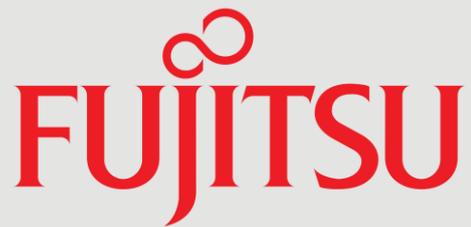




Inside Track Research Note

In association with



Enterprise Storage Architectures

Is it only about scale up or scale out?

August 2015

About this Inside Track

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

Enterprise storage arrays have always been designed with very clear objectives in mind, namely to deliver reliable storage platforms that could operate for many years with very high availability and good consistency of service quality.

With limited ability to upgrade capacity and performance without service interruption, organisations were often compelled to acquire arrays that had far more storage than was required at system installation time.

In a nutshell

Enterprise scale up storage arrays have traditionally been expensive to acquire and complex to manage. While effective, scale up systems are not without their challenges. For example, you may well recognise the issue of storage capacity paid for up front that is still unused three, four or even five years later. To tackle this and other challenges inherent in scale up storage systems, vendors have developed new offerings designed to meet the challenges of today. In particular 'software defined scale out' storage architectures have recently gained high visibility, but don't fall into the trap of thinking they are the answer to life, the universe and everything. This paper looks at the pros and cons of both scale up and scale out storage solutions and asks if there might be a third storage architecture that combines the best features of both approaches without bringing along too many of their limitations.

A history refresher

Traditional enterprise storage arrays have always been designed with very clear objectives in mind, namely to deliver reliable storage platforms that could operate for many years with very high availability and good consistency of service quality. In addition they were built with 'enterprise class' data protection capabilities to ensure data could be safely accessed over the lifetime of the platform. Most were built using 'scale up' architectures.

Such systems usually came bundled with an expensive price tag and required highly skilled IT staff to keep them functioning effectively. They were also notoriously difficult to size cost-effectively as few organisations had an accurate idea of just how much data they would need to hold over the extended lifetime of the storage. As a consequence some systems were very poorly utilised in terms of storage capacity usage, making them less cost-efficient than desired.

Alternatively if they were initially under scoped they could quickly run into trouble when it became necessary to upgrade storage capacity. In addition it could be difficult, if not impossible, to add required functionality as business needs changed.

Today new approaches to storage are grabbing the headlines, especially with scale out systems such as Ceph and Swift attracting considerable interest. Some commentators even go so far as to say that enterprise scale up storage systems, are dead or dying. The reality is that enterprise storage is still evolving and this paper looks to see if there is still life in the old dog.

Traditional enterprise scale up challenges

The acquisition of new enterprise storage platforms has always been a complex undertaking. A major issue concerned the fact that selecting an array of the appropriate storage capacity could be problematic. With limited ability to upgrade capacity and performance without service interruption, organisations were often compelled to acquire arrays that had far more storage than was required at system installation time.

With scale up arrays this over-provisioning was undertaken partly to ensure that future capacity growth needs could be fulfilled, but also reflected that in many organisations obtaining funding for additional capacity at a later date could be a

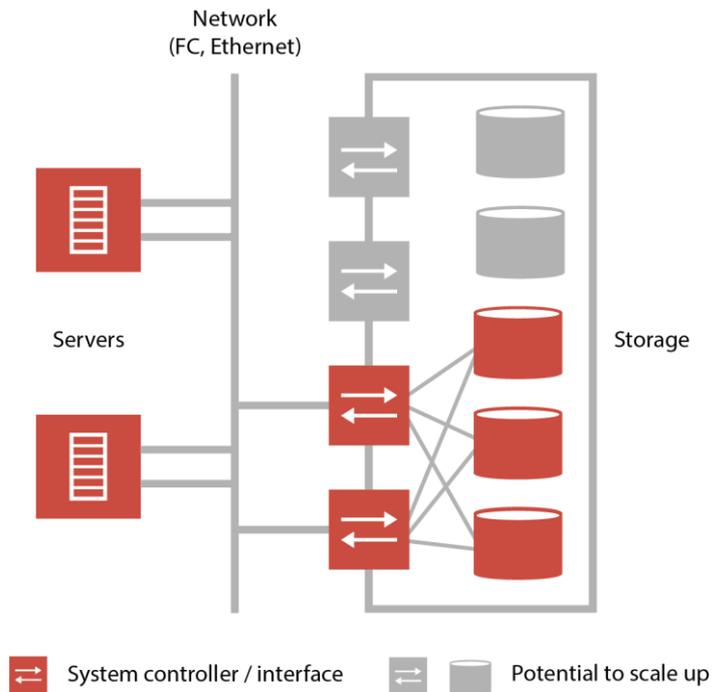
With the cost of storage falling every year in terms of €/Gb, \$/Gb or £/Gb, overprovisioning on day one made little financial accounting sense.

