IT Professionals' Guide to a Smarter Storage Environment:

Solutions for Managing Data Growth and Controlling Costs

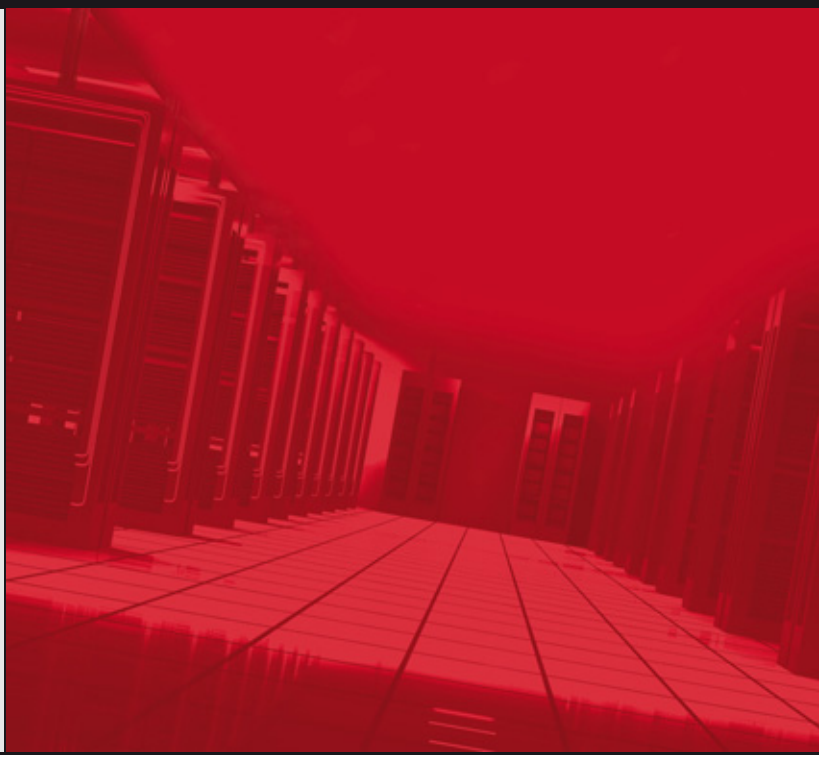

IT agility. Your way.

Data growth will bleed you dry.

Keeping up with data storage requirements is a constant drain on your budget. Only a smart storage tiering strategy will stop the bleeding—simplifying data management to save you up to 50 percent in storage costs.

www.f5.com/solutions/storage

IT agility. Your way.

IT Professionals' Guide to a **Smarter Storage Environment:**

Solutions for Managing Data Growth and Controlling Costs

The storage landscape continues to evolve. Faced with such an array of storage ideas and options, how does one begin to make sense of such complexity?

If you're reading this, then the chances are you're about to gain some useful insights into how best to create a Smarter Storage Environment within your business.

You'll discover that understanding the business requirements that need to be met by your storage infrastructure is the first step in the process. This is followed by the knowledge that, although you can't achieve everything you want overnight, you do need to decide in which order your current pain points should be addressed.

And if there's one gem to take out of the following pages, surely it's the quote from Dave Allerton, IT Director at RHWL Architects: "Implementing ARX from F5 Networks has made my life easier and saved RHWL more than £243,000 a year."

In times like these, few of us can afford to ignore those kinds of savings!

Phil Alsop
Editor SNS Europe & SNS UK,
DataCentre Solutions (DCS)

contents

IT agility. Your way.

Tiers of Joy

Phil Alsop, Editor of SNS Europe, raises the issue of data management with Kirby Wadsworth, VP Marketing, F5.

The mantra "doing more with less" has become a popular one in the storage industry over the past year or so. There is a not-so-perfect storm of global financial chaos, ever-growing data volumes, shrinking IT budgets and emerging storage technologies that end users know they must address. Recognition of this status quo is one thing, but knowing what to do next, quite another.

To date, many of the attempts to deal with the key problem of exploding data volumes and the attendant higher storage costs have been technology driven. However successful such an approach might appear, it does not recognize the underlying issue – that data needs to be better managed. Technology will play a part in such an optimized data management strategy, but it should not be the starting point.

