

Mediaflux® Universal Data System

**Converged Data Management, Data Orchestration,
Multi-Protocol Access, and Storage in One Platform**

Revolutionary Cost Savings

**Manages the Entire Data Lifecycle – On-Premises
and Cloud, with Globally Distributed Access**

Immense Scalability

**Underlying Storage Can be From any Vendor or
Multiple Vendors - Breaking Vendor Lock-In**

Limitations of Clustered File Systems

Clustered file systems have been a mainstay in the data storage landscape for years. Their primary function has been to store data across multiple servers, ensuring that the system can be accessed and scaled as one unit. However, their capabilities have traditionally been limited to just that – storing files. They lack the sophisticated tools and features required for comprehensive data management and orchestration. In a world where data is more than just bytes and bits, but a valuable asset that drives decision-making and innovation, this limitation is significant.

The Arcitecta Mediaflux Journey

Until recently, Arcitecta Mediaflux was built on top of third party file and object storage systems from other vendors. This approach, while functional, was not without its challenges. The exponential growth in data generation has shifted the demands of our customers and highlighted the need for more affordable and flexible data management solutions. The costs associated with existing enterprise storage systems were becoming increasingly unsustainable as data volumes surged.

Introducing the Universal Data System

Recognizing these challenges, Arcitecta has introduced the Universal Data System. This innovative solution offers clustered storage capabilities without the need for any third party software. Whether you're using block storage from one vendor or multiple vendors, the Universal Data System can seamlessly integrate and manage them all. This flexibility is a game-changer in the industry.

Revolutionary Cost Savings

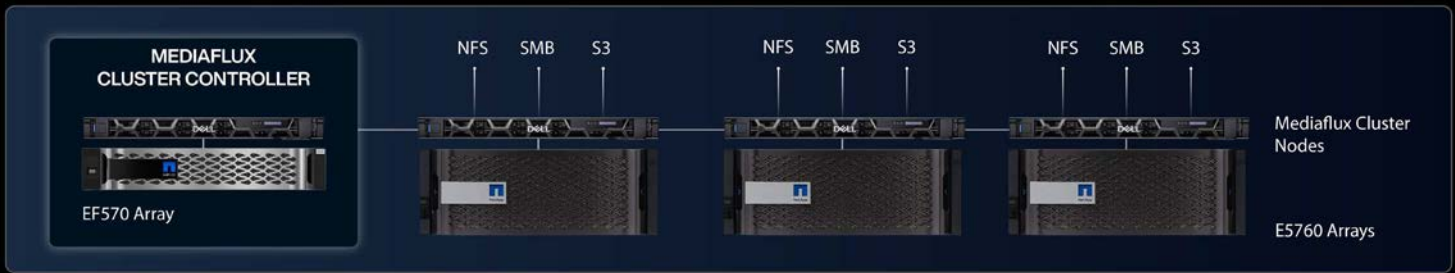
Traditional clustered and distributed file systems come with a hefty price tag. The Universal Data System is a response to this pain point, allowing customers to bypass third party storage file systems, and constructing a storage ecosystem tailored to their needs, resulting in significant cost savings. Moreover, they benefit from the advanced capabilities of the Mediaflux data management platform without incurring additional expense.

A Universal Approach to Data Management

The Universal Data System is not just about storage; it's about managing data throughout its entire lifecycle. From the moment data is acquired, to its cataloging, transformation, dissemination, preservation, and eventual storage – whether on-premise, in the cloud, or distributed across the globe – everything can be managed within a single system. This holistic approach ensures that data is not only stored but is also accessible, usable, and valuable.

Immense Scalability at an Affordable Price

Mediaflux licensing is decoupled from the data stored. This means organizations can affordably scale their storage needs to hundreds of petabytes, accommodating hundreds of billions of files without the financial strain typically associated with such vast capacities. This groundbreaking model not only ensures that organizations can grow their data infrastructure as their need evolve but also guarantees that they can do so without the looming concern of escalating costs.



Mediaflux Universal Data System with NetApp E-Series Block Storage

NetApp E-Series Block Storage Example

As shown above, NetApp E-Series arrays can connect to a Mediaflux server using various protocols such as SAS, InfiniBand, Fibre Channel, iSCSI, or NVMe over Fabrics, depending on the specific configuration and requirements. The Mediaflux Cluster Controller houses the Mediaflux XODB metadata database and runs on a flash-based NetApp EF570 array for maximum performance. Mediaflux creates an EXT4 file system across each E-Series array. EXT4 was chosen for its reliability and performance and is a natural choice for the Mediaflux data management application.

Mediaflux unifies multiple E-Series array servers, presenting them as a single logical NFS/SMB/S3 endpoint with multiple IP addresses. Mediaflux selects nodes for I/O based on its internal load balancer. This load-balancer evaluates the workload assigned to each node, factoring in write affinity needs to maintain E-series block optimization. Mediaflux employs various algorithms to distribute the load across the E-Series cluster, which allocates requests based on a node's response time, ensuring resilience against potential node issues.

The Basis of a National Compute Infrastructure

Australian National University's National Compute Infrastructure presents a compelling case study for the capabilities of Mediaflux Universal Data System. The lead site has 4PiB of NetApp block storage managed by Mediaflux Universal Data System, which also tiers Spectra Logic Black Pearl connected tape libraries. With the integration of Mediaflux Universal Data System, the National Compute Infrastructure can incorporate a variety of storage solutions at varying price points, optimizing their storage costs. Universal Data System also extends its reach to external researchers and collaborators across Australia. The ambitious project, once fully realized, will manage an impressive 140PiB of data, showcasing the system's immense scalability.

Managing the Entire Data Life Cycle

Mediaflux has emerged as a groundbreaking standard bearer in data management, providing a holistic view of data irrespective of its storage system or geographical location. With its multi-protocol support, including NFS, SMB, and S3, Mediaflux ensures the effortless integration of new storage technologies, thereby amplifying data accessibility. Its intelligent data placement feature further optimizes storage efficiency by automatically tiering data based on usage and access patterns.

Among its myriad capabilities, Mediaflux provides disaster recovery replication, ensuring data safety and continuity. Its extensive metadata capabilities not only improve data discovery but also foster collaboration and knowledge sharing. The platform's global distributed access ensures data can be retrieved from any location, facilitating international collaboration. Supporting this is Mediaflux's integrated highspeed WAN file transfer capabilities, ensuring rapid and secure data transmission across vast distances. Its multifactor authentication and file versioning features underscore its commitment to data security and integrity, making it an indispensable tool for data-centric organizations and workflows.

The Future of Data Management and Storage

Mediaflux Universal Data System is the culmination of Arcitecta's vision for the future of data management. By merging world-class data management, orchestration, multiprotocol access, and storage into one cohesive platform, Arcitecta is setting the standard for the industry. It's not just about storing data; it's about making data more accessible, manageable, and valuable than ever before.