complex and time consuming process. With the cost of storage falling every year in terms of €/Gb, \$/Gb or £/Gb, overprovisioning on day one made little financial accounting sense.

But ultimately how far the storage could grow was a challenge, as absolute capacity was constrained by the number of shelves holding disks that could be fitted into the array's cabinet, while performance was limited by the ability of the controllers to handle data flow to and from the disks (Figure 1).

Figure 1

Scale up storage architectures combine hardware and software components into a tightly integrated system



In addition many enterprise storage arrays had little capacity to deliver storage with different service levels to meet the needs of multiple applications. Thus it could also be the case that arrays were acquired to meet the needs of the most performance stringent workloads meaning that other applications were being serviced on more expensive / performant solutions than might otherwise be required.

The capabilities embedded in scale up enterprise arrays typically developed very slowly, if at all.

A final factor that also influenced the approach of 'buying the best you can afford' is the simple fact that the capabilities embedded in scale up enterprise arrays typically developed very slowly, if at all. When new capabilities were developed on other storage devices, enterprise platforms could take years to update their own management tools. In some instances the only way to obtain some functionality was to migrate all data to an entirely new platform, a task still today seen to be complex and risky.

Some new storage solutions have now reached a level of maturity whereby they are ready for mainstream enterprise usage.

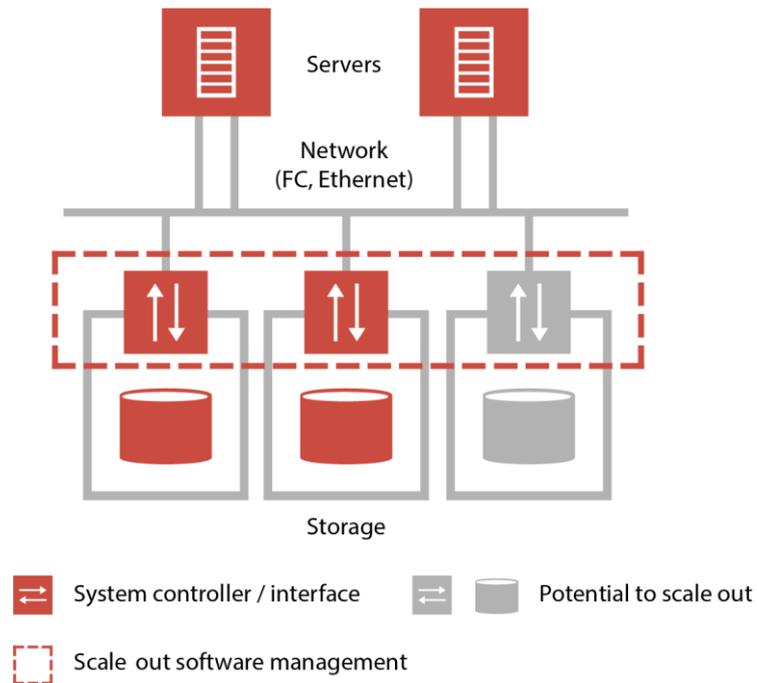
Storage evolution – software defined scale out architectures

In the face of such technical challenges and the requirement of modern businesses to be able to deploy new services rapidly as business conditions change, the storage industry has been actively developing new solutions. Some of these have now reached a level of maturity whereby they are ready for mainstream enterprise usage rather than being constrained to niche technical edge cases. One such solution that has received a lot of attention is 'software defined scale out storage'.

