

The future of storage is now.
Consider a two tier storage strategy.



Sponsored by Scality
March 2016

TABLE OF CONTENTS

TABLE OF CONTENTS.....	2
EXECUTIVE SUMMARY	3
LOCAL AND DISTRIBUTED PERFORMANCE	3
THE ROLE OF SOFTWARE-DEFINED STORAGE	4
TARGETING A TWO TIER STRATEGY	5
UNDERSTANDING THE MARKET	6
A LOOK AT THE MARKET LANDSCAPE	6
BUILDING THE RIGHT STORAGE STRATEGY	8
BOTTOM LINE	10
JUKU	11
WHY JUKU	11
AUTHOR.....	11

EXECUTIVE SUMMARY

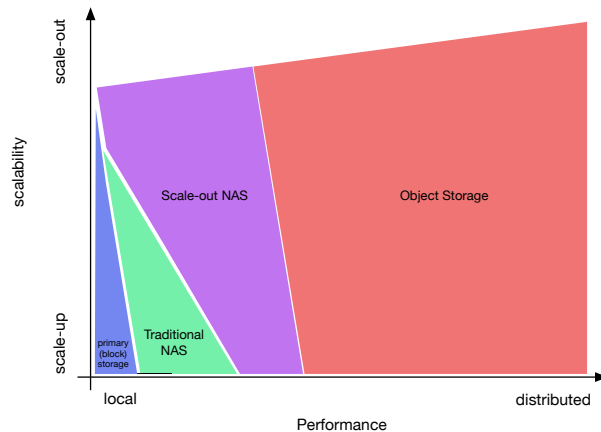
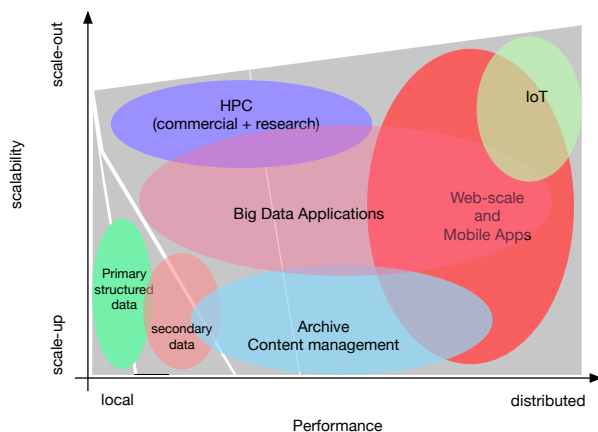
The enterprise data storage landscape has been rapidly changing in the last few years. Demand for more storage resources is unstoppable but, contrary to the past, enterprises and ISPs are experiencing a strong diversification in the type and quality of resources needed to satisfy each single different user or specific business need.

Data growth and workload diversity are placing a major challenge to traditional infrastructures, and are revealing all their limits and constraints in terms of agility and, more in general, TCO.

New requirements are being heavily driven by different workloads and types of data that were quite uncommon just a few years ago. It's not only the growing amount of data to manage or faster access speed to serve more transactions in the same time. Cloud and mobile computing, as well as new access patterns imposed by web applications, global data access and user behaviors are impacting the entire infrastructure stack and ask for more agility and flexibility.

LOCAL AND DISTRIBUTED PERFORMANCE

Over time, demand for local performance, which can be described with high IOPS and latency sensitive data and workloads running in between the datacenter, has been joined by distributed performance needs that are associated to high throughput, parallelism and distributed data accessed from anywhere and on any device. And even though high IOPS demand can also be a characteristic of distributed performance workloads, this is always associated with higher latency due to the nature of design of distributed systems and consistency of network connectivity involved.



Data growth and workload diversity is challenging traditional infrastructures in a major way and are revealing all their limits and constraints in terms of agility and, more in general, TCO.

The quantity of primary data, usually stored on monolithic or modular scale-up systems (i.e. structured data, such as DBs), is growing much less than any other form of data. A visible market trend is showing that server-based storage is growing (+10% Y/Y) while traditional external disk storage systems are seeing a

3.9% decline in the same period of time¹. Even though some of this change is due to an increment in large datacenter expenditure from public cloud service providers², part of it comes from newly designed scale-out software-defined infrastructures for large and mid-sized organizations.

