

# Xcalar Provides Advanced Data Analytics on NetApp

Highly competitive enterprises are continually looking for ways to maximize and accelerate the value that they can derive from their data. Because of the growing number of data sources, many customers have large volumes of data on NetApp storage. This data can provide powerful, actionable business insights when coupled with a self-service big data analytics platform for discovery.

## Understanding the Value of Your Data

NetApp and Xcalar have partnered to provide customers with a powerful, self-service solution for business analysts, data scientists, and line-of-business managers to understand and discover insights from their NetApp data. With Xcalar's True Data in Place™ technology, NetApp users can explore all of their business data without copying or moving the data. Designed with a familiar spreadsheet-like paradigm, Xcalar enables all NetApp users to build models that ask questions and create business value at scale, directly and interactively, without requiring programming skills. Xcalar leverages all the resiliency, availability, and scalability built into NetApp storage and extends it with a simple, powerful, and easy-to-use data analytics solution.

## Making Discovery as Easy as Searching

Xcalar Design, an intuitive web-based user interface for viewing and modeling data, provides a familiar spreadsheet-like environment. NetApp users can quickly discover meaning from their data by pointing and clicking through their data. Xcalar Design represents the sequence of user actions as dataflow graphs with full data lineage and auditability. Xcalar users can leverage these dataflow graphs to compile sophisticated algorithms, which can be deployed in production on the Xcalar Compute Engine without assistance from programmers or IT.

## True Data in Place™

Xcalar's patented True Data in Place™ technology keeps all data in its original form on the NetApp storage platform. This retains all the data management, reliability, and efficiency expected from NetApp and relied upon by its users. Processed data and analytics result sets can be written directly back to the user's NetApp storage. All data, including dataflows, system logs, and user-defined functions, is stored on NetApp, providing a true separation of storage and compute for scaling resources to efficiently meet analytics needs.

## HIGHLIGHTS

- True Data in Place allows users to reference source data without copying it into a new format.
- Xcalar is agnostic to storage type and format so users can directly access data on NetApp® FAS, StorageGRID Webscale, and Hadoop HDFS.
- Performance and capacity scale linearly with the number of nodes.
- Scale storage independently of compute to accommodate petabyte-sized data lakes.
- Deploy on the premises or in the cloud across multiple locations with no vendor lock-in.

## SOLUTION COMPONENTS

- ONTAP®
- SnapMirror®
- StorageGRID Webscale
- Xcalar Data Platform

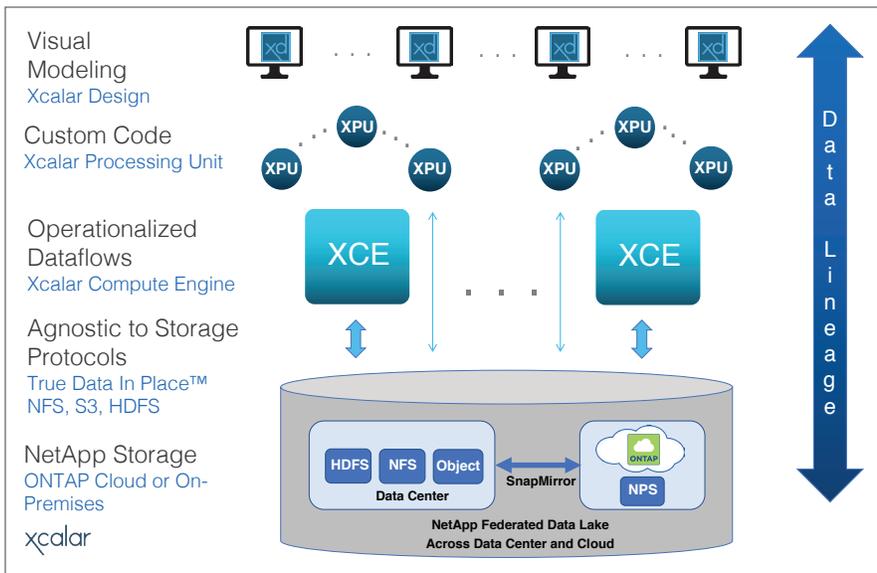


DIAGRAM 1: NETAPP XCALAR SYSTEM ARCHITECTURE

## NetApp Data Management

NetApp's joint solution with Xcalar fully leverages the capabilities of ONTAP, the industry's leading enterprise data management software, combining simplicity and flexibility with powerful data management capabilities and storage efficiency. With ONTAP deployed as NetApp Private Storage, users can build a hybrid cloud that is the foundation of a Data Fabric that spans flash, disk, and cloud. Users can also seamlessly manage their data as it flows to where it is needed most, to help make the best possible decisions for the user's organization. For example, a user can deploy one Xcalar cluster in its test and development storage environment to model Xcalar dataflows, then operationalize by seamlessly deploying dataflows in its production

storage environment on a separate Xcalar cluster. Take advantage of NetApp SnapMirror, Snapshot™, and cloning to maximize efficiency and eliminate redundant data copies. RAID DP® and advanced storage software provide storage efficiency and resiliency without the cost of double- or triple-mirroring the underlying data.

## Direct Data Access

Xcalar is agnostic to the type of storage on which the data resides. Xcalar supports cloud object storage, including Amazon S3, as well as SAN, HDFS, and NAS. Therefore, from a single Xcalar deployment, the user can directly access NetApp storage using NFS, NetApp StorageGRID Webscale data using Amazon S3, or Hadoop storage

backed by NetApp E-Series, without configuring gateways or other intermediary steps. Xcalar interprets data in any format, which enables the modeling of dataflows in minutes, without the need for a separate ETL process or data copying.

## Cloud Deployment or On-Premises Options

Xcalar architecture separates compute from storage. This enables great flexibility for on-premises, cloud, and hybrid cloud deployments. Here are some of the salient features:

- Xcalar clusters are elastic and can be resized to meet unique service-level objectives.
- Xcalar runs on all the major cloud platforms, so that users are not locked into one particular cloud vendor. This allows the user to choose a cloud provider based on the provider's cost, performance, and services. Xcalar is deployed on Microsoft Azure, Amazon Web Services, and Google Cloud and is also available on premises.
- Users can deploy Xcalar on premises with local access to their NetApp storage, with the option to deploy to the cloud with minimal disruption.
- Xcalar dataflow graphs are transferable and can be moved to other cloud platforms or on-premises environments, without disruption to the user's NetApp storage deployment.
- Multiple Xcalar instances running different workloads can access the same files on shared storage.

### ABOUT Xcalar

**Xcalar is an open, interoperable, and extensible analytics platform that delivers the fastest time-to-insight on the market with unprecedented simplicity, speed, and scale. Its patented technologies provide a complete and powerful suite of products for data engineering, data science, scale-out data warehousing, and ad-hoc analytics. Xcalar's enterprise grade software scales to hundreds of nodes, thousands of users, and petabytes of data, for both on-cloud and on-premises platforms.**

### ABOUT NETAPP

**NetApp is the data authority for hybrid cloud. We provide a full range of hybrid cloud data services that simplify management of applications and data across cloud and on-premises environments to accelerate digital transformation. Together with our partners, we empower global organizations to unleash the full potential of their data to expand customer touchpoints, foster greater innovation and optimize their operations. For more information, visit [www.netapp.com](http://www.netapp.com). #DataDriven**