


Accelerating Storage Tiering Whitepaper



Performance-Aware Storage Tiering with Virtual Instruments

Today's monolithic storage systems are designed to deliver massive scalability—enabling systems that scale to hundreds of terabytes of storage and tens of millions of IOPS, supporting hundreds of application servers. But many enterprise applications do not need the power and cost of these systems. That's why most enterprise IT departments are moving many applications to secondary tiers by what is generally referred to as "static storage tiering". Static tiering forces IT staff to make best guesses about performance even though the penalty for a bad decision can be a multi-week long migration project, best case, or really angry users, worst case. This is why many storage professionals are reluctant to move to secondary tiers even though the cost advantages can be huge. A more scientific analysis is required to lessen the risk.

With their capacity and performance analysis tools, storage vendors offer tiering automation. Automation is provided by performance and capacity measurement and trending, plus storage virtualization techniques which ease the movement of data. This paper describes how Virtual Instruments VirtualWisdom complements and adds critical intelligence to these storage vendor tools to significantly improve re-tiering decisions and accelerate realization of lower storage costs.

VirtualWisdom Adds:

- **Drastically reduced risk of performance problems due to tiering decisions.** Provides end-to-end data path visibility tracking for every I/O from every initiator to every storage system in continuous real time. Adds the one key metric, Exchange Completion Time, that definitively measures the effect of tiering
- **Single pane of glass for SAN health** Provides all the required information on a single dashboard to clearly show whether an application performance problem is with the SAN or not, for real-time root cause analysis
- **Historical trending** Provides historical performance, error, alerts and event trending information for end-to-end application I/O conversations
- **Quick resolution of performance or component issues** Alerts based on trends or thresholds quickly enable IT administrators to pinpoint performance issues and intermittent failures and take the appropriate action(s) to correct the problem

Background


Companies undertaking a tiered disk strategy are doing so to drive down operating (OPEX) and capital expense (CAPEX) costs of their storage infrastructure. However, many of these initiatives fall woefully short of expectations because the business units and storage managers are reluctant to aggressively pursue the use of lower tiers of storage due to the lack of metrics needed to baseline the performance requirements of the applications and to assess the impact of storage tier changes.

In addition, poor I/O performance is often attributed to the storage arrays when in fact other bottlenecks in the storage infrastructure (i.e. congestion and/or misconfiguration) are the root cause. Due to insufficient monitoring of storage transactions these issues are not detected or corrected, resulting in more applications being assigned to upper storage tiers in an attempt to “fix the problem” even when the problem is not the array!

VirtualWisdom from Virtual Instruments provides storage managers with the ability to measure and monitor storage I/O performance, thereby ensuring that applications are assigned to the most cost-effective storage tier while meeting the performance requirements of the business. In addition, VirtualWisdom identifies other conditions within the SAN infrastructure that ultimately limit the ability of the storage arrays to deliver the desired performance. With VirtualWisdom, the net effect is a balanced and properly configured SAN infrastructure that fully supports tiered storage initiatives and ensures the businesses’ storage needs are met in the most cost-effective manner.

Storage Tiers

The storage array features that are typically used to define “storage tiers” include: performance (e.g. amount of cache, speed of disks), availability (e.g. RAID levels, active-active controllers) and business continuity functionality (e.g. snapshots, replication, etc.). In most cases the IT storage architecture team develops a set of standard storage tiers associated with specific storage technologies. From a performance perspective, the general assumption is made that that Tier 1 storage has better performance than Tier 2 and Tier 2 has better performance than Tier 3. This assumption is based on the vendor’s specifications (e.g. supported number of IOPs and bandwidth) and features such as the types of disks used (FC versus SAS versus SATA) and number of array controllers and cache. The following table shows a typical tiered storage definition.



