FUJITSU

White Paper

The new Storage Zoo – which storage animals fit to your IT ecosystem?

Contents

The storage impact of a digitalized world

The digitalization of the world will have a big impact on storage. More data sources create more data, which needs to be gathered, transported, processed and stored. Analysts predict a growth of data volumes by a factor of 10 in the next three years. This is generating increased storage requirements in terms of speed as well as faster and greater scalability. In order to contain the costs of data growth the costs per terabyte need to be reduced significantly. Since storage systems have to host larger data volumes, customers require longer lifecycles, as the migration of hundreds of terabytes or even petabytes of data is a complex task. Furthermore, it is necessary to align the scalability of storage servers and networks more tightly in order to avoid bottlenecks generated by this new flood of data.

More data → **More storage requirements** → **More storage species**

- will be gathered
- will be stored
- will be analyzed
- will be processed
- will be transported
- will be online

- Higher speed
- Faster, greater scalability
- Less costs per terabyte
- Extended storage lifecycles
- Aligning storage, server, network scalability

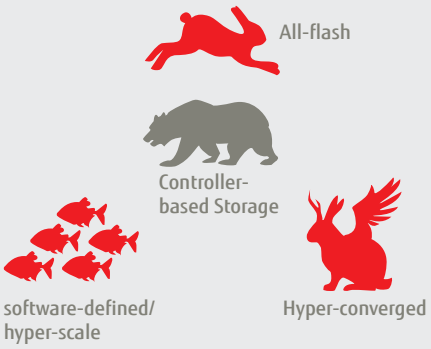

Not all of these new requirements can probably be solved with a single storage architecture. This is the reason why we can observe the emergence of new storage species like all-flash systems, software-defined and hyper-scale architectures or hyper-converged approaches, which are added to the classic controller-based storage designs. In a nutshell, the storage zoo is filling-up with new species. In the following we will take a closer look at these new storage animals but also at how innovations of classic controller-based systems help to serve new requirements.

Rapid like a rabbit
All-flash storage

The objectives of all-flash storage systems are to reduce the latency which is created by hard disks. Flash-optimized storage systems are used for areas like business analytics, virtual desktops, performance-demanding virtual machines, OLTP environments with high frequency transactions and so on. Here all-flash storage systems can deliver great benefits, however they also produce some new challenges. In particular purpose-built designs are typically on a high cost level. They often come with completely diverse management software allowing no management integration with disk systems. Furthermore, they often have only basic high-availability capabilities.

IT'S WORTH TO COMPARE

Comparing architectures you will essentially find three types of all-flash storage systems. There are purpose-built systems using proprietary flash modules. Secondly, there are flash-optimized storage systems with a design which has been derived from standard storage architectures. The third approach uses classic storage systems, fully equipped with SSDs. Surprisingly, in real-life system comparisons the response times of particular systems are not necessarily dependent on their architectural category. For example, in official SPC benchmarks Fujitsu could demonstrate a superior performance with its SSD-based all-flash system ETERNUS DX200F compared with purpose-built flash systems from competitors. Even a standard ETERNUS DX600 hybrid system, fully equipped with SSDs, could outperform specialized all-flash designs. Thus, comparing official benchmarks with the required response time can save a lot of money as quite often it is not necessary to invest in specialization.

Model	SPC-1 IOPS™ **	Response time (ms)
ETERNUS DX200 S3	200,500.95	0.63
ETERNUS DX600 S3	320,206.35	0.61

* Results are current as of July 25, 2014 and available at: http://www.storageperformance.org/results/benchmark_results_spc1_active/#fujitsu_spc1
** Input/Output performance per second

REAL-LIFE EXAMPLE

The German market research company GKL offers price information for vendors of consumer goods. The company scans the price tags in supermarkets in several regions and stores and gathers all this data in big databases. In order to provide a good service level to their customers the company needed a fast storage system with the lowest response times. After some testing they decided in favor of an ETERNUS DX600 system with SSDs, as it could provide sustainable performance in the range of response times between 0.1 and 1 ms. A great example of how well-engineered, classic storage systems can reduce the variety of your storage zoo whilst offering an excellent price-performance ratio.

