DDN®

WHITEPAPER

WOS®
COMPREHENSIVE OBJECT STORAGE

Collaborate • Distribute • Archive

EXECUTIVE SUMMARY

The industry has come to understand that an alternative storage methodology is required to efficiently and securely store the exabytes of unstructured information we generate every day. Object storage is destined to be this new storage paradigm, featuring unlimited scalability, unbeatable reliability and the lowest cost per terabyte. Over the past decade, several generations and flavors of object storage have been introduced, each one of them designed to solve a particular unstructured data challenge. Most of these platforms were designed for archive applications, as they support large volumes of data, with relatively low performance requirements.

But, data archives are just the tip of the unstructured data iceberg. To solve the unstructured data challenge, we need flexible object storage platforms that support the entire data life cycle and integrate with a wider range of applications: genomics research, worldwide collaboration, file sync & share, social media applications, content distribution and many more.

In essence, there are 5 key requirements customers will need in order to define the architecture of their object storage infrastructures: scalability, accessibility, efficiency, reliability and performance. While the requirements can affect each other (for example, increasing performance may require more copies of data to be stored), organizations should always be able to tune their architecture to meet all requirements, without having to compromise on others.

WOS is the only platform that offers full flexibility in data protection schemes and enables performance optimization to comply with all data or application requirements. In the title, "Compreshensive Object Storage" refers to the versatility of the platform, which offers a complete set of protocols and which is available as an integrated plug 'n play system, as an archive appliance and on OCP-certified Hyve platforms.

WOS was designed as a single storage solution for all your unstructured data needs. In this paper, we will investigate what the requirements are to deploy object storage on a bigger scale and illustrate how WOS enables customers to build storage infrastructures that can be tuned to meet the requirements of any data set, any application and any storage architecture.

TABLE OF CONTENTS

THE CHALLENGE OF UNSTRUCTURED DATA

APPLICATIONS ARE OPENING **A WHOLE NEW RANGE OF OPPORTUNITIES,** BUT ARE ALSO CREATING **BIG CHALLENGES.**

In the 1980s and '90s, the most valuable digital data was transactional data, such as database records that are created and accessed through database applications. This led to the success of large database companies, rich database applications and data mining tools. The relatively slow capacity growth of structured data never presented significant challenges to database platforms or the underlying, typically block-based (SAN), storage infrastructures. Life was simple.

For unstructured data (data not stored in a database), things have been quite different. With the advent of MS Office® in the 1990s, unstructured data would become much more important than it had ever been before. Halfway through '90s, every office worker had a desktop computer with a host of applications. E-mail allowed us to send slides, spreadsheets and text documents to anyone we wanted to share information with and storage consumption went through the roof. Enterprises would soon be challenged to build shared file storage infrastructures and backup and archiving became another challenge. Tiered storage was born. Data storage became both hot and cool. Over the next 2 decades, we would see plenty of innovations to manage fast-growing unstructured data sets and the file storage (NAS) industry skyrocketed.

While enterprise databases are still the central nervous system for many large corporate organizations, unstructured information has also become the lifeblood for them and the rest of the world. Think of audio and video files in the media and entertainment industry, genomics data in research centers, satellite and drone footage for defense - or the exabytes of data we are all generating, uploading and viewing through scale-out web applications, mobile devices, etc. These applications are opening a whole new range of opportunities, but are also creating big challenges, as traditional file storage was not designed to scale to these levels of capacity growth, quantity of users and performance demands.

The industry has come to understand that an alternative storage methodology is required to efficiently and securely store the exabytes of unstructured information we generate every day. Object storage is destined to be this new storage paradigm, featuring unlimited scalability, unbeatable reliability and the lowest cost per terabyte. Over the past decade, several generations and flavors of object storage have been introduced, each one of them designed to solve a particular unstructured data challenge. Most of these platforms aim at archive applications as those are designed for large volumes of data, with relatively low performance requirements.

But data archives are just the tip of the unstructured data iceberg. To solve the unstructured data challenge we need flexible object storage platforms that support the entire data life cycle and integrate with a wider range of applications: genomics research, worldwide collaboration, file sync & share, social media applications and many more. In this paper, we will investigate what the requirements are to deploy object storage on a bigger scale and illustrate how WOS enables customers to build storage infrastructures that can be tuned to meet the requirements of any data set, any application and any storage architecture.

