

WHITE PAPER

Approaches To Storage Cost Management In The New Economy

Sponsored by: Infortrend Technologies Inc

Simon Piff

June 2009

IDC OPINION

Amid the economic crisis brought about by the collapse of a number of financial institutions in 2008, organizations of all sizes across the globe are facing a major challenge when it comes to the management of their data. In spite of the reduced economic activity that the recession has created, data storage growth has not slowed significantly; indeed in some sectors, the economic collapse has increased the amount of data being created, and yet IT budgets are by no means growing.

The complexities of managing existing data continue to add to the challenges of provisioning for new data. This is emerging as a major issue for organizations and the previously accepted practices for provisioning new storage capacity need to be thoroughly reviewed with a more strategic intent, if IT departments expect to operate a holistically efficient datacenter environment.

In December 2008, IDC published a report titled *Worldwide Storage Top Ten Predictions: Grappling With Content Growth in a Contracting Economy (Doc #216026)* which pointed out that IT spending growth rates would decline from 5.1% to just 2.6% in 2009. In the January 2009 report titled *Economic Crisis Response: Worldwide Disk Storage Systems 2008-2012 Forecast Update (Doc#216074)*, IDC forecast that, for the first time in many years, worldwide storage revenues would experience negative growth of 1.7%. So there is a dichotomy that organizations need to address: the ongoing growth of data juxtaposed to the decreases in IT, and specifically storage, expenditure.

Data growth is not a new issue. Over the past years, a number of solutions have been introduced to the market to help address the explosive growth in data. However, for many IT departments, the ongoing decline in the cost of disk storage or dollar per gigabyte (\$/GB) has meant that the easier solution is simply to add more disk space to address data growth. This approach has been perceived to be the fast, easy and cheap option for many organizations; however, the market dynamics are changing and the acquisition of storage space is no longer where the major cost concerns are being felt. More and more organizations have deployed a variety of storage solutions from multiple vendors across different technologies — be they Fibre Channel or iSCSI connected storage area networks (SAN) or point network-attached storage (NAS) systems to cater for the rapid growth in file-based data. In addition to this complexity, many new servers were deployed with a high degree of internal storage built up over the years. This organic growth of data storage systems is now beginning to emerge as a management nightmare for many organizations.

IN THIS WHITE PAPER

In this White Paper, IDC discusses the reasons for ongoing data growth within organizations and examines a number of technology approaches that can help IT departments more effectively manage the costs associated with provisioning storage.

Understanding what drives the growth of data, the challenges inherent with existing practices for provisioning new storage and the economic impact of the low utilizations rates for storage found in most enterprises will help businesses in assessing the right approach to their next storage acquisition, in light of the change in economic dynamics and the looming threat of a storage management meltdown.

IT departments have several primary concerns about storage infrastructure:

- Ensuring application and data availability
- Limiting downtime, both scheduled and unscheduled
- Ensuring data security and integrity
- Reducing the initial acquisition and ongoing costs associated with storage
- Managing storage resources and their utilization levels
- Addressing the complexities within the storage environment

Each of these issues alone is of great concern to IT departments and, if not handled correctly, can lead to challenges for the business. Since many IT departments are concerned with all of these issues, we can begin to understand the many challenges faced when it is time to provision new storage capacity, be it for a new application or to cater for the growth in data from existing applications or users.

SITUATION OVERVIEW

Explosive Data Growth and Under-Utilized Capacity

Let us consider how new applications are provisioned in many environments. To begin, the business unit defines the pain point and refers to it to source a suitable solution that will address the requirements of the business. This solution is then provisioned to a service level defined by the business. From a storage perspective, the major concerns are data availability, integrity and security. In order to ensure availability, IT organizations first decide on the amount of disk space that will be allocated to a specific application. This allocation includes both the current existing storage demands, plus an estimation of future storage capacity requirements.

Herein lies the first area of inefficiency, since these estimates are simply that, estimated, there is no guarantee this excess capacity that has been allocated will get used. To compound this issue, in order to ensure that the application performs as well as can be expected and has sufficient capacity for data growth, it is not an uncommon practice for capacity to be over-allocated in the first instance. This is by no means a unique practice across the world. Consider the impact on the IT department if an application freezes or ceases to operate due to insufficient storage capacity. The impact on business would be huge and the corresponding impact on the individuals responsible would likely be quite severe, and so it is no surprise that over-provisioning of capacity is not simply a case of individuals being over-cautious but, in

many cases, is a systematic process that is part of the IT standard operating procedure. As new applications are added, this over-provisioning continues, and in a manner that locks out the unused disk space from other devices and so the growth of unused and, more importantly, unusable disk space has emerged in recent months as a cost overhead within IT that now needs to be addressed.

