




Inside Track Research Note

In association with



Storage Quality of Service Management

The automation imperative

January 2015

In a nutshell

About this Inside Track

The insights presented in this document are derived from independent research conducted by Freeform Dynamics. Inputs include in-depth discussions on the latest technology developments with IT vendors and service providers, along with intelligence gathered from mainstream enterprises during broader market studies.

If your organisation is similar to its peers, the growth in data volumes you are experiencing will be high across the majority of application types.

Most are playing the consolidation game and increasingly turning to shared storage solutions.

Organisations are increasingly turning to multi-tiered shared storage environments to deal with growing data volumes, changing business requirements, and IT delivery models. Managing Quality of Service (QoS) at scale, however, requires a highly automated approach to enable continuous 'lights out' optimisation.

Breaking down the silos

Not so long ago, storage was largely procured and implemented on an application-by-application basis. Today things are different. Fuelled by developments in virtualisation and cloud computing, it's now normal to think of storage as a shared resource. The need for dedicated storage hasn't gone away, but now it's the exception rather than the rule in relation to new application requirements.

The trouble is, however, that historical investments and infrastructure have a habit of hanging around. Over the years, you may have accumulated multiple types and generations of equipment. If this is the case you will be living with a fragmented storage landscape and the high management overhead that comes with it. You'll also be suffering from under-utilisation of assets due to reuse constraints, and poor returns from incremental investments as capacity is over-provisioned in each silo individually to provide adequate headroom for demand fluctuation and growth.

And if your organisation is similar to its peers, the growth in data volumes you are experiencing will be high across the majority of application types. Add this to the constant stream of new application requirements in areas such as workforce collaboration, digital customer engagement and business analytics, and your challenges are going to get worse.

In terms of business impact, this is significant because it isn't just about direct expense and overhead; fragmented and disjointed systems are notoriously fragile and both difficult and costly to change. You can therefore add business constraint and risk into the mix, as the organisation's ability to respond to market threats and opportunities is inhibited. If you think this sounds a little dramatic, take a minute to consider how almost any business change initiative nowadays will have a significant electronic data component, and therefore storage-related implications.

The shared storage nirvana

Against the above background, it's not surprising that most are playing the consolidation game and increasingly turning to shared storage solutions. Some are further along in making the shift than others, but the trend is clear, and for good reason. Shared environments allow you to tackle many of the problems we have been discussing, with more efficient use of assets, and easier administration.

Pooling of storage also goes hand-in-hand with the move to dynamic private cloud architectures and the implementation of modern delivery practices such as DevOps. The latter, for example, relies on rapid provisioning and deprovisioning of resources to support iterative software development, testing and deployment. The last thing you need when implementing flexible and responsive modern methods is a reliance on an elongated physical asset procurement cycle in order to get anything done.

One of the decisions you have traditionally had to make when procuring and implementing storage on a dedicated basis is the class of equipment to buy.

The effectiveness of an application-specific manual-intensive approach is undermined when you shift to a highly shared storage environment.

As demands fluctuate, there will be periods in which applications end up competing with each other for available resources.

On the surface, the shared storage approach therefore makes many historical challenges go away, and sets the business up well for the future. But does it really lead to a state of nirvana?

A new set of management challenges

One of the decisions you have traditionally had to make when procuring and implementing storage on a dedicated basis is the class of equipment to buy. At the simplest level, a high throughput business critical application would require fast resilient storage, while a slower moving back office administrative system might operate perfectly acceptably on cheaper commodity devices.

Of course in practice it has been a bit more complicated than this when you take things like the age and status of data into account. Even in the most demanding transaction processing system, for example, older records tend to be relatively static and infrequently accessed. Storage tiering has been a popular way to handle this, based on schemes such as the offloading of data onto lower class devices when it reaches a certain age or state.

In a dedicated storage environment, the point at which migration between tiers takes place is typically based on the logic and usage patterns associated with the specific application concerned. Furthermore, it has generally been possible to manage things through periodic batch movement of data using manual administration processes.

However, the effectiveness of an application-specific manual-intensive approach is undermined when you shift to a highly shared storage environment. While a storage pool could be based on multiple classes or tiers of equipment, managing the placement and migration of data when many applications are involved soon becomes an unwieldy task.

To begin with, it's extremely hard to formulate a workable set of data placement and migration policies and rules when applications each have different information lifecycles, access patterns and service level requirements. Maintaining the logic thereafter against the backdrop of relentless growth and ever-evolving business requirements can then become prohibitively complex.

Quality of Service (QoS) implications

The management challenges we have highlighted can easily result in negative quality of service implications due to the difficulties of keeping data placement in line with throughput, latency and availability requirements.