THE PROMISE OF OBJECT STORAGE

The promise of object storage is simple: to enable customers to build a highly reliable, infinitely scalable and most efficient storage pool for all their unstructured data needs. For this to be true, object storage platforms need to meet all of the object storage requirements and provide a set of tunable parameters – yet few platforms available today live up to that promise.

Object storage platforms need to scale granularly, but efficiently, and provide full flexibility with respect to data protection and performance. Then, they can be used for any mix of large and small objects and be tuned for high IOPS, high throughput or both. Such

5 KEY REQUIREMENTS FOR SCALE-OUT STORAGE: SCALABILITY, ACCESSIBILITY, EFFICIENCY, RELIABILITY AND PERFORMANCE.

platforms need to be hardware-independent, support single or multi-data center architectures and meet the most stringent TCO objectives, without heavy trade-offs.

In essence, there are 5 key requirements customers will need in order to define the architecture of their scale-out storage infrastructures: scalability, accessibility, efficiency, reliability and performance. Each of these requirements has a number of sub-requirements, called requirement dimensions. Both the requirements, and their dimensions, can affect each other. For example, increasing performance may require more copies of data to be stored – but organizations should always be able to tune their architecture to meet their key requirements, without having to compromise on others.

Object storage was purposefully designed for very large volumes of unstructured data, with unlimited scalability as the ultimate objective. Inherently, storage platforms can be scaled in three dimensions: the total volume of storage, the number of objects and the number of sites. The number of objects tends to be a particular challenge when objects are very small or the ratio of small vs. big objects cannot be accurately predicted. Applications may not have high scalability requirements for all three dimensions initially, but over time those requirements can change due to external elements, so it's important to build in all three scalability dimensions from the start. For example, a new web application may run just fine over two sites when it is launched locally in the US, but when the service expands overseas, the built-in capability to seamlessly add more sites will save a lot of time, resources and money.

One of the key benefits of true object storage is the absence of file systems. This also creates a challenge: the accessibility of the data. Typically, object storage is accessed through applications, which use application protocol interfaces (API's) to interact directly with the backend storage pool. Several attempts have been made to have the industry agree on a standard object interface, with mixed success. The Amazon® S3 and Openstack® Swift API's are currently seeing the widest adoption, but it still remains to be seen if one of them will ever become a true "standard", and an open API. Therefore, it is important for object storage platforms to provide wide support for multiple protocols and applications, including file system gateways to integrate with legacy applications.

The most difficult requirement of all, and coincidentally the only one the CFO (who signs off the storage budget) really cares about is total cost of ownership (TCO). As a result, it is almost impossible to find an object storage platform that does not promise to be the most cost efficient solution on the market. And many platforms "will" be the most cost efficient, at least in one particular configuration, for a very specific use case and for a very limited range of parameters. What is more difficult is to find the platform that offers the overall best efficiency, including infrastructure cost (raw storage, data center, etc.), management cost (how many people are required to manage the infrastructure) and bandwidth consumption. The latter is especially ignored by the distributed erasure coding1 solution providers, who enable their customers to build relatively reliable (durability only – read further) infrastructures with low storage overhead, by distributing data chunks over multiple data centers. The problem with those solutions is that they heavily rely on the WAN and generate a lot of bandwidth traffic for data rebuilds.

Central in any storage architecture is reliability. Whether you are building a low-cost archive or a high-performance storage cloud, reliability is key. But there is a lot more to data reliability than most storage providers like to admit. Data reliability dimensions are: availability (the time your data is instantly available for access); durability (the guarantee that your data will not be corrupted); and integrity (the assurance that your data will remain unchanged and cannot be tampered with). Security is a feature inherent to data integrity. Only when reliability is provided in all three dimensions, will you achieve true data reliability. Most platforms provide acceptable reliability grades on one, sometimes two of the above.

Storage performance is measured in throughput, IOPS and latency. Throughput is the total volume of data that is transferred within a time frame. For object storage, this is typically measured in MB/s or GB/s.

IOPS2 refers to input/output operations per second. Historically, IOPS-intensive environments (think databases) would be deployed on block storage, so IOPS referred to the number of blocks that could be written to and read from a disk drive. Most of today's object storage platforms are not optimized for small objects, so IOPS tends to be neglected in performance conversations. But, when you are storing massive volumes of small objects, IOPS needs to be added to the performance equation.