RECOMMENDATIONS

- Most use cases can be served with well designed hybrid systems equipped with SSDs
- Well designed hybrid storage systems can cope with all-flash systems
- Use performance tests to challenge the need of all-flash systems. Keep your storage zoo small!
- Be careful with all-flash systems with always-on dedup
- Check high-availability/quality of service functions, avoid management islands

More information: White papers created by the British analyst firm "Freeform Dynamics"
- The Promise and Practicalities of Flash Storage **(Link)**
- Everyone is talking about dedupe **(Link)**

Swarm intelligence
Hyper-scale, Software-Defined Storage (SDS)

About hyper-scale, software-defined storage there is currently a great deal of hype in the industry. One rough point of reference to decide whether controller-based or SDS systems should be used is a distinction between whether the usage scenario needs more performance-optimized storage versus more low-cost capacity-optimized storage, for which performance does not play a decisive role. For scenarios like storing lots of unstructured data, historical data which needs to be processed in big data environments or large online archives, the software-defined architecture may provide better cost efficiency.

IS SOFTWARE-DEFINED STORAGE THE SILVER BULLET?

SDS decouples the data management logic from the underlying hardware All higher management functions like data placement, data redundancy, data recovery, high availability and optimization are part of the software layer which runs in parallel on x86 servers. This software can typically distribute data between different nodes, allowing you to easily scale capacity and performance with each added node, or to exchange nodes during operation when the hardware reaches end-of-life. So the storage system can theoretically "live" forever. Thus, storage migration projects can be avoided or at least the frequency can be greatly reduced. As industry-standard server hardware is used and various software-defined storage platforms are available as an open-source product, there is savings potential in the overall purchase costs. However, a build-your-own-storage approach requires substantial testing and evaluation, adds complexity to maintenance and lifecycle management, and may generate significant operational efforts. All that can create a lot of hidden costs and risks resulting in higher TCO than anticipated.

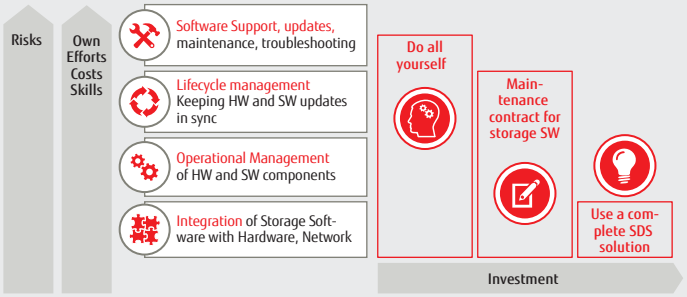

BEST APPROACH: OPEN SOURCE BASED SDS IN A TURNKEY SOLUTION

Fujitsu has introduced with ETERNUS CD10000 a complete enterprise-ready solution. The appliance runs Ceph, probably the most capable platform on the market, supported by a large community of developers. ETERNUS CD10000 is delivered as an appliance with the right-sized hardware, integrated management console, and end-to-end maintenance and support services for the whole solution. In its current version the system can scale between 200 TB and 50 PB. So it is the right platform to manage unpredictable and aggressive data growth dimensions without any practical limits in a very cost efficient way. With this approach customers can fully benefit from SDS and open source without burdening the IT team with high operational efforts and risks.

RECOMMENDATIONS

- SDS can achieve high cost efficiency in storing large amounts of data
- Open source based SDS reduces lock-in, increases viability and has the best cost saving potential
- Build your own storage may generate significant operational efforts, risks and costs
- Look for pre-built solutions with end-to-end support and maintenance over the entire lifecycle

More information: White papers created by the British analyst firm "Freeform Dynamics"
- Hyper Scale Data Management **(Link)**
- Ceph as an Enabler of Growth and Scalability **(Link)**

The cross-breed species
Storage in hyper-converged IT

Another new approach of storing data are hyper-converged infrastructures (HCI), also called hyper-converged IT. In IT terms a hyper-converged solution typically consists of clustered building blocks with integrated network, storage and computing functions. All these functions are controlled by an overarching software stack that orchestrates all the resources. This can be, for example, hypervisors with add-on functionalities, meaning they are no longer just virtualizing servers but they also deliver data management and network control functions.