In software defined scale out storage systems sophisticated management software is used to combine diverse physical storage resources together into a single logical entity. In essence the physical storage resources such as CPU / Memory and a variety of storage devices such as disk or flash / SSD are turned into pooled virtual resources by the scale out software and a network, usually Ethernet or gigabit Ethernet. Many scale out storage systems utilise software defined storage architectures (Figure 2).

Figure 2

Scale out storage architectures are based on the idea of adding in hardware components as and when required using software solutions to glue everything together as a single logical entity.



A major benefit of clustered scale out storage platforms is their ability to scale performance and capacity independently.

A major benefit of clustered scale out storage platforms is their ability to scale performance and capacity independently. The addition of either storage capacity (usually in the form of disks / flash / storage arrays) or processing power (CPU, I/O and memory) takes place as and when needed. If you need more storage but performance is OK, just add more disks or SSDs. If capacity is OK but performance is an issue due to a CPU, memory or I/O bottleneck just add another server to the solution.

The only limits on storage capacity and the number of CPU/Memory resources added to a system come from the scale out management software rather than any physical limitations.

In effect, scale out storage systems can grow from very small to very, very large at the pace you need, not that at which your vendor wishes to sell you storage.

This capability completely sidesteps one of the major challenges with traditional scale up storage, namely having to acquire all, or at least very considerable chunks, of capacity up front, even if it is unlikely the usage of the array will get close to the capacity acquired for some time. In effect, scale out storage systems can grow from very small to very, very large at the pace you need, not that at which your vendor wishes to sell you storage.

Another factor that is becoming more important is that scale out storage solutions also tend to have new functionality added to its data management and protection capabilities quite rapidly. In addition the distributed nature of its architecture can allow for software upgrades, and even, when required, hardware upgrades or fault resolution to be handled non-disruptively.

Scale up versus software defined scale out

It is clear that the clustered scale out storage architecture has attractive capabilities, but as with nearly all IT there are advantages and disadvantages with any approach. Let's take a high level look at the pros and cons of each of these storage architectures.

Benefits of scale up architectures include consistent latency and performance as data volumes expand, plus they are well suited for most general enterprise workloads.

Scale up architectures

At a high level the advantages inherent in scale up storage arrays can be summarised:

- Simple capacity growth within array
- Consistent latency as data stored increases up to predictable capacity
- Consistent performance dependent on storage medium (HD, flash etc.)
- Simple networking requirements beyond array
- No additional networking gear required as array scales up
- Specialised controllers to ensure data integrity and performance
- Well suited for most general enterprise workloads
- Procured from and supported by a single vendor

While their disadvantages can include:

- Limited flexibility to modify performance / capacity requirements
- Power of array fixed by maximum number of controllers
- Capacity fixed by number of controllers multiplied by number of disks
- Fixed maximum capacity within each array
- Performance may degrade as utilisation of storage grows beyond 79 / 80%
- New storage management functionality dependent of supplier upgrades resulting in long wait times for new capabilities

Scale out architectures

Again, at a high level the advantages inherent in scale out storage arrays can be summarised:

- Ability to scale capacity and performance independently
- Ability to utilise standard x86 / x64 based equipment
- Wide range of storage devices supported, often including JBOD
- Potential capability to upgrade / refresh hardware without service interruption, thereby elongating usable life of the system and minimising the need for future migrations
- Storage management features are independent of any one hardware device
- Total storage capacity and performance limited only by management software and network constraints
- No single point of failure

While the disadvantages potentially inherent in scale out may include:

Disadvantages of scale up can include a limited ability to modify performance of the array, and storage management upgrades can be slow to be delivered.

Advantages of scale out architectures can include the ability to utilise a wide range of storage devices, plus storage management features are independent of any one hardware device.

Disadvantages of scale out architectures can include the potential for latency to grow as storage volumes expand, plus setup and management of solutions can be complex.

- Potential for latency growth as storage volumes / number of processing nodes increases
- Network infrastructure (switches, cables etc.) linking clustered scale out kit can become complex as cluster grows
- Set up and management can be complex unless solution acquired from single supplier
- Support may depend on multiple suppliers unless entire solution acquired from single supplier
- Use of general IT equipment may not deliver best performance and consistency as system scales out

Of course this begs the question, isn't there any way of getting the best of both worlds?