Latency is the delay between the client system initiating a request (to read or write data) and the execution of that request. A common way to think about this is for a disk write is the time it takes to confirm that the last byte of an object has been written - and in the case of a read, the time for the first byte to be received by the requester. The key latency influencers in object storage systems are both disk and network latency. In the case of the disk, latency is dominated by seeking the head to the correct location and the characteristics of the network the latency of which is often dominated by the WAN. Most object storage platforms do not allow for latency optimization, which is the reason they do not mention latency in their performance conversations.

WOS, THE MOST COMPREHENSIVE OBJECT STORAGE PLATFORM

As data repositories started to grow in the 1990s, tiered storage architectures became common practice to reduce the total cost of storage. Tiers are typically based on the requirements as defined in the previous chapter: reliability, performance, scalability, accessibility and efficiency. But, even with all the available automated software tools, tiered storage infrastructures can be very complex to manage and the cost savings realized are often marginal.

Also, for an increasing number of environments, tiering is just not practical. Think, for example of scale out web applications where users can store images, video files and other data. It is very difficult to predict what type of data they will store most, what the large to small files ratio will be and the number of views from respective geographic locations, etc. Web scale applications require scale-out storage repositories.

WOS was designed as a single storage solution for all your unstructured data needs. We have carefully analyzed the storage requirements and built a platform with unprecedented tuning options. This enables customers to give priority to the requirements that are most important for their data - or to their customers, in all dimensions.

ARCHITECTURE: CHOOSE ANY NUMBER OF SITES

PLATFORM: CHOOSE INTEGRATED DDN WOS HIGH-PERFORMANCE OR ARCHIVE APPLIANCE OR OCP INSPIRED HARDWARE

DATA PROTECTION: CHOOSE FROM A WIDE RANGE OF PROTECTION SCHEMES

INTEGRATION: CHOOSE NATIVELY INTEGRATED APPLICATIONS, DIRECT INTEGRATION THROUGH NATIVE OBJECT API'S OR FILE GATEWAYS

WOS is the only platform that offers full flexibility in data protection schemes and enables performance optimization to comply with data, application or SLA requirements. Comprehensive Object Storage refers to the versatility of the WOS platform, which offers a complete set of protocols and is available as an integrated plug 'n play DDN WOS7000 appliance, a cost-efficient WOS Archive appliance or integrated on third-party OCP-inspired appliances.

WOS BALANCES PERFORMANCE AND PLATTER UTILIZATION WHILE MAINTAINING LOW SYSTEM OVERHEAD.

WOS is enabled through the underlying NoFS architecture and the unique use of WOS Object Identifiers. Most other object storage platforms still have a POSIX file system layer on the disk level. WOS, however, was designed as a true object storage platform: a flat, single layer, address structure where objects are stored in a manner in which disk operations are minimized. This results in as little as a single-disk-operation in the case of reads, and as few as two operations in the case of writes. (Contrast that with the 8-10 operations required in a POSIX file system, which adds additional performance and network overhead, slowing down access and adding cost). The on-disk layout used in WOS carefully balances performance, platter utilization3 , all while keeping overhead for the system very low.

Through a set of parameters, customers can tune their platform to the requirements that are important for their users and data:

- WOS offers full flexibility to storage architects. WOS can be configured as a single-site, dual site setup, or deployed across a number of sites. For collaboration and data distribution environments, in comparison to other systems, a higher number of remote sites can actually reduce (rather than increase) the latency to access data in WOS.
- WOS is available preinstalled on the high-density, high-performance WOS7000 storage appliance, as a WOS Archive appliance and on OCP-ready appliance hardware (by Hyve) for scale-out web applications.
- A unique choice of data protection schemes allows customers to optimize toward higher data protection, more I/O performance or latency optimization for distributed environments. For archival use cases, through the new Global Object Assure4, WOS can act as a very low cost active archive.
- WOS provides the most complete choice of interfaces, including a set of native (REST) API's, file access interfaces. Additionally, WOS can be provided with pre-integrated applications such as iRODS® for data management, ASG® Atempo for archival purposes or Ctera® for secure global file sharing.

WOS NO FILE SYSTEM, PURE OBJECT STORAGE

WOS is the only object storage platform that was specifically designed to solve the unstructured data challenges and which leverages the immutable nature of this data. The NoFS architecture enables the WOS "Comprehensive Object Storage" paradigm, (which is truly unique, as opposed to other platforms that are typically optimized for a limited number of requirements and use cases).

