Issue 1

Active Archive Solutions to Manage File-based Data Growth

Solutions to the problem of file-based data growth must protect data, while being affordable and easy to integrate into existing environments.

Welcome

According to virtually every study, analysis, pundit or perspective, the growth of file-based enterprise data is skyrocketing. Gartner expects ECB disk storage vendors to continue innovating at a pace that supports annual terabyte growth rates of approximately 55% with declines in cost per terabyte remaining within the 30% range during the period from 2011 through 2015. This translates in a growth from the 2010 number of 11.8 million terabytes sold to the projected volume of 107.5 million terabytes in 2015.

This growth rate does not include drives inside laptops or desktop computers. It doesn't include drives that might be used in a myriad of other devices and technologies. This growth directly represents the increase in storage infrastructure for business, new file-oriented applications used in both enterprise and technical computing, disk-based backup and archive deployments and expanded server and desktop virtualization projects.

Gartner is but one to make this prediction. Virtually any IT manager will tell you the same thing. All data is exploding exponentially and in the process causing primary storage and backup infrastructures to grow massively and expensively to keep pace.

But there is a companion problem that goes along with the issue of data and infrastructure growth. The problem is that even though file-based data are filling up ever-larger disk silos, the utilization of those files does not necessarily increase at the same rate. In other words, people may be creating more and more files, but they are still using them only a few at a time. My own hard drive has tens of thousands of presentations, documents, photos, emails, and other files, the large majority of which I have not touched in months or years. And yet, I want them available at all times for when I need them.

Translate this to an enterprise and the problem becomes astounding. Because the problem moves from the realm of personal preference: (I ***want*** my files available all the time) to business necessity: (my business ***needs*** to have access to its data at all times).

This newsletter contains market analysis as well as solutions focused on building affordable, usable active archives that can work within existing file-based data environments.

Floyd Christofferson
Director, Storage Product Marketing
SGI

Featuring research from
Gartner

Addressing the Problem of Inactive Data Filling up Expensive Active Disk Storage

The study found that most of the files in active disk arrays will seldom if ever be reopened or changed.

In a 2008 study at the University of California, Santa Cruz funded by the National Science Foundation, an active storage pool of 22TB used by 1500 employees in business and technical workflows was analyzed for utilization of network file system workloads. In other words, they studied usage patterns for the type of workloads used in virtually every enterprise in the world.

What did they find? With the continual growth of cheaper and larger storage infrastructures, files are being kept an order of magnitude longer than was found in previous studies. What's more, these files are rarely reopened; 95 percent are reopened fewer than five times. Over 60% of file reopens occur within a minute of the first open. Over 76% of files are never opened by more than one client, and of those that are opened by others, 90% of file sharing is read only. Finally, even though the files contain valuable information that must be preserved, most are not re-opened once they are closed.

The net conclusions of the study were that most of the files sitting in the active production disk arrays will seldom if ever be reopened or changed. And yet, just like the files on my laptop that I haven't touched in years, business users have a very difficult time removing files from active storage. It is simply too difficult to do so in a way that keeps the files easily accessible and protected. Because even though the files may be seldom used, they still have great value for compliance, legal, reference or intellectual property reasons.

And so datacenter disk infrastructures keep growing at an astronomical rate to keep up with the burgeoning growth of file-based data.

This problem is compounded by the cost of that growth. And not just in the acquisition of new disk arrays. As those active production disk systems grow, so must the backup infrastructure needed to protect them. No wonder deduplication schemes for backup are hot commodities as a way of reducing the compound effect of runaway data growth.

Let's not forget the cost of adding data center space, the cost in electricity and cooling of the data centers for disk drives that are always on and continually generating heat for data that is seldom being used.

Translated into real numbers, a leading network attached disk-based storage array uses about $20 per TB annually for power and cooling. For a two-petabyte system, this translates into $49,000 per year in power and cooling costs alone, at typical U.S. utility rates. This does not even consider the cost per square foot of data center space or the operational cost for managing that transactional storage.

Yes, the data is available at all times for users to access. But at what cost? There has to be a better way.

The IT director for a major pharmaceutical company recently complained about the lack of utilization of the very large and very expensive tape-based archive system they had installed over a year ago. "We needed to free up our production disk to reduce our backup times and avoid growing the production disk infrastructure," he said.

"But after a year, our production disk is still growing and the archive is largely unused," he said. The problem, he explained, was that users strongly resisted using the archive because of the difficulty and time required to get files back once they had been put there.

"Even if they rarely use source material for a project, for example, when they need it, they need it right away," he said. "Deadlines and revenues are often tied to the speed in which they can access those archives. So we keep growing our disk infrastructure and the archive sits unused."

Inactive Data Sitting on Active Disk

The solution is as simple as it is often been beguiling. The key is to create an 'active archive' where archived data is available in an 'online' state for easy access, where the data is protected for extremely long-term retention, and where the operational costs are extremely low. If any of those three elements is missing, archive strategies tend to fall flat.

The problem is that most archive solutions only address part of the issue. When backup (protection of active data) gets confused with archive (retention of inactive data), data paralysis occurs. Backup and restore times become impossible to manage because they must encompass both inactive and active data. Seldom-accessed data becomes hard to find. Operating costs skyrocket when additional production disks are needed just to keep up with the relentless growth of data.

What's worse, the excessive growth of data stored on production disk can be a contributing factor to data becoming segmented into incompatible silos. This makes collaboration between different departments areas either impossible or at best a manual process prone to error and wasted effort.

Users do their work with files but are forced to work within file systems. The job of a proactive data management strategy, an active archive strategy, is to let the users focus on their work, and not waste time, infrastructure, or energy on just setting up to do their work. That is what the tool kit that comprises active archiving solutions enables. Namely, a way for IT managers to keep data accessible, affordable, and protected without requiring the users to add work to their day understanding how it is done or where it is, or what steps need to be taken to get there.

Key Concepts: Distinguishing Between Backup and Archive

At the heart of resolving these problems is sorting out an integrated approach to both archive and backup. This makes data accessibility the priority to users while still allowing for cost containment and data protection priorities for IT managers.

The problem is that backup and archive are often confused. Not necessarily in concept, but in day-to-day practice.

FIGURE 1 Unmanaged Data Growth

Source: SGI

As file data continuously grows, production disk capacity must be increased, placing an increasing burden on backup. This often leads to an overflow, and an unmanaged archive.

The process often goes something like this:

As data continues to grow, primary or production disk arrays fill up and must be expanded. As noted above, this is typically a mixture of active data with older, seldom-used data. Backup is needed to provide protection for primary disk, and as the overall data volume grows, the backup windows grow as well, even when deduplication strategies are employed on the backup.

Whether it is due to the inability to even get backups done within the available time windows or because the backup environment simply fills up, IT managers are often left with no choice but to take excess backup data and put it on a shelf as an 'archive'. As unwise as even they know it is, they are often left with little choice but to do that or

continue to grow their environment. The problem is that this 'archive' is not a true archive at all, and is typically unmanaged. Data with a high time value is mixed in with low-value inactive data.