Array Characteristic	Typical Spinning Disk-based Storage Tiers		
	Tier 1	Tier 2	Tier 3
Usable capacity	> 1PB	600TB	300TB
Disk drive size	300 GB, 500GB, 1TB	1TB	2TB
Disk drive connectivity	Fibre Channel	SAS	SATA
Disk drive speed	15K RPMs	10K RPMs	7.2K RPMs
Front end FC ports	128	16	4
Array cache	512GB	64GB	16GB
Replication	Sync and async	Sync only	Batch
Snapshots	Full point-in-time mirror copy-on-write	Copy-on-write	None
Thin provisioning	Yes	No	No
Deduplication	No	No	Yes

Business Application Tiers

On the business side, business applications are assigned a specific “application tier” based on the value of the application/data to the business. For example, a revenue-generating application is classified as “Tier 1” while email is classified as “Tier 2” and a marketing data warehouse may be classified as “Tier 3”. Additionally, flash / SSD-based “Tier 0” is being deployed for extremely latency-sensitive applications. Each application tier also tends to be associated with a set of common “data management/data protection” requirements such as recovery time object (RTO) and recovery point objective (RPO). These requirements drive specific characteristics of the storage that is to be used such as: level of RAID protection, replication requirements and backup and recovery needs. For example, a Tier 1 application might require RAID-6 Fibre Channel storage with synchronous replication and local snapshots taken every six hours.

Rarely is there an explicit performance requirement associated with the business applications. Most often an implicit relationship is made between each application tier and the performance requirements of applications assigned to each tier such as: Tier 1 application requires Tier 1 storage performance (i.e. “the best” performance), Tier 2 application only requires Tier 2 application performance and Tier 3 only requires Tier 3 storage performance.

Assigning Tiered Applications to Storage Tiers

This assignment of storage tiers to each application tier typically occurs without any real analysis (measured or estimated) of the actual performance requirements of the application. So, while an application may be classified as Tier 1 from a business perspective, the performance requirements of that application may only necessitate Tier 2 (or possibly Tier 3) storage. The opposite may also be true; for example a Tier 2 application should utilize Tier 1 storage because of its performance requirements.

However, the evolution of storage array technology is now making it possible to better align the applications within a business with the most cost-effective storage technology. Some of these storage technology innovations include:

- Solid state memory disks
- Multiple physical disk technologies (e.g. Fibre Channel and SAS) within a single storage array
- Improved performance (IOPs, bandwidth) and data protection features (e.g. asynchronous replication, copy-on-write snapshots) of Tier 2 storage arrays
- Virtualization of physical storage (e.g. IBM SVC, HDS USP-V)

So, now it’s possible for example, to take a Tier 1 application and assign it to Tier 2 storage since the Tier 2 storage technology often has the functionality to meet the business requirements (e.g. replication, snapshots and RAID protection) of the application. However, there is still often a risk of unfulfilled performance requirements. Will the Tier 2 storage be able to satisfy the performance requirements of the application? The only way to guarantee performance is to measure the effect of I/O performance on the application, and establish an SLA based on this metric.

More Effective Tiered Storage

While tiered storage has great potential for reducing storage OPEX and CAPEX costs, most companies fail to realize these potential cost savings. There are several factors that contribute to this failure including:

1. There is a heavy reliance on the storage subsystem performance specifications provided by the storage vendors rather than real, measured performance

2. There is a significant level of reluctance to moving off higher tier storage technology by both the Lines of Business as well IT management because it represents a risk, especially as it relates to performance and
3. The storage arrays themselves, regardless of which tier is being used, often are not constraining performance; the constraint often exists in another part of the SAN infrastructure and cannot be overcome by even the fastest performing arrays!