Why NoFS? File systems are designed to enable multiple users to collaborate on shared data sets without corrupting the data. But the complex technology that enables things like file locking, make file-based storage platforms expensive at scale. As over 80% of unstructured data that is stored will never be changed again, DDN recognized it is more optimal to build a pure object storage platform that allows you to make changes by saving a new object.

Traditional file systems require at least three layers of software constructs to execute a file operation. To allow file amendment by multiple users, they must maintain complex lock structures with OPEN and CLOSE semantics. These lock structures must be distributed coherently to all of the servers used for access. This complexity increases exponentially at scale.

On the disk level, data is stored in available blocks that are gathered and apportioned based simply on availability at the time

WOS OFFERS A **FLEXIBLE DATA PROTECTION MECHANISM** TO MEET **PERFORMANCE, RELIABILITY** AND **EFFICIENCY REQUIREMENTS.**

the data is written. As data is placed based on random block availability, this creates fragmentation- especially in environments with large volumes of unstructured data, different data sizes and various data types. As a result, using a file system designed for amendable data to store immutable data constitutes an inappropriate and wasteful use of bandwidth and compute resources. This highly inefficient practice results in the requirement for a great deal of additional hardware and network resources to achieve data distribution goals.

The WOS NoFS architecture solves the file system challenges and provides a number of features that enable the WOS benefits:

- WOS is a flat, single layer, address structure where objects are stored in an on-disk format, such that disk operations are minimized and performance is maximized.
- WOS groups objects of the same or similar size in addressable containers. This approach avoids unconstrained fragmentation of the data on disk.
- WOS generates object names in the form of OIDs (Object Identifiers). This means the data structures required to retrieve an object do not require accessing inefficient file system metadata.
- As an object is stored in WOS, a checksum is calculated, stored and checked several times along the data path. In this way, WOS can provide end-to-end data integrity.
- Each object, as it is stored in WOS, is associated with a "policy". The Policy defines how WOS will protect that object. For example, you can define the number of replicas and how that object's data is protected within a WOS Node. The policy also defines the locality of the data. This allows WOS to control where the data physically resides (in a distributed environment) for both performance and compliance reasons.
- Unlike the many complex commands of file systems, the WOS NoFS semantics have been simplified to PUT, GET, and DELETE.

The end result is that a WOS storage cluster is an automated, high performance data distribution and protection mechanism that functions without continuous external management - and can be used to establish an inherently unbreakable, distributed and balanced cloud service.

WOS DATA PROTECTION

Different data types and sizes require different protection schemes, as do different applications. For example, web applications with users around the globe will require high throughput and/or high IOPS, but also low latency. Active archives, on the other hand, may just require decent throughput, as latency is measured in seconds, not milliseconds. WOS offers a unique choice of data protection mechanisms, that enable customers to optimize their infrastructure to meet the performance, reliability and efficiency requirements of their applications. Data protection mechanisms are policy based and can be defined as granularly as on a per-object basis5.

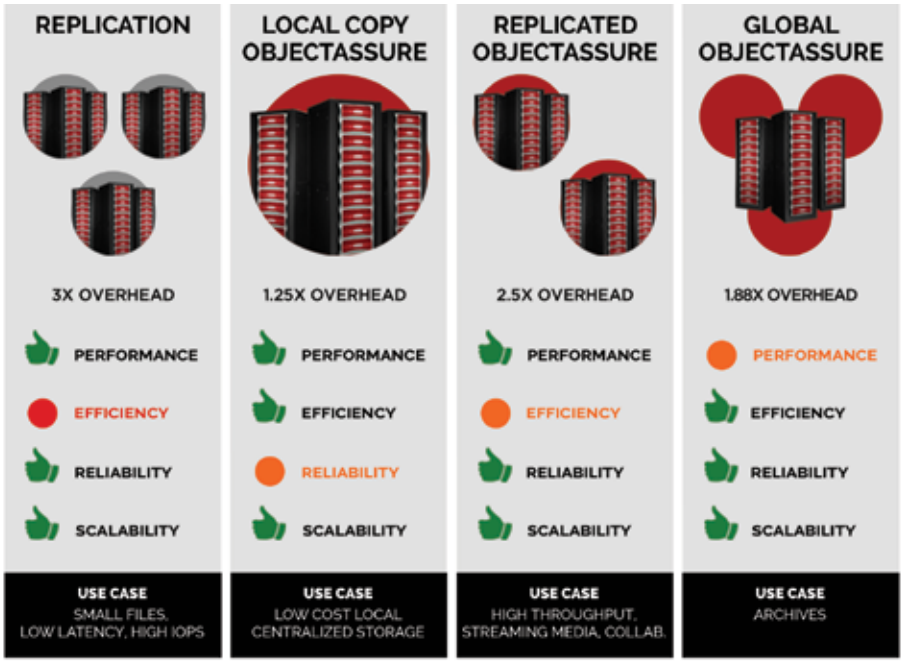

WOS TAKES A MORE **INNOVATIVE APPROACH** TOWARDS **REPLICATION** THAT RESULTS IN **HIGHER, MORE RELIABLE THROUGHPUT**.