A third approach

With both scale up and scale out storage solutions offering advantages and challenges, some vendors have recently been developing solutions that try to combine the best of both worlds while limiting the disadvantages of each.

One approach adopted by some vendors has been to build storage that on the surface appears to be a 'scale up' array but internally uses a 'clustered scale out' architecture. These often make use of a multi-controller RAID architecture built as a clustered scale out array which can offer a combination of the benefits of both architectures:

- Storage solution is supplied by a single vendor as a single order unit
- Entry capacity can be small but with the ability to scale up as workloads increase, thus avoiding under-utilisation of storage and unnecessary spend
- Choice of hard drives, flash / SSDs as required
- Solution built using specialist components rather than generic hardware, e.g. controllers and optimised networking to link all storage drives
- Latency and performance designed to be consistent as storage capacity grows
- Single management tool for entire storage

More sophisticated solutions may also offer additional capabilities now intimately associated with 'enterprise class' storage solutions.

- Automated management capabilities
- Service level management using different storage tiers
- Very high availability with no single point of failure
- A range of data protection / replication capabilities both internally and to external storage resources
- High Availability / Disaster Recovery via synchronous data replication to a second clustered scale out array

One clear potential down side of the clustered scale out array architecture is that the system still has a physical limit on its ability to expand, in particular in terms of how many storage disks and controllers can be accommodated before it becomes necessary to migrate to another platform.

Some systems offer a combination of the benefits of both scale up and scale out architectures.

One clear potential down side of the clustered scale out array architecture is that the system still has a physical limit on its ability to expand.

Beyond technical issues it is worthwhile evaluating the ability of potential storage suppliers to provide solutions appropriate to needs rather than promoting only one offer to every use case.

In addition the pre-set components of the controllers, storage disks and internal networking can limit the total flexibility of performance in the platform. This may not be important for many services but could be an issue if there are specialist applications, for example, those that demand the very lowest latency possible.

Beyond these technical issues, it is worthwhile evaluating a few additional points with your prospective supplier.

- Do they provide consultancy on solution selection taking full account of your existing start point and the needs of your organisation?
- Do they provide financing to enable storage to be acquired on a 'pay as you grow' basis to avoid unnecessary overprovisioning of storage capacity?
- Do they offer assurance of service quality / performance / latency as data storage grows?
- Can they provide guidance on the maximum utilisation rate of storage capacity before performance is impacted?
- Do they offer guidance / services around data migration from your existing infrastructure to the new platform?
- Is the service offering robust enough for your needs?

As these systems are designed for 'enterprise' class usage without some of the financial challenges long associated with scale up arrays, if the suppliers with whom you are speaking cannot give good answers to all of these questions, you might want to look for someone else to help you move your storage forwards.

The bottom line

Enterprise storage is one of the most important assets in your IT Infrastructure, but unless something has gone seriously wrong it is also likely to be one of the least visible and most underappreciated. However, the absolute importance it has as the repository of business critical information can also make those charged with its operation very cautious in making major changes to the enterprise storage landscape.

After all, history tells us that the entire migration process from one enterprise storage platform to another is time consuming, complex and fraught with business risk. Moving data between systems has become more straightforward, but the continuous ramping up of business expectations makes change difficult to evaluate. Storage professionals want to be able to trust the platform they are acquiring which often means turning to vendors with a demonstrable 'track record'. But such vendors must also be able to show they have solutions that meet modern business requirements.

Storage professionals want to be able to trust the platform they are acquiring which often means turning to vendors with a demonstrable 'track record'.

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

Fujitsu's ETERNUS storage portfolio balances storage capacity, performance and costs for the complete lifecycle of data - from production, business analytics and big data to backup and long-term archiving.

The comprehensive ETERNUS storage family comprises

- all-flash systems
- hybrid disk and flash storage arrays
- hyper-scale and software-defined storage solutions
- backup appliances
- tape libraries

ETERNUS excels in "Business-centric Storage" by synchronizing storage resources with business priorities whilst reducing costs.

Based on ETERNUS products, leading technology from our partners and comprehensive services from Fujitsu we deliver complete storage solutions.

For more information, please see 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.