THE ROLE OF SOFTWARE-DEFINED STORAGE

Software-defined and commodity hardware are part of a huge paradigm shift in the way medium and large size infrastructures are designed, and this is also moving the needle from scale-up to scale-out architectures:

- Commodity hardware is now capable of great performance at reasonable prices while interconnect technology, at the base of scale-out systems design, now offers microsecond latencies and high throughputs.
- Distributed systems, like NoSQL-based back-ends for example, allow vendors and users to build rich sets of modern features, adding analytics functionalities while granting unmatched scalability.

Only a small percentage of data (between 10 and 20%) needs local performance, all the rest is unarguably better served by infrastructures capable of a distributed performance.

Today, only a small percentage of data (between 10 and 20%) needs local performance³, all the rest is much better served by infrastructures capable of distributed performance. The latter is also the kind of data that sees the highest growth, and this is why scalability is a key factor.

	Type of data	Performance defined by	Common design	Interfaces	Media	Efficiency defined by
First tier	Mostly structured	IOPS Low latency Local access	Scale-up Scale-out ⁴	Block File (limited)	Flash Memory NVRAM	Data footprint reduction (deduplication, compression) Data services Hypervisor integration
Second tier	Mostly unstructured	Throughput Parallelism Distributed access	Scale-out	API File Scale-out FS	HDD SSD (metadata)	Data protection (embedded replication, erasure coding) Policy based automation

¹ Source: IDC quarterly storage tracker Q215-Q214 comparison

² Source: IDC WW storage for public and private cloud 2014-2018 forecast

³ Source: Latency vs. Capacity Storage Projections 2012-2026, Wikibon

⁴ Limited scale-out design, usually up to 8 node maximum

Accordingly to IDC findings in recent market research⁵, the sum of server-based and ODM (Original Design Manufacturer) storage systems is already outpacing external storage in terms of capacity installed, and will be doubling it by 2016. With a tendency that sees prices going down to an aggregated average (ODM+internal) of \$0.05 per GB against \$0.20 per GB of external storage systems by 2019, there is no doubt that aiming at having a sustainable infrastructure in the coming years means leveraging software-defined scale-out storage technology based on commodity hardware design for the largest part of data under management.

TARGETING A TWO TIER STRATEGY

In order to achieve the best in terms of TCO and service levels, a two-tier strategy is fundamental.⁶ Primary data will be stored in very high performance Flash based systems that emphasize block protocols and usually a more traditional design or an integrated hyper-converged architecture.

On the other side, large scale-out repositories, capable of simultaneously serving file and object protocols, will serve less latency sensitive workloads.

Such a strategy will improve infrastructure agility and manageability allowing an IT organization to best serve a wider range of workloads and implement cloud storage services while sustaining data growth with an overall better infrastructure TCO.

⁵ Source: IDC WW Enterprise storage systems forecast, 2014-2019

⁶ Source: <http://juku.it/en/storage-changing-thinking-new-strategy/>

UNDERSTANDING THE MARKET

Following a similar market study done by Jerome Lecat (Scality's CEO) last year⁷, defining market trends of FOBS (File and Object Storage combined) will help to easily identify new opportunities and the right strategy to build sustainable infrastructures for the next decade.

Traditional market segmentation is no longer enough to identify end user needs: workload diversity and infrastructure scale are changing the way different solutions are evaluated by customers.

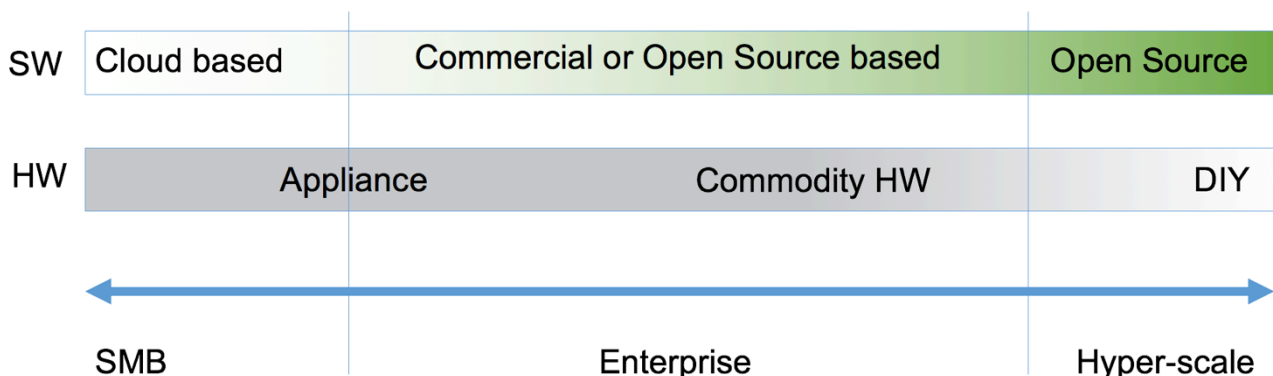
A LOOK AT THE MARKET LANDSCAPE

There are some important considerations, coming from market observation, analyst researches and end user interviews, that are fundamental to understanding what is really happening in the current market:

