

Dell Storage and Arcitecta Mediaflux® Multi-Site: A Single Global Namespace for Data

Geo Distribution is Driving Opportunity

Content production is a team effort that leverages the creativity and energy of increasingly globally distributed people. Whether collaborating between New York and Chicago or Los Angeles and Tokyo, distributed teams face hurdles accessing shared data and resources in real-time, or near real-time, across different locations, leading to collaboration and productivity challenges and inefficiencies.

The need for a global namespace arises from the requirement for teams to access and interact with data seamlessly, regardless of its physical location. When projects and workflows are distributed, additional issues arise as latency increases. Consider, for example, a visual effects company that shares project work between geographically distributed offices, including outsourcing work to sub-contractors. To establish “follow the sun” pipelines, data needs to move across networks efficiently and transparently arrive when and where it is required, no matter where the expertise, talent, or computer resource is located.

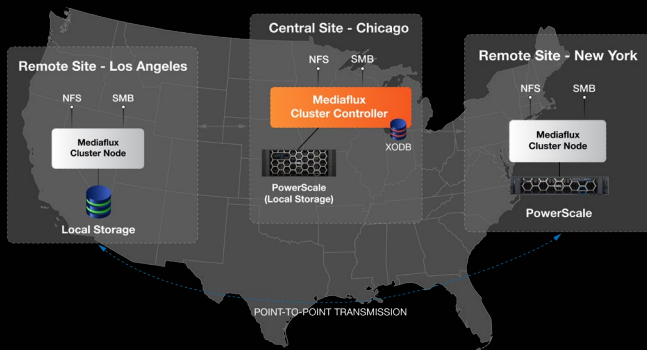
A Single View of Data

Mediaflux, Dell PowerScale (file storage), and/or Dell ECS (object storage) are optimal for managing and storing extensive data across dispersed locations. Mediaflux, installed at each site and linked to either PowerScale or ECS/ObjectScale local storage, offers users a unified view of distributed data through a global or federated namespace. Operating efficiently either in-band or out-of-band, Mediaflux boasts a high-performance XODB database for seamless scaling, alongside an open API for easy integration into any workflow. With its robust security features, Mediaflux ensures data visibility and access is restricted subject to permissions. Combined with PowerScale, the world's most cyber-secure scale-out file array¹, your data has unparalleled protection, making it the ideal solution for safeguarding your critical data assets. This solution streamlines data management, metadata curation, and facilitates data movement, ideal for organizations operating globally. Paired with Dell PowerScale, Mediaflux enables global data accessibility, fostering collaboration and accelerating project completion by leveraging multi-site capabilities.

Depending on the distance between teams, there are two options:

Global File System (GFS) – Figure 1 illustrates how locating a unified mount point less than 2,000 miles (about twice the distance between Chicago and New York City) from either team location allows data added in one location to instantly propagate to others, ensuring seamless collaboration worldwide.

Federated File System (FFS) – Over greater distances, such as between Chicago and London (Figure 2), teams can be connected by a federation where metadata and data are asynchronously transmitted between sites. Each site will still see the same data – within 100s of milliseconds – yet a given file can only be created or modified by one site at a time.



Global File System (GFS)

A global file system deployment consists of a central metadata server and local edge nodes at each site that present a local file system and accessibility to data from all sites. Each satellite may have its local storage for caching and increased I/O at that site. All file writes occur locally. The system also allows for in-place (distributed) data analysis and metadata extraction.

The central system – the cluster controller – houses the master metadata database, XODB. It can also store duplicate copies of data for backup or local access. The global file system ensures uniform and authorized access across all locations. When a file from another site is accessed, it is transparently retrieved and cached locally.

Federated File System (FFS)

A federated file system deployment enables connectivity between high-latency, geographically distant, locations such as Sydney and London. In this setup, teams are interconnected via a federation that uses asynchronous transmission of metadata and data between sites. Despite the distance, each site maintains visibility to data in near real-time (hundreds of milliseconds).

This federated approach is well-suited to environments where files are created at one site and transmitted to another site for access and computation.



The Mediaflux Difference

Mediaflux offers a solution for a true global namespace by placing the data management layer in the data path, between storage and user layers (in-band). This architecture allows users to access the same data concurrently through different protocols via a single mount point or share, eliminating the need to track data storage locations.

There are several advantages to this:

- **Real Time Synchronization** – All user I/O passes through Mediaflux in real-time, enabling data routing to or from the appropriate storage via policy, thus facilitating active storage management.
- **Workflow Optimization** – Policy driven data management is a major advantage of configuring Mediaflux in-band. It enables data to be moved automatically across distributed environments based on workflow policies or triggers, which means that far fewer people can be more efficient at managing large amounts of data.
- **Storage Optimization** – Mediaflux sits in front of the storage layers, thus the organization can change the backing technology to any storage mix because users access data via a global namespace. As a result, organizations become free to optimize storage using any desired technology at any point in its lifecycle without disrupting end-users' path to find data. Mediaflux can even automate these storage optimizations via policy.

High Speed Transfers are Imperative

Mediaflux utilizes asset metadata to efficiently eliminate redundant content transfers and enhance the movement of data by employing parallelized transfers across latent networks. It excels in optimizing the transfer of your data, regardless of file size or type, and seamlessly scales by consolidating smaller assets or breaking down larger assets into optimal chunk sizes to align with the underlying compute or network infrastructure. Additionally, Mediaflux may be utilized as a transfer node for edge locations acting as an edge cache for data being centralized or utilized in the cloud.

Furthermore, when content is encrypted for transmission, which is often the preferred setting for sensitive and confidential data, Mediaflux achieves a multithreaded transfer speed that is nearly twice as fast as alternative solutions.

Arcitecta + Dell Technologies: Better Together

It shouldn't matter where these data workflows occur – joint solutions from Arcitecta and Dell Technologies deliver data where it's needed at the right time. Arcitecta's pioneering metadata and data orchestration tools coupled with Dell Technologies' powerful, industry-trusted infrastructure enable a global distributed edge that stays simple and performant, no matter the complexity of your workflows.

About Arcitecta

Arcitecta has been crafting highly advanced platforms since 1998. Today, Arcitecta is transforming data management and backup with Mediaflux, a rich end-to-end data fabric that simplifies data-intensive workflows in petabyte-scale environments to improve business and research outcomes. Mediaflux unifies data management processes into a single platform, simplifying the administration of big data and allowing world leaders to solve some of the most challenging problems on the planet.

About Dell Technologies

Dell Technologies is a trusted leader in media and entertainment storage for content creators, broadcasters, and content delivery providers. Dell Technologies storage forms the foundation of a simple and future-proof infrastructure, giving you the agility to transform business operations, and the flexibility to adapt to new media workflows. Through innovative products and services, Dell Technologies accelerates the creation and monetization of media, helping media professionals store, manage, protect, and analyze their most valuable digital media.

¹ Based on Dell analysis comparing cyber-security software capabilities offered for Dell PowerScale vs. competitive products.