Is Data Essential?

When examining file-based data, for example, it quickly becomes apparent that there is a very simple 80:20 rule – 20 percent of data can be classified as frequently accessed; 80 percent, however, is accessed infrequently, if at all. Historically, with no good method for determining in which category a particular file resided, all data was treated the same, typically stored on a single tier of high performance disk and backed up every night whether necessary or not. With file-based data growing at a compound annual growth rate (CAGR) of approximately 80 percent, such a solution is simply not sustainable.

Now technology has made it possible to simply and automatically determine a particular file's access characteristics. What's more, having recognized the difference between active and inactive business information, technology separates frequently accessed information from moribund, automatically moving less frequently accessed data to a second, far less expensive, tier of storage.

Although the data is now shared across multiple storage tiers, intelligent file virtualization allows end users to see only one pool of

storage. If they access a file that has been unused for six months, it is served up by the tier 2 storage device in such a way that the user never realizes it was moved to the second tier for lack of use. This movement of active and inactive files is carried out automatically. The operator need only set the policy - how many days a file remains inactive before migration occurs, for instance.

Shrink the Backup Window

What's more, the files residing on tier two, by default, have not changed. Therefore, it is not necessary to back them up on a regular basis. With some 80 percent of the data remaining virtually unchanged over time, only the active 20 percent needs to be regularly backed up. The reduction in cost and time to back up is obviously significant.

Due to the risk of overflowing and the subsequent loss of application availability, traditional single tier NAS systems are typically run at a relatively low utilization rate, sometimes as low as 40 percent. In a tiered environment, operators can be confident that tier one storage will constantly be offloaded to tier 2, thus offering a low risk of capacity overruns. Therefore, tiered storage systems can run much higher utilization rates, perhaps as high as 90 percent. This improves cost effectiveness and reduces the necessity for purchasing more tier one storage.

Further savings accrue from this automation approach as (manual) storage administration costs are reduced – a 19 percent reduction is not uncommon.

Elsewhere in this supplement you can read about the technologies that underpin tiered storage, as well as about some of the customers who have benefitted. Properly implemented, tiered storage is a game-changing approach to many of today's storage challenges. For example, many large organizations contemplating consolidation of a host of SAN-attached, Windows-based file servers would be facing, not uncommonly, nine months of data migration, with the accompanying outages and angst. F5 is confident that, using ARX technology as part of a tiered storage strategy, such a migration project can be completed in two and a half weeks, with an ROI of four to five months on a multi-million dollar solution.

Time to swap tears of frustration for tiers of joy?

TechValidate

"A Global 500 banking company reduced their backup time from 34 hours to two hours after deploying their F5 ARX, a 94% reduction." **TVID: 0F4-DD0-DEB**
(See page 11 to learn more about the TechValidate Realtime Research Network.)

All Data is not Created Equal

By Mark Govan, Storage Consultant, F5

Enterprises, faced with enormous and growing volumes of data washing over them, have a storage problem. The storage team, who is tasked to manage this data cost-effectively, add more storage. What they often end up with is a siloed storage infrastructure with no consideration for different types of data or applications, a costly setup tailored for the highest common denominator.

When environments grow, more servers and storage are added, management of the data becomes more complex... and ultimately you spend time and money to protect and back up data in the same way. This is untenable. Data needs to be looked at in a different way, from the standpoint that different values can be apportioned to data depending on its characteristics. In short - a tiered approach. Most stored data is not critical and therefore does not warrant expensive tier 1 storage. Older data is generally less relevant and gets changed less. In most companies, over two-thirds of files haven't been altered in the last three months or more. If not all data has equal relevance to the business, then not all data need reside on the same type of storage. There are wildly varying performance, availability and cost points associated with different types of storage, and thus tremendous efficiencies and cost savings open to IT teams.

Reducing Complexity and Cost

The challenge in being able to realize these cost savings lies in the ability to identify different types of data, to place them on appropriate tiers of storage, and to manage this relationship over time.

All of this needs to be achieved in an automated manner without impacting client access to the data or increasing management costs.

The SIGNAL IDUNA (opposite) and RHWL Architects (pages 12-13) case studies are great examples of how adopting this approach results in significant benefits to the business itself.

