

SCALE-OUT DATA LAKE SOLUTION BUILT ON DELL EMC ISILON

Simple, scalable, efficient, and secure

ESSENTIALS

A data lake based on Dell EMC Isilon provides

- CAPEX savings: With 80 percent utilization, enterprise IT will benefit by simplifying the infrastructure on an Isilon scale-out data lake.
- OPEX savings: A single volume, single file system reduces management costs and better meets future needs.
- Elimination of storage silos: Consolidating data into a single store eliminates delays and lead times to ingest, load, move, and transfer data in and out of systems.
- Regulatory compliance: Meet compliance, governance, and regulatory requirements like SEC Rule 17-a4, PCI DSS, and HIPPA by utilizing support for data confidentiality, integrity, and availability along with self-encrypting drives (SED's).
- Multiprotocol access: These mechanisms reduce the risks to IT projects and added IT infrastructure.

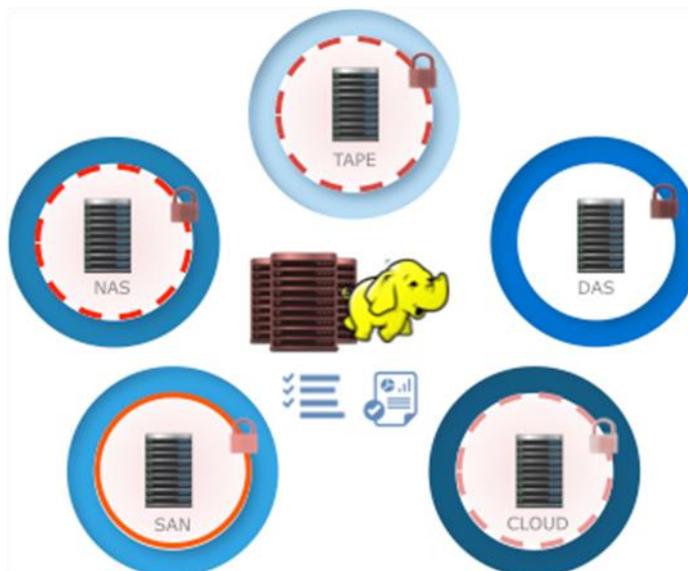
Dell EMC® Isilon® scale-out network-attached storage (NAS) is a simple and scalable platform to build a scale-out data lake and persist enterprise files of all sizes that scale from terabytes to petabytes in a single cluster. It enables you to consolidate storage silos, improve storage utilization, reduce costs, while providing you a future proofed platform to run today and tomorrow's workloads.

THE DATA CHALLENGE

Organizations are faced with the challenge of handling large quantities of unstructured data that is growing rapidly, is needed by more employees for business operations, and must meet stringent security and compliance mandates. With the growing popularity of Hadoop as an analytics tool of choice, businesses are

- Growing data needs combined with islands of storage that are difficult to manage
- Enabling access to more data and the results of analyses for a larger diverse audience
- Grappling with inconsistent security policies as datasets grow
- Experiencing longer latencies to move data from the sources to the analytics systems due to the size and complexity of access mechanisms

These challenges, along with access difficulties, are preventing business insights from driving business action.



DATA LAKE—ONE CONSOLIDATED STORE

Enterprise data growth is best addressed by consolidating storage into one single central repository of persistent data—a data lake—that simplifies the IT architecture, is efficient, and scales as business needs change. Organizations also need the peace of mind offered by a consistent application of security policies and access to the same data from a variety of sources and clients without an additional translation layer. The data lake future proofs your storage by enabling you to run traditional and next generation applications.

SIMPLE, EFFICIENT, AND MASSIVE SCALE-OUT STORAGE

Isilon clusters simplify storage by combining the file system, volume manager, and data protection into the Dell EMC Isilon OneFS® operating system. Through the clustered use of Dell EMC Isilon high-performance S-Series and X-Series nodes, high-capacity NL-Series and high density HD-Series nodes, a single Isilon cluster can contain a mix of tiers that provide the best economics, throughput, or IOs per second into the petabyte range. With over 80 percent storage utilization, Isilon clusters need less raw capacity than most storage systems, which can be further improved by 35 percent through Dell EMC Isilon SmartDedupe. Compared to traditional direct-attached storage (DAS) Hadoop, it can do so at a third of the storage capacity while providing more protection.

Consolidating your unstructured data on Isilon results in greater efficiency, simplified management, and cost savings.



DATA LAKE 2.0 STRATEGY

The new data lake 2.0 strategy expands the data lake to extend from the data center to the enterprise edge locations and to the cloud. With IsilonSD Edge software defined storage, enterprises can extend the data lake to the edge and consolidate their data to the core data center. Extending the analytics to the edge and to the core

enables enterprises to have a more holistic view of the data and to extract maximum business value.

With Isilon CloudPools software, the data lake can be extended to provide virtually limitless capacity without adding any complexity to store or manage the data. Cloud integration is seamless to users and applications.

CONFIDENTIALITY, INTEGRITY, AND AVAILABILITY

The OneFS operating system's features support the three aspects of data security, namely:

- Confidentiality of data is maintained by proper authentication, authorization, and access controls through support of Active Directory, Lightweight Directory Access Protocol (LDAP) and, access control lists (ACLs).
- Integrity of data against accidental, premature, or malicious alteration or deletion is achieved through Dell EMC Isilon SmartLock® for data immutability
- File system auditing, which provides the chain of custody information.
- Self-Encrypting Drives (SEDs) for increased security
- Data availability through Dell EMC Isilon FlexProtect™ technology creates an N-way, redundant fabric that scales as nodes are added to the cluster, providing 100 percent data availability, even with four simultaneous node failures. Dell EMC Isilon SnapshotIQ™ and SyncIQ® together provide local and remote data replication and disaster recovery.

MULTIPROTOCOL ACCESS

Isilon natively supports Server Message Block (SMB), network file system (NFS), FTP, HTTP, Network Data Management Protocol (NDMP), Swift, and Hadoop Distributed File System (HDFS) at the protocol layer, simplifying and speeding up access. This enables applications to store and retrieve data using their native format without additional translational overhead.

BENEFITS

A scale-out data lake built on Isilon can help your organization reduce costs, decrease storage complexity, and comply with corporate governance or regulatory mandates. The data lake 2.0 strategy extends the data lake from the data center to the enterprise edge and to your choice of public or private cloud options.

With the growing popularity of Hadoop as the Big Data analytics platform, Isilon helps speed time to insights and reduces the risks associated with deploying new systems or extending existing ones as business needs change.

TAKE THE NEXT STEP

Contact your Dell EMC sales representative or authorized reseller to learn more about how the Scale-Out Data Lake solution built on Dell EMC Isilon can benefit your organization

Also see our solutions in the Dell EMC Store at <https://store.emc.com/isilon>.

