Nine Steps to Effective Networked Unified Storage Consolidation

Marc Staimer, President & CDS of Dragon Slayer Consulting
marcstaimer@comcast.net 503-579-3763

Introduction

Is there an IT storage administrator left on the planet that has not heard about why he or she should consolidate their storage? More importantly, how networked unified storage consolidation makes it so much easier for the storage admin to manage the ever-swelling amounts of storage while simplifying the data protection? If so, there are not very many. And candidly, networked unified storage consolidation benefits are unquestionably and quantifiably real.

Of course, most vendors, system integrators, resellers, trade press, and pundits paint networked unified storage consolidation as a storage admin panacea. Repeatedly, there is the unstated assumption (now and again explicitly stated) that the networked unified storage consolidation process is simple and easy. In some cases, it ***definitely*** is. More often than not, it is neither simple nor easy. It can in fact be excruciatingly painful.

Many IT organizations have experienced frustratingly long networked unified storage consolidation projects. Projects that had been forecast to take hours or days fell far short of the mark. Many of the networked unified storage consolidation projects extended into weeks or months. Vendors set expectations about performance, savings, costs, ease of data migration, flash cutovers, ongoing manageability, and data protection just did not match objective reality. They overcommitted and underperformed, leaving a sour taste in the mouths of those IT organizations.

How could this have happened and more importantly, continue to happen? Simply put, vendors and administrators alike did not know what they did not know. Perhaps it can be explained by the very old anecdote about a Rabbi, Priest, and Minister fishing together on a small rowboat in the middle of a lake.

> *While fishing, the Priest suddenly realized he had left his best lure back on the dock. Rather than asking his friends to row back, he jumped out of the boat, ran across the water to the dock, retrieved his lure, then ran back across the water and got back into the rowboat. The Minister was unphased by the Priest's miraculous ability to walk on water, whereas the Rabbi was completely flustered. Noticing that the Minister said nothing, the Rabbi decided to keep his own consul and inquire about it later. A short while later the Minister accidentally dropped all their bait out of the boat. Feeling guilty, he jumped out of the boat, ran across the water, raced up to boathouse, bought more bait, ran back across the water, and climbed into the boat. The Rabbi was flabbergasted and too stunned to speak. He thought about what he had just witnessed and decided that he was just as pious as his friends. He figured if they could walk on water, so could he. With mock exasperation, the Rabbi exclaimed he had forgotten his water bottle. Before his friends could say anything, he jumped out of the boat into the cold water of the lake and immediately began flailing and splashing about. The Priest turned to the Minister and said: "I think we should have told him where the rocks were."*

The purpose of this white paper is to point out where those networked unified storage consolidation "rocks" are located, demonstrate how to get optimum value out of networked storage consolidation, and finally how to avoid the "gotchas" that are so maddeningly frustrating.

Quick Review – The Benefits of Networked Unified Storage Consolidation

Networked storage consolidation is designed to significantly reduce the escalating cost of storage systems. It accomplishes this task by enabling storage resources to be pooled in a dedicated storage network. Networked storage resources can be shared far more effectively with heterogeneous servers, which increases storage utilization. Increased storage utilization means more storage targets can be consolidated into fewer larger storage targets. Fewer storage targets reduce both capital expenses and operational expenses. Reduced capital expense is a direct result of less acquired storage targets including all the associated hardware and software components. Reduced operating expenses come from having to manage less storage systems, fewer support contracts, and higher storage admin productivity in managing more storage with fewer people.

Unified storage is the combination of network-attached storage (a.k.a. file-based storage or NAS) and block-based storage (a.k.a. SAN storage) providing the best of both environments. NAS is typically deployed for its simplicity and shareability. SAN is typically deployed for higher performance and scalability. Some applications are optimized for SAN targets (Microsoft Exchange and various DBMS systems). Other applications are optimized for NAS targets (CAD/CAM, file based data, unstructured data, and various DBMS systems). Unified storage is optimized for both. Having a unified target provides quantifiable benefits. Unified networked storage leads to much higher storage utilization resulting in reduced storage target acquisitions. Managing fewer storage targets reduces operating expenses from improved storage admin productivity, less maintenance, less service, less power, less cooling, and less real estate.

Networked unified storage consolidation should then provide the best of both, shouldn't it? Success as with most things in life, is dependent on multiple factors.

