

Object storage: the unified storage infrastructure for the cloud era

Prepared for DDN
April 2016

TABLE OF CONTENTS

TABLE OF CONTENTS	2
EXECUTIVE SUMMARY	3
NEXT GENERATION UNIFIED STORAGE INFRASTRUCTURES	4
THE UNIFIED STORAGE FOR THE MODERN ENTERPRISE	5
A SEAMLESS MIGRATION PATH	5
PERFORMANCE OF FILE ACCESS ON OBJECT STORAGE	6
WHY DDN WOS?	7
BOTTOM LINE	8
JUKU	9
WHY JUKU	9
AUTHOR	9

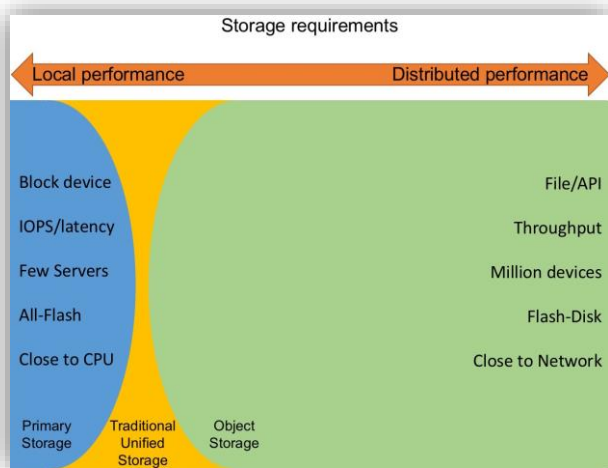
EXECUTIVE SUMMARY

Enterprise storage needs are quickly changing and traditional unified storage systems are no longer sufficient to cope with the demand generated by a multitude of different devices and locations. Newer data access patterns and cloud-based applications have gone through major changes as well, contributing to an important shift in the necessary requirements to implement modern storage infrastructures.

Newer data access patterns and cloud-based applications are contributing to an important shift in the necessary requirements to implement modern storage infrastructures.

Distributed organizations, with users accessing data freely from desktops and mobile devices, are the new standard. But unfortunately, traditional file access protocols (like SMB or NFS) were designed to work on stable and low latency local networks, and they are inadequate to serve data across the internet or WANs. On the other hand, we have modern primary storage systems which are aiming at latency-sensitive workloads, leveraging flash memory and block level access to overcome any latency issue that could compromise performance consistency. This storage category is also very expensive and even less adequate for remote access.

Traditional unified storage systems, which were considered “good enough” to serve both structured and unstructured data in the past, are now less appreciated by end user because of their inability to cope with the extreme expansion in capacity requirements, workloads and access methods of different data types.



Due to scalability limitation and protocol constraints, the number of use cases for traditional unified storage systems is quickly shrinking, especially when it comes to provisioning storage for remote users. In fact, major concerns for end users come from infrastructure sustainability and security. For example, storage consolidation in large data centers accessed remotely from state-less gateways/caches dramatically simplifies management, monitoring, security as well as backup and DR operations while improving overall infrastructure efficiency. And it's not made possible with legacy architectures like the one usually adopted on traditional unified storage systems.

NEXT GENERATION UNIFIED STORAGE INFRASTRUCTURES

By leveraging an object storage back-end and a new set of protocols at the front-end, including HTTP-based protocols like Amazon S3 API, it is now possible to build next generation storage infrastructures with unmatched scalability and flexibility.

This kind of solution overcomes all the limits of traditional architectures while providing an overall better TCO and multi-protocol access to support more use cases, types of applications, and legacy solutions. It also provides the necessary characteristics for the development of several innovative services capable of sharing the same data, regardless of location or device. Actually, this could be considered a new class of storage system, where scalability and performance are as important as multi-protocol access.

