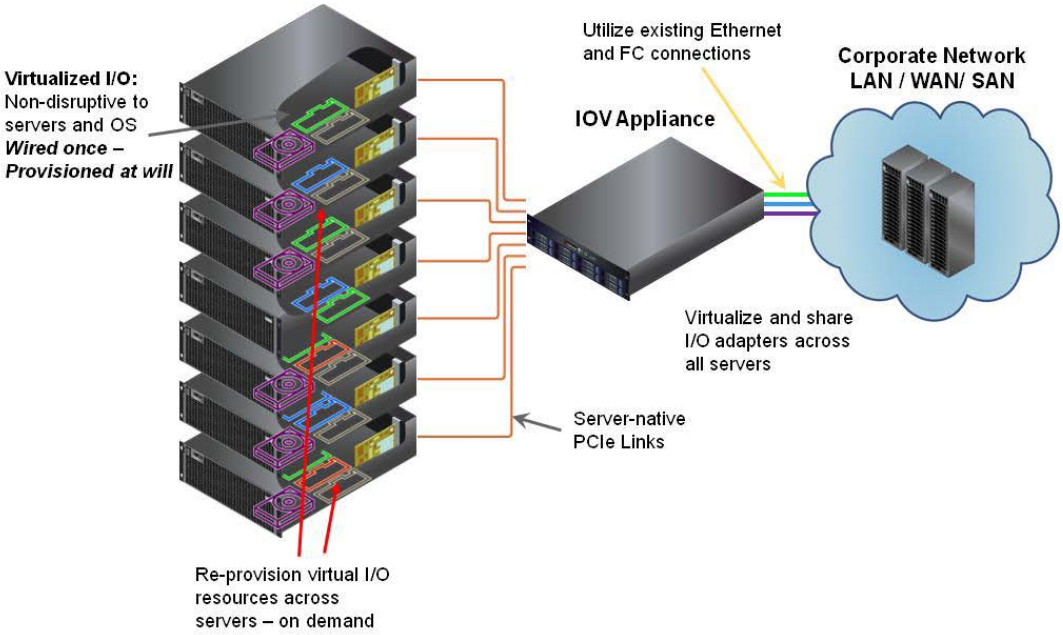

Figure 5

‘IOV allows fewer HBAs to provide the same or better levels of connectivity’

In the case of the FC HBA, it is now replaced with a passive low power, low profile PCIe extension card which carries traffic from the PCIe bus in the server via a PCIe extension cable to the IOV appliance at native bus speeds (PCIe Gen 2 x4 functions at 20 GigaTransfers per Second or 16 Gbps in raw PCIe bandwidth). In addition redundancy is possible to further increase reliability. Utilizing a virtual management console, it is possible for the server administrator to provision and re-provision virtual I/O resources across servers on demand.

As we have already seen, utilization of installed HBA resources in storage network environments is typically highly inefficient, with up to 40% completely unused. I/O Virtualization allows fewer of these expensive FC HBAs to provide the same or even higher levels of utilization.

Myth #3 - It's Difficult to manage and troubleshoot

'Averaged metrics fail to give a true insight making FC SAN troubleshooting problematic'

The advent of virtualization and the growing nature of SAN fabrics and heterogeneous storage arrays within environments, plus the added abstraction that these new layers bring has led to FC SAN troubleshooting being deemed problematic. Furthermore it can take hours if not days to resolve issues, especially if the cause or problem is something difficult to locate, for example, a flapping HBA, a damaged FC cable, an oily finger tainted SFP or even worse simply a FC cable that is bent too far at an angle leading to added latency and consequent poor application response time and performance issues.

What makes identifying such hard to locate and problematic issues even more difficult is that the host based monitoring and SRM tools available depend on averaged metrics that fail to pinpoint intermittent issues or latency down to the millisecond. Also with such device specific views, true I/O latency can be missed when these issues occur intermittently or only during peak periods and therefore beyond the scope of the 'averages'. In Figure 6 for example, this 1 minute interval graph would be reported by a host based or SRM tool as an averaged out 5MB/s. The highlighted intermittent peaks of 20MB/s would be missed leading to a failure in correctly diagnosing the root cause of the problem. It is due to this lack of insight that troubleshooting inherently takes so long and that FC OPEX costs are consequently considered so high.

