

Scality Ring & Primary Data DataSphere

Primary Data's machine learning lowers the cost of managing cold data.

Storing cold data on high-performance storage is expensive. Since identifying and archiving cold data is a challenge, IT commonly overprovisions primary storage capacity. This strategy doesn't scale with the need to manage the growth and cost of unstructured data. Scality RING and DataSphere completely change the game, with machine learning powered data management and object storage.

Scality and Primary Data have partnered to help customers successfully adopt hybrid-cloud strategies for the enterprise, improving performance while lowering the TCO of their entire storage infrastructure, using machine learning data management to identify cold and dark data, and moving it non-disruptively into scalable, cost-effective object storage.

IT knows they have up to 90% cold and dark data in their data centers but storing this type of data on primary storage is expensive. Since identifying and archiving cold data is a challenge, IT commonly overprovisions primary storage, a strategy that doesn't scale for performance or cost. This is where Scality and Primary Data change the game.

Primary Data's intelligent data management platform, DataSphere, automatically and non-disruptively moves data to the most appropriate storage to meet user-defined business objectives. DataSphere employs machine learning to build intelligence and automation to help enterprises manage data across their hybrid-cloud. The DataSphere metadata engine adds awareness between applications and infrastructure, virtualizing data by mounting all storage in a global namespace. This makes heterogeneous data stores such as Scality RING file and object storage simultaneously available to all applications without requiring any changes to application workloads.

What's In It for You?



Optimize the entire storage infrastructure

by reducing the need to overprovision resources. DataSphere's machine learning metadata engine monitors data and application performance to identify hot, cold, and dark data and evaluate all data against defined objectives every 15 seconds.



Break down storage silos

with a global namespace, making all data available to any application or user regardless of where the data is located. This virtualization of data pools enables heterogeneous environments to deliver the performance of NAS with the economics of object storage.



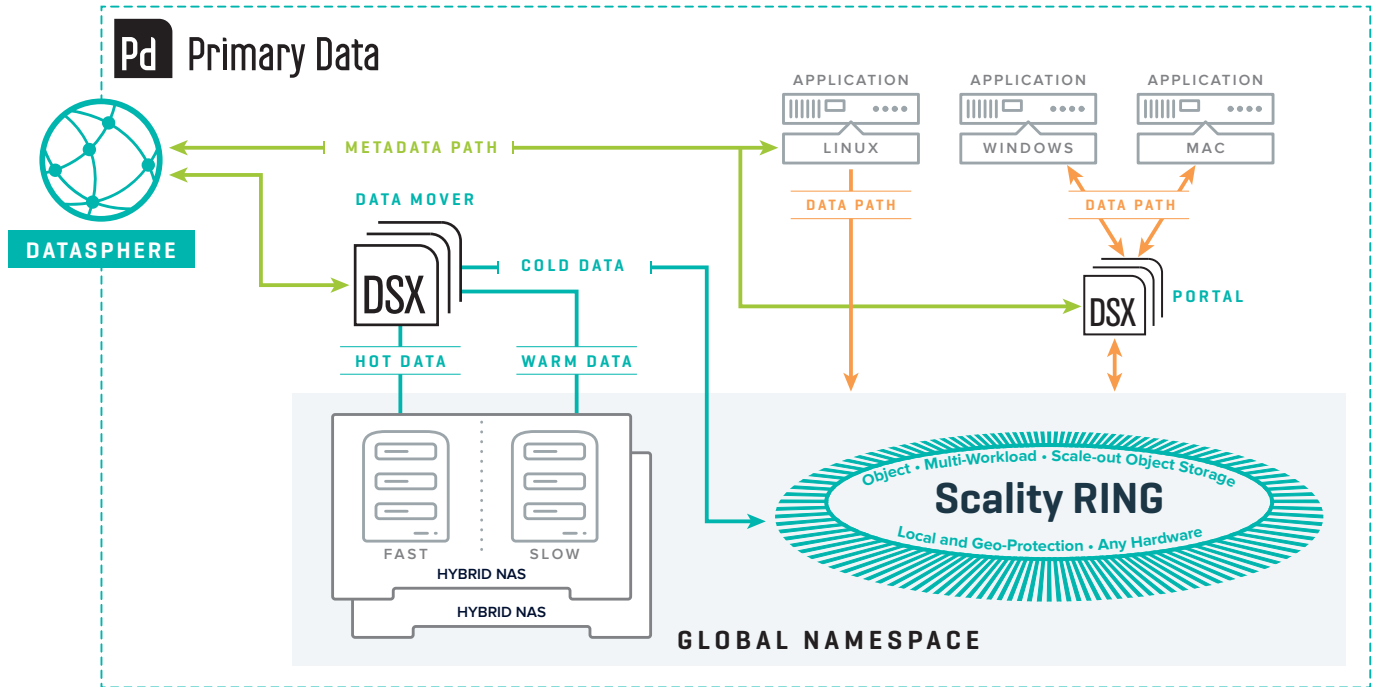
Automate the management of unstructured data

with non-disruptive auto-tiering of live data. DataSphere eliminates the need for caching multiple copies of the same data across the infrastructure, while using insights from machine learning to intelligently moving the right data in and out of S3 object storage without any disruption or reconfiguration of application workloads.



Pay as you grow

with a software-only solution. DataSphere is licensed per instance and is not based on capacity. It can be deployed either on bare-metal or in a VM. The DataSphere data mover technology will scale-up and -down to adapt to changing dynamics within an environment.



About Primary Data

Primary Data develops intelligence and automation software for enterprise data management across on-premises IT infrastructure and into the cloud. Its DataSphere platform combines metadata management and machine learning to move the right data to the right place at the right time across a global namespace, automatically and without application disruption. DataSphere makes heterogeneous data stores simultaneously available to all applications, enabling enterprises operating at petabyte scale to easily manage billions of files, automate data migration, integrate the cloud, and scale out NAS performance while getting the most value out of infrastructure investments on a per-client, per-file basis. To learn more, visit us at PrimaryData.com, follow us on Facebook.com/PdDataSphere, or Twitter at [@Primary_Data](https://Twitter.com/Primary_Data).

About Scality

Scality is a pioneering innovator of software-defined, multi-cloud data storage at petabyte scale. Recognized as a leader in distributed file and object storage by Gartner and IDC, Scality assures data control and freedom to manage data across clouds. Our products scale on-demand, non-disruptively, and drive lower cost for today's leading enterprise companies.

Follow us on Twitter [@scality](https://Twitter.com/scality) and visit us at www.scality.com to learn more