REPLICATION

Replication is the most straightforward form of data protection – the storing of multiple copies6 of the objects on different disks or systems, so data can be recovered from a spare copy in case one is lost. Most object storage platforms make very simple one-to-one copies of objects. WOS takes a more innovative approach, resulting in higher data throughput, that enables faster rebuilds and higher reliability.

Traditional replication techniques store each copy of the data to a separate disk. This approach slows the write throughput down to what one disk can ingest. When a disk becomes inoperable, you rely on just one system to rebuild the data. WOS is fundamentally a scale out system: even though you might be storing what the client sees as a single object, any object larger than 1MB is split into smaller objects (1MB each), which are then written to different disks. This way, the system can engage multiple disks and nodes to store a single large object: a 100MB object could be stored to 100 disks in parallel. WOS Replication provides two major benefits:

- As multiple pieces of hardware are processing your transaction, WOS provides higher throughput when storing data.
- Because large objects are stored as fragments across multiple disks and nodes, rebuilds

ObjectAssure PROVIDES **HIGHER DURABILITY** AND **AVAILABILITY** AT A **MUCH LOWER COST.**

can be done fully in parallel: the recovery of a disk or node is done in a many-to-many fashion (versus one-to-one rebuilds).

Replication is a popular choice for data sets with small objects (as erasure coding often creates too much storage and processing overhead), and for distributed environments, like CDNs (where low latency is more important than low overhead). WOS provides both synchronous and asynchronous replication, offering the customer the flexibility to optimize toward performance and reliability.

SYNCHRONOUS REPLICATION

The object is sent to multiple storage targets (nodes) synchronously; the client application receives a write confirmation when all required copies have been stored. The benefit of synchronous replication is that the object is immediately protected according to the storage policy, including site failure protection.

ASYNCHRONOUS REPLICATION

The object is sent to only one storage target, but a temporary spare copy is made on a different disk. This temporary copy is deleted when another copy has been made on a different target. The benefit of asynchronous replication is performance, as the data will be stored with the lowest latency. The temporary copy ensures the object is protected against disk failure during the time that is needed to replicate the data to another target.

ObjectAssure

ObjectAssure® is DDN's implementation of erasure coding. WOS is the only platform that provides local, replicated and globally distributed erasure coding. This provides customers with a lot more flexibility than traditional dispersed erasure encoding techniques, where customers have very limited choice in number of sites or data protection levels.

WOS ObjectAssure requires less storage capacity than replication but still provides very high data reliability. The biggest benefit of DDN's erasure coding vs. other technologies is that ObjectAssure does not abuse the WAN, as rebuilds are done locally without data traffic over the WAN. This provides higher durability and availability, at a much lower cost.

LOCAL COPY ObjectAssure

Reduce your storage overhead while maximizing durability for single-site deployments with local ObjectAssure. This approach maximizes efficiency, while still providing the highest possible local data durability. Local ObjectAssure is often used as an alternative for traditional RAID for cost-efficient hot archives. Often, an additional copy is stored on tape.

REPLICATED ObjectAssure

Protect your data against site failure with Replicated ObjectAssure. To improve availability without affecting the WAN cost, a copy of the data is made to a second site. Rebuilds due to disk or node failure are done without generating WAN bandwidth traffic. Replicated ObjectAssure is typically selected for larger objects in a multi-site setup, e.g. for content delivery networks or scale-out web applications.

GLOBAL ObjectAssure

Global ObjectAssure is similar to distributed erasure coding, providing the lowest overhead for multi-site deployments, but without the bandwidth dependency. Global ObjectAssure is basically a double implementation of ObjectAssure with local and global (cross-zone) data protection. In case of failures, data rebuilds will always first happen locally, which provides much faster rebuilds and thus higher durability.