It comes down to rolling up the sleeves, digging in, and completing an effective nine step detailed networked unified storage consolidation plan. The first step identifies and establishes the current storage performance (IOPS and throughput), plus capacity requirements (by application), as well as those for the life of the networked unified storage. Step two calculates the current and future storage network infrastructure requirements such as technology, bandwidth, ports, data protection specifications, (snapshot, incremental snapshot, synchronous and or asynchronous mirroring, etc.). Step three establishes the organizations service and support requirements. Step four ascertains the staff skills, their training requirements, and maximum individual load productivity capabilities in managing networked unified storage.. Step five establishes the budget for the networked unified storage consolidation (both capital and operating budgets). The budget is then compared with the total cost of ownership of each of the vendors' solutions. Step six puts all of the meticulously established requirements into a "weighted" criteria matrix. This matrix is then used to ascertain the vendors and products that best meet the requirements while winnowing the list down to two finalists. Step seven is a realistic clearly defined test of the top two vendors in a bakeoff preferably using production data. Step eight negotiates with both (yes, both) vendors to come up with the best price/performance contract that best meets the requirements. Step nine establishes the detailed implementation plan for installation, data migration, testing, and production cutover by application.

This nine-step plan may seem like too much work. However, the payback in reduced headaches, saved time, lower costs, and much decreased back-end work is in the end, well worth the trouble.

Step 1: Networked Unified Storage Consolidation Capacity & Performance Requirements

Many IT people focus on capacity as their primary concern. Capacity is important; however, the latest high density Fibre Channel, SAS, and SATA disk drives have made it relatively simple to find storage systems that have capacity in the hundreds of terabytes. A bigger issue is how and when that capacity must be purchased and allocated, which has direct correlations to both performance and the total cost of ownership (TCO).

Larger capacity drives means fewer disk heads per terabyte, resulting in lower IOPS and throughput. Lower capacity drives means more disk heads per terabyte and higher RPMs, resulting in higher IOPS and throughput. The latter case means higher performance and higher TCO.

When the networked unified storage system's capacity is purchased also directly affects TCO. Disk drives have historically decreased an average of 10% per TB per quarter. If both the current and future networked unified storage capacity must all be acquired up front it will be purchased at the highest cost per TB. Also, new software technologies such as "thin provisioning" allow current and near future capacity to be purchased now while future capacity can be added as needed, on-demand, allowing that capacity to be purchased at the lowest possible cost.

Thin provisioning tells the attached file system it has the maximum defined capacity requested. In reality, it has much less and only adds physical capacity when policy-based thresholds are met. When physical capacity is added, it is done so transparently to the file system.

Performance is a much more important requirement factor than capacity and much more difficult to measure outside of testing. It requires more understanding, knowledge, planning, and foresight. Mostly, it means not being dazzled by technology and instead, staying focused on the specific application requirements.

There are numerous factors that affect performance, including, but not limited to: the number, technology type (SCSI, FC, SAS I, SATA I, SATA II), and speed (RPM) of hard disk drives (HDDs); networked unified storage target internal architecture; bandwidth through the networked unified storage target; efficiency of the ports IO; number of ports; amount and speed of the cache; processing MIPS, the NFS and CIFS read and write processing capabilities of the networked unified storage target; and software efficiency. The sum or end result of all these factors is far more important than any individual specific technology..

Many vendors attempt to point at their CPU clock rate, the latest x86 chip, a special ASIC, or FPGA to show how "fast" their system is. The truth is a bit more complex. It is not the individual components that determine performance or efficiency; it is sum of all the parts. There is one factor that has proven to provide both noticeable and quantifiable system performance increases. It comes from the amount of the network unified storage operations that are performed in hardware (ASICs or FPGAs) versus the amount in software. This is especially noticeable when the system is under "load".

One technology that significantly affects the performance characteristics of a centralized storage system is the disk type. The advent of SATA technology has provided lower cost higher capacity drives. Unfortunately, there is no such thing as a free lunch. SATA's greater capacity at a lower cost were not designed to sustain the duty cycles, higher throughput and IOPS characteristics of higher end drives such as SCSI, FC, and SAS. This is especially noticeable for small files and random workloads. The trick is to utilize SATA technology as tier 2 or archival storage where those characteristics are far less important. This dramatically improves the networked unified storage system price performance.

