



POINT IN TIME

POWERED BY MEDIAFLUX®

Continuous inline data protection
for trillions of datums.

Nobody likes to lose things

How do you secure one hundred petabytes of data? Moreover, how do you find a single file in repositories that large?

It's estimated there will be nearly 100 zettabytes of data in the world by the end of 2022. With the scale of data continually growing, making it secure and resilient is becoming harder to achieve. How would you backup petabytes of data in the hundreds? You probably wouldn't with traditional backup, and that's precarious. Vulnerabilities scale with data growth: corruption, malware, accidental deletion, mysteries. Furthermore, the time it takes to find lost data with traditional backup systems increases with the amount of backup data stored. IT departments are constantly pulled into the task of data recovery. Data resilience for trillions of datums, and instant, self-serve data recovery is not possible with backup as we know it.

Traditional backup was a solution for millions of files.

But we're in an age of billions of files. Traditional backup works by scanning a file system to find and create copies of new and changed files. The problem is scanning takes longer as the number of files grows – so much so that it's becoming impossible to complete scans within a reasonable time frame. They usually run during the night when systems are likely to be less volatile. The process occurs at set intervals, which means any change before the next scan will be lost if there's a system failure. Traditional backup cannot and does not meet the objective of zero data loss.

Recovering data in petabyte-sized repositories is time intensive.

The process of recovery is not what it should be – it's tedious and slow. When someone wants to recover data, they will typically ask an IT administrator for help to recover it. The administrator will then ask them for the path and names of the missing files, along with the date and time they existed – many people will not remember those details exactly, and so begins a process where different backup sets are restored one after another and inspected until the missing or damaged files are found. That process can take hours, days, or longer, to recover data – a process that is inefficient and costly. This inefficiency compounds when there are many files to find and restore.

Mediaflux Point in Time. The Backup Revolution.

Mediaflux Point in Time is software-defined storage that virtualizes existing storage infrastructure to provide continuous data protection. It presents as NFS or SMB – both protocols can be used concurrently on the same data – in addition to a host of other protocols that includes sFTP, DICOM, S3 and more. Every structural and data change is captured by Mediaflux Point in Time and recorded in Mediaflux's high performance database XODB®.

Mediaflux Point in Time achieves maximum data resilience at scale.

Mediaflux Point in Time is designed for the scale and complexity of today's data demands – billions of files and hundreds of petabytes of data.

Point in Time achieves Recovery Point Objectives (RPO) of near zero (typically within milliseconds).

Every change that happens in the data path is captured in Mediaflux's high-performance database XODB®. Every metadata change has an RPO of zero. Data changes are immediately and continuously replicated – the RPO will depend on the time it takes to transmit to alternate storage determined by network speeds and available bandwidth. When recovery is required, every point in time is there.

Mediaflux Point in Time achieves Recovery Time Objectives (RTO) of zero.

With an intuitive GUI, Point in Time gives users (with the correct permissions) the ability to easily find and access their data across all points in time, eliminating the need for IT intervention. Find files amongst billions or more in milliseconds, no matter when they existed, at any point in time using wildcard searches. Mediaflux Point in Time instantly changes the shape of the file system to appear as it was at the selected points in time, achieving an RTO of zero.



Point in Time secures data at scale, rolls back ransomware, and makes data recovery expedient.

Point in Time ensures data resilience at scale.

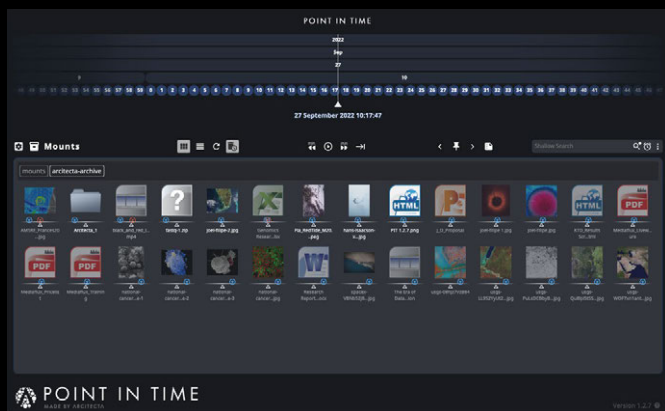
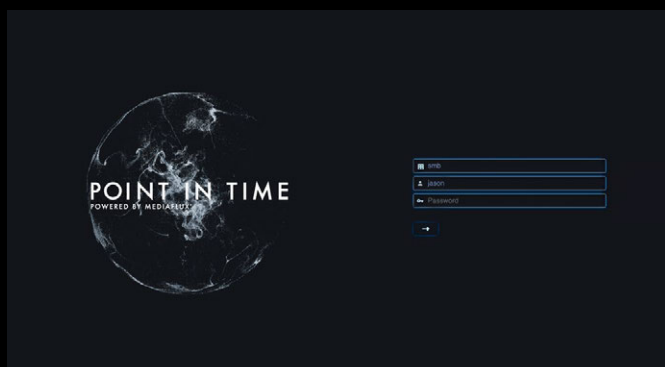
By uniting the file system and the power of XODB®, a database for trillions of datums, Mediaflux Point in Time provides inline data protection that actively records every significant change that occurs for every file, in real-time. The detailed metadata trail provides provenance for forensic analysis and reconstruction.

Roll back ransomware attacks.

Mediaflux Point in Time's real-time recording of changes means that whenever a ransomware attack occurs, enterprises are already protected – no matter the scale of the attack, simply rewind to unwind the damage. No decryption keys needed. Unwind ransomware events that happen over a long period with a single operation.

Find what you need, fast.

The user-friendly interface empowers end-users to find and recover things quickly. Simply scroll the time-slider to