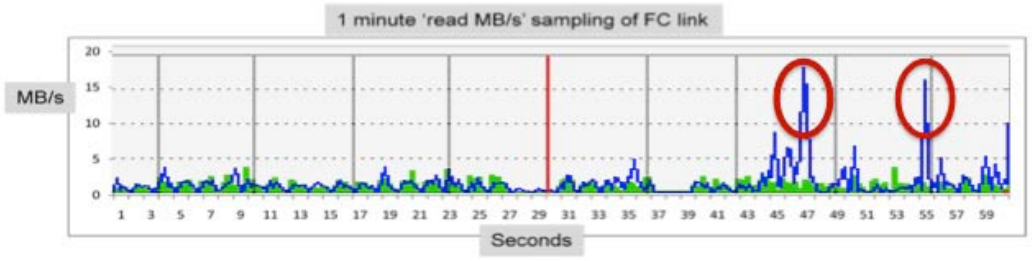

Figure 6

To resolve this it's vital to have a platform that reports in real time and on what is termed 'Exchange Completion Time'. By being able to measure the Exchange Completion Time, administrators will be empowered with the knowledge of the exact amount of time it takes for a complete Exchange i.e. the time that all FC Frames and Sequences have passed from the host to the Storage Port to the LUN and acknowledged back to the host for any given transaction.

Deploying the Virtual Instruments architecture as shown in Figure 7, every single FC Frame can be measured in real time via the deployment of TAPs that non-intrusively split the FC light to the ProbeFCX analyzer. Coupled with the metrics gathered from vCenter via the ProbeVM software, and the ProbeV software which takes via SNMP from the SAN switches, the TAPs and the ProbeFCX provide a comprehensive, real-time monitoring system that is out-of-band, causes no latency and is able to report on application I/O latency down to the millisecond.

'Pinpointing physical layer issues enables the proactive remediation of FC SAN issues'

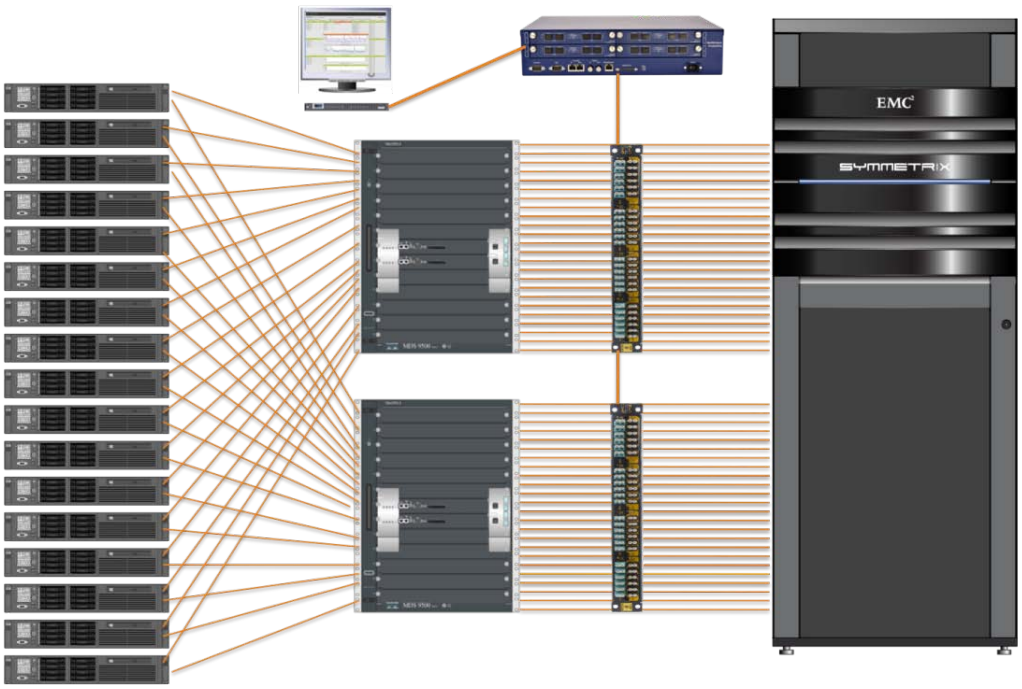

Figure 7

Such instrumentation also enables the identification of physical layer issues in the FC SAN environment such as a bent or damaged cable or even a flapping HBA or SFP. This leads to proactive remediation of such issues before they significantly affect application performance. Additionally slow draining devices can be quickly detected by monitoring any significant increases in pending and minimum pending exchanges. By optimizing the Queue Depth settings, which are the most common cause of credit issues, the problems caused by slow draining devices can be quickly mitigated. In this context the ProbeFCX enables users to see the true impact of Queue Depth settings and overall latency in the environment.

Additionally using historical playback, application performance during different periods can also be monitored and base-lined. So for example looking at an application's Exchange Completion Time as displayed in figure 8, a Storage Administrator can quickly identify from these red peaks that latency has been increasing slowly over the last few hours. This granularity would be missed by the application owners due to their dependence on averaged metrics.