All these factors can be addressed if the SAN infrastructure is properly monitored and performance metrics collected from the environment and analyzed prior to making these tiered storage decisions. For a performance-based tiered storage strategy the metric that is fundamental to a deployment meeting the needs of the business is I/O response time measured at the host. In the Open Systems world, where SCSI is the upper layer storage protocol of choice, I/O response is often referred to as SCSI Exchange Completion Time (or ECT). This is where Virtual Instruments' VirtualWisdom SAN Performance Probe technology comes in. SAN Performance Probe deciphers the SCSI exchanges contained within the Fibre Channel frames exchanged between the host's HBA and the storage array LUN. From this information, VirtualWisdom is able to compute the overall ECT and instantly report on this response time metric as well as maintain a historical record of this performance. This is key information when developing a performance-based tiered storage model based on an application's business requirements AND I/O performance SLAs. With this data, the storage engineering and operations teams can now align an application's real business requirements with the storage technology that meets those needs. For example, if an application is moved from Tier 1 to Tier 2 storage, the impact on response time can be measured in real time and compared to the historical performance and defined SLA.

The keys to value realization of a tiered disk storage strategy are:

1. Ensuring the business applications (and associated data) are matched to the most cost-effective storage technology that meets the needs of the business. Because performance is one of the top factors in making the decision regarding what storage tier to use, it is imperative that performance (I/O response time) be measured and tracked in both real-time and historically. Virtual Instruments' VirtualWisdom technology is the only solution that can provide these performance metrics automatically on a consistent and continuous basis.
2. Ensuring other components within the SAN infrastructure do not constrain performance either due to congestion or misconfiguration. Using VirtualWisdom to monitor the SAN infrastructure will allow the storage operations team to pre-emptively detect these SAN issues before they negatively impact performance and SLA targets are missed.

VirtualWisdom can be used to look at performance (in real-time and historically) in many ways including: read and write performance, overall performance for a single server, performance for a single LUN on a single server, performance of a group of servers (e.g. all Exchange mail servers), average response time, peak response time, over different time ranges (by minute, by hour, by day, by week, by month). Therefore, SLAs can be defined that are associated with as narrow or as wide a span of time and servers as is desired. For example an SLA can be established that says that the average write response time for all mail servers over a 7-day period must be less than 35ms.

Storage Vendor Tools

With storage vendor tools, LUN level activity is monitored, and LUNs or even "sub" LUNs are dynamically relocated based on perceived performance needs, moving "hot" LUNs to higher-performance drives and "cold" LUNs to higher-capacity SATA drives. Users can provide for dedicated resource memory, enabling more predictable tiered storage performance; and they can prioritize application workloads, allowing higher priority applications to get faster response times. Performance and capacity monitoring tools collect and store historical array performance and

capacity data, providing insufficient insight for more optimized tiering decisions, which will be explained in the next section.

Further, typical logic of vendor tiering is based on the frequency of data access. The rationale is that data that is frequently accessed is more valuable so it should reside on higher storage tiers and infrequently accessed data should be moved to lower tiers. This simplistic algorithm ignores the performance requirements and implications of tiering.

Most SAN management products are focused on capacity optimization, but capacity cannot be expanded without a comprehensive view of how it affects performance and end-user response time.

Impact of the SAN on tiering decisions

In the days of direct-attached storage, when the I/O performance bottleneck was disk transfer rate, these storage vendor tools would be sufficient to determine where to place data for an optimum balance of performance and costs. But today, in the world of shared storage, typically via SANs, it's much too simplistic to look just at the difference in external transfer rate between a Fibre Channel disk and a SATA disk. An analogy we often use is one most of us can relate to - traveling to work in a typical metropolitan-area traffic jam. Having the fastest car in the world is typically not going to reduce your commute time. Knowing the readings on your car's tachometer is useful data about how the car is performing, but seeing the speedometer (miles per hour) in real-time is much more relevant to your commute time. The speedometer reflects the current traffic conditions, weather conditions, how many lanes are available, etc. This analogy is almost certain to be representative of your SAN. But even the speedometer is an imperfect tool.