Also, when archive data is taken from backup data it becomes haphazard and incomplete. The data, which may have significant time value, is offline to users and is often irretrievable without significant cost, effort and time.

This is a common problem across industries that are otherwise extremely careful with their data production and protection. Indeed, it is because of the high volume of data and the difficulty in managing the distinction between high-value and low-value data that many IT managers are left with little option but to keep everything, which merely compounds the problem.

The solution is to create a clear distinction between backup and archive, and to decouple the needs of data protection from that of data retention.

Backup strategies should be for short-term production data. They are to protect what is done in the short term in case of catastrophic failures.

Archive or data retention strategies on the other hand are long-term by nature. Disaster recovery protection is still needed for this data, but does not need to be done within the tight time window needed by backup.

Building an Active Archive Strategy

An active archive is one in which all data is always available in an 'online' state all the time. An 'online' state does not mean that it is taking up expensive primary disk capacity. In the context of an active archive, and 'online' means that the data is available in an environment that is immediately and easily accessible to users, that is not drawing power, and one in which the data is protected for very long retention.

In fact, when properly applied, an active archive strategy significantly reduces the overall storage and data management costs while at the same time increasing efficiencies and the ability of users to access all data. Using an active archive strategy, costs are contained because the production disk does not need to grow very often. Inactive data that still has retention value is moved into an archive tier storage that, although 'online' and visible to the user, is typically in a powered-down state using MAID (Massive Array of Idle Disks) technology that completely removes power from the array. These archives, while still available to users, can be managed with very different disaster recovery techniques, and at a fraction of the operational costs of conventional disk-based file stores.

Thus, data lives where it is most efficient. Online, active data is contained only on the primary disk arrays. Seldom used data can be moved to an archive tier, either programmatically or manually, where it remains available but it in low-cost

Backup	Archive
Copies data	Moves data
Supports operations and recovery	Supports business and compliance
Supports availability	Supports operational efficiencies.
Short term in nature	Long-term in nature
Data typically overwritten	Data typically secured, not overwritten
Poor solution for regulatory compliance	Ideal solution for regulatory compliance
No historic relevance	For historic information
Not easily searched	Easily Searched

Source: SGI

FIGURE 2 Managed data growth with an active archive strategy

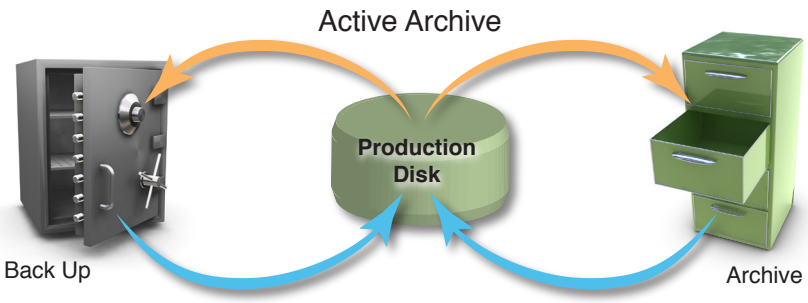

Implementing an active archive strategy means that both production disk sizes and backup capacities are managed, and in this way the growth of active data storage is contained. The key is having an easily-accessible, file-based archive to provide a low-cost, long-term data preservation target.

Source: SGI

protected state. Since data growth is not managed by increasing production disk, backup of those arrays does not mushroom either. Backup is reserved only for active data, keeping costs down while reducing backup and recovery times.

Implementing an Active Archive

There are numerous tools that can be employed to implement an active archive strategy. These will vary by industry, by use case and workflow. Not all are needed in every circumstance. In fact, what is most important in devising a strategy is to approach data growth proactively from this whole-system perspective rather than reactively by throwing more disk capacity to solve short-term problems. As we have seen, such short-term solutions compound the overall problems, always leading to higher costs and risk.

Some tools to consider:

- Digital Asset Management Solutions:

  A key problem in determining whether data is active or not is in the strategy used to classify it. This problem is compounded when production data is distributed across multiple silos.

  Leading digital asset management solutions such as LiveArc™ from SGI enable content to be indexed automatically in multiple ways as it is created and modified. Users can search for data, administrators can easily set policies to determine which data should remain on production disk and which can migrate to second or third tier storage.

Another key benefit of a digital asset management platform like LiveArc is the ability to bridge multiple namespaces, or data silos, to provide a global view across all storage, data and metadata types. In this way, IT managers have complete control over their environment and can implement back-end changes without impacting the ability of users to have 'always-online' access to their data. Users don't need to know or care where the data actually is in the hierarchy of storage infrastructure because it is always visible to them within the management interface on their desktop.

- Hierarchical Storage Management (Tier Virtualization):

  Another key practice that can aid in developing an active archive is to virtualize tiers of storage through the use of hierarchical storage management solution, such as SGI® DMF (Data Migration Facility). DMF enables multiple tiers of disk and tape to appear to the users as one large aggregated volume even though the data is actually distributed across multiple storage types.

  For example, production disk is typically higher performing (and thus higher cost) disk, according to the needs of the industry. But as we've seen, if only a fraction of the data is active the result is that expensive disk is used to house inactive data.

With a tier virtualization solution like DMF, the expensive high performance disk is linked with 'nearline' or cheaper, capacity disk. This in turn can sit in front of a MAID solution which powers disks down and/or a tape library.

The beauty of a tier virtualization solution is that all the data appears to the user to be online in the expensive production disk all the time. But in reality, even though the file appears to be right where the user put it in the file system, it has been automatically migrated to lower cost disk. The result is dramatic operational savings without the need for users to wonder where their content is.

With a solution like DMF, the rules for when or if the inactive data migrates can be automated by policies, such as file type, time since it was last accessed, etc. In addition, since DMF can manage multiple copies of the same file, backup becomes optimized to a significantly smaller amount of data.

- Low power mass storage using MAID (Massive Array of Idle Disks):

A MAID solution is another significant tool in creating an active archive. With true MAID platforms like SGI® COPAN™ and ArcFiniti™ platforms, whole sections of the disk array are selectively powered down until a user accesses that data. In the case of SGI COPAN and ArcFiniti, the power utilization of the archive can be strictly controlled within a predetermined power budget. This dramatically reduces operation costs and the power and cooling requirements of the data center. Built with technology based upon Intel® Xeon® processor E7, COPAN and ArcFiniti have the added advantage of much higher performance and proactive data protection. SGI COPAN 400 MAID and VTL solutions are best of breed in this category, providing extremely high density and low power options. See "Key Technology for Solving the Problem of Persistent Data Retention" in this newsletter for more information.

Protecting the Data that is Your Business

An active archive strategy requires planning and tools, but when done properly can dramatically reduce the overall costs of managing a growing pool of data. More importantly, by de-coupling production disk from other tiers of storage, single points of failure are virtually eliminated. Individual components can be upgraded or changed without impacting overall utilization for users. Scalability becomes an asset in this scenario, not a headache.

