

FEDERATION & REPLICATION

A Rich Suite of Tools to Enable Collaboration and Protection of Distributed Data

Mediaflux includes key capabilities at its core to enable global collaboration and policy-based disaster recovery across multiple locations of a distributed infrastructure.

Leveraging the power of metadata, Mediaflux Federation and Replication services are a powerful set of collaboration tools to ensure that data can be accessed when and where it needs to be, in the most efficient way possible, across multiple uses and user applications within an organisation.


FEDERATION & REPLICATION

As data volumes and use cases in an organisation increase in size and complexity, it becomes increasingly challenging to ensure the relevant data can be found and made fully accessible to those who need it, with complete security and in a cost-efficient way across a distributed infrastructure.


Whether it be network latency or the sheer volume of data, simply duplicating the entire data set or infrastructure is often not possible, nor is it affordable. Additionally, collaboration with remote offices or partners must be managed seamlessly, and often with subsets of data. This must be done in ways that do not add complexity or risk.

Mediaflux Federation and Replication services provide the power to extend data management policies and workflows across distributed locations and multiple simultaneous use cases, all with the same granular control that is a hallmark of all Mediaflux capabilities.

- Users can perform searches or execute other services across multiple locations in parallel, and receive a consolidated result.
- Collaboration can be tuned to allow complete or partial access by remote partners, deployed users, or field offices, according to role-based authentication permissions.
- Policies can be established to automate replication of some or all data between sites to enable disaster recovery scenarios or to automate complex workflows.



Mediaflux Federation enables different Mediaflux instances to act in concert whether they are in another building or around the world. In this way, different sites with unique content can be virtually combined to enable secure global collaboration.



Mediaflux Replication is used when data and/or metadata needs to be synchronised between sites. This can be applied to all or part of the data under management, according to policies and access privileges.

MEDIAFLUX FEDERATION

Mediaflux Federation services enable multiple Mediaflux nodes in different locations to act in concert, so that actions may be executed across any or all of them in parallel from a single command. Unlike clustering, where multiple Mediaflux servers in a single location are bound together to parallelise local operations, Federation enables distributed operations across multiple sites, whether across a campus, or in other parts of the world.

MEDIAFLUX REPLICATION

Mediaflux Replication services support the sharing of data and metadata between multiple Mediaflux servers, whether the servers are separated by a few metres or located on different continents. Typically a Mediaflux server is placed wherever important data is located, and replication services are used to enable convenient access to data throughout an organisation.

KEY REPLICATION & FEDERATION FEATURES AND BENEFITS

- Apply any Mediaflux service or query in parallel across multiple locations.
- Granular role-based access limits which data a user may see, query or use.
- Mediaflux servers can participate in multiple federations without the need for a central control server.
- Replication policies can be applied globally, or on subsets of data or triggers.
- Metadata and data can be replicated independently, either uni-directionally or bi-directionally.
- Replication can be synchronous or asynchronous operations, to enable deployed or highly latent uses for distributed data.

A FEDERATION OF TRUSTED NODES


Data in each location is owned by the Mediaflux server storing it. There is no need for a central server to arbitrate data access across a federation, so there is no uncertainty over data ownership.

When a service is executed on a remote server, the user's authentication is linked to the unique identity of the originating server. This allows for extensive configuration over remote access rights to metadata and data.

PEERS

Mediaflux servers can be configured as a network of connected peers, and laid out to reflect the data topography of an organisation.

The architecture is extremely flexible, such that each peer can have multiple replication policies for all or parts of the data. This enables easy integration of partners or complex workflows into a shared environment.



INTELLIGENT REPLICATION

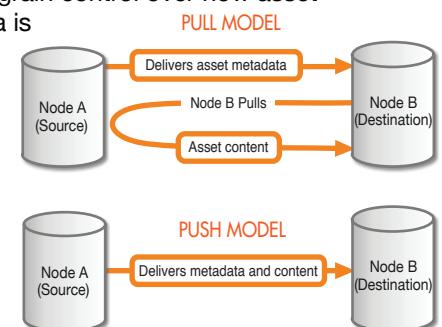
Replication within Mediaflux can be a regularly scheduled event for disaster recovery, or the result of a work flow or other event trigger. It can be the result of a simple or arbitrarily complex database query. It can be based on derived metadata or on direct content attributes. Restoration takes existing metadata and data into account and only merges in required versions. Replication may be "by reference", where metadata is replicated according to the schedule, and data is replicated on demand when requested.

SYNCHRONOUS AND ASYNCHRONOUS

Mediaflux supports both synchronous and asynchronous links to optimise scalability between large numbers of nodes. This also enables remote deployments or peers with limited bandwidth.

Mediaflux provides fine-grain control over how asset content and/or metadata is transferred between nodes.

Air-gap replication also enables content to be shipped physically if need be, to catch up with the metadata and accommodate difficult network scenarios.



ABOUT MEDIAFLUX

Mediaflux is a rich, platform-independent system for curating, managing and controlling huge amounts of information/data in all forms and life-cycle phases for historical, contemporary and future needs and across any type of storage technology. For over a decade Mediaflux has enabled customers to bridge data and metadata incompatibilities, as well as to virtualise diverse storage types to accelerate time to discovery and reduce data management complexity so they can realise the maximum value from their data and infrastructure.