From the end user point of view:

- Traditional market segmentation (SAN, NAS, Object, ...) is no longer enough to identify end user needs: workload diversity and infrastructure scale are changing the way different solutions are evaluated by customers.
- Thanks to cloud computing, the difference between scale-up and scale-out solutions is now clearly understood and reluctance against to the latter (because of potential complexity issues), is no longer a problem, even for the most conservative enterprise environments.
- It's not only about storage, Software-Defined solutions are getting a lot of attention at every level. Adoption rates of Software-defined Networking and hyper-converged infrastructure solutions are also quickly increasing too. Hyper-convergence in particular is rapidly moving from specific use cases (like VDI for example) to a wider set of business applications.
- Private and Hybrid cloud infrastructures are becoming more common in large and medium sized organizations. End users are now more inclined to invest in human resources (developers and SysAdmins) than in the past.

Software-defined storage adoption model (per organization size)



⁷ The disruptive transition to software-defined: Sizing the opportunity <http://www.scality.com/futureofstorage-marketstudy/>

- Large ISPs and public cloud providers prefer software-defined DIY solutions based on open source software and, where this is possible, cost effective open-source hardware. On the other hand enterprises will likely continue to leverage commercial software to achieve similar results. Smaller organizations will continue to prefer appliances and commercial software.
- Mobile computing and geographically dispersed clients are demanding new forms of storage (e.g. Sync & Share and distributed NAS for remote and branch offices - ROBO), which are not feasible with traditional storage architectures.
- As a consequence of the previous point, demand for object-based storage solutions accessed via standard APIs is growing rather quickly even in smaller organizations. It's not unusual to find end users looking at object storage for on-premises implementations starting in the order of 100TB.
- Many end users are relatively happy with traditional storage solutions already in place, but complain about the high overall TCO.

From the vendor point of view:

- There is a proliferation of new startups working on storage platforms with an object storage back-end, some of them don't even expose RESTful APIs at the moment, but they all have similar basic characteristics like a scale-out design and a software-defined approach.
- The number of solutions supporting object storage APIs (primarily Amazon S3 API) is growing at an incredible pace, and now counting more than 4,000 different products.
- There is an interesting Market consolidation happening. Three startups developing object storage systems have been acquired in the last 18 months (and the number of acquisitions is much higher if we consider all the operations that have involved storage products designed to work with an object store at the backend).
- Now, all primary storage vendors have an object storage solution in their product line up, or an agreement to resell a third party solution.

BUILDING THE RIGHT STORAGE STRATEGY

Traditional enterprise storage does not meet the need of modern applications. Looking at an iconic, and legacy, product like EMC VMAX for example, it's plain to see that it can neither address local nor distributed performance as necessary. Taking no action to replace legacy storage systems for private cloud or any large scale deployment will severely impact the overall infrastructure TCO, putting competitiveness and agility of the entire IT organization at risk.

If \$/GB of a traditional storage system is too high when compared to a modern All-Flash array, it's much worse when the comparison is made with a large capacity scale-out distributed system such as an object store.

High availability of this type of system is still very high but most end users are not really interested in 6-nine availability capabilities. For most, especially those running a cloud infrastructure, 5-nines are more than enough and they are much more interested in performance, latency predictability and TCO. And All-Flash arrays can easily scale to higher performance with a much better overall

efficiency of any legacy disk-based product. In fact, systems like EMC VMAX fail behind in the comparison. It turns out that not only the \$/GB that is worse than an All-Flash array but also \$/IOPS and IOPS/GB, which in this case are even more important than \$/GB, fall very far behind in the comparison.

If \$/GB of a VMAX-like system, once considered a highly scalable system, is too high when compared to a modern All-Flash array (like EMC XtremIO for example), it's gets even worse when the comparison is made with a large capacity scale-out distributed system such as an object store. The latter is now targeting a TCA of cents/GB and its scalability is far beyond any enterprise array possibilities. At the same time, what was historically considered "good" in terms of Disaster Recovery and replication features, is greatly inferior when compared to the capabilities and reliability of the embedded automated multisite data protection schemes offered by an object store that can easily be configured to achieve unmatched data availability, durability and resiliency across continents.