Integrity Issues

The second area of concern for IT is integrity. This can be defined in a number of ways but for our illustration, let us leave it as ensuring the data is actually resident within the storage environment. One of the more common solutions for addressing this is the backup environment. This simply creates a copy of the data that is being stored, ideally within a less expensive disk storage system environment, but in a manner that can provide for simple and relatively fast restoration should the primary data storage environment fail.

So now we have at least two copies of the same data. We say "at least" because this is where data proliferation can begin to get out of control. In the case of an application that is transactional and much of the data is system-generated, there is little opportunity for identical duplicates of the primary data being created; however, when there is a chance for user intervention, there is always the opportunity that multiple copies of information can exist. One example could be a customer relationship management system where one member of the account team creates a company identity, but the next user is unable to easily locate it and creates another entry. A duplicate entry, if you like, of the same details. Whilst this type of situations does occur, the impact on data growth can be minimal since in such an example the updating of account activity and contacts is likely to only occur in one instance of the customer record, mostly.

Where duplicate data growth can get out of hand is in the area of file-based data and it is here where for some organizations the economic recession is actually driving increased levels of data growth. Consider the example of an individual within the marketing department. The marketing process usually begins with data analysis to understand what is happening in the markets and, ideally, to predict where market opportunities exist for a product or service. Spreadsheet analysis is a data-intensive activity that also results in the creation of new data in the form of outcomes and forecasts. For many organizations, these outcomes are then rewritten into some form of presentation document which, in today's environment, can be a data-rich as well as media-rich document that consequently increases the file size. Whilst a presentation may well address the needs of a large audience, there are departments that will want to also review a more in-depth word processing document that explains much of the thought processes behind the decisions being presented. Each of these processes creates new data files, files that will require storing. For many organizations the initial work is stored locally on a PC or laptop hard drive. This data is then further backed up to the shared storage infrastructure as part of standard IT procedures.

Where the data growth has most impact is how the items are shared. Email is the most commonly used tool for sharing file-based data and we can therefore see how each of the documents created by the marketing department are then replicated within the email environment. This continues across the organization as many users will make copies of these documents onto their own local hard drives, which are subsequently backed up by the standard backup processes of IT. In the process of document creation and editing, there is highly likely to be multiple — but only

marginally dissimilar — copies of the same files being stored, and emailed and stored again across the organization.

In this example, we can see how a single process within a business can produce many copies of the same document across all the interested parties that are in touch with the process and, of course, all this data will also need to be backed up.

Now we add in to the mix the situation brought about by the current recession and demands of the business to find growth in declining markets or save costs, and the data creation cycle is increased.

Datacenter Consolidation Trends

Another factor that is contributing to the closer inspections of storage infrastructure is the ongoing drive for datacenter consolidation, driven by multiple factors including server consolidation and virtualization and, in certain parts of the world, a need to reduce the electricity demands due to either a desire to reduce costs or in some regions an inadequate supply to existing facilities. The resulting datacenter consolidation programs are driving an increased demand for SANs. At the same time, IT administrators are required to consolidate storage requirements for multiple purposes and pools of shared storage across multiple diverse applications. Clearly there are going to be some applications that will still require dedicated storage (e.g., core banking systems or telco billing systems) where the data availability, integrity, security and growth is a critical part of the core mission of the business), although in such environments there is also an opportunity to critically examine past practices and determine if indeed these are still the optimal method for provisioning storage now that the economics of IT have changed.

FUTURE OUTLOOK

Having stated that data is growing in an almost uncontrollable manner across many different types of applications within the enterprise, the question remains: What can be done to minimize the challenges for storage administrators, bearing in mind the current economic situation?