TechValidate

"A large enterprise electronics company confirmed that it will defer purchasing 100 TBs of storage in the next 12 months as a result of deploying their F5 ARX." **TVID: 4E7-703-BF2**
(See page 11 to learn more about the TechValidate Realtime Research Network.)

Financial Services Case Study: SIGNAL IDUNA

SIGNAL IDUNA
Versicherungen und Finanzen

"Thanks to F5's ARX, we have been able to free up a massive amount of online storage space, greatly improving performance."

Klaus Krüger, Systems Engineering Project Manager, SIGNAL IDUNA Group

The SIGNAL IDUNA Group is one of Germany's ten largest insurance companies. Operating in Germany, Hungary, Poland, and Switzerland, the group employs a workforce of around 8,000 and is represented by 27,000 insurance agents in the marketplace.

Rapid Data Growth

Given exponentially growing volumes of data, Klaus Krüger, Systems Engineering Project Manager, and his team had serious concerns about the extremely slow access and data backup speeds of an online storage system based on a four-node, Microsoft Windows 2003 cluster. A detailed investigation revealed that there was a considerable volume of seldom-used files that took up a disproportionately large amount of online storage space.

Apart from the high storage costs per megabyte, this severely impacted performance. An initial remedy was found in a hierarchical storage management (HSM) solution, but a desire to do without the file "stubs" that this system left behind meant that SIGNAL IDUNA turned to F5's ARX file virtualization device.

Lower Capital Expenditure

"The first thing that actually impressed us was the fact that ARX is a network switch—a combination of high-performance hardware and integrated software," says Krüger. "It neither generates stub files nor requires that agents are installed on servers or terminals. It is fast, efficient, and 100 percent transparent for our users—and that's precisely what we want."

ARX automates the swapping out of non-critical business data to more cost-effective mass storage devices. This not only improves performance, but also lowers capital expenditure significantly.

What's more, the ARX solution downsizes the backup windows and lowers infrastructure costs, because it reduces the size of the redundant data regularly saved.

Reducing Costs and Improving Efficiency for File-based Storage

IDC
Analyze the Future

In findings that yet again hammer home the extent of how file storage demands are growing, IDC research in Europe and the US found that companies are increasing their file-based storage by 40 percent to 120 percent a year, and place a high priority on boosting the efficiency and reliability of their management processes for file-based information.

The majority of new storage capacity in data centers was driven by unstructured, file-based data, and this growth is projected to accelerate.

By 2012, over 75 percent of new storage capacity shipped will be dedicated to the storage, organization, and protection of files.

This growth is being driven by the expanding set of applications that companies use to compete in an increasingly difficult business environment.

Organizations rely on email, collaboration tools, and Web sites to communicate and conduct business.

They collect, store, and analyse more information than ever before. Records, design documents, videos, and other types of unstructured data are being digitized to increase efficiency, offer new services, and comply with evolving government regulations.

Shifting Business Priorities

This has a number of implications for organizations. At the top of the chart for storage teams is the need to boost the efficiency and reliability

of management processes for file-based information. In turn, the unique use and retention demands associated with file-based data drives them to examine new technologies that have the twin benefits of optimizing storage utilization while reducing data protection and business continuity burdens. End users need to look at different ways of solving problems because key weaknesses in existing management practices for file-based storage are being exposed. Many companies deploy large numbers of dispersed, underutilized file servers in support of different workgroups and geographically remote sites. Others have consolidated file-based storage on larger NAS systems but still struggle to deal with rapid data growth and rapid technology changes.

There are direct economic impacts as a result of these practices:

- The use of dispersed file servers leads to severe underutilization of storage capacity, often at levels below 15 percent.
- In addition, enterprises require an increasing allocation of backup resources to meet higher business continuity standards.
- IDC's research found that actual backups consisted primarily of older, unchanging files that have already been backed up many times or files that are somewhat less than critical, e.g., employee MP3 files. Old or low-priority files often account for 80–90 percent of all files being backed up and are responsible for driving back-up times and costs up.
- Efforts to migrate and consolidate data from even a limited number of file servers onto larger,

more centralized NAS systems can often take six months to a year.

