

# The future of storage is now.

## Market landscape and opportunities.

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.

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.

As a consequence the entire storage industry is living one of the most interesting transformations ever. Software-defined and Server-based storage solutions are quickly growing in popularity among end users, putting even more pressure on traditional vendors, and forcing them to consolidate their position on the market with acquisitions and reorganizations aimed at improve efficiency and cut costs.

Looking at what has happened in the last six months alone, it is not difficult to get a clear picture:

- Dell is in the process of acquiring EMC with the hope to improve overall efficiency, cut costs, and optimize its product portfolio.
- IBM has acquired Cleversafe to improve its cloud services and build hybrid cloud storage infrastructures for its customers.
- Pure Storage, after the IPO, is now the second largest independent storage vendor in terms of market capitalization and has recently added a new product to cover file/object storage needs.
- Netapp, which has invested for years in the concept of unified storage systems, has completely revised its strategy. Now it is heavily working on in its object storage system and has acquired Solidfire.
- HPE, which has concentrated all its efforts for years in 3PAR arrays, has recently made a strategic investment in Scality.

Now, all primary vendors have a stronger proposition for both primary and secondary storage. Trying to cope with the increasing demand in storage expressed by end users, both in terms of latency-sensitive and capacity-driven workloads.

Once again, it is interesting to note that most competitive capacity driven solutions fall in the category of software-defined storage, based on commodity x86 servers and capable of serving multiple file/object protocols at the same time.

# UNDERSTANDING THE MARKET

Following a similar market study done by Jerome Lecat (Scality's CEO) last year<sup>1</sup>, defining size and 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.

## 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.
- 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.
- Multi protocol support is a basic requirement to manage legacy and modern applications.
- 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:

---

<sup>1</sup> The disruptive transition to software-defined: Sizing the opportunity <http://www.scality.com/futureofstorage-marketstudy/>

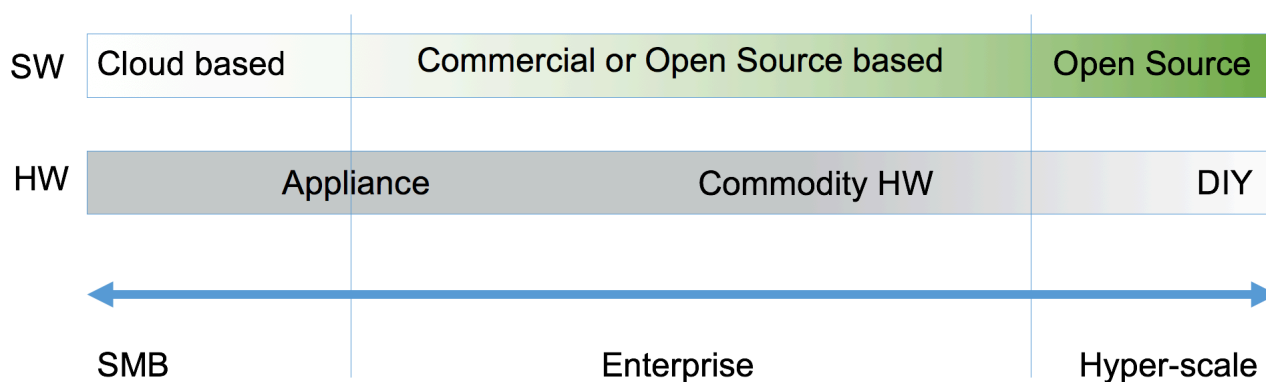
- There is a proliferation of new startups working on storage platforms with an object storage backend, some of them don't even expose RESTful APIs at the moment, but they all have similar basic characteristics like a scale-out design and a software-defined approach.
- The number of solutions supporting 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).
- Most successful vendors in capacity-driven systems are those who are capable of exposing multiple file/object protocols at the same time.
- Now, all primary storage vendors have an object storage solution in their product line up, or an agreement to resell a third party solution.