In practice these are the reasons why we are already seeing a substantial decline in traditional enterprise storage system revenues.

Eventually traditional storage systems will be in the middle of an expanding fork that sees:

- All-Flash arrays, and hyper-converged systems, that are pushing towards very low latency and new PCI interconnect fabric protocols (like NVMe and NVMe-fabric) to bring data as close as possible to the CPU with the best of local performance and the lowest latency.


- Large scale-out, and highly automated, capacity storage systems spread out in multiple datacenters capable of serving distributed performance, less latency sensitive, and high throughput, applications and workloads ranging from archiving to mobile computing, IoT, Big Data and more.

The most viable strategy for large IT organizations is to build a two tier infrastructure. The fast tier will leverage Flash memory arrays alongside software-defined solutions (like hyper-convergence). In the second, and slower tier, scale-out systems with an object-store back-end will be the preference and they'll be based on very cheap commodity hardware and software-defined solutions, supporting both File (for legacy workloads) and Object protocols.

BOTTOM LINE

Object-based storage is the only option today to address large multi-petabyte environments with ease. FOBS is already a \$21 billion market already and growing rapidly.

Storage is in great demand with end users but specialized infrastructures are needed for different data types and workloads. The two tier strategy proposed in these pages, with All-Flash for the primary tier and All-Object for the rest of the data under management is the optimum way to get the best in terms of performance, capacity and TCO.



Enterprises are beginning to build data lakes and new challenges are also coming up, like IoT for example. Huge repositories are needed to store data coming from multiple globally dispersed sources that will eventually be analyzed later, if not in real time. This requires massive object streaming throughput, reliability and data availability across geographies, all features that can be found only in modern object-based storage platforms.

Storage is in great demand with end users but specialized infrastructures are needed for different data types and workloads.

Scality RING is one of the few products on the market that has already proven its potential with massive installations on the order of hundreds of Petabytes. This kind of installation will become more and more common from now on. The software-defined approach and multi-protocol support make this product a particularly

interesting solution for designing storage infrastructures of any size, capable of serving either legacy or next generation applications.

JUKU

WHY JUKU

Jukus are Japanese specialized cram schools and our philosophy is the same. Not to replace the traditional information channels, but to help those who make decisions for their IT environments, to inform and discuss the technological side that we know better: IT infrastructure virtualization, cloud computing and storage.

Unlike the past, today those who live in IT should look around themselves: things are changing rapidly and there is the need to stay informed, learn quickly and to support important decisions, but how? Through our support, our ideas, the result of our daily interaction that we have globally on the web and social networking with vendors, analysts, bloggers, journalists and consultants. But our work doesn't stop there, the comparison and the search is global, but the sharing and application of our ideas must be local and that is where our daily experience, with companies rooted in local areas, becomes essential to provide a sincere and helpful vision. That's why we have chosen: "think global, act local" as a payoff for Juku.

AUTHOR



Enrico Signoretti, Analyst, trusted advisor and passionate blogger (not necessarily in that order). Having been immersed into IT environments for over 20 years, his career began with Assembler in the second half of the 80's before moving on to UNIX platforms (but always with the Mac at heart) until now when he joined the "Cloudland". During these years his job has changed from deep technical roles to management and customer relationship management. In 2012 he founded Juku consulting SRL, a new consultancy and advisory firm highly focused on supporting end users, vendors and third parties in the development of their IT

infrastructure strategies. He is constantly keeping a vigil eye on how the market evolves is constantly on the lookout for new ideas and innovative solutions. You can find Enrico's social profiles here: <http://about.me/esignoretti>

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources Juku Consulting srl (Juku) considers to be reliable but is not warranted by Juku. This publication may contain opinions of Juku, which are subject to change from time to time. This publication is covered by [Creative Commons License \(CC BY 4.0\)](#): Licensees may cite, copy, distribute, display and perform the work and make derivative works based on this paper only if Enrico Signoretti and Juku consulting are credited. The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors. Juku consulting srl has a consulting relationship with Scality. This paper was commissioned by Scality. No employees at the firm hold any equity positions with Scality. Should you have any questions, please contact Juku consulting srl (info@juku.it - <http://jukuconsulting.com>).