Non-disruptive Migration

There is a clear need, then, for companies to manage files throughout their lifecycle, from creation to archive to possible destruction. The inability to automatically, but non-disruptively, migrate information from high-performance storage to lower-cost storage leads to excessive spending on systems that cost three to four times as much to acquire and administer. Network-based file virtualization is a key technology for addressing these issues and introducing more manageable policies for storage teams facing the environment and issues described above. They address the requirement of allowing organizations to deploy (and consistently manage files across) a wide range of disk storage tiers with different performance, capacity, availability and cost characteristics, without user disruption and without requiring wholesale displacement of existing assets.

Virtualization, as we know, is a major trend throughout the IT industry, and the storage market is no exception. When it comes to file storage virtualization, however, the focus is different. For most IT managers, virtualization means consolidating many x86-based applications onto a single server.

In the storage market, it's about allowing storage managers to manage multiple storage systems (e.g., SAN attached arrays, NAS systems, or file servers) as a common pool of capacity.

Common Capacity Pool

This allows IT managers to non-disruptively migrate information between storage, boost the utilization of installed capacity, implement tiered storage policies to manage costs, and better scale data protection processes.

As such, file storage virtualization is an effective answer for enterprises struggling to better manage their file-based information assets.

ADAPTED FROM THE F5-SPONSORED IDC WHITE PAPER "THE ECONOMIC IMPACT OF FILE VIRTUALISATION: REDUCING COSTS AND IMPROVING EFFICIENCY FOR FILE-BASED STORAGE," RICHARD L. VILLARS, APRIL 2009, Doc. ID #217787

The Realtime Research Network

You'll have noticed frequent references to TechValidate and TechFacts throughout this booklet. TechValidate is a research network that facilitates the **peer-to-peer** sharing of information about how technologies are used in the real world. The power of the network is based on the insights of thousands of verified IT and business professionals who contribute information about the specific products and services they use on a day-to-day basis.

F5 has used TechValidate to survey how our data management products are used by those verified professionals who are also F5 customers. All of the TechFacts you see throughout these pages are independently verified by TechValidate based on anonymous feedback from their realtime research network.

Why You Should Join TechValidate

Any registered member of TechValidate can browse **thousands of research reports, deployment statistics, and case studies** taken from the actual experiences of their peers in the network. **Over 41,000 IT and Business Professionals like you have joined TechValidate to share their technology experiences with peers.**

Realtime Information

Unlike traditional market and industry research, **TechValidate streams its peer-based findings in real time** to its research network. TechValidate's research is never stale or out of date with shifting market sentiment or deployment patterns. As a result, you can stay abreast of the latest market and technology trends without having to wait for a new analyst report to be written.

High Data Quality

TechValidate's exclusive research data is compiled from Web-based surveys conducted directly by TechValidate. To preserve the authenticity and veracity of the data, TechValidate maintains **stringent verification controls** over who can contribute information that includes confirming their organization affiliation and cross-checking that their organization has purchased the specific product or service. Furthermore, TechValidate guarantees the anonymity of all participants so that professionals can share freely and openly.
www.techvalidate.com

Architectural Case Study: RHWL ARCHITECTS

RHWL Architects generates large, technical documents embedded with images and diagrams. These huge files can place a considerable strain on servers and rapidly consume available storage capacity. The company found itself struggling to meet the needs of the users for additional capacity while attempting to manage its storage budget. Users were often disrupted as IT shuffled files and archived files offline to salvage enough storage capacity to operate.

RHWL's IT infrastructure consists of multiple servers running Microsoft Windows applications and a range of specialized Computer-Aided Design (CAD) and image manipulation packages. In addition, the company has a Storage Area Network (SAN), which holds the majority of its file storage.

Emergency Data Archiving

As the company grew and started to work on more and more projects, it began to face a series of challenges. Storage space on their SAN solution continually ran out, resulting in the continuous need to perform emergency archiving of data that, ideally, should be online at all times. Soon, these "emergencies" became an almost weekly event.

This cost the business a lot of money in terms of wasted IT time, wasted employee time and project delays. Additionally, RHWL performs a weekly backup of all of its files, with smaller incremental backups occurring during the week. This process caused problems as well – it took up so many tapes and so much time that the IT team was struggling to accomplish the process during the course of a weekend.