Performance is primarily measured as input output per second (IOPS) and throughput (MBps). Published IOPS are the most manipulated statistic in storage. It can refer to NFS IOPS, CIFS IOPS, block IOPS, read IOPS, write IOPS, or a mix of read and write IOPS (variable percentage), with different file and block sizes. Then there are the pseudo-standardized IOPS tests (SPC-1 for block storage IOPS and SPECsfs IOPS for NFS IOPS). Even these standardized IOPS tests can be manipulated by the number of drives, drive technology, RAID configuration, packet size, number and type of storage processors, cache, etc. In spite of this, the SPEC tests are the only apples-to-apples comparison of network storage. It can be used, but ensure that the system listed in the SPECsfs results is the same system configuration being proposed, (e.g. all functions are turned on and the configuration is a standard configuration not a one off.)

These numbers should only be used as relativistic comparisons against requirements. Look at the example of consolidating 40 Windows file servers. Typical Windows NFS IOPS ranges from

4,000 to 6,000 IOPS. Throughput ranges from 200 to 300 Mbps. To effectively consolidate these 40 Windows servers into a single networked unified storage target would require a that system to have a peak performance capability of 200,000 NFS IOPS (40 times the range midpoint) and approximately 10 Gbps (again the range midpoint) in throughput to meet the current "peak" requirements. This is most likely a worst-case scenario. It is far more likely that a lower IOPS and throughput number is required because it is extremely doubtful that all 40 Windows servers would ever run at max performance concurrently.

Performance and capacity are key elements in determining how much can be consolidated on a specific networked unified storage target system. Factors such as high availability, projects separation/isolation, and location also may have a bearing on the amount of consolidation possible. It will ultimately require testing to determine how many consolidated targets will be required. As a politician once remarked: "trust but verify."

One important networked unified storage consolidation pitfall is to assume the same unified performance by application. The e-mail and OLTP databases will require much higher transactional IO performance then the account payables application. The CAD/CAM files will require much higher throughput than the sales presentations. The streaming of video and pod casts will also require much higher throughout than the downloading of reports.

The key is to select networked unified consolidated targets flexible enough to accommodate all levels of performance (IOPS, throughput, and capacity) without compromise or undue costs.

Step 2: Network Infrastructure Requirements

Network infrastructure requirements can be a bit tricky and usually come down to three things:

1. Storage network technology, type, and speed (bandwidth)
    - TCP/IP on Ethernet for file storage
        - 10/100Mbps, 1Gbps, 10Gbps
    - iSCSI on Ethernet for block storage
        - 100Mbps, 1Gbps, 10Gbps
    - Fibre Channel (FC) for block storage
        - 1/2/4 Gbps edge ports and 1/2/4/10Gbps interswitch links
2. The amount of storage network oversubscription
    - Number of switch ports and bandwidth within the switched storage networks
    - Number of ports and bandwidth into the networked unified storage targets
    - Amount of bandwidth from the application servers versus the amount of bandwidth through the storage network
3. IT organization skill set/bias (The skill set bias influences which technology has the least organizational resistance and highest potential productivity.)
    - TCP/IP
    - Ethernet
    - Fibre Channel

Many IT organizations find storage network oversubscription a bit vexing. It is pretty easy to create a storage network bottleneck, causing the application server performance to decrease.

Fibre Channel is the most difficult because it is a deterministic (the routes between given pairs of nodes are pre-programmed or determined) layer 2 channel protocol-based network. One major factor to consider with Fibre Channel storage networks is change management. It is an offline process that requires significant detailed analysis, planning, simulation, and experience.

TCP/IP, iSCSI, and Ethernet are easier to minimize network oversubscription and congenstion because they are layer 3 non-deterministic routed protocols. It is straightforward to make changes or add storage network bandwidth that increases throughput and eliminates network bottlenecks. Change management for TCP/IP and iSCSI storage networks on Ethernet are simple online processes requiring little preparation, planning, or simulation.

The other oversubscription bottleneck is often the networked unified storage target itself. Often the unified storage technology meets the consolidation capcity requirements: however, the IOPS or bandwidth for that technology does not. The end result is an oversubscription bottleneck. Oversubscription by itself is not necessarily something to be avoided. In fact, when calculated correctly, it is to be embraced, because it reduces storage network over-engineering and cost. As previously explained with the Microsoft Windows file server consolidation, it will be quite the rare occasion (if ever) when all application servers will be hitting the networked unified storage targets concurrently. This means oversubscription can (and should) be employed to a degree while still meeting the consolidation requirements.