'Know the I/O latency of Mission Critical Applications down to the millisecond by measuring Exchange Completion Times.'

Figure 8

A corresponding time aligned view as can be seen in Figure 9, quickly identifies the root cause being a large amount of SCSI reservation conflicts, Pending Exchanges and cancelled transactions that have gradually built up. Via a customized alarm the Storage Engineer would proactively identify and remediate such issues before they caused any problems on the application layer leading to significant time-savings in FC SAN troubleshooting and consequent OPEX costs.

Figure 9

Myth #4 - It's too costly

It is well known that FC SAN is expensive to deploy with SAN switches, storage arrays and their many front end port and processors. The first fundamental question that must be asked is how many storage ports and processors are actually required?

'How many HBAs, switch and storage ports do you actually need?'

With the ability to measure Exchange Completion Times, end users can accurately determine the right amount of storage processors and storage ports they actually need by aligning it to their application I/O demands. This can lead to FC end users consolidating and reducing their storage port acquisitions considerably while concurrently optimizing their performance.

For example as can be seen in the scenario in Figure 10, we have the metrics of 2 applications that are currently on 2 separate storage ports. Prior to consolidating we can use the modeling feature to see how these two applications would perform once consolidated onto the same storage port. Using the live data and historical playback feature we can see how these applications would also perform during peak times, ensuring consolidation and optimization without risking SLAs.

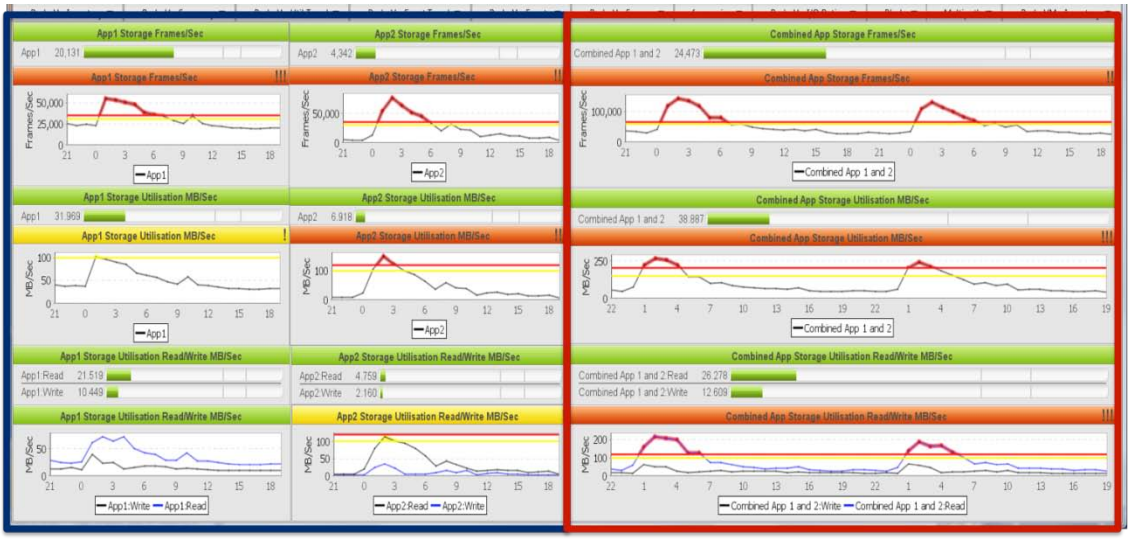

Figure 10

Another mistaken belief, that FC is expensive to deploy, is due to the high cost associated with over provisioned and under utilized SAN Switch ports and HBAs. Provisioning I/O is time and resource intensive involving "physical touching" of servers by up to 5 times:

1. Installing the FC interface
2. Provisioning the network ports

3. Terminating cable connections
4. Configuring the Hypervisor for the new I/O
5. Configuring the server OS for the new I/O

This leads to the second fundamental question - How many HBAs and Switch ports do you actually need? I/O Virtualization consolidates and virtualizes traditional server network and storage I/O connectivity for both physical and virtualized server environments and simplifies management of the resulting FC SAN. Implementing an I/O Virtualization solutiondelivers:

'Lower TCO through a reduction of I/O port count (8:1)'

- Simplification in bandwidth provisioning to servers through utilization of a smaller number high powered virtualized HBAs.
- Lower TCO through a reduction of I/O port count (8:1) and streamlining of the cabling infrastructure.
- Simplification of I/O provisioning, by "wire once - provision at will" and migration of server personalities between physical serversallows for "move, adds, & changes" to become more dynamic/flexible.