### CALCULATING THE MARKET SIZE

Traditional storage vendors are seeing an overall flat or negative growth<sup>2</sup> with a steady decline in revenues for legacy storage systems. The latest reports<sup>3</sup> on EMC, for example, have shown that VMAX, VNX and Data Domain product lines have all experienced declines between 3 to 15 per cent in the last four quarters. In the same period of time, software-defined and scale-out solutions from the same vendors have grown rapidly<sup>4</sup>. This is in line with what is usually seen in the field, with end users moving away from expensive and rigid storage solutions towards hyper-converged and All-Flash systems for primary data.

All analyst forecasts made between the last 6 to 12 months anticipate about an overall substantial increase in scale-out File- and Object-based storage solutions in the next few years. For 2015, IDC<sup>5</sup> expects a total revenue for scale-out systems in the range of \$B21.6 (the equivalent of 82.9 exabytes of capacity shipped). Numbers that are expected to grow up to \$B35.6 and EB249.1 by 2018, and showing a CAGR between 2013 and 2018 of 21.5% and 47.9% respectively! The lion's share from now until 2018 is being taken by

### Software-defined storage adoption model (per organization size)



object-based solutions (which also includes SaaS and Cloud NAS). Considering the same time frame (2013-

<sup>2</sup> Source: The Register, [http://www.theregister.co.uk/2015/10/13/dell\\_emc\\_prelim\\_results\\_q3\\_fy2015/](http://www.theregister.co.uk/2015/10/13/dell_emc_prelim_results_q3_fy2015/)

<sup>3</sup> Source: The Register, [http://www.channelregister.co.uk/2015/10/15/emc\\_trad\\_storage\\_decline\\_accelerating/](http://www.channelregister.co.uk/2015/10/15/emc_trad_storage_decline_accelerating/)

<sup>4</sup> Source: Stifel EMC financial analysis (EMC 3Q15 review), <https://goo.gl/NFisU8>

<sup>5</sup> Source: IDC, Worldwide File- and Object-Based Storage 2014–2018 Forecast (#251626)

2018), file-based scale-up solutions and single storage servers will be seeing a negative revenue growth, while the total capacity shipped will show a less impressive CAGR of 24%. On a similar wavelength a note, recently published by IDC<sup>6</sup>, that is showing an important increase in enterprise storage spending for server-based and hyper-scale storage infrastructures in the first quarter of 2015.

	Revenues 2015	Revenues 2018	CAGR % 2013-2018
<b>Scale-out storage (includes scale-out NAS + object)</b>	\$B 21.6 (82.9 EB)	\$B 35.6 (249.1 EB)	21.5% (47.9%)
<b>Scale-up storage (includes single servers + appliances)</b>	\$B 9.5 (14.6 EB)	\$B 8.2 (24.5 EB)	-4.5% (24.0%)
<b>Storage Services</b>	\$B 37.5	\$B 45.3	6.2%

These trends confirm, once again, that traditional storage solutions are less attractive than in the past to end users, who are now looking for better TCO and scalability.

The numbers above do not include storage services. But they are forecasted for a CAGR of 6.2% from 2013 and 2018<sup>7</sup>. By unpacking this number, it is interesting to note that expenditure for integration services is in line with Management, Support and Consulting registered in the past years. This is important because it proves that moving away from traditional storage systems to software-defined infrastructures will allow to manage much more storage capacity while maintaining integration and management costs under control.

Last but not least comes \$/GB ratio. In fact, due to the strong consolidation that is possible with scale-out systems, \$/GB will continue to decrease and remain in a range of one fourth to one third of the \$/GB forecasted for scale-up systems.

<sup>6</sup> Source: IDC, press release Jun/5/2015 <http://goo.gl/BQA23c>

<sup>7</sup> Source: IDC, Worldwide Storage Services 2014–2018 Forecast (#252135)

## SIZING THE OPPORTUNITY

It's clear that data storage, as well as all IT infrastructure, is experiencing great change. Agility requested by modern business needs has opened the door to an all out rethinking of IT organizations and infrastructure design.

The two major paradigm shifts that we are seeing, strictly related to each other, are the adoption of Software-defined infrastructures and cloud computing. In fact, looking at the modern datacenter, the best way to deploy private and public cloud is through software-defined solutions, and this is particularly true for larger infrastructures.

According to Gartner<sup>8</sup>, between 2015 and 2018, enterprise (101-500 racks) and large (500+ rack) datacenters will grow more than smaller datacenters in terms of end user storage expense. This is due to various factors:

- 1) Smaller organizations will spend less for on-premises infrastructures by leveraging more public cloud computing in the form of IaaS or SaaS. In most cases non-primary storage needs will be moved to the cloud (for example, backup, sync & share)
- 2) Large size organizations will continue to invest and build their own private cloud infrastructures.
- 3) Hyper-scale xSPs will see the largest increase, since they will be providing resources to small end users. xSPs are the type of end user that will benefit the most from scale-out software-defined storage infrastructures because the higher consolidation ratio and better TCO as well.

DC size	2015	2016	2017	2018	CAGR 13-18
<b>0-25 racks</b>	3.191.106.206,40 €	3.207.874.795,57 €	3.264.993.666,30 €	3.357.782.504,89 €	1%
<b>26-100 racks</b>	3.702.211.676,98 €	3.745.988.540,06 €	3.877.197.572,46 €	4.030.505.219,96 €	2%
<b>101-500 racks</b>	13.349.022.531,30 €	13.919.981.020,79 €	14.598.791.013,07 €	15.329.848.739,71 €	4%
<b>501+ racks</b>	9.322.775.986,90 €	10.047.597.987,36 €	10.641.348.145,58 €	11.222.492.255,96 €	5%

From another research, published by IDC<sup>9</sup> about public and private cloud storage revenue, it's interesting to note that both public and private cloud storage are expected to experience a fast growth rate in the next few years, confirming once again the trend already described. Large datacenters (large enterprises and cloud providers) will absorb most of the storage spending and the sum of public and private cloud (considering only on-premises installations) will reach a revenue of \$B19,2 with a CAGR 15-18 of 8%. This means that most of the datacenter expenditure for storage between now and 2018 will be allocated to cloud infrastructures.

		2015	2016	2017	2018	CAGR 15-18
<b>Public Cloud</b>	SaaS	6,3	7,0	7,6	8,3	7%
	Other	4,4	5,0	5,5	6,2	9%

<sup>8</sup> Source: Gartner, Forecast: Datacenter Worldwide 2010-2018 2Q14

<sup>9</sup> Source: IDC, WW Storage for public and private cloud (Nov 2014)

<b>Private Cloud</b>	On-prem	3,6	4,2	4,7	5,1	9%
	Off-prem	4,9	5,8	6,7	7,7	12%
<b>Public (SaaS+Other) + Private On-prem</b>		14,2	16,2	17,7	19,6	8%

By 2018 the total of \$B19.2 reported in the table above will count for a capacity of 303.66 Exabytes. About 80-90% of that capacity (and around 50% of the total revenue) will most likely be stored in large scale-out software-defined infrastructures optimized for distributed performance, less latency sensitive, workloads. The rest of the data (between 10 and 20%) will be stored in All-Flash arrays and hyper-converged systems configured to serve local and high latency sensitive workloads.

## BOTTOM LINE

Object-based storage is the only option today to address large multi-petabyte environments with ease. FOBS is already a \$B21 market already and growing rapidly.

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

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.

Among startups, Scality has one of the few products on the market that has already proven its potential with massive installations in the order of hundreds of petabytes. The kind of installation that will become more and more common from now on. The software-defined approach and multi-protocol support make Scality's product a particularly interesting solution for designing storage infrastructures of any size, capable of serving either legacy or next generation applications.

Among the larger vendors, IBM picked up a strong solution with Cleversafe. Pure's new offering, still in beta, is an interesting addition to their all flash array product. Both of them further validate the need and demand for secondary storage, and how technology that was once considered futuristic is being adopted right now.

## 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](mailto:info@juku.it) - <http://jukuconsulting.com>).