Additional articles in this newsletter outline some of the best-of-breed technologies available in the SGI® InfiniteStorage ecosystem to implement such a strategy.

Source: SGI

Extreme Data Growth is Shifting to File-Based Unstructured Data

The ArcFiniti network-accessible active archive solution to reduce the pressure on production storage.

As noted in the previous article, the compound annual growth rate of controller-based disk infrastructures is very strong. New estimates show that this trend is increasing. But what the research shows is that the shift dramatically tilting towards file-based data. For example, according to the Gartner Forecast Analysis: External Controller-Based Disk Storage, Worldwide, 2011-2015, 1Q11, which is included in this newsletter, "The worldwide NAS market continued its strong sales momentum in 4Q10 with a year-over-year growth rate of 36.4%." According to the report: "Many new vertical-industry applications (such as those for Web 2.0, life science and manufacturing design) generate more unstructured file data than structured data."

The shift to file-based data is also reflected in the decline of old-generation disk archiving systems. According to the report, "The fortunes of this segment will likely continue to decline because of the number of file-based systems that are now available to support the disk archiving market using industry standard interfaces, such as NFS and CIFS".

So, even though data rates are growing exponentially, the dramatic growth of file-based unstructured data means, as noted in the previous article, increasing volumes of inactive data are being held in costly active network-attached storage silos.

The Next Generation Active Archive Solution

SGI's newly released ArcFiniti platform directly targets this problem. As a network-accessible dedicated archive solution, ArcFiniti brings the best of SGI's powerful data management software tools and scalable hardware together for the first time in a fully integrated archive solution aimed specifically at providing an easily accessible solution to the problem of unparalleled growth of unstructured file-based data.

As a next-generation disk-based archive solution, ArcFiniti combines high-performance network access with built-in long-term data integrity advantages over tape. The ArcFiniti solution is designed not only to house archived data, but to protect it for the long term. With features such as proactive health monitoring of hardware and data integrity checking, ArcFiniti ensures that archival data is accessible and valid for a dramatically longer lifecycle than traditional disk or tape solutions. If ArcFiniti detects any potential disk drive problems, it proactively migrates data and verifies data integrity while alerting administrators to replace the faulty part. This dramatically reduces the need to take the system offline for costly RAID rebuilds due to disk errors.

Virtualized Tiers

As a fully integrated solution, ArcFiniti is designed to be easy to deploy and easy to use. Users see the entire system as one unified online archive pool. But to provide data protection and performance, ArcFiniti virtualizes a high-performance Primary Cache with an Archive Tier based upon SGI's award-winning MAID technology. This virtualization is managed automatically in the background by an archive policy engine. All files are always in an available online state to users, with ArcFiniti ensuring that archive content is protected for long-term retention on the most cost-effective storage tier.

Mitigate Risk, Reduce Costs, Improve Productivity

Built with technology based upon Intel® Xeon® processor E7, ArcFiniti is a complete solution that brings to the general IT market an easy-to-deploy package featuring the

The shift to file-based data is also reflected in the decline of old-generation disk archiving systems.

powerful suite of hardware and software tools that SGI has used to power some of the world's largest archives. With ArcFiniti you have a revolutionary alternative to costly and outdated disk-based archive solutions and to slower, less reliable tape solutions.

ArcFiniti gives you the best of both worlds—the reliability, easy management, and quick data access of disk—all at a fraction of the cost of transactional storage. Plus you get all the advantages of unprecedented reliability, scalability, efficiency and low power and cooling consumption.

Model Number	A156	A234	A468	A936	A1404
Max Usable Capacity	178 TB	256 TB	490 TB	970 TB	1.44 PB
Primary Disk Cache	10TB* 22TB - Opt	10TB* 22TB - Opt	10TB* 22TB - Opt	22TB* 34TB - Opt	34TB* 46TB - Opt
Power Managed Archive Tier (Usable)	156 TB	234 TB	468 TB	936 TB	1.4 PB
GigE Ports	2	2	2	4	6
10GigE Ports	2	2	2	4	6

ArcFiniti is available in 5 bundled configurations, modularly upgradable to a maximum of 1.4PB of usable archive space in a single cabinet. The primary cache has two size options available, to accommodate variable usage requirements.

Source: SGI

From the Gartner Files:

Forecast Analysis: External Controller-Based Disk Storage, Worldwide, 2011-2015, 1Q11 Update

On the strength of a strong 4Q10 close and a modestly improving economic outlook in most of the world, Gartner is increasing its external controller-based (ECB) disk storage forecast for 2011 through 2015 by just under 2% per annum. This rather modest adjustment produces a five-year compound annual growth rate of 8.7%. However, year-over-year growth rates have been moderated by a tenth of a point from our 4Q10 update forecast, reflecting the uncertainty presented by geopolitical events in the Middle East.

Key Findings

- Led by network-attached storage (NAS), ECB disk storage vendor revenue exceeded its 2008 high-water mark by over $1.4 billion in 2010.
- The broadening availability of technologies that improve utilization efficiency is slowing the annual growth rate of raw terabytes shipped into the ECB disk storage market to 55.5%.
- Gartner predicts that annual price-per-terabyte declines will hover around 30% for ECB disk storage during the 2011-through-2015 forecast period.
- The Asia/Pacific region (sans Japan) and the Americas (North America and Latin America) are most likely to increase market share at the expense of EMEA and Japan.
- With 2010 revenue falling another 16.6% on a year-over-year basis, special-purpose disk archiving storage systems with proprietary host interfaces continue to lose share to file-based systems that are available to support the disk archiving market.

ASSUMPTIONS

The 2011-through-2015 ECB disk storage market forecast is based on several assumptions outlined in the following paragraphs.

State of the Regional Economies

Gross domestic product (GDP) expectations directly influence investments in IT infrastructures in general and in ECB disk storage infrastructures in particular.

While real GDP estimates vary by region, this forecast is predicated on the expectation that global real GDP growth will range between 3.5% and 3.8% over the forecast period of 2011 through 2015.

Server and Client Virtualization Deployments

Gartner predicts that the percentage of virtualized processors shipped will increase from 12.5% in 2010 to more than 47% in 2015, providing further growth stimulus for ECB disk storage.

Server and client virtualization deployments require shared ECB disk storage to optimize infrastructure efficiency and flexibility. Shared ECB disk storage enables numerous virtual machines to be cost-effectively supported. Booting from shared ECB disk storage reduces capital expenditures, improves server mobility, and automates boot images and application deployment. Using local and remote replication functions based on shared ECB disk storage enables users to modernize data backup processes.

ECB Disk Storage Refresh Cycles

From 2011 forward, users will generally return to refreshing their ECB disk storage infrastructure every three to five years.