Data growth will not change much, so there is little point in attempting to implement policies that will restrict data growth (although there is an argument that the implementation of policies that permit chargeback for IT services to the departments that use them may indeed provide certain "data hogs" to reconsider their approach to the creation of data, chargeback is however a challenging practice that is only just emerging in many markets). But whatever strategy businesses choose to adopt, they should take into account the following concepts:

- Integration with existing systems
- Simplified management
- Ability to maximize storage optimization

Within this decision-making process, there are a number of technologies that can be considered that will ensure these three tenets are observed. These technologies include but are not limited to:

- Thin provisioning

☒ Storage virtualization

☒ De-duplication

Regardless of the technology deployed, the objective for new deployments has to ensure that the business demands for data can be met whilst also serving the requirements for the new datacenter. In order to better manage storage cost, the idea of having all the disk storage space viewed as a single, shared pool of storage will help increase the utilization. Increased utilization in turn will lead to lower management and support costs, and less of an overhead for IT.

The "shared-pool" approach can help IT in its task of intelligently managing the consolidated storage requirements across the organization by providing a "storage-on-demand" approach to provisioning new space. Instead of dedicating specific arrays to specific applications or departments which, as we have seen, leads to redundancies in the storage and data as well as over-provisioning and a long term growth in unusable disk space, IT can now apportion parts of the shared storage to specific requirements, and growth can be handled in a more natural manner. Provisioning can be managed in a shorter timeframe and, assuming that the capacity can be expanded upon demand and the available resources are carefully monitored, IT will be able to experience a much higher utilization rate of existing disks. At the same time, this approach to disk storage provisioning has the added value of reducing both valuable space within the datacenter, but also has the capability to reduce the associated costs for power and cooling, since this shared resource will take the place of multiple disparate resources that are likely to be resident within independent chassis and require independent and incremental power and cooling.

This approach to deploying storage also fits well with the new paradigm of highly virtualized datacenters. With the ongoing move towards server consolidation and application virtualization, IT is no longer provisioning individual server and storage solutions for specific applications. It is now possible to run multiple applications within the same physical environment, by the implementation of server virtualization technology. This has also led to the growth in SANs as the preferred type of storage for virtualized environments. Therefore, the ability of IT to provision a shared pool of disks to multiple applications has moved from being a strategic opportunity to becoming an essential method of provisioning storage.

IT DEPARTMENT CHALLENGES

One of the biggest challenges facing IT departments is the realization that operating costs are now consuming so much of the IT budget that there is very little funding to develop new strategically-aligned solutions for the business. The economic environment has compounded this and the likelihood of new funding becoming available is highly unlikely, leaving IT in a situation where data growth is unchecked whilst budgets are in decline. In order to offset this, the concept of highly virtualized datacenters has become a reality for many organizations, but this in turn has led to a new set of management issues.

Prior to the extensive use of virtualization within the server environment, applications were deployed in dedicated hardware with dedicated storage and networking. The advantage of this approach is that each application can be managed in isolation of the others, providing IT with a defined environment to monitor for specific applications and, arguably, an easier environment to troubleshoot should any application cease to

function within normally acceptable parameters. The move to a highly virtualized environment now means that the management and troubleshooting of specific applications has to be managed within a system that is also running multiple other, and possibly non-related, applications. This can make the process of problem solving more challenging since it relies primarily upon the management software to inform of issues and suggest solutions, and whilst it is possible to restart specific individual virtual environments, the hardware itself cannot be shut down since it is required to keep the other virtual applications up and running (of course, hardware has matured sufficiently that server hardware failures are exceedingly rare and the need to power down a server is significantly less likely in today's environment than was in the past). If we add into this complex management environment a shared pool of storage that is accessed (in terms of both read and write) by multiple applications, then all of a sudden the storage administrator has been forcibly moved out of his or her comfort zone and into an environment where he or she had not previously had to deal with. Therefore, over and above the technological accomplishments that are possible through thin provisioning and storage virtualization, there is a social engineering aspect that will need to be addressed to ensure that storage administrators are suitably skilled in the implementation, provisioning and management of these new storage environments. Since these individuals are expected to contribute to the recommendations concerning the types of storage being purchased by the organization, and unless their fears are suitably addressed and allayed, then the financial benefits that can be obtained by organizations that deploy shared storage cannot be realized.

INFORTREND TECHNOLOGIES ESVA SERIES

To address many of these issues that are emerging for storage administrators Infortrend has launched a new range of disk storage systems that can help address these concerns. The new range of products, named Enterprise Scalable Virtualized Architecture (ESVA) offers five new products in two broad categories.