A better, longer-term storage solution was clearly needed. The obvious choice of expanding SAN capacity would have required spending several hundred thousand pounds - money that the firm simply did not have.

As a result, they began to look at automated tiered storage solutions, with the goals for the new solution being to add more space, reduce the cost of storage per gigabyte, and gain a solution that would be much more scalable to support future growth.

"Recently…implementing ARX from F5 Networks has made my life easier and saved RHWL more than £243,000 a year"

Dave Allerton, IT Director RHWL Architects

Reduced Backup Times

The F5 ARX Series allows companies to define their own policies for how to manage stored data across servers, normally based on age, type, and frequency of access characteristics. RHWL could therefore adjust its data archive to suit its business needs and make changes to it, even on the fly if necessary. Analysis of the company's files revealed that most of them hadn't been touched in over six months. To address this, the IT team was able to set up the F5 ARX to migrate different files to different platforms, according to how recently they had been modified.

With its new solution in place, RHWL no longer has to perform lengthy and costly backups at the end of each week, dropping weekend backup times from fourteen hours to just one and one half hours, creating savings in both time and the cost of tapes.

As part of the total solution, RHWL purchased one new lower-cost storage array for tier 2 storage. However, most of the additional capacity that the company has acquired has come from making better use of its existing server capacity. The ARX solution maximizes utilization and organizes files, so that RHWL doesn't have to constantly migrate files to archives.

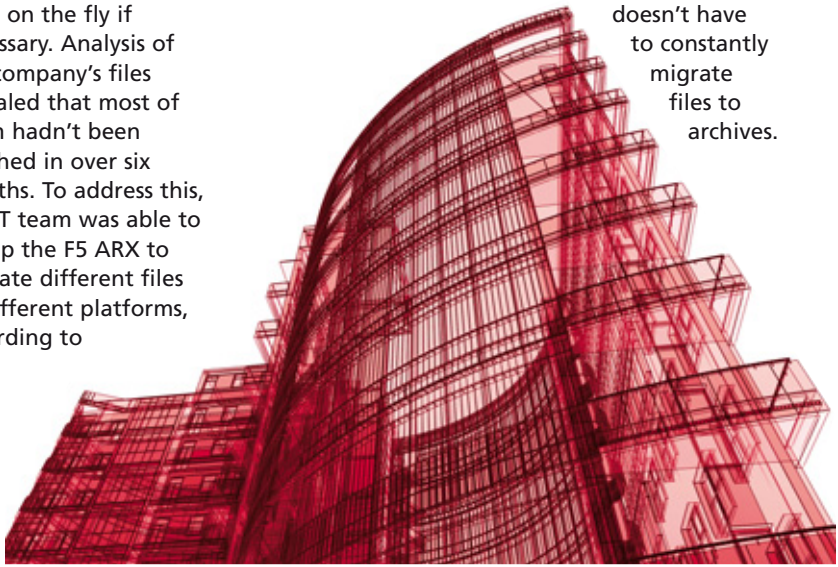

Breaking the Ties that Bind

Nigel Burmeister, Director, Product Marketing, F5

Data growth has skyrocketed. Infrastructures are increasingly complex. Stringent regulations require data be preserved for long periods of time. And, as always, lean operational budgets have become even leaner. What's the solution?

File virtualization reduces the complexity associated with managing file data, which in turn addresses many of these challenges. Removing the static mapping between client and storage resources breaks the ties that bind your storage infrastructure, making it easy to migrate data between servers, to provision new storage, and to share storage capacity, all without disruption to users or applications.

Global Namespace

The key to file virtualization is the Global Namespace, which is a logical representation of file content irrespective of where the content actually resides. When a file storage environment is virtualized, clients access the Global Namespace, rather than the physical file servers and NAS devices themselves. This dramatically reduces the number of mount points that exist. And, since the client always retains the same logical drive mapping, IT never needs to perform laborious reconfigurations.