Figure 11 shows the simplified environment of a typical I/O Virtualization deployment when compared to a "traditional" deployment.

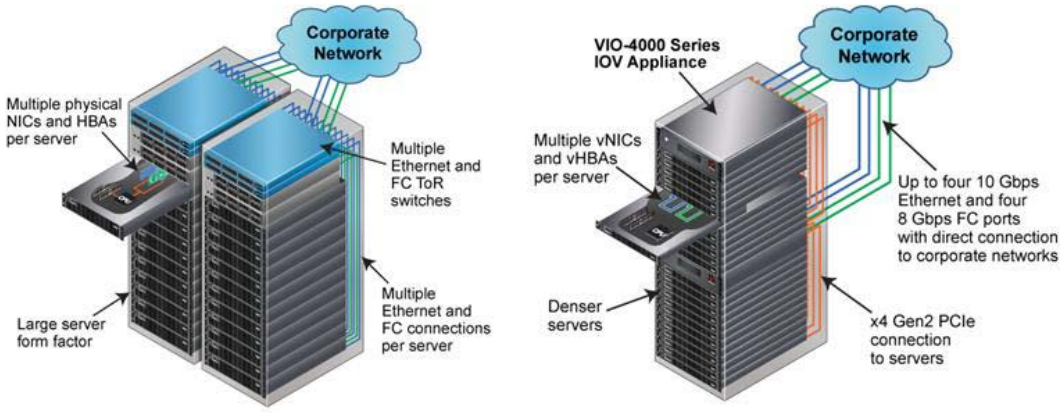

Figure 11

By measuring the actual utilization requirements and correctly balancing your resources, deploying an IOV appliance can potentially consolidate the SAN switch ports requirement down to an 8:1 ratio. Furthermore end users that have previously purchased 96 port Storage Arrays have found that by measuring Exchange Completion Times and hence aligning their application I/O demands to correctly, actually get better performance with 32 port Arrays. This brings significant cost savings, room for more scalability and a vast improvement in the utilization of FC resources. An example of such a consolidated set up can be seen in Figure 12.

'Consolidation of expensive FC resources can actually lead to optimization of performance'

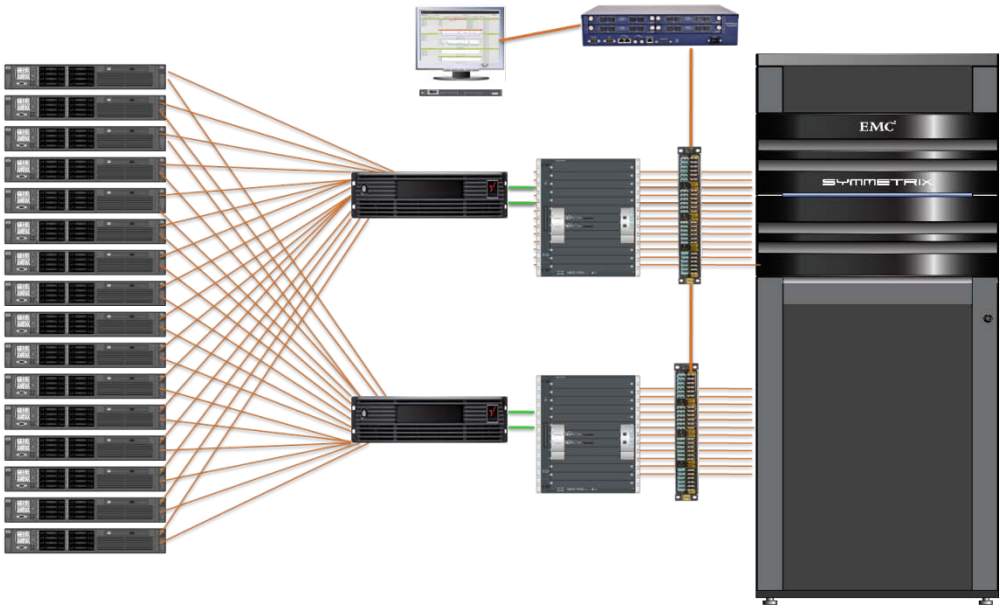

Figure 12

Such consolidation not only eases FC SAN management but also troubleshooting due to the reduction of physical devices and potential points of failure. To further optimize such a set up it is important to understand the utilization metrics of not only each switch port but also the Switch Blades and their corresponding ASICs. For example Figure 13 shows metrics that will allow SAN and server engineers during either post or pre consolidation to correctly allocate the right SAN ports for their servers without risking traffic congestion and performance problems. This action would counter many of the problems faced as a result of the SAN switch's inherent oversubscription rates. Furthermore as such metrics are recorded they can be played back and referenced each time a new virtual or physical server is provisioned, thus enabling optimized performance during busy periods.