Rather than pay the annual maintenance fee after the warranty or prepaid maintenance period expires, users typically elect to refresh their ECB disk storage systems. This enables them to take advantage of the latest technological innovations that have a positive impact on total cost of ownership. In addition, improvements in ECB disk storage capacity and performance scalability provide the opportunity for storage infrastructure consolidation, thereby simplifying management of the storage infrastructure.

Technology Advancements Support Cost-Effective Growth

The more cost-effective ECB disk storage becomes, the more data users want to store online for ready retrieval.

Indeed, the availability of high-capacity, low-cost SATA disk drives has made it possible for many users to migrate away from cumbersome tape infrastructures for data backup/recovery. New data reduction techniques, such as data deduplication, are encouraging users to deploy low-cost and energy-efficient ECB disk storage systems to support archiving applications. Continued advancements in hard-disk drives (HDDs), solid-state drives (SSDs), input/output bus, CPUs, storage networking, global file systems, data reduction, storage provisioning and tiered-storage technologies are driving forces enabling enormous amounts of digital data to be cost-effectively stored online in ECB disk storage systems. Gartner expects ECB disk storage vendors to continue innovating at a pace that supports annual terabyte growth rates of approximately 55% with declines in cost per terabyte remaining within the 30% range during the period from 2011 through 2015.

File-Access Outpaces Block-Access ECB Disk Storage

Gartner projects that the deployment of file-access storage will outpace block-access ECB disk storage during this forecast period.

Storage consolidation is shifting more file storage from general-purpose file servers to NAS systems. An increasing number of users are adopting NFS as the storage protocol for their VMware environment for simplified management. Large-scale audio and digital content, and Web 2.0 applications, generate more unstructured data than structured data, resulting in more shipments of scalable NAS solutions. With support of industry standard protocols and deduplication technologies, stand-alone NAS continues to be a popular disk-based recovery and archiving target.

Rational Vendor Pricing Strategies

The average cost-per-terabyte reductions and average selling price (ASP) assumptions used to build this forecast are based on vendors continuing historic pricing practices.

Vendor pricing strategies have traditionally followed HDD technology trends, resulting in typical year-over-year cost-per-terabyte declines in the 30%-to-35% range. However, future cost-per-terabyte declines are expected to moderate to the lower boundary of this range as more ECB disk storage systems are increasingly deployed with a mix of SSDs, high-performance Fibre Channel and Serial Attached SCSI (SAS) HDDs and low-cost, high-capacity Serial Advanced Technology Attachment (SATA) HDDs.

Impact of Public Cloud Deployments

The forecast assumes and takes into account that public cloud deployments based on ECB disk storage infrastructures will gain increasing traction during the forecast period.

Public clouds based on ECB disk storage infrastructures promise to offer an enterprise-worthy storage infrastructure that, with colocated computing and networking infrastructures, provides IT executives a cost-effective alternative to in-house IT storage infrastructures, also known as private clouds. Therefore, public clouds based on ECB disk storage infrastructures represent a growth opportunity for storage vendors that have products that are easy to scale, provide multitenancy protection and incorporate space-saving utilization features.

ANALYSIS

Total ECB Disk Storage Revenue: 4Q10 Update vs. 3Q10 Update

In 2010, the total ECB disk storage system market grew by 18.1% year over year and by 7.9% over 2008, indicating that the market has fully recovered from the 2009 dip and resumed its growth phase. NAS experienced strong growth of 36.2% year over year, while the block-based ECB storage market grew at a respectable 14.9%, offsetting the sharp decline of 16.6% in the special-purpose disk archiving storage system market. As a result, Gartner has revised its annual ECB disk storage forecast slightly upward to 8.2% growth in 2011 from 8.1% in the last iteration. Figure 1 illustrates the forecast update changes compared with the last iteration.

Block-Access ECB Market in More Detail

4Q10 block-access vendor revenue came in strong, pushing year-over-year revenue growth for 2010 to 14.9%. Compared with block-access ECB disk storage historical year-over-year vendor revenue growth patterns, the 14.9% increase over 2009 may appear irrationally high. However, it is important to note that the 2010 block-access ECB disk storage vendor revenue results represent only

a 3.4% increase over 2008, the year before the global recession took its toll.

In spite of uncertain global economic projections, our 2011-through-2015 block-access ECB disk storage vendor revenue forecasts reflect a belief that the economies in North America, Latin America and Asia/Pacific will strengthen while the Japanese economy struggles and the Pan-European economy remains murky. Indeed, as evidenced by stronger GDP estimates for China and India, Gartner predicts that the Asia/Pacific region will increase its share of the block-access ECB disk storage market by almost six percentage points from 2010, reaching 19.8% by 2015.

Even though Gartner surveys indicate that IT capital expenditure budgets remain tight over the five-year forecast period, there is sufficient evidence to support the expectation that executives responsible for IT infrastructures will continue to invest in block-access ECB disk storage to take advantage of the new technologies that produce improved utilization and reduced operating expenditure. Moreover, the expanded deployment of virtualized server and desktop host infrastructures in smaller organizations further supports the demand for block-access ECB disk storage. In fact, new server and desktop hosting projects reinforce Gartner's prediction that modular block-access ECB disk storage systems will represent more than 73% of the total block-access ECB disk storage market, up from 65.1% in 2009. Even though Gartner does not forecast ECB vendor performance, we expect eight vendors — Dell, EMC, Fujitsu, HP, Hitachi/Hitachi Data Systems, IBM, NetApp and Oracle — to collectively increase their share of the total block-access ECB disk storage market as they acquire smaller competitors, expand their channel marketing clout, and hone their private and public cloud infrastructure strategies.

NAS ECB Market in More Detail

The worldwide NAS market continued its strong sales momentum in 4Q10 with a year-over-year growth rate of 36.4%. For the entire year of 2010, the NAS market grew 36.2% compared with 2009, ahead of our last quarter's forecast of 33.5%. Growth was especially strong in the Americas region with a 48% annual growth rate in 2010, whereas Japan showed up as the slowest-growing region with 13% annual growth. As 2010 was a strong recovery year with many customers having refreshed their systems or added new systems, we expect the annual growth rate will slow down and be normalized to the 20% range in 2011 and beyond.