Policy-based Automation

A Global Namespace offers significant benefits, but the truly compelling advantages of file virtualization become clear when policy-based automation is added to the mix. Consider, for example:

- **Migration**

Moving files among different file storage devices can be a labor-intensive, disruptive process that demands downtime. Automation allows large upgrades, server consolidations, platform changes and any other project that requires the movement

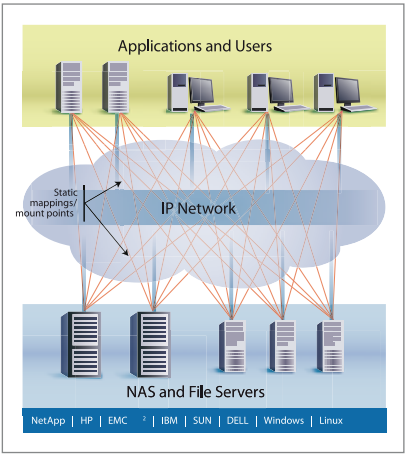

Static Infrastructure

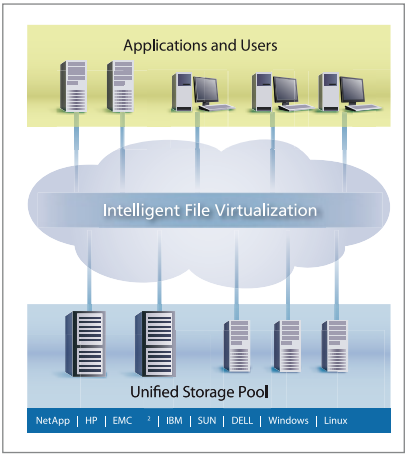

Flexible, Dynamic Infrastructure

of data to be performed simply and quickly, even during normal business hours.

- **Storage Tiering**
Automating the movement and placement of data on appropriate devices according to the value of the data offers significant cost and backup benefits, as noted in the earlier sections of this booklet.

In addition, it is easy to integrate technologies like SSD, SATA, deduplication, and even cloud storage, from any vendor, as your business needs and budget dictate.

- **Capacity Balancing**
Consolidating existing storage resources into a "virtual file server" helps you maximize your existing resources and eliminate bottlenecks. Automating the process makes it transparent to users and applications and frees IT resources for other tasks.

As organizations around the world are proving every day, breaking the ties that bind their client and storage resources goes a long way toward relieving the strain of data management.

TechValidate

"A large enterprise media and entertainment company reduced its data migration time by 70% using their F5 ARX." **TVID: A48-EA0-364**
(See page 11 to learn more about the TechValidate Realtime Research Network.)

Internet Shopping Case Study: shopzilla

Shopzilla is dedicated to helping shoppers around the world find the best price on the right products. Its patent-pending technology empowers shoppers with the ability to quickly and easily compare prices on more than 30 million products offered by 92,000 stores.

Shopzilla's IT organization was being challenged by the complexity of physically managing rapid storage growth while minimizing downtime. Specifically, the group faced a huge data migration issue.

Manual Data Migration

When Shopzilla's file servers reached end of life or a new platform was introduced, IT was faced with a laborious data migration process. Not only would the data have to be manually moved from the old platforms to the new, but the applications themselves would have to be "rewired to understand the movement from the old to the new," according to Burzin Engineer, VP of Infrastructure Technology at Shopzilla.

This "rewiring" of the applications involved modifying the application source code to recognize the new storage location. While this process worked, it could take up to four months to do a migration.

The process consumed significant IT resources, increased the potential for bugs and would impact customer-facing applications. As a result, upgrades were often delayed, and Shopzilla was forced to pay exorbitant support costs to maintain its older gear.

"Having a device like F5 ARX which masks all of that complexity helps significantly,"
Burzin Engineer, VP of Infrastructure Technology

To address the problem, Shopzilla installed a pair of F5 ARX6000 intelligent file virtualization devices to create a global namespace and virtualize its file storage environment.

In its first migration project following implementation, Shopzilla seamlessly phased out an older storage array by automatically migrating a 4TB volume in just three days, without IT staff intervention and without impacting application operations. "Having a device like F5 ARX which masks all of that complexity helps significantly," Engineer said.

No Staff Intervention Necessary

Now, each time a migration is required, the IT staff simply establishes the policy and lets the F5 ARX device run quietly in the background. Migrations continue to be completed in days rather than weeks or months. The policy-driven data management provided by F5 ARX has also enabled Shopzilla to automate many other tasks that had previously been performed manually.

