

WANdisco Fusion[®]

Product Overview

Consistent Data Everywhere

Keeping data consistent in a distributed environment is a massive challenge. WANdisco Fusion, an enterprise-class software platform, solves the exponentially growing challenge of keeping unstructured data available across diverse IT environments regardless of geographic location, architecture, or cloud storage provider. Used by enterprises worldwide, our technology is based on a high-performance coordination engine called DConE, which uses consensus to keep Hadoop and object store data accessible, accurate, and consistent in different locations across any environment—on-premises, hybrid-cloud, multi-region cloud, or multi-cloud. WANdisco Fusion is a foundation for a modern cloud data strategy—a LiveData strategy—because it prevents data disasters, de-risks data migration to the cloud, and simplifies hybrid and multi-cloud data management.

What Is A LiveData Strategy?

Make your data globally accessible and consistent everywhere with a LiveData strategy. Alleviating the challenges of siloed data, a LiveData strategy ensures that enterprise data stays accurate, accessible, and consistent across your global IT environment. With a LiveData strategy, every user and every application has always-available data, regardless of geographic location, data platform architecture, or cloud storage provider.

WANdisco Fusion Overview

- Coordinates changes to data across heterogeneous storage systems achieving consensus and consistency at scale.
- Minimizes latency by allowing applications to use data locally, while replicating changes consistently with other replicas.
- Eliminates points of failure by ensuring data remain available even when systems fail.
- Self-healing in the event of an outage with automated recovery.
- Multi-directional replication capability allows active workloads in all locations.
- Works with Hadoop file systems, object storage, Hive data, and security metadata.
- Works across any network topology—LAN, WAN, public and private clouds.
- Leverages a common administration and management interface.

Are your replication capabilities keeping pace with your data volume growth—and your business needs?

Guaranteeing the safety and accessibility of your business data is vital. If you work in a data-powered business, you need 24/7 access to that data to keep operations moving smoothly and stay competitive. However, as data infrastructures grow more distributed, many companies rely on complex hybrid and cloud environments, often from a combination of vendors and platforms, in multiple regions. Enterprises need new technology to enable continuous data availability and consistency in a mixed, multi-cloud IT environment.

What is DConE and how is it different?

WANdisco Fusion is powered by DConE, which is a high-performance, coordination engine able to work across wide-area networks. The technology uses distributed consensus to guarantee data consistency.

Traditional approaches to data replication are batch-based, do not guarantee data consistency, and cannot operate over wide-area networks or the public internet. Unlike other technologies which move data from one location to another, WANdisco Fusion uses DConE to coordinate distributed changes to data, enabling shared access to common data sets. The technology works by applying a mathematically-proven approach to consensus which works regardless of the distance between data sources or types of data stores.

Data Consistency In A Multi-cloud World

WANdisco Fusion platform enables any hybrid and multi-cloud strategy. Enterprises are recognizing the benefits of a multi-cloud strategy: greater flexibility and functionality, stronger resilience and disaster recovery, and improved business performance. Different applications have distinct needs, and different cloud products have distinct strengths and weaknesses. As a result, each department within an organization may be using a different cloud platform for their applications. Even though they may not have deliberately planned it, most enterprises are finding themselves working with several cloud vendors. Managing a complex network of multiple cloud vendors and instances, spread across many regions, creates big data management challenges including global data consistency, and data availability during data migrations

Cloud and distributed architectures require fundamentally different replication technology

Business Benefits

- **Business continuity in a distributed environment:** A new approach to large scale data replication ensures data consistency and availability, even during large scale data migrations / cloud migrations.
- **Global data resiliency:** Continuous replication prevents data disasters.
- **Always available data:** Users and applications can access and change data from any location.
- **Improved application performance:** Eliminates latency, points of failure, and availability challenges associated with geo-distributed data regardless of distance.
- **Ability to meet strict data SLAs:** Near-zero RPO/RTO removes data gaps inherent with point-in-time backups or snapshots, and enables a long-term strategy for meeting regulation compliance requirements.
- **Reduced data gravity challenges:** Avoids cloud vendor lock-in and simplifies management of multiple cloud storage providers.
- **Self-healing:** Eliminates the need for administrators to respond to system level failures with automated recovery.

LiveData Use Cases

Disaster prevention for any environment

Disaster recovery and high availability via near-zero RTO and RPO ensures business continuity and data SLA compliance across on-premises, hybrid and multi-region, multi-cloud environments.

Cloud migration without downtime

Zero-disruption data migrations allow users to continue working on data—even for large-scale migrations.

Data lake consolidation

Synchronize multiple data lakes, unify data silos, and maximize your data lake investment and line of business data utilization.

Hybrid-cloud enablement

Keeps on-premises and cloud environments continuously in sync capturing every change wherever it occurs.

Protection from cloud vendor over-reliance

Eliminate application downtime from the failure of an object storage service by having strongly consistent replicas of objects in another cloud or location (and avoid vendor lock-in).

Use the right cloud for the right job

Optimize your multi-cloud strategy and business performance by taking advantage of each cloud's unique or best-performing capabilities (and avoid vendor lock-in).

Multi-cloud data operations enablement

Guarantee global, continuous data accuracy to applications in different clouds across geo-distributed, multi-region, multi-cloud architectures.