FIGURE 1 ECB Disk Storage Vendor Revenue Forecast, Worldwide, 2011-2015

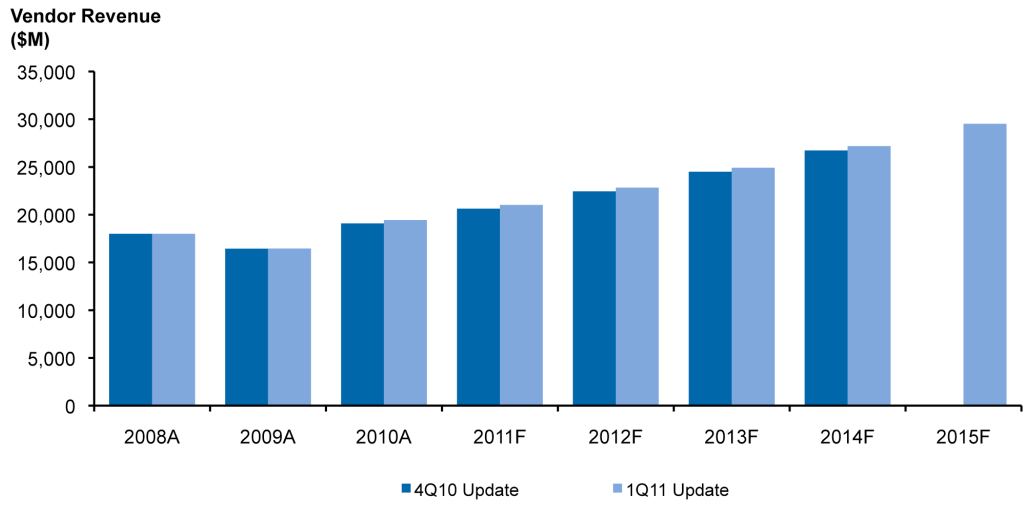

A = actual
F = forecast
Source: Gartner (March 2011)

From the vendor's perspective, market leader NetApp grew by 43% in 2010, gaining back revenue market share slightly to 34% of the total market. EMC's total NAS system revenue, including the bulk of Data Domain and Isilon's fourth-quarter 2010 revenue, grew by 93% in 2010, gaining almost 10 points of market share to 28% in 2010. Together, NetApp and EMC pulled ahead of the competition and increased their combined market share to 62% in 2010, compared with 52% in 2009. We expect the two market leaders to continue their strong growth in the file-access NAS market, while many other NAS vendors will continue to grow based on their competitive advantages. The overall growth drivers remain unchanged, including the increased adoption of NAS/storage area network (SAN) unified storage in the data centers and a large amount of unstructured data generated by new applications, as well as by backup with deduplication and archiving with retention.

Special-Purpose Disk Archiving Storage System ECB Market in More Detail

The special-purpose disk archiving storage system market segment appears to be on a death spiral. Results from 4Q10 continue to be poor: down 17.5% compared with the same time period in 2009, and down 16.6% for all of 2010. This segment is dominated by the EMC Centera systems, and their revenue fall-off represents the major reason for the drop in total vendor revenue in this segment. The fortunes of this segment will likely continue to decline because of the number of file-based systems that are now available to support the disk archiving market using industry standard interfaces, such as NFS and CIFS.

MARKET MODEL

Figure 2 presents the ECB disk storage market forecast model.

FIGURE 2 ECB Disk Storage Market Forecast Model

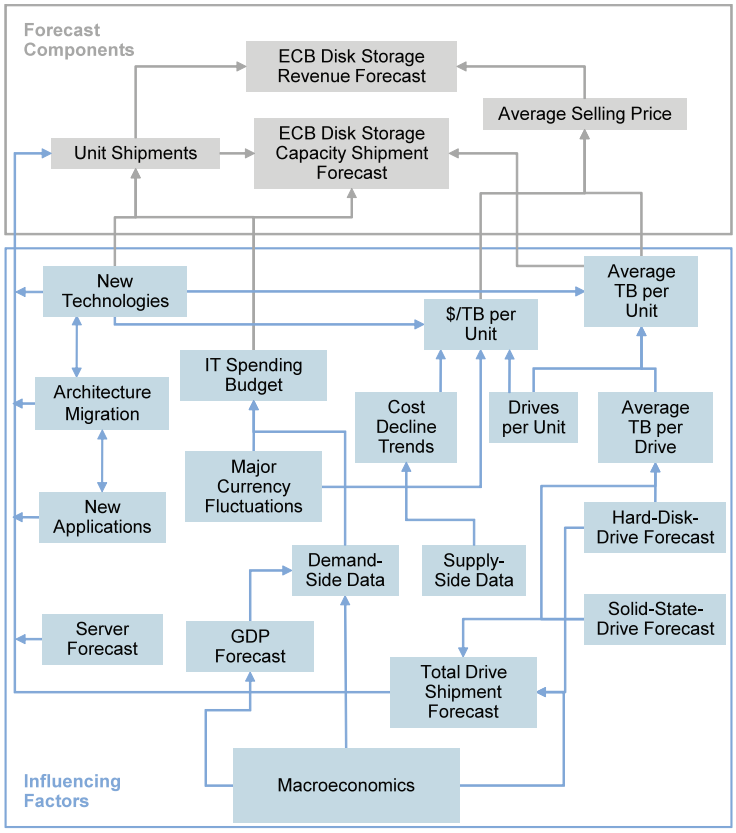

TB = terabyte
Source: Gartner (March 2011)

ECB Disk Storage Revenue Forecast

The ECB disk storage market is defined as the total net revenue measured in U.S. dollars for vendor-branded ECB disk storage systems. Vendor revenue market share credit is given to the vendor brand under which the system is sold to the distribution channel or the end user (via a direct sales force). For cobranded systems, vendor revenue credit is given to the primary or lead brand.

Forecast Components

The ECB disk storage market model consists of three fundamental components.

ECB Disk Storage Capacity Shipment Forecast

This component estimates the total raw terabytes shipped by all vendors in the ECB disk storage market. It is based on Gartner's trend analysis and assumption of the average raw capacity per storage system. It's the product of the average raw capacity per unit and the total unit shipments.

Unit Shipments

This component estimates the total ECB disk storage system unit shipments. Dual-controller configurations associated with scale-up modular systems are counted as one unit. In a failover clustered configuration with no cluster file system, each node in the cluster is counted as one unit. In a scale-out cluster where a single storage pool is an aggregate of multiple nodes, each cluster is counted as one unit.

Average Selling Price

This component estimates the average vendor selling price for an ECB disk storage system sold through all channels, expressed in U.S. dollars. With direct channels, the vendor ASP is the equivalent of the average user purchase price. However, with indirect channels, the vendor ASP is typically lower than the average user purchase price because of channel margins. The ASP for a scale-out cluster is the estimated ASP of the average size of the cluster.

Influencing Factors

Average Terabytes per Unit

This influencing factor estimates the average raw storage capacity in terabytes within a storage unit.

Dollars/Terabyte per Unit

This influencing factor estimates the average cost per terabyte of raw capacity shipped in a storage unit.

HDD Forecast

This influencing factor provides drive capacity growth rate input to the average raw capacity per drive in the ECB disk storage forecast. Recognizing that HDDs represent 40% to 80% of an ECB disk storage system, the Gartner enterprise HDD forecast is leveraged to make necessary adjustments to the capacity per HDD and the number of HDDs per ECB disk storage unit.

SSD Forecast

This influencing factor provides drive capacity growth input to the average raw capacity per drive in the ECB disk storage forecast. Gartner expects SSD technology to change the way ECB disk storage systems are configured. Accordingly, the Gartner SSD forecast is leveraged to help produce capacity and unit estimates for ECB disk storage units.