DDN WOS®, a robust object storage, available in both software and appliance models, can be leveraged to implement next generation storage infrastructures where S3, Swift, NFS and SMB protocols are exposed and can be used for a comprehensive set of use cases. Different configurations can serve installations of any size, starting at a few hundred Terabytes to several Petabytes, and any level of performance, thanks to the broad lineup including software-based-storage and hardware appliances. Thanks to the protocols supported it enables different applications like Sync and Share, cloud collaboration as well as local and remote NAS access.

THE UNIFIED STORAGE FOR THE MODERN ENTERPRISE

Contrary to traditional unified storage systems, which expose file and block protocols, usually on dated file-based architectures, Object-based scale-out solutions offer much more coherency and flexibility with a set of protocols focused on unstructured data workloads and cloud applications.

Object-based scale-out solutions offer much more coherency and flexibility with a set of protocols focused on unstructured data workloads and cloud applications.

Modern object storage platforms, like DDN WOS®, allow to access same data volumes through different protocols concurrently. This is why you can consider them as unified storage. Depending on each single implementation, RESTful APIs (S3 or Swift, for example), are usually presented alongside proprietary API sets and NFS and SMB are available as well. S3 enables cloud-based applications with SMB and NFS allowing to manage legacy workloads and traditional clients. This capability is a sizeable enabler and freedom of choice is a key factor for the enterprise and the ISP/CSP.

A SEAMLESS MIGRATION PATH

Adding file access protocols to object storage broadens the number of use cases and helps to speed up adoption in legacy environments:

- Multiple large File-based archive repositories, can be easily migrated and consolidated to an object store - taking advantage of a considerably better \$/GB and vastly improving data protection and disaster recovery capabilities;
- Most advanced object-stores implement sophisticated automation and policy engines, moving the needle from hundreds of Terabyte to several Petabytes managed per SysAdmin;
- The ability to access the same data object simultaneously from a file system and APIs lays the foundation for a seamless migration path from legacy to cloud-native applications. For example, data can be ingested through a file interface and globally accessed via APIs.

API based access is much more common than in the past, however. The list of S3 and Swift compatible solutions is getting longer by the day, and native applications and gateways get the best out of object storage:

- Private Sync and Share is the most common example and is in very high demand among enterprises. It's not only about the improved TCO, when compared to public cloud, but also because it's much easier to enforce specific security policies in order to have a complete control over data and how it is accessed.

- Remote stateless NAS gateways are similarly successful because they provide the same user experience of a local file server but with the agility, level of protection, easy scalability, and economics of a cloud back-end.
- Several backup software solutions now support S3 repositories to store backup archives. In fact, object stores can easily have a tremendous throughput and the best durability while their distributed architecture makes back-end procedures (like e-vaulting) transparent and even more secure with a low \$/GB. And if the object store supports multi-tiering automation capabilities and tape, \$/GB will fall further.



All these features and use cases can be combined onto a single object-based unified storage infrastructure allowing a massive consolidation of all secondary and non latency-sensitive data and workloads.

PERFORMANCE OF FILE ACCESS ON OBJECT STORAGE

Most object storage systems are not designed with performance as their primary goal, especially when data is accessed through file interfaces, but this is rapidly changing and there are several solutions available to mitigate performance issues:

- DDN, as well as some other vendors, is now providing powerful ingest (i.e. GRIDScaler, EXAScaler) and caching solutions (i.e. SFX, IME) which can leverage high performance file access and Flash memory on the front end to speed up operations;
- Flash can also be part of the object store which is especially useful when objects are particularly small;
- Most flexible products, like DDN WOS®, can be configured to manage different kinds of workloads efficiently, giving the end user plenty of choice without having to trade too much performance for capacity.

Contrary to the past, when object storage was usually considered a very durable but slow storage system, good for cold data only, today performance is fundamental at every scale. In the specific case of DDN, a leader in the HPC market, performance of WOS-based infrastructures has already been proven in the field by the end users with different applications and protocols.

WHY DDN WOS?

DDN WOS was initially designed to meet requirements of web-scale applications, addressing performance and capacity at scale and, in the latest versions, a lot of new features have been added to meet traditional enterprise needs as well. The result is a next generation object storage based solution that is well integrated with the rest of the DDN product lineup, which thanks to native and third party solutions, can be the foundation of a modern enterprise unified storage infrastructure.