MOST SUITABLE FOR VIRTUALIZED ENVIRONMENTS

The strength of the hyper-converged IT approach is a deep integration of storage in virtualized environments, the possibility of fast and simple scalability and potentially typically lower purchase costs. This point is highly dependent on the licensing scheme and costs of the orchestration software.

Available are middleware offerings from a variety of vendors or from open source communities. A prominent example is VMware VSAN. VMware has extended its hypervisor functionality with storage capabilities. VSAN aggregates internal disks of servers as well as disks of attached JBODS into one logical and clustered data store. With its latest version it includes a rich quality of service functionalities, comprehensive data management capabilities and several high-availability services to prevent data loss in case of disk and node failures. This is an ideal technology for data services in larger and growing VMware environments.

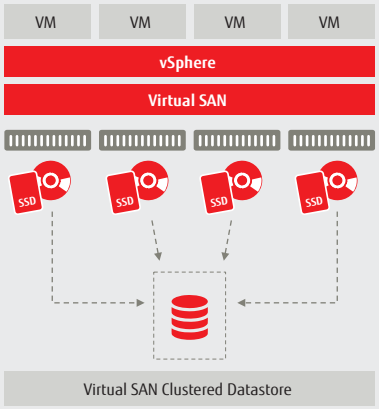

INTEGRATED SYSTEMS REDUCE THE TCO OF HYPER-CONVERGED IT

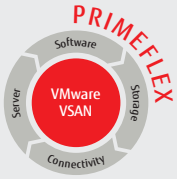

Fujitsu has developed PRIMEFLEX for VMware VSAN, to support customers in reducing the implementation and operational costs. This complete solution integrates all hardware and software components in one right-sized infrastructure. Thus, customers can benefit from a short implementation time. Furthermore, Fujitsu offers support and maintenance for the entire solution along its complete lifecycle. This reduces the TCO of hyper-converged solutions through higher efficiencies in operational aspects.

The UK customer "Pea Soup" for example, deployed PRIMEFLEX for VMware VSAN to serve a group of customers in the area of media streaming by delivering fast-scaling virtual infrastructures. With this solution they can offer the right compute and storage capacity as needed without doing too much pre-investment themselves, and they were able to reduce their internal operational and maintenance efforts.

RECOMMENDATIONS

- Alternative storage approach for midsize and large data centers, greenfield installations, larger remote sites, data center modernization projects, for highly virtualized workloads
- Integrated systems reduce implementation and lifecycle efforts
- Typically storage in hyper-converged solutions coexist alongside controller-based storage, thus creating an additional storage silo
- Cost comparisons with dedicated storage work only on a TCO basis
- Include backup into your solution design right from the beginning

More information:
- PRIMEFLEX strategic briefing **(Link)**
- PRIMEFLEX for VMware VSAN **(Link)**

The agile bear
Evolution of controller-based storage

Faster response times, greater scalability, reduced storage TCO or stronger integration in virtual IT: In the latest version of its controller-based ETERNUS DX storage line Fujitsu has embedded a lot of innovations that serve all these new requirements. ETERNUS DX is a fully compatible family of controller-based storage systems ranging from cost-efficient economy storage to highly scalable enterprise storage arrays for data centers. By using the same hardware components and the same management software across the line, ETERNUS DX helps to reduce storage complexity resulting in higher operational efficiency and thus contributing to reduced storage TCO.

HIGHLY VERSATILE STORAGE FAMILY

Faster response times: ETERNUS DX belongs to the performance-leading storage systems in the market and achieves superior response times. Performance-hungry scenarios can be served without the need to add another storage silo. For example, the systems use the latest processor technologies with multicore and multithreading capabilities. The internal operating system has been optimized to make full use of them. They support large amounts of DRAM caches and can enhance the cash capacities by using pools of SSDs.

Automation reduces TCO: Very often customers invest in bigger systems than they actually needed from a capacity point of view or they use discrete storage systems for different workloads in order to avoid performance conflicts. This is why Fujitsu has developed a technology called automated quality of service management. This allows you to prioritize the data performance of the system according to business priorities. Operational efforts can be reduced and system utilization can be increased to 90%.