Total Drive Shipment Forecast

This influencing factor, which is part of the HDD forecast and the SSD forecast, provides input to the unit shipment forecast in the ECB disk storage market.

New Technologies

This influencing factor provides input to all ECB disk storage forecast components. The march of technology innovations, including advancements in HDDs, SSDs and CPU performance — coupled with thin provisioning, automated data tiering, data compression and data deduplication, as well as host interface protocols such as 10 Gbps Ethernet and Fibre Channel over Ethernet (FCoE) — directly and indirectly impact average terabytes per unit, dollars per terabyte per unit, unit shipments and ASP. In every case, these technologies enable users to place more data online while reducing storage system cost. They help expand the total available market for ECB disk storage systems by enabling the deployment of new horizontal applications, such as virtual tape libraries, disk-to-disk backup/recovery, active archiving and "cloud" storage services, as well as vertical applications that generate a tremendous amount of digital data.

New Applications

This influencing factor provides input to the ECB disk storage unit shipment forecast. New applications refer to new server or client applications that generate more digital data output that needs to be stored. Many new vertical-industry applications (such as those for Web 2.0,

life science and manufacturing design) generate more unstructured file data than structured data. We also see a growing trend of using NAS to support horizontal applications, such as VMware and Oracle.

Architecture Migration

This influencing factor provides input to the ECB disk storage unit shipment forecast. Server and storage architecture migration, often triggered by new technologies, new applications and emerging business needs, impacts our forecasts for the overall ECB disk storage market, as well as different segments of the market. For example, the server virtualization trend has often triggered the redesign and migration in storage architecture, driving new storage unit growth, especially for NAS and SAN storage systems, as networked storage is required for many advanced features in the virtual environment. In addition, the emerging business needs for Web 2.0, cloud storage services and client virtualization will likely trigger deployments of new ECB disk storage solutions based on a scale-out architecture.

Demand-Side Data

This influencing factor provides input that impacts the IT spending budget. Data from periodic Gartner surveys among storage users and ongoing user inquiries provide vital demand-side insights regarding terabyte growth expectations and future application deployments, as well as plans to adopt new storage technologies. All this data is used in preparing the future-year forecast for ECB disk storage systems.

Real GDP provides a high-level backdrop in terms of regional projections.

Supply-Side Data

This influencing factor provides input that helps to determine cost decline trends. Gartner participates in quarterly financial calls and other briefings conducted by storage vendors to gain a sense of how they see the ECB disk storage market unfolding. In addition, we have periodic discussions with various contacts within the storage vendor community forming Delphi positions and associated assumptions regarding how the ECB disk storage market is expected to change during the current forecast cycle.

IT Spending Budget

This influencing factor provides input to the ECB disk storage unit shipment forecast. The Gartner IT spending forecast update provides the global economic backdrop necessary to establish an ECB disk storage forecast that is in alignment with overall and regional IT spending trends and other economic projections. Fundamentally, the Gartner IT spending forecast refers to how much budget user organizations will likely allocate to spend on information technologies such as ECB disk storage systems.

Major Currency Fluctuations

This influencing factor provides input into the IT spending budget and dollars/terabyte per unit factors. Major exchange currency fluctuations may impact regional IT spending as imported IT products may become more or less expensive for the region.

Server Forecast

This influencing factor provides input into the ECB disk storage unit shipment forecast, as the dependency of servers (physical or virtual) on ECB disk storage (instead of their internal storage) is increasing. Gartner's server (physical or virtual) forecasts are used to provide directional guidance to make sure that the block-access ECB disk storage system forecast follows the future deployment expectations associated with operating system type and with server technology.

Connections

The following paragraphs illustrate the connections between various forecast components and forecast influencing factors.

Macroeconomics > GDP Forecast > IT Spending Budget > Demand-Side Data

2009 taught us that the state of the global economy has the most profound impact of all the influencing factors on the ECB disk storage system market. When the global economies are good, users are willing to invest in their ECB disk storage infrastructures, and vendors are able to expand R&D initiatives and production facilities. When the global economy suffers as it did in 2009, demand recedes and vendor production abates accordingly.

New Technology > Architecture Migration > New Applications > Unit Shipments

Over the course of the past decade, we have seen new technologies have a direct impact on expanding the scope of the ECB disk storage market, the enabling of new applications including virtual tape libraries and Disk 2 Disk backup, and the development of new ECB disk storage

ACRONYM KEY AND GLOSSARY TERMS

ASP	average selling price	**NAS**	network-attached storage
ECB	external controller-based	**SAN**	storage area network
GDP	gross domestic product	**SSD**	solid-state drive
HDD	hard-disk drive	**TB**	terabyte

architectures, such as scale-out and unified storage platforms. Considered as a whole, these developments have led to the increase in annual ECB disk storage unit shipments.

Supply-Side Data > Cost Decline Trends > Dollars/Terabyte per Unit > Average Selling Price

Vendors are constantly developing new technologies that increase ECB disk storage system productivity and efficiency. These developments influence cost trends, which enable vendors to reduce the dollars/terabyte per unit while at the same time retaining appropriate ASPs that reflect the value of the offering.

SSD Forecast > HDD Forecast > Average Terabytes per Drive

SSD and HDD vendors are continually enhancing capacity, performance and cost characteristics of these fundamental components of ECB disk storage systems. The projections from these companion Gartner forecast documents are used to establish the estimate for average terabytes per drive.

Drives per Unit > Dollars/Terabyte per Unit > Average Terabytes per Unit

Increasing the capacity of individual HDDs is often insufficient to meet the necessary ECB disk storage unit requirements. Accordingly, the number of HDDs per unit continues to increase, which, of course, reduces dollars/terabyte per unit and increases the average terabytes per unit.

Unit Shipments > ECB Disk Storage Capacity Shipment Forecast > ECB Disk Storage Revenue Forecast

Unit shipments have a direct impact on both the ECB disk storage capacity shipment forecast and the ECB disk storage revenue forecast. Unit shipments have increased each year as the scope of the ECB disk storage market expanded, enabling, except for 2009, year-over-year ECB disk storage revenue growth.

ASP > ECB Disk Storage Revenue Forecast

The ASP multiplied by unit shipments produces the ECB disk storage revenue forecast. ASPs fluctuate based on the introduction of new technologies and lower-cost entry-level ECB disk storage systems, which increases the number of units sold.

This document is published in the following Market Insights:

Storage Worldwide

Gartner Dataquest G00211733, Roger W. Cox, Pushan Rinnen, 16 March 2011

Key Technology for Solving the Problem of Persistent Data Retention

Overcoming the limitations of both conventional RAID arrays and Tape Libraries in assuring affordable long-term data integrity with the core technology used in SGI COPAN and ArcFiniti platforms.

Techniques for ensuring that persistent data archived today will really be available when it is needed in the future.