In the SAN, you are guaranteed to have bottlenecks somewhere. At best, most SAN and storage management products report IOPS or MB/s, which are primitive measures of true storage performance. By far, the best measure of SAN performance is the effect on application response time for every transaction. Going back to our automobile analogy, simply looking at IOPS or MB/s is like looking at an automobile speedometer, and guessing how long it takes to go to the market and back for a loaf of bread. And with most commonly used tools, it's even worse than that. Legacy performance tools might look at a number of trips to the market and tell you what your average speed was, and they won't tell you if one of the trips took 50% longer than the others, or if one of the trips resulted in a crash. Using VirtualWisdom to look at application latency is like having a stopwatch and reporting on exactly how long each trip takes to get to the market and return with the bread. And with a latency measurement you would know, in real time, if one of those trips is taking 50% longer than normal. Virtual Instruments gives you the complete picture, not just the tachometer or speedometer reading. Storage vendor tools most often only look for bottlenecks within the storage arrays; they have no view into SAN-wide performance information or transmission problems. Those that do report on application latency cannot break out the effect of the SAN from the effect of the host.

In our experience, there are much bigger factors to consider when moving data from tier to tier, factors that only come through a comprehensive understanding of the complete I/O path. Congestion in the fabric can cause response time to be 100 times (or more) slower. Not having HBA queue depths set properly can have a 10 times response time difference. If the disk drive doesn't know that a sequential request is actually sequential because of the way it is configured, it could have an 8 times or greater response time effect. Errors on the links can cause sporadic impact that can vary from having no effect to having response times 3,000 times as long. While the storage vendor tools do a superb and necessary job of measuring the array-based metrics, they are not designed to measure the actual effect of the tiering change on the SAN, and the one metric that matters – application latency. When the tiering decision is based largely on performance, and not array reliability or availability, you need a solution that will track I/O conversations across the SAN infrastructure, not just within a specific device.




Virtual Instruments' VirtualWisdom

Virtual Instruments is the missing puzzle piece of the tiering value proposition. Drive speed is only one factor in achieving the desired application performance for a tier and is often the smallest, most expensive performance differentiator. VirtualWisdom tracks I/O conversations to provide both real-time monitoring and trend analysis of overall performance between enterprise applications and all components of the Fibre Channel SAN – the host HBA, switches, virtualization appliances, storage ports, and LUNs. Combined with the storage vendor tools with their emphasis on storage system metrics, VirtualWisdom provides the critical missing data to enable IT organizations to more confidently use lower cost storage to provide higher tier performance. Virtual Instruments adds application latency data and other information to properly plan and optimize the environment. This is combined with alerting capabilities to prevent user impact as demands change, enabling IT to realize huge CAPEX improvements.

Beyond adding full I/O path awareness to let IT make more intelligent performance-based tiering decisions, VirtualWisdom also provides deep, physical-layer knowledge of the fabric to enable actual problem avoidance, avoid costly fabric over-provisioning, accelerate troubleshooting, improve VMware server utilization, and help accelerate the move to a cloud computing model.

Summary

Effective tiered disk storage initiatives must accommodate both the business and performance requirements of the application data. In addition, visibility into the complete server-to-storage data path is required in order to ensure that bottlenecks are not introduced that render the performance features of the storage arrays ineffective. VirtualWisdom is the only technology available today that provides this rich suite of performance and health monitoring of a Fibre Channel SAN, including linking the performance metrics with virtual machines in a VMware environment. With VirtualWisdom, businesses can now unleash their tiered storage initiatives and finally realize the full financial benefits that have eluded them for so long. When an organization applies a more aggressive tiering strategy based on actual measured performance data, not only are capital costs reduced, but also maintenance, power, cooling, and floor space costs are all reduced by 50% or more.

	Corporate Headquarters 25 Metro Drive Suite 400 San Jose, CA 95110 Phone: 408-579-4000 Fax: 408-579-4001	Sales sales@virtualinstruments.com Phone: 408-579-4081	Support support@virtualinstruments.com
--	---	---	--

©2011 Virtual Instruments. All rights reserved. Features and specifications are subject to change without notice. VirtualWisdom, Virtual Instruments, SANInsight are trademarks or registered trademarks in the United States and/or in other countries. All other brands, products, or service names are or may be trademarks or servicemarks of, and are used to identify, products or services of their respective owners. 09/11