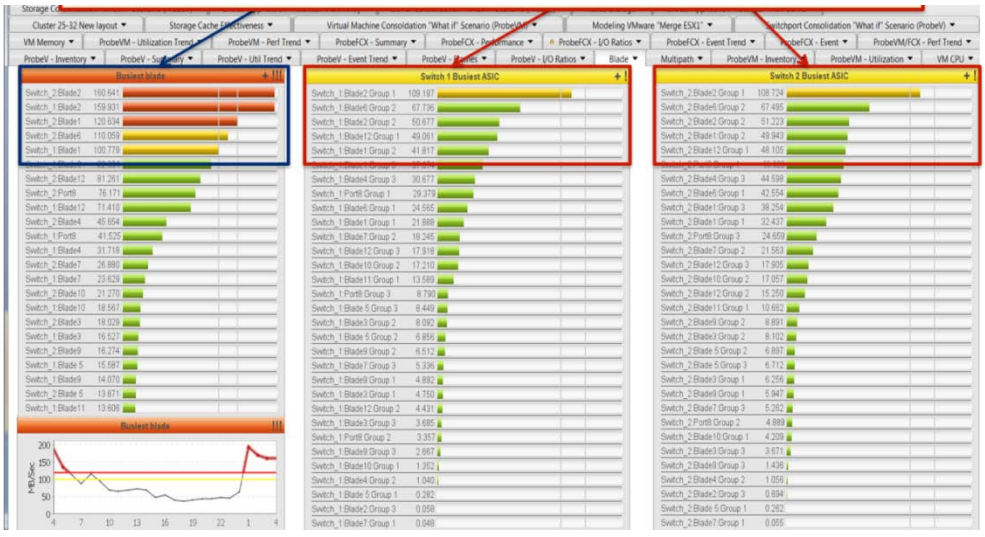

Figure 13

Summary

Without considering and proactively monitoring for slow draining devices and physical layer issues in your FC SAN it is wrong to assume that the problem lies with 'there not being enough pipe.' Indeed once a FC SAN has been de-risked and is healthy, a clear insight into the utilization of HBAs, SAN switch and Storage ports still shows that the mere balancing out of resources can easily meet the highest of demands.

By using IOV technology HBA resources can be pooled into a dynamic model that share I/O resources in accordance to demand enabling better usage and balance of an already deployed environment. Furthermore the reduction of physical devices that this brings simplifies the provisioning of FC resources and quickens deployment as well as increase the scalability of an existent infrastructure. This allows end users to gain far more from what was initially considered a costly investment into FC.

Additionally monitoring your FC SAN environment in a comprehensive manner that incorporates the SAN fabric and provides metrics such as the Exchange Completion Time rapidly changes FC SAN troubleshooting from a reactive to proactive exercise reaping significant OPEX savings. It also enables Server, Storage and Application admin to have a common language of 'response times' thus eliminating any potential silos. With the knowledge of application I/O latency down to the millisecond, FC SAN Storage administrators can quickly be transformed from the initial point of blame to the initial point of resolution.

The consolidation that IOV technology and SAN I/O monitoring bring to a FC SAN infrastructure not only optimizes performance and utilization of resources but also brings significant CAPEX and OPEX savings. This ensures that the extreme waste of resources that are currently deployed in the majority of FC environments are put to use enabling a far greater ROI and an eradication of the idea that FC is too costly.

About Virtensys

Virtensys' award-winning VIO-4000 series I/O virtualization appliances are changing the way IT managers manage and deploy I/O resources to standard rackmount-based servers, enabling a more agile, scalable and dynamic data center. Built upon Virtensys' patented PCIe I/O virtualization technology, the VIO-4000 product line consolidates and virtualizes traditional server network and storage I/O connectivity for both physical and virtualized server environments. Server administrators can now simply "wire-once" and then remotely manage and provision their I/O resources as needed. For more information, visit **www.virtensys.com** or **www.facebook.com/virtensys**

About Virtual Instruments

Virtual Instruments is a SAN and Virtual Infrastructure Optimization company that offers the award-winning VirtualWisdom and SANInsight virtual infrastructure optimization solutions. They offer a combination of software and hardware products that provide comprehensive, real-time I/O instrumentation and measurement capabilities that enable IT managers to optimize the performance, utilization, and availability of virtual and physical IT infrastructure by analyzing SAN I/O traffic data. For more information visit **www.virtualinstruments.com**