Most people assign a value to digital data based on their perception of its immediate importance rather than when they might need it at some future date. As a result, the vast majority of data is stored in a state of limited usefulness due to how and where it is housed, how it was initially identified or tagged, and how easily it can be retrieved.

For the purposes of this paper, we will distinguish between "transactional" data, which is data that is actively being created, changed, and accessed by an application, and "persistent" data, which has been created and is being retained in an unchanged state for potential future use. When an event happens in which users have a critical need to locate and retrieve specific portions of this persistent data, they may be severely limited in how quickly and accurately they can access it. Yet, the value of the data may be even more important at that point than when it was originally created.

With the onset of government-mandated regulations surrounding the retention and retrieval of corporate data, it has become critical to know where this persistent data is, to be able to immediately retrieve precisely what is needed, and to be completely sure that the data is still intact even as it is locked away in the archive. The ability to do this can make the difference between winning or losing lawsuits, saving lives, and ensuring that years of irreplaceable work is protected.

Because corporations are increasingly aware of the role and value of persistent data in the enterprise, they are taking a more careful look at storage technologies and methodologies to ensure it is protected. This includes extending the definitions of data availability and data integrity. These were formerly centered around "mission critical" online transactional data, and were generally stored on the most expensive and reliable storage systems.

In this newsletter, we discuss the techniques for ensuring that persistent data in the enterprise today will really be available when it is needed in the future.

The Role of Persistent Data in the Enterprise

In corporations today, the four basic roles and usage models for persistent data are:

- Backup/Recovery;
- Business Continuity/Disaster Recovery (DR);
- Active Archive; and
- Passive Archive, including Compliance and Deep Archive.

These roles and usage models have both common and conflicting requirements, as detailed in Table 1.

In the past, companies have used data tape to store almost all of their persistent enterprise data, due to its low capital cost. While tape may serve well for some environments and helps for portability off site, it can also cause problems with long retrieval times and difficulty with the data integrity requirements of all four usage models above. As a result, the industry has been moving to low-cost disk solutions. The problem is that conventional disk also has limitations for long-term storage, in particular with the requirements for longer retention periods and low costs outlined in Table I.

This is especially true for compliance data, for which guaranteed integrity is not only a business necessity but a legal requirement. In that situation it is critical that companies be able to find the exact copies of specifically-requested data as well as to verify that it has not been changed from its original state.

Tape is still appropriate for storing many types of Passive Archive data, especially when it needs to be stored offsite (Deep Archive). This is because tape is inherently portable, tape is inexpensive, and access time is often not critical. However, enterprises are finding out the hard way that

Table 1. Storage Requirements and Usage Models for Persistent Data

Requirements	Backup / Recovery	Business Continuity / DR	Active Archive	Passive Archive
Application Path	Backup Software	Must Be Restored	Direct	Must Be Restored
Access Frequency	Seldom	Rarely	Seldom	Rarely
Retention Period	Days -- Months	Days -- Months	Years	Decades
Retrieval Time	Minutes -- Hours	Hours	Seconds	Days
Retrieval Objects	Files -- Volumes	Files -- Volumes	Files	Files
Data Integrity	Very High	Very High	Very High	High -- Absolute
Retrieval Accuracy	Last Good Copy	Process Recoverable	Very High	Exact, Inclusive
Other Copies	Multiple	Few	Few	Few -- None
Storage Cost	Low	Very Low	Low	Low -- Very Low

Source: SGI

because tape is so portable, it can easily be misplaced, stolen, damaged or lost.

On an industry wide level, what has been desperately needed is a new persistent data storage platform that combines the best attributes of both disk and tape.

Why MAID Is the Right Platform for Storing Persistent Data

Addressing this need is MAID (Massive Array of Idle Disks), an idea rooted in common sense. MAID technology operates under the premise that when persistent data is stored on disk it is seldom actually accessed. So why reduce the service lifetime of those disks and needlessly consume more power by spinning all of them constantly? The basic MAID concept is to turn on only those particular disks that are needed for a specific action, and then turn them off again when they are not needed. As implemented in the industry-leading SGI COPAN MAID platform and integrated into the archive tier of the SGI ArcFiniti, the entire system can be tuned to a strict power budget, to ensure stability of the data as well as of the operational costs.

In that sense, it is similar to how you use your car. It is turned off when not needed, which extends the service life of the engine and tires, plus saves on gasoline.

While this strategy has limitations for online transactional data where data access requirements are measured in milliseconds, as part of a tiered architecture or an Active Archive of persistent data MAID technology properly implemented solves the biggest liabilities of both disk and tape-based systems.

A key problem with any long-term storage strategy, however, is to ensure that the data remains intact on whatever storage medium it is placed. When a book or file is put on the shelf for long-term archive, the environment is carefully maintained to ensure it is protected. Too much humidity and the pages can deteriorate and even rot. Magnetic disk or tape media has a similar problem and is subject to 'bit rot', which is when the magnetic substrate deteriorates over time. This can lead to the gradual degradation of data integrity in any storage environment, even when it is just sitting there unused.

This is one of the great failings of tape as a long-term storage medium for the four usage models in Table I. That is, users don't really know if the data is intact until it is accessed. For most companies, it is too labor intensive and expensive to constantly keep mounting the tapes and checking on them. Doing so would also cause additional wear on the magnetic tape itself. For all companies, finding out the data is gone at the very moment when they are trying to access it is not an acceptable result.

This can lead to the gradual degradation of data integrity in any storage environment, even when it is just sitting there unused.

Thus, when MAID technology was first introduced commercially in the COPAN platform in 2004, engineers recognized that they had multiple challenges to overcome to make long-term, cost-effective digital data preservation absolutely secure for all of the major use cases. Specifically, the system had to:

1. Quickly power on and off a selected set of disk drives in a manner that assures only the right set of drives are on when needed, and off when not needed – all within a power budget that satisfies multiple simultaneous requests to a large array of drives;

2. Make sure that the data is still intact when it is needed, even though that data had not been accessed in a long time;

3. Ensure that the drives are healthy when they are actually called on to perform;

4. Have enough intelligence and bandwidth within a large array of mostly idle disks to deliver the needed performance from the particular drives that are actually being accessed;

5. Pulling all these requirements together into a robust system architecture that also provides high-level functionality to meet the rest of the primary data preservation requirements in Table I, such as intelligent identification and search capability for data stored on disks that are usually off; and

6. Deliver all of this at the lowest possible cost in an extremely dense, enterprise-class package to be competitive with existing tape solutions that is easy to service and maintain, while capable of providing decades of operational life.

To meet these criteria, the COPAN MAID platform was built using three key technologies:

1. Three-tiered architecture, with connectivity and computing intelligence at each tier to dynamically distribute the bandwidth and processing capabilities where and when needed;

2. Power-Managed RAID® software, for full RAID 5 data integrity protection and fault tolerance within a predetermined and controlled power budget; and

3. Disk Aerobics® software, which ensures healthy drives and solid data protection while proactively moving data off faulty or dying drives.