Tight integration in virtualized IT: All our ETERNUS DX systems support virtual volumes or "VVOL" – a fairly new technology from VMware. ETERNUS DX systems can create virtual volumes which can be directly managed out of VMware, thus relieving administrators from mapping VMDK files to LUNs. Furthermore, the automated quality of service management function of ETERNUS DX supports the definition of service levels regarding different performance needs of particular virtual machines.

GREAT STORAGE SCALABILITY

Should the capacity or the performance of a particular ETERNUS DX model be exhausted, it is possible to upgrade the system to a higher class and so helping it to overcome limits. In the enterprise class ETERNUS DX systems support the scale-out of storage capacity by adding additional storage controllers to one system.

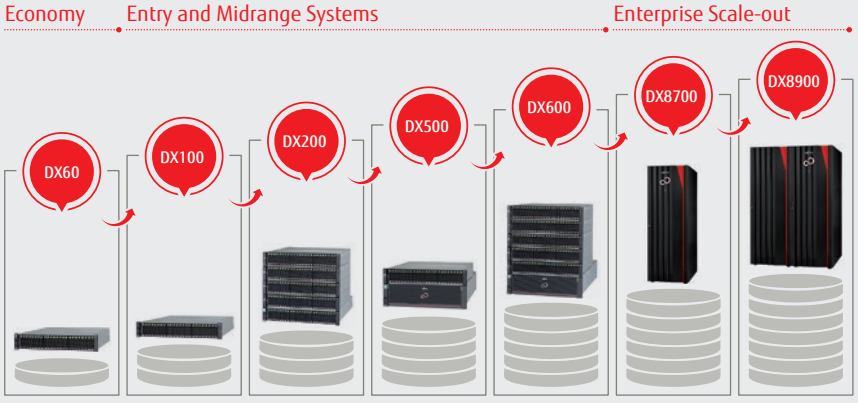

RECOMMENDATIONS

- **Well engineered controller-based storage systems can cover a lot of new storage requirements helping to keep your storage zoo lean.**
- **Look for performance architectures with officially published benchmark, and automated quality of service functions enabling large savings in investment and operations**

More information: White papers created by the British analyst firm "Freeform Dynamics"
- Enterprise Storage Architectures **(Link)**
- Storage Quality of Service Management **(Link)**
- The Impact of VMware VVOLs on Storage **(Link)**

Summary

As new storage requirements have triggered the emergence of new storage species, the key question in order to decide which storage animal you should look for is the following: If you need faster response times, greater scalability, reduced storage TCO and stronger integration in virtual IT, is it necessary to deploy many point solutions? Or can these new requirements be covered with the help of classic general propose storage.

- We recommend keeping your storage zoo as lean as possible to contain complexity and to reduce investments as well as operational efforts.
- Do performance tests and look for official benchmarks to justify the usage of all-flash systems. Many performance needs can be covered with general purpose storage like ETERNUS DX.
- Software-defined storage shows its strengths in areas of large and fast scalability needs, in particular for large amounts of unstructured data which do not require fast access performance.
- Storage in hyper-converged approaches currently works best for large-scale server virtualization scenarios. Here they can unfold their particular strengths.
- Look for complete solutions in the hyper-converged and software-defined storage area in order to keep implementation and lifecycle efforts low.

Fujitsu has experience and offerings in all storage categories and in helping you to find the right solution for the particular storage needs. So let's talk about your future storage zoo!




Business-Centric Data Center – the right solution for your business needs

On premise · Hybrid IT · In the Fujitsu cloud

Managed by yourself · Managed by Fujitsu

Robust IT · Bimodal IT · Fast IT

Business-Centric Integrated Systems

Business-Centric Storage

Business-Centric Computing

Contact
Fujitsu Technology Solutions GmbH
Mies-van-der-Rohe-Straße 8
80807 Munich

© 2016 Fujitsu – All rights reserved, including intellectual property rights. Technical data subject to modifications and delivery subject to availability. Any liability that the data and illustrations are complete, actual or correct is excluded. Designations may be trademarks and/or copyrights of the respective manufacturer, the use of which by third parties for their own purposes may infringe the rights of such owner.