

KEY BENEFITS

- Data Migration without Application Impact
- Accelerate Performance with Existing Storage
- Seamless Cloud Integration
- vSphere Integration with VASA
- Flexible Storage Reporting

DataSphere is an enterprise metadata engine that delivers live data mobility without application interruption across enterprise infrastructure and the cloud. With DataSphere, enterprises can automate data migration, easily adopt cloud storage, scale performance in parallel, and automate VM storage management to maximize efficiency in virtualized environments.

By separating how an application logically views data, DataSphere allows the data to be freed from where it is physically stored. Freeing the data from its physical location maximizes the value of your existing infrastructure and helps enterprises easily integrate with the cloud. Metadata analytics provide insight on data activity at the file level, giving IT the visibility needed to take full control over the management of their storage resources.

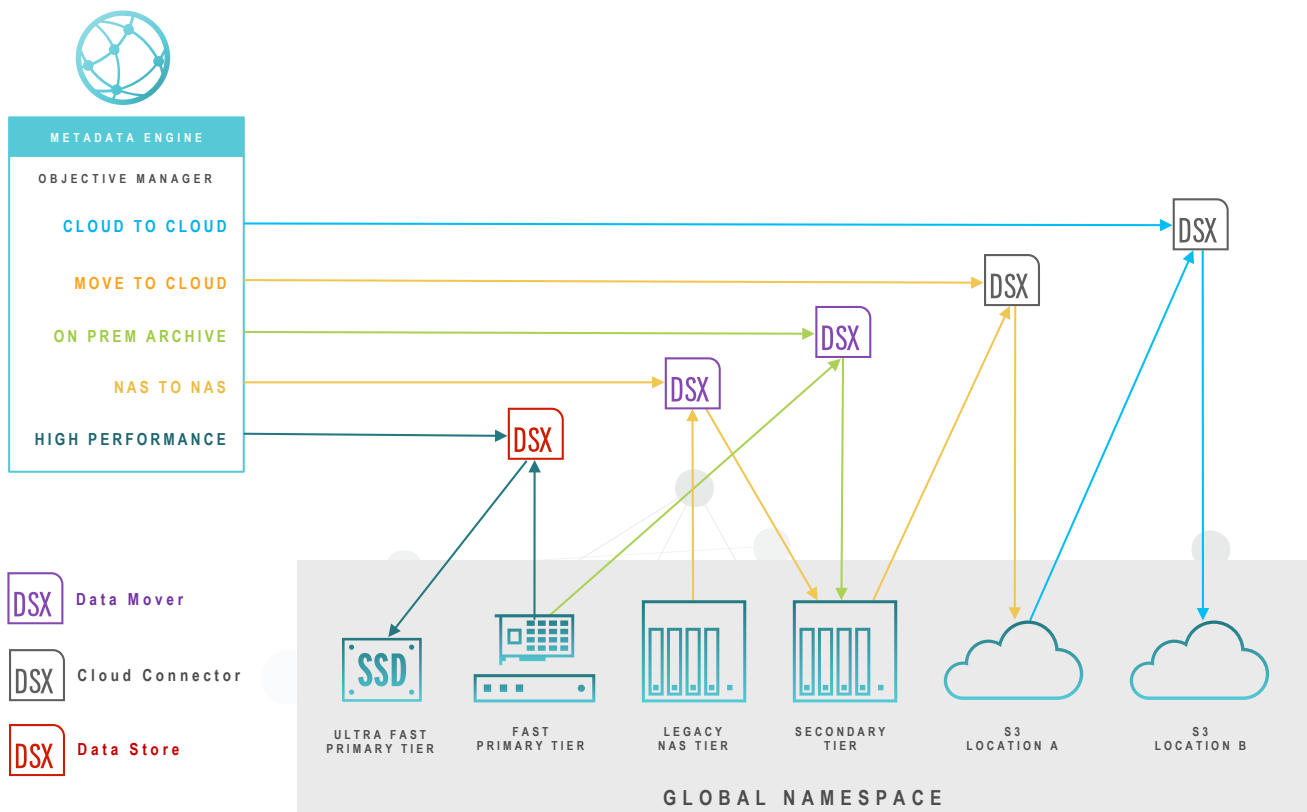


Figure 1. DataSphere makes real-time automated decisions for data placement, moving data without disruption to overcome or prevent outages, and to maintain alignment with service level agreements or objectives.

The storage-and vendor-agnostic DataSphere architecture unites different types of storage into a global namespace and automatically places data on the most appropriate storage resource to meet business and IT objectives across performance, protection and price. This helps enterprises overcome performance bottlenecks, integrate with the cloud for savings and active archival, and easily adopt new resources from any vendor to achieve unprecedented improvements in performance, efficiency and scalability.