Next come the data protection requirements. Query all of the application administrators on their RPO and RTO requirements. Establish the tradeoffs between protection and cost. Calculate whether SATA drives (tier 2 storage) are adequate or is FC, SCSI, or SAS drives necessary? Make sure the data protection requirements cover business continuation, disaster recovery, regulatory compliance (may require checking with the organization's legal counsel), WAN bandwidth requirements, etc. Technologies to consider include snapshot, high frequency incremental snapshot (a.k.a. small aperture snapshot), synchronous, and asynchronous mirroring.

Not all applications and data sets have the equal data protection needs. This calls for a system flexible enough to provide tiered levels of RPO and RTO while matching each application's requirements at the correlating price points without compromise.

The key take away from step 2, is to engineer the storage network infrastructure to meet current requirements while being flexible and scalable enough to meet future requirements.

Step 3: Service & Support Requirements

Service and support requirements tie directly back to:

- Number of people and workload demand (also calculated as headcount per capacity);
- Scalability (number of systems required) and simplicity of the networked unified storage manageability;
- System reliability (application disruptions and system downtime for storage tasks, maintenance, or problems);
- Organizational knowledge/experience skill level as described in the next section.

Greater system management sophistication, scalability, and serviceability (translates into simpler and easier) reduce the task and knowledge load on those providing the service. An effective networked unified storage system that's easy to set up, operate, scale, and manage will ultimately minimize administrator headaches and unhappy end users.

Step 4: Staff skills, Training, Training Requirements, & Management Difficulty

This step requires inventorying IT staff storage and storage network skills, experience, and training. This inventory provides a starting point in determining what skills need to be acquired, trained, or outsourced to accomplish the networked unified storage consolidation project. It also establishes cost parameters, which must be taken into account in the final analysis.

Specialized skill sets require more training, more costs, more salary, a more limited pool of workers to draw from, and higher turnover as these highly trained workers are recruited away. Generalized skill sets require less training, lower cost, less salary, a larger pool of workers to draw from, and lower turnover.

The IT organizational skill set will have a significant influence on both the storage network and networked unified storage target technology. Those with file system, TCP/IP, and Ethernet knowledge/experience skills will lean towards NAS and, to a lesser extent, iSCSI SAN storage. Those with LUN management and Fibre Channel knowledge/experience skills will lean towards block-based SAN storage. Combinations of knowledge/experience skills can and will occur, which means the networked unified storage consolidation project can be a mix of technologies.

Lesser knowledge/experience skills require networked unified storage systems with more expertise built-into the management interfaces. The management interface must provide much greater automation and intuitive person machine interfaces.

Greater knowledge/experience skills require less sophisticated networked unified storage system management interfaces. This allows the unified storage system to be more dependent on the administrator expertise than the management system.

Networked unified storage management scalability is a very important factor in management simplicity. If the system does not scale well, it creates more systems to manage and more management tasks. NAS is an excellent example. Even if the NAS system scales in capacity, ports, and bandwidth, it needs to scale in the size of the file system. Larger file systems mean a lot less management. If the file system is limited to 16 terabytes, it requires a lot more time managing the multiple file systems with mounts, data protection, load balancing, and data management to spread large data sets across multiple systems. In contrast, NAS systems that have larger file systems of 256 terabytes require as much as 16X less management than those with 16 terabyte file systems. File systems of 512 terabytes require as much as 32X less management. These larger file systems allow large data sets to be centrally stored, with single or fewer mount points, and less end-user confusion looking for the location of their data.

The key take away point is that management simplicity is a direct result of management scalability – how much more cost, and complexity will it take to double the capacity? Will it require additional file servers or can a faster node handle growth in capacity or performance keeping management costs relatively flat.

Step 4: Service & Support Requirements

The service and support requirements tie directly back to three things:

- Organizational knowledge/experience skill level as described in the previous section;
- Number of people and workload demand;
- Scalability and simplicity of the networked unified storage manageability;
- The reliability of the system.

Greater system management sophistication, scalability, and serviceability (translates into simpler and easier) reduce the task and knowledge load on those providing the service.

Step 5: Networked Unified Storage Budget Requirements

Establishing the networked unified storage budget must include capital and operating expenses. This includes the acquisition of the networked unified storage targets, ongoing capacity, bandwidth, power, cooling, floor space, rack space, and interface costs. Additionally, it includes all system software, storage network infrastructure, (cables, connectors, HBAs or NICs, switch ports, etc.), along with the installation, maintenance, professional services, and all other operational costs. The budget needs to be realistic while providing a realistic TCO range for the networked unified storage consolidation project.

Factors to consider when determining a budget and the project ROI are:

- Current cost of support contracts for the high number of storage nodes currently in place.
- Expected upcoming purchases to cover growing performance or capacity demands.