Lower Cost Structure

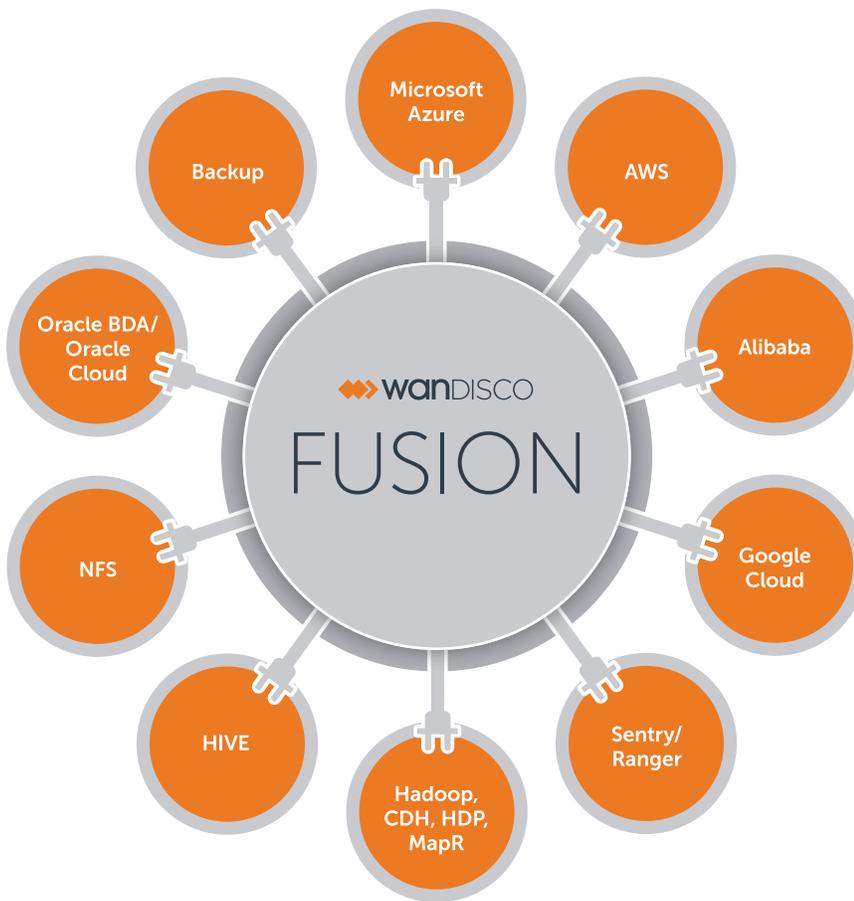
Fully utilize hardware previously reserved for backup and recovery, and improve your infrastructure ROI with more capacity. By freeing up capacity, your entire data infrastructure is actively leveraged to run your business. You no longer need to dedicate one location to just storing a copy of your primary system's data, but can take full advantage of every location for running active workloads at the same time. Multi-directional replication allows active workloads in all locations.

The WANdisco Fusion Platform Plugins

WANdisco Fusion’s plug-in architecture is highly flexible, enabling multi-Hadoop distribution, hybrid cloud, and multi-cloud strategies.

The extensible platform plugins work on popular data sources, physical devices, and Hadoop components, e.g., Ranger, Hive, and Sentry, ensuring compatibility with multiple systems.

The platform plugins enable operations in unlimited regions in AWS, Azure, and major cloud vendors, or a mix of on-premises and cloud S3-compatible storage services.



Features

- Flexible selective replication defined at a per-folder level or per-data set.
- Compatibility with common security protocols—Kerberos, SSL/TLS, LDAP.
- Quickly compare and validate data consistency between zones and replicated environments.
- Replicate encrypted at-rest data between different platforms.
- Network traffic is between WANdisco Fusion servers so that application environments can stay isolated.
- UI-controlled throttling of network bandwidth utilization.
- API support of server administration and automation.
- Data is replicated and stored in cloud native formats for direct access by applications and services.
- Run applications in multiple, geographically separate locations with local endpoints in each.

Supported Environments

HADOOP DISTRIBUTIONS

Cloudera CDH
Hortonworks Data Platform
MapR
IBM BigInsights®
Microsoft Azure HDInsight®

CLOUD PROVIDERS

Amazon Web Services
Microsoft Azure
Google Cloud
Oracle® Cloud
Alibaba Cloud

OPERATING SYSTEMS

Centos
Oracle Linux
RHEL
SLES
Ubuntu

STORAGE SYSTEMS

HDFS and other Hadoop-compatible file systems
EMRFS
MapRFS
ADLS Gen1 and Gen2
WASB
Google Cloud Storage
Local and NFS mounted file systems
NetApp ONTAP
OpenStack Swift
AWS S3
AWS Snowball
Azure Blob Storage
IBM Cloud Object Storage
Dell EMC Elastic Cloud Storage
Ceph Object Storage
Dell Virtustream Storage Cloud
Oracle Cloud Object Storage
Oracle Cloud Object Storage Classic
Alibaba Cloud Object Storage Service
Pure FlashBlade Object Storage
HGST (Western Digital)
ActiveScale Object Storage
Minio Gateway for Azure Blob Storage
Scality Zenko CloudServer

About WANdisco

WANdisco is the LiveData company that empowers enterprises to revolutionize their IT infrastructure with its groundbreaking distributed coordination engine (DConE) in the WANdisco Fusion platform, enabling companies to generate hyperscale economics with the same IT budget — across multiple development environments, data centers, and cloud providers. WANdisco Fusion powers hundreds of the Global 2000, including Cisco Systems, Allianz, AMD, Juniper, Morgan Stanley and more. With significant OEM relationships with IBM and Dell EMC and go-to-market partnerships with Amazon Web Services, Cisco, Microsoft Azure, Google Cloud, Oracle, Alibaba and other industry titans – WANdisco is igniting a LiveData movement worldwide.



5000 Executive Parkway, Suite 270
San Ramon, California 94583

Talk to one of our specialists today

EMEA +44 114 303 9985
APAC +61 2 8211 0620
ALL OTHER +1 925 380 1728
US +1 877 926 3472

Join us online to access our extensive [resource library](#).

Follow us to stay in touch