As a result, the company has experienced both direct savings, by eliminating the need for high-cost support on older equipment, and indirect savings, by freeing critical IT resources to focus on other business-critical initiatives.

These important management and operational efficiencies will ensure the IT infrastructure continues to deliver the unsurpassed online user experience Shopzilla customers have come to expect.

Clouds Hang over the Storage Landscape

Rick Gillett, VP of Data Systems Architecture, F5

As mentioned previously in this guide, there's no doubt that many end users are in danger of putting the cart before the horse in their quest to achieve an optimized storage environment. Headline technologies, such as solid state disks (SSDs) and, "The Cloud" can become dangerous obsessions. They must be implemented as soon as possible, no matter how well they match or don't match the business requirements of the company.

Wrong. Unless an effective, and realistic, storage management strategy is established up front, new technology investment of this nature is likely to be money poorly spent. SSDs and clouds are not some kind of universal panacea, but a couple of new options to help implement such a storage tiering approach.

For example, F5's intelligent data management brings with it a range of benefits that has the potential to make a more immediate and significant impact on a business than any cloud-based approach. ARX technology offers capacity balancing (improved applications performance, increased productivity, elimination of bottlenecks); data replication (entire file systems or individual files via an automated policy); data migration (automated file movement between heterogeneous appliances without reconfiguration or any impact on the user); and, most critically, storage tiering (automated data movement and placement through the storage network).

Supporters of cloud solutions tout benefits such as economies of scale, speed of deployment and better management. These benefits are already appreciated by end users who have implemented F5-based virtualization!

Sophisticated Storage

Where new technologies do bear investigation is in adding sophistication to a basic two-level tiered storage strategy. Returning to the 80:20 inactive/active file-based data model, it would seem sensible to break down these two types of data into further subsets. It might be that, of the 20 percent of high-value, highly available data, only 30 percent of this needs to be available 24x7, while the remaining 70 percent is not quite so critical. Depending on the file type and usage, it could well be that a Tier 0, higher even than Tier 1, is established, with this data residing on SSDs.

Similarly, the 80 percent of inactive legacy data could well be subdivided according to how much it will need to be referenced, creating another tier at this "lower" end.

Alongside this tiered storage refinement in terms of frequency of data usage, there is the opportunity to consider whether putting any of the data into a private, or public, cloud would bring further management efficiencies. Does that 80 percent of inactive data need to be kept within the company? Or is it used so infrequently, or has it become so relatively unimportant, that it could be handed over to an outside service provider to store and manage, just so long as it is viewable, and available, within an organization's virtual storage environment?

Or maybe there's a case for implementing a DR strategy that involves one or more service provider?

Know the Value of Data

Matching the value of your business data with the platform it sits on is the key foundation of a successful storage tiering approach. Everything else that follows – F5's ARX technology, deduplication, SSDs, clouds and even convergence – will then provide significant business benefits. Fail to establish this foundation, and so-called disruptive technologies will remain just that.

"A Global 500 pharmaceuticals company saved $200,000 of Capital Expenditures (CAPEX) costs by deploying their F5 ARX." **TVID: 249-296-E9A**
(See page 11 to learn more about the TechValidate Realtime Research Network.)

F5 Networks
North America
e : info@f5.com
USA
t: (206) 272-5555

F5 Networks EMEA
e : emeainfo@f5.com

UK
t: (+44) 01932 582 000

Spain
t: (+34) 934671309

Italy
t: (+39) 02 62 03 2180

Finland
t: (+358) 20 75 28 680

Germany
t: (+49) 0 89 94 383 0

The Netherlands
t: (+31) 20 658 63 50

Belgium
t: (+32) 16 39 47 60

Denmark
m: (+45) 2877 4333

France
t: (+33) 1 41 44 89 50

Norway
t: (+47) 99480071

Israel
t: (+972) 3 769 1500

Kingdom of Saudi Arabia
t: (+966) 56 008 7660

Dubai
t: (+971) 4 361 6654

Sweden
t: (+46) 0 70 551 61 00

More Contact Information
www.f5.com/about/contact

IT agility. Your way.