- And the cost to manage the sprawling storage solution vs. a centralized unified storage infrastructure, (which will usually count towards a deferred headcount or hiring in another area such as applications or network administration.)

NOTE: If there is already an existing budget for the file and storage systems being consolidated, that budget can be used as a starting point for the networked unified storage consolidation project. The TCO will be significantly less than the consolidated file and storage systems' TCO over the same period of time. This savings will typically provide payback in less than three years.

If starting out with a limited set of consolidation goals, the project's price per terabyte will often be higher as the cost of the unified storage servers will not be amortized across a larger pool of initial capacity. As capacity increases for additional consolidation efforts, costs will continue to decline providing an improved ROI over the life of the project. It is often better to bite the bullet and get the right sized unified storage solution upfront that meets both current and future requirements, than trying to save initial money with a "lesser" system. This ends up requiring the purchase of additional systems, causing a repeat of the very problem that consolidation is attempting to solve. A better choice is picking a networked unified storage system that's modular and easily scaled up. This alleviates the tight budget, minimizing the initial investment while providing for future additional capacity purchases over time (which is a major part of the system TCO).

Step 6: Weighted Criteria Matrix

This step provides a critical, non-emotional, objective tool for picking the correct networked unified storage consolidation solution. It codifies the requirements already established in the previous steps while adding other aspects such as financial health of the vendor; number of installed similar consolidation projects; number of references of said projects; previous track record if any, etc. It then weights each and every requirement.

This is a good tool to combine all of the project's requirements for capacity, IOPS, throughput, data protection, manageability, serviceability, and training. It allows all of the requirements to be ranked and prioritized. So if synchronous mirroring is not important and asynchronous is, weight it appropriately. If IOPS has a higher priority than throughput, note it here. Or if thin provisioning is incredibly important because it eliminates application disruptions when increasing addressable storage, note it here. Then to simplify the decision process, they should be separated into three groups:

- Must haves or "can't live without;"
- Important and "can reluctantly live without;"
- Luxuries or "nice-to-haves."

It is essential to realize that must have requirements may not necessarily be technology issues, they can be comfort issues, financial issues, or support issues. Again, each requirement must be weighted as to the importance within the final decision. One rule of thumb in picking a no-regrets solution is that it meets 100% of the must haves, 75% of the "can reluctantly live without" requirements, and 50% of the luxuries. A weighted criteria matrix eliminates much of the bias inserted into the decision process.

Requirements/Criteria	Weight (1-5)	Vendor 1		Vendor 2		Vendor 3	
		Relative Value 1=lowest, 10=highest					
		Value	Total	Value	Total	Value	Total
Functional Requirements							
Performance scalability	4						
Capacity scalability to "X"	3						
IOPS scalability to "Y" IOPS	5						
Thin provisioning	4						
Throughput	1						
Amount of over supscription	3						
File system size	4						
10G Ethernet TCP/IP & iSCSI	3						
1G Ethernet TCP/IP & iSCSI	1						
4G FC	1						
Data Protection Requirements							
Snapshot	4						
Incremental Snapshot	4						
Continuous Snapshot	3						
Synchronous mirroring	2						
Asynchronous mirroring	5						
Flexibility to assign differently by app	5						
Manageability Requirements							
Manageability scalability	4						
Headcount required	5						
Reliability	1						
Serviceability	1						
3 Yr (complete) TCO	5						
Productivity multipliers	4						
Support capabilities	2						
Financial Strength	3						
# of installed production accounts	1						
Product roadmap	2						
References	3						
Total							

Table 1: Example of a Weighted Criteria Matrix

Step 7: Realistic Testing of Top Two Vendors From the Matrix in a Bakeoff

This might seem like common sense; however, as Voltaire once said, "common sense is not so common", and many IT organizations do not take this step. This is in spite of the fact it is the only step that provides direct empirical evidence of a networked unified storage system to meet the requirements in a controlled setting. This verifies or disproves vendor claims and clarifies the networked unified storage systems' ability to meet the organization's requirements. Every environment is different. Getting the networked unified storage system in-house will ensure that it delivers as promised. If this is not possible due to time frame, budget, space, etc… an option may be to bring the organization's reference data to the vendor's lab. Many are set up to emulate the specified environment enough to get a performance characteristics with sample data set or at least go through general test scenarios. Again if this is not possible, a final option is to at least ensure that the vendor has reference customers that are in similar businesses and ask to talk to them. Published benchmarks (not just marketing numbers) and case studies can also help.