HOW TO TRANSFORM YOUR ENTERPRISE WITH DATASHERE

KEY BENEFITS

- **Seamlessly add** single or multiple cloud storage tiers
- **Move data to the cloud and back** without disrupting application access
- **Save costs** with automatic data deduplication and compression
- **Free existing storage capacity** by moving cold data to the cloud

Easily Adopt a Low Cost, Limitless Storage Cloud Tier

Over the last few years, the broad adoption of cloud technology has dramatically disrupted the way that enterprises store their data. The vast potential to scale resources on-demand increases operational agility, changing how enterprises tier their storage.

DataSphere helps enterprises reduce the total cost of ownership (TCO) for managing and storing data by providing the ability to define and account for data movement based on price-to-performance targets. With DataSphere, you can pair local NFS storage for hot data, with cloud or on-premises object storage as a low-cost, lower performance, highly reliable storage resource to store cold data, manage snapshots, ensure data governance.

For customers looking to fully embrace a pay-as-you-go, on-demand consumption model, this approach can be financed through leasing options available from many storage vendors and connected to a cloud capacity storage tier.

KEY BENEFITS

- **Non-disruptively** change, add, upgrade or tier the storage serving your data
- **No application impact** while migrating data between storage systems
- **Gain flexibility and save** budget by leasing storage

Transform Data Migration From an IT Headache to a Business Opportunity

Storage migrations typically take months, consuming a large portion of IT's budget and resources. Migration plans require multiple steps to minimize disruption or downtime, often restricting access to data by halting applications, manually copying data to the new storage and reconfiguring, and then restarting applications while hoping that everything went according to plan.

DataSphere allows IT to ease or eliminate the common issues associated with data migrations. Once DataSphere is in place, organizations are no longer faced with the headache of planning and performing multiple steps to minimize disruption or downtime when migrating storage systems.

Admins can select the specific storage capabilities required to meet business needs. These objectives can be applied to single files, directories, or shares to provide unprecedented control. DataSphere analyzes if objectives are being met and will automatically redistribute data between different storage devices to meet performance, cost or reliability requirements.

Instead of a source of stress, Data migrations can be an opportunity for IT to add value by optimizing the infrastructure to meet changing data demands.

KEY BENEFITS

- **Overcome performance bottlenecks** using existing storage
- **Improve application performance** by moving data to your fastest storage tier without disruption
- **Increase clustered NAS performance** with automated, non-disruptive load balancing at file granularity
- **Accelerate metadata operations**
- **Automatic capacity balancing** across multiple vendors without application disruption

The New NAS: Parallel Performance, Power and Scale

Network Attached Storage (NAS) offers a rich environment for file-based management, which is often preferred by enterprise IT.

With DataSphere, admins can deliver higher performance with new forms of scale-out NAS storage solutions built from existing NAS deployments. DataSphere makes it possible to combine NAS arrays from different vendors for cost savings and agility, create logical storage tiers for improved capacity efficiency, define performance tiers for increased application throughput, automate demands using objectives, upgrade storage without disruption and leverage the use of the cloud today – seamlessly and without changing applications. DataSphere expands architectural storage choices to meet both IT's budget constraints and the application demands of the business.

KEY BENEFITS

- **Adopt VMware's Virtual Volumes** using existing storage
- **Implement VVol's policy-based storage management** across multiple storage types and tiers
- **Simplify administration** with a universal VASA provider and a single scale-out VVol datastore
- **Ensure compliance** to VM Storage Policies automatically as workloads change

Automated Storage Policy Management for Virtualized Environments

VM Storage Policies ensure VMs get initially provisioned to suitable storage, but they cannot automatically adapt to environmental changes. While admins get alerts when VMs fall out of compliance, they must still find where the VMDK is located and manually redistribute workloads through Storage vMotion, sometimes guessing as to whether or not the new distribution will result in different hot spots.

DataSphere can automatically redistribute workloads when a VM falls out of compliance, and can do so intelligently. For example, it might decide to move a less critical VMDK that isn't struggling to protect the I/O performance of a mission-critical VM. It can also automatically archive inactive VMDKs to deliver significant cost savings on recovered capacity, without admin intervention required.

DATASPHERE AND DSX TECHNICAL SPECIFICATIONS

Namespace	<ul style="list-style-type: none">• Unlimited size and number of files	<ul style="list-style-type: none">• Unlimited number of exported shares
Client Protocols	<ul style="list-style-type: none">• NFS v3• NFS v4.2	<ul style="list-style-type: none">• SMB 2.1• SMB 3.x
Clients	<ul style="list-style-type: none">• Linux• Windows• Mac	<ul style="list-style-type: none">• ESX• UNIX• BSD
Storage Protocols	<ul style="list-style-type: none">• NFS v3• S3	<ul style="list-style-type: none">• Server-attached storage• DAS (SSD, HDD)• SAN (VMware only)
Storage Vendors	<ul style="list-style-type: none">• NetApp• Isilon	<ul style="list-style-type: none">• Generic NFS storage• Any S3 compatible cloud / object storage

NOTE: For the latest Hardware Compatibility List, please contact your Primary Data sales team.

WANT TO LEARN MORE? Connect with us at: DeepDive@primarydata.com