The three-tiered architecture allows Power-Managed RAID, I/O operations, and application processing to all be done in parallel using separate processing resources. As a result, the new SGI COPAN 400 MAID system can deliver wire-speed read/write performance across eight 8 Gbps Fibre Channel connections (6,400 MB/second sustained) with any combination of up to 25% of 896 SATA or nearline SAS disk drives operational at one time. With the current 3 TB nearline SAS drives, this provides an industry-leading 2.7 PB of raw capacity or of 1.9 PB useable persistent data storage in a single COPAN 400 cabinet. In addition, the SGI MAID platforms features two power options, enabling either 24% or 50% of the disk to be spun up at a time for a given shelf. This gives IT managers the option of tuning performance and while also setting pre-determined power usage parameters requirements to the balance that best suits their requirements.

In day-to-day operations, idle drives are automatically powered on in four-drive RAID 5 sets within seconds of receiving an access request. They are arranged physically so that there is no vibrational interaction or local heating issues. Once on, they respond to all other I/O requests at normal millisecond disk speeds just like conventional RAID arrays.

All of the drives are hard-mounted with vibrational damping, in easily accessible pull-out canisters for scheduled replacement. They are also actively monitored by individual thermal sensors and kept below 45°C, which assures that they always operate well within the drive's maximum reliability range. SGI purposely kept the number of drives in the RAID sets small in order to reduce the probability of a drive member failing. This also minimizes the time needed to rebuild the RAID set if one drive should suddenly fail.

A key high-availability feature of the COPAN MAID three-tier architecture is that it offers multiple replication options. Not only is export to tape available for creating a physical tape cartridge from a virtual tape cartridge, but also replication between COPAN 400 VTLs can be done over a WAN to avoid the cost and security risks inherent

in the physical transport of tapes. These options offer very fast disaster recovery capability, as well as the ability to consolidate backups from geographically-dispersed sites into a centralized backup facility.

The third key technical innovation, Disk Aerobics software, continuously monitors all operating drives within the system and periodically turns on and monitors each idle drive to ensure that the entire system is healthy and operating within design limits. If Disk Aerobics detects that a particular drive is showing signs of possible failure, it proactively replaces that drive with a spare and retires it from the system. As a result, Disk Aerobics substantially reduces the probability of a RAID set rebuild due to a failed drive, and virtually eliminates the likelihood that a second drive in the RAID set would fail before the RAID set could be restored.

Thus, the combination of these architectural principles in the COPAN 400 and ArcFiniti to provide the assurance of online data integrity for long-term persistent data for the four key categories of persistent data preservation.

Implementing Disk Aerobics Software

Disk Aerobics is a technology that enables the system to:

1. Actively monitor disk health and environmental data;

2. Periodically exercise idle drives to ensure good health;

3. Proactively replace degrading or end-of-life drives; and

4. Retire replaced drives and maintain spares.

The active monitoring of disk health and environmental data involves analyzing a selected subset of the internally generated SMART (Self-Monitoring, Analysis, and Reporting Technology) data from the drives, as well as checking temperature readings for each drive, canister temperatures, fan temperatures, and voltages. This is all kept in a database, and the system employs heuristic techniques to detect if a drive is becoming marginal or if environmental conditions are moving out of limits.

If any of these parameters falls outside specified limits or if a drive fails to return a SMART status, the data on that drive is immediately backed up to a spare and the drive is proactively replaced in its RAID set and then retired. Because the backup is done directly between the suspect drive and the spare drive without affecting the other members of the RAID set, this is accomplished far more quickly than the time it takes to rebuild an entire RAID set following a failure.

In addition, the SGI MAID systems use a number of proprietary techniques that anticipate potential data path problems or drive failures before they happen by monitoring I/O timeouts, analyzing CRC errors, and reconstructing bad blocks. Along with SMART data analysis and proactive replacement, the net result is that the vulnerability window for a second drive failure is reduced to milliseconds from what would normally be many hours as is typically needed to rebuild a RAID set if it were allowed to simply fail. Thus, SGI MAID technology together with Disk Aerobics anticipates potential failures, instead of reacting to failures as always-on RAID systems must do.

Disk Aerobics also periodically exercises the idle drives to ensure that they are healthy. To do this, the system sets aside a fixed amount of the power management budget for idle drives so they can each be turned on at least once a month and run through the same set of diagnostics and proactive replacement techniques as the active drives. All of this is accomplished as background tasks, with no interruption of normal operational performance.

So how long can disk drives be left powered off before head-disk stiction, corrosion, or some other non-operating problem begins to cause an additional reliability risk? Although industry answers to this question vary, the emerging consensus is that spinning drives up monthly is adequate to keep them healthy. Even with this assurance, the COPAN 400 monitors and analyzes each drive's start up time, the amount of electricity it's motor draws and other key indicators as another proactive measure to forestall failure. If any of these parameters begins to increase beyond acceptable limits, the drive is immediately tagged as a candidate for proactive replacement and then retired by the normal procedure. The COPAN 400 also proactively replaces and retires any drive that has exceeded a fixed percentage of the specified power-on hours or start / stops for that drive

These technical innovations provide the assurance of online data integrity for long-term persistent data.

model, even if the other operational parameters still all look good. When it comes to data integrity, there is zero margin for error.

To handle the proactive replacement needs of any SGI MAID product, the system currently uses 40 spare drives (5 per shelf), which should result in scheduled service or maintenance action to replace retired drives with fresh spares no more often than every 12 to 18 months, even with 896 drives in a fully-loaded COPAN 400 MAID system. Allocating spares and keeping track of the number remaining is also a key part of the Disk Aerobics software. If the number of spares has been reduced to less than a pre-determined amount, the COPAN 400 or ArcFiniti systems will 'phone home' and request an immediate service call to replenish the full inventory of spares.

SGI MAID technology dramatically improves drive reliability

We have been tracking the reliability of the MAID products in the field since the first COPAN system was shipped in 2004. The actual AFR of the over 10,000 SATA drives installed in customer environments was at 0.42%. This is equivalent to an MTBF of over 3 million hours. Thus, the measured failures for the drives in the COPAN MAID systems was less than 25% of the failures for conventional SATA drives in published studies; and the COPAN study was tested on a sample size over twice as large. As a result, the projected service life of the drives should improve by at least five times over typical always-spinning drive storage systems. The accumulated field test time for these drives is now over 70 million hours. In addition, there have been no incidents of lost or unavailable data. This is truly enterprise-class reliability and data availability.

And now with the introduction of even more reliable 3TB nearline SAS drives in the COPAN and ArcFiniti, the long-term protection of data is enhanced even more.

The combination of Disk Aerobics and SGI MAID technologies in the COPAN and ArcFiniti platforms is truly the best solution for storing persistent data. The combination of enterprise-class reliability and data availability, along with superior disk drive performance at a cost comparable to tape, makes these solutions ideal for all four of the persistent enterprise data usage models.

Source: SGI