Global ObjectAssure is typically used in environments that consist of more than two sites. The benefit is that your data is protected against site failure (+additional nodes) at a much lower overhead than Replicated ObjectAssure. Global ObjectAssure enables customers to build very low-cost infrastructures for very large data sets. With less than 1.9 storage expansion, you can still obtain the highest availability. This scheme is typically used for archival scenarios where low overhead and high reliability are more important than low latency.

USE CASES

WOS CAN EASILY BE **INTEGRATED** FOR **CUSTOM APPLICATIONS**

WOS enables customers to build the perfect storage platform for any data type or size. The platform can be tuned to support any application that needs to store and access large volumes of unstructured data. Policies can be set per object, so a single WOS infrastructure can be used to support several applications. For example, it is perfectly possible to use WOS for a combined sync and share, DR and archiving solution - or to build a scale-out web application that will be generating billions of small files, but which also needs to provide high throughput to stream high-resolution movies.

DDN's WOS is extended with a mix of in-house developed solutions, integrated applications from third party ISV's and reference architectures. With its rich choice of API's and file system gateways, WOS can easily be integrated for Custom Applications as well. WOS currently supports a wide variety of use cases with out-of–the-box integrated solutions:

SYNC AND SHARE

Automated Sync and Share applications enable users to securely upload documents to the cloud, synchronize files and devices, and easily share information with others. This is one of the more popular applications that utilize WOS, leveraging the latency-aware and data placement techniques that are unique to the platform. WOS Sync & Share comes as a pre-integrated partner solution, from companies like Ctera® and OwnCloud®.

VIDEO STREAMING

Enable Video on Demand, Cloud DVR or other video streaming services for residential or corporate users. WOS provides high-throughput, low latency video delivery streaming for geographically distributed viewers. WOS Video Streaming can be deployed as a custom solution (integrated with API's or file system gateways) or as a pre-integrated solution using the technology from partners like Arris®, a global innovator in cable, video and broadband technology.

CONTENT DELIVERY NETWORK

Leverage WOS to build your own Content Delivery Network for worldwide distribution of massive volumes of data with high throughput and low latency. The unique latency-aware technology in WOS, combined with the flexibility to optimize for small and large file performance, make WOS the perfect CDN storage origin. DDN has engaged with several partners to build CDN architectures scaling to as many as 60 origin storage sites.

WORLDWIDE COLLABORATION

Store assets in a globally distributed storage cloud to enable collaboration between distributed teams and integrate with your favorite workflow suites or file sync and share clients. WOS is the only platform that enables integration with (and federation of) parallel file systems. Leading research institutions and universities around the world leverage WOS to build global collaboration libraries, enabling more efficient workflows and quicker time to result/discovery.

VIDEO POST-PRODUCTION

WOS enables collaborative editing for distributed teams in video post-production environments. Integrate WOS with your favorite post-production workflow suites, directly integrated through REST or by leveraging the integrated CIFS and NFS file system gateways. WOS can be optimized to replace several tiers of the storage infrastructure with one, scalable and easy to manage storage pool. DDN has over 600 customers in the media and entertainment industry and a rich partner ecosystem for video post-production workflows.

ACTIVE ARCHIVES

Many providers have been promoting disk storage solutions as an alternative to tape to build "Active" Archives, but few are able to provide the cost-efficiency that is required to build petabyte-scale repositories. WOS enables customers to monetize their data and build highly reliable, scale-out archive infrastructures at the lowest TCO. WOS provides instant access to all archived assets and integrates with popular archival platforms such as CommVault®, ASG® Atempo and iRODS®. The WOS-GRIDScaler Bridge also enables archive solutions for parallel file systems. DDN is a member of the Active Archive Alliance.

ABOUT DDN®

DDN is the world leader in high-performance and massively scalable data storage. We empower Enterprises to extract maximum value from their information while accelerating bottom line results. By leveraging the best talent in the industry, DDN® storage and Professional Services help content rich and high-growth IT environments achieve peak levels of systems scalability, performance, efficiency and simplicity.

Across traditional and commercial high performance computing, our customers rely on DDN solutions to solve the most demanding big data problems in cloud, online content and social networking, security and intelligence, life sciences, finance, energy and media production.

Deployed in thousands of mission critical environments worldwide, DDN engineers proven solutions for the most scalable data centers and delivers you the ultimate competitive advantage needed to succeed in today's information-driven Enterprise.

DDN®

BOTTOM LINE RESULTS. ACCELERATED.

SALES@DDN.COM | +1.800.837.2298

v2 (8/16)