Beyond this, there is then the problem of contention. All storage setups have finite limits when it comes to I/O and network bandwidth, and as demands fluctuate, there will be periods in which applications end up competing with each other for available resources. Unless this is dealt with effectively, the impact of such contention can be both unpredictable and undesirable. You could easily end up with a situation in which a business critical application is bottlenecking because a less important workload is hogging a key resource.

One way around this is to size the system to cope with the theoretical aggregate peak requirement across all applications, but you then end up simply recreating the old

Shared storage environments can easily create as many problems as they solve when you try to implement them at scale.

It's about focusing on the 'what', and using technology level automation to deal with the 'how' of QoS management.

Modern tools and techniques go a long way to delivering on the promise of storage pooling, while avoiding the common drawbacks.

utilisation and efficiency issues. And anyway, the headroom you would need when dynamic cloud architectures and/or DevOps-style delivery are in the mix is going to be ever changing anyway. This is because a fundamental business requirement when designing modern infrastructure is an ability to introduce new applications and workloads safely and efficiently at short notice.

So, it looks as if shared storage environments can easily create as many problems as they solve when you try to implement them at scale. How do we get around this?

Rethinking the problem

The trick is to get away from trying to instantiate application and information level logic as data placement and migration rules at a storage level, and take a more empirical approach instead. It's about focusing on the 'what', and using technology level automation to deal with the 'how' of QoS management.

The whole process begins with considering the quality of service requirements of each application. These can be defined in terms of minimum throughput, response times, availability and other relevant metrics. As part of this, you either implicitly or explicitly define the relative priorities of individual applications, so when contention occurs, 'winners' can be determined objectively based on genuine business requirements, and the activity of 'losers' throttled accordingly.

And when it comes to the distribution of data across storage tiers, this is continually optimised based on historical and current activity rather than theoretical projections and rule sets. Access frequency and performance in relation to all data is monitored on an ongoing basis, with reference to quality of service settings. Movement of data between tiers then takes place in an automated manner.

Through this kind of mechanism, 'hot' data with a high quality of service requirement will rapidly end up in the highest performing tier, regardless of whether the intensity of access was anticipated or not. At the same time, less frequently accessed data will be automatically relegated down the tiers, though only as far as the predefined quality of service parameters allow.

But it's also important that the logic works in all directions. Lower priority data may be automatically migrated to a higher tier based on increased access, but not so high that it begins to interfere with a critical application. As a simple example, you would not want heavy reporting activity against historical data to interfere with live transaction processing taking place against the same storage pool.

The bottom line

The kind of advanced automation capability described above is available on the market today and effectively provides continuous 'lights out' optimisation of a multi-tiered environment.

Solutions in this space might not totally negate the need for finely tuned application-specific landscapes in exceptional scenarios. Pretty much any organisation of any size, however, can be enabled to manage quality of service effectively across the majority of their application portfolio in a shared storage environment, regardless of the level of expertise in place. As such, modern tools and techniques go a long way to delivering on the promise of storage pooling, while avoiding the common drawbacks.

About Freeform Dynamics

Freeform Dynamics is an IT industry analyst firm. Through our research and insights, we aim to help busy IT and business professionals get up to speed on the latest technology developments, and make better informed investment decisions.

For more information, and access to our library of free research, please visit www.freeformdynamics.com.

About Fujitsu

Fujitsu is the leading Japanese information and communication technology (ICT) company offering a full range of technology products, solutions and services. Approximately 162,000 Fujitsu people support customers in more than 100 countries. We use our experience and the power of ICT to shape the future of society with our customers.

Business-centric storage

Under the ETERNUS storage solution brand, Fujitsu provides disk systems and data protection appliances, which allow customers to flexibly manage their increasing data volumes at a lower cost of growth, and to benefit from a very reliable architecture and radically simplified operation.

Storage Management Software

The ETERNUS SF storage management software is a uniform management solution that works across the storage portfolio, from entry-level systems to high-end configurations, including virtualised server infrastructures. All essential operations for resource management, monitoring, reporting, tiering, performance management, disaster resilience and business continuity are integrated.

For more information, please see <http://www.fujitsu.com/eternus>.

Terms of Use

This document is Copyright 2015 Freeform Dynamics Ltd. It may be freely duplicated and distributed in its entirety on an individual one to one basis, either electronically or in hard copy form. It may not, however, be disassembled or modified in any way as part of the duplication process. Hosting of the entire report for download and/or mass distribution by any means is prohibited unless express permission is obtained from Freeform Dynamics Ltd or Fujitsu. The contents contained herein are provided for your general information and use only, and neither Freeform Dynamics Ltd nor any third party provide any warranty or guarantee as to its suitability for any particular purpose.