Step 8: Negotiate a Contract with Both of the Top Two Vendors

Lowest TCO is only one criterion in weighted criteria matrix. Negotiating with both vendors ensures that the vendors are engaged and working to solve the system requirements at the best cost for value. The evaluation will determine the best technical fit. Contract negotiations look at the system costs versus the budget. If the right solution does not fit the budget, negotiating with two vendors provides some leverage, but the right solution is not always the lowest bidder so trade-offs on technical vs. price may have to be made.

Step 9: Detailed Implementation Plan

This step is the most crucial to making certain this is a positive networked unified storage consolidation experience. Planning the work and working the plan is cliché because it is true. Thorough planning by application for installation, data migration, testing, application expectation setting (performance, timing, etc.) and production cutover means few surprises and happy application owners.

After implementation, perform a post mortem to determine what went well and what should be changed to improve the experience next time.

Where to Start

Although the applications being consolidated may not be performance focused, unified storage consolidation requires a culmination of many applications, which can drive up performance and capacity. Thus looking at a leader in the high performance NAS space makes sense. Two excellent networked unified storage consolidation systems to seriously consider are BlueArc's Titan 2100 and Titan 2200. These hardware-based systems have tremendous scalability in performance, capacity, ports, bandwidth, and manageability. Titans have consistently been selected as the right choice for networked unified storage consolidation projects and been deployed to consolidate dozens of DAS servers or multiple NAS filers in a variety of markets. The systems have a clustered file system, thin provisioning (both of which allow a pay-as-you-grow approach), intuitive management, and block storage virtualization. These functions further increase the scalability while reducing the workloads allowing reduced complexity and management costs.

The solution has a back-end SAN that is virtualized behind the Titan server, a FPGA hardware-based implementation that provides industry leading high speed access to this virtualized storage via standard protocols (NFS, CIFS, iSCSI and NDMP). This allows dynamic online change management of the storage, critical for a unified storage solution. Using standard protocols also eliminates the need to have HBA and software on all the application servers as it can sustain SAN like performance and scalability with ease of management and standard file sharing ability built into client server operating systems.

	Titan 2100	Titan 2200
Scalability / System		
SPECsfs Ops/sec. *	Up to 75,000	Up to 100,000
Throughput	Up to 5Gbps	Up to 10Gbps
Capacity Scalability	256TB	512TB
File System	256TB	256TB
Global name space	Yes	Yes
Gig Ethernet Ports	Six	Six
10Gbe Clustering Ports	Two	Two
Fibre Channel Ports	Four 4Gb	Four 4Gb
Manageability	Four 4Gb	Four 4Gb
N-Way Clustering	Yes	Yes
Protocol Support		
NFS	Yes	Yes
CIFS	Yes	Yes
iSCSI	Yes	Yes
Data Protection Capabilities		
Snapshot	Yes	Yes
Incremental Snapshot	Yes	Yes
Synchronous Mirroring	Yes	Yes
Asynchronous Mirroring	Yes	Yes
Mixed Drive Support		
FC	Yes	Yes
SATA II	Yes	Yes
Manageability		
Intuitive, Automated, Simple	Yes	Yes
Thin Provisioning	Yes	Yes
Storage Pools	Yes	Yes
Flexible Modules	Yes	Yes

* For more details about the SPECsfs benchmark and a current list of validated results, please visit www.spec.org/sfs97r1/results/sfs97r1.html.

Table 2: BlueArc Titan

For more information go to: www.bluearc.com.

Summary

Networked unified storage consolidation will reduce escalating storage and management costs. To have a "positive" networked unified storage consolidation experience requires the discipline of the three Ps (planning, preparation, and process.)

Success requires focus on:

- Organizational and application requirements
  - Performance scalability (IOPS & throughput), capacity scalability, manageability scalability,
  - Storage network technology, infrastructure, & bandwidth
  - Organizational knowledge/experience skills
  - Serviceability, support, simplicity, and functionality
- Total cost of ownership
- Manageability
- Objective criteria matrix
- Implementation plan & review.

A good place to start is with the BlueArc Titan series. The Titan 2100 and 2200 have proven to be the market pace setters for networked unified storage consolidation. They have set the bar for networked unified storage system performance scalability, capacity scalability, functionality, reliability, ease of implementation, ease of operations, and ease of management.

When initiating that next storage consolidation project, make sure the Titan series is on the list.