Serving the market with Fibre Channel connectivity are:

- ☒ ESVA-F20. A 16 to 112TB system based on up to sixteen 7,200rpm 3.5 inch SATA drives for high-capacity demand workloads in a 3U enclosure providing four 8Gb Fibre Channel ports
- ☒ ESVA-F40. 1.8 to 7TB system based upon up to twelve 10,000 rpm 2.5 inch SAS drives providing high throughput and the best IOPS performance for its size in a 1U enclosure providing four 4Gb Fibre Channel ports
- ☒ ESVA-F60. Up to 33.6TB systems based on up to sixteen 15,000 rpm 3.5 inch SAS drives providing mission-critical performance in a 3U enclosure with four 8Gb Fibre Channel ports

Complementing this group are two offerings with iSCSI connectivity. Each provides expansion up to 64 drives and connects through eight 1Gb iSCSI ports in a 3U enclosure:

- ☒ ESVA-E20. Up to 64TB in sixteen 7,200 rpm 3.5 inch SATA drives
- ☒ ESVA-E60. Up to 19.2 TB in sixteen 15,000 rpm 3.5 inch SAS drives.

As indicated within the Infortrend naming conventions, virtualized scalable architecture is one of the key technology attributes of this new range of systems. Inherent within these systems is the ability to scale on demand, be it for performance requirement, whereby load-balancing technology is embedded within the infrastructure of these systems, or for more efficient utilization, where thin provisioning is enabled across the system allowing administrators to add more capacity to an array; users can add up to 12 subsystems in a single virtual pool, effectively permitting 12 x 112Tb in a single pool – all of which can then be dynamically allocated at the time of writing data, thereby avoiding the costly approach of pre-allocating disk storage to applications.

Infortrend's Challenges/Opportunities

The primary challenge for Infortrend is that this is a move into a new market area for the company. Whilst their pedigree within the storage business with respect to shipping RAID technology cannot be argued with since they have some of the most technically demanding organizations as their customers, this move to deliver entire new storage systems is a step beyond their RAID controller core competence. However, this may also be seen as a natural progression for a company that has prided itself, justifiably, on delivering new technologies to the markets time and time again over recent years.

The second issue for Infortrend is brand recognition. In this highly competitive market, the storage vendor landscape has consolidated to only a handful of mega vendors and only a few of these are pure-play storage vendors (as opposed to IT hardware vendors). Infortrend has managed to secure some high-profile customer wins which should help with increasing their brand awareness, and the company has a proven ability to focus on product development. If they can execute on their marketing and sales campaigns as effectively, then they should be able to overcome this challenge with relative ease.

Finally Infortrend has stated clearly that these products are primarily targeting the mid-market organizations. The issue here is that while these customers require a strong channel play, they also demand enterprise level support. This conundrum has not been a huge issue for Infortrend's channel since the company already offers a mature support infrastructure either through its channel partners or directly to its customers. Extending this technological knowledge to disk systems should present no major barriers to success, depending upon the speed with which training programs can be rolled out.

Clearly, there are also opportunities within the market for this product lineup. With the major competitors in this space having relatively recently been acquired by larger organizations, there are still questions with regards to the ability for Infortrend's main competition to effectively execute in light of changes within the channel structures and partnerships. Whilst backed by significantly larger vendors, Infortrend's competitors no doubt have the ability to create markets; however, if Infortrend can execute within these markets more effectively, bearing in mind their channel has not had to undergo such drastic changes, then there is every opportunity for the company to gain significant share within the markets where they are strong.

CONCLUSION

The decision-making process currently used by end-user organizations around disk storage systems will have to undergo a change if organizations are going to avoid storage management chaos. The currently held belief that, because the cost of storage in terms of \$/GB is in decline, has led to storage decisions being mainly tactical with little thought given to a larger strategic vision of internal disk storage infrastructure.

The current appetite for datacenter and server consolidation and virtualizations presents a valuable opportunity for organizations of all sizes to rethink their approach to storage buying decisions in light of the continued growth in data verses the probable decline in available funding.

Infortrend's new product lineup addresses a number of areas that mid-market organizations will be able to take advantage of, and indeed some larger enterprises may benefit from putting this product range through their due diligence process when buying new systems.

The essential storage feature requirements for most businesses include being able to provide storage "on-demand", scale as the data grows, deliver thin provisioning, and essentially being able to scale with the organization whilst not adding to the levels of poor utilization. In this respect, Infortrend should be well placed to address the needs of many mid-market businesses globally.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2009 IDC. Reproduction without written permission is completely forbidden.