Its clever back-end design is based on some key features like, for example, the absence of an underlying file system to store objects (NoFS) that boosts space efficiency up to 99% while improving latency at the same time.

The number of supported protocols include S3, Swift, WOSREST as well as CIFS, SMB (2.1, 2.2) and NFS (v3, v4) file protocols. An optional bridge allows to integrate external tape libraries and DDN GRIDScaler or EXAScaler appliances through automated tiering policies to cover an even larger number of workloads and use cases.

The overall efficiency of a WOS-based infrastructure is granted by its flexibility:

- Its layout is organized in relatively small clusters, so that any change (or fail) to a single node will not impact the performance of the whole infrastructure.
- Erasure coding can be configured globally or locally and various forms of replication are possible, allowing the user to choose the best option based on performance and capacity optimization.
- Flash memory can be used to lower latency and accelerate access to small objects.

The customer can choose between hardware and software appliances for any kind of installation. In fact, the software solution can be the right choice for the user who wants to standardize its datacenter on a specific hardware vendor or type (like HP or Dell for example), as well as for those smaller organizations that need to test out Object Storage on their virtualized environment. On the other hand, DDN offers a complete line of fully tested and supported hardware appliances for any capacity or workload.



BOTTOM LINE

A storage infrastructure with the characteristics described in these pages can easily manage up to 80% of modern enterprise storage needs.

It offers a remarkably high consolidation ratio and a much lower TCO than traditional storage systems. At the same time, thanks to native protocols and additional gateways/applications, it can serve both legacy (local NFS/SMB) and private cloud storage applications (S3- or Swift-based applications, like Sync and Share and remote office NAS for example).

New use cases for object storage are complementary to traditional applications and can now be served by the same infrastructure offering a remarkably high consolidation ratio.

These new use cases are complementary to traditional object storage applications and can be served by the same infrastructure, including backup targets, active archives, content distribution, cloud collaboration and any other form of cloud storage concurrently. Contrary to what happens with traditional storage systems, with object storage, thanks to its multi-tenancy capabilities, the rule is always: the bigger, the better.

DDN WOS has already proven its capability, in the field at scale, both in terms of performance and capacity with cloud applications, and now it can also be proposed as a solid solution to building enterprise unified storage infrastructures for large enterprises and ISPs thanks to:

- Data access flexibility, with APIs and files;
- Performance and scalability, thanks to its modern design;
- A unique architecture;
- Hardware and software appliances;
- Integration with the rest of DDN products to cover end-to-end complex scenarios, including IoT and Big Data Analytics.

In a two tier scenario that includes all-flash storage for primary workloads characterized by low latency and a high number of IOPS, an object storage based infrastructure could work out very well and manage capacity and less latency sensitive data-driven 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 decision makers in their IT environments, to inform and to discuss the technological side that we know better: IT infrastructure virtualization, cloud computing and storage.

Unlike the past, today those who live in the IT environment need to be aware of their surroundings: things are changing rapidly and there is a need to be constantly updated, to learn to adapt quickly and to support important decisions - but how? Through our support, our ideas, the result of our daily global interaction 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 in providing an honest and productive vision. That's why we have chosen: "think global, act local" as a payoff for Juku.

AUTHOR



Enrico Signoretti is an analyst, trusted advisor and passionate blogger (not necessarily in that order). He has been immersed in 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 until now when he joined the "Cloudland". During these years his job has changed from highly technical roles to management and customer relationship management. In 2012 he founded Juku consulting SRL, a new consultancy and advisory firm deeply focused on supporting end users, vendors and third parties in the development of their IT infrastructure strategies. He keeps a vigil eye on how the market evolves and 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 DDN. This paper was commissioned by DDN. No employees at the firm hold any equity positions with DDN. Should you have any questions, please contact Juku consulting srl (info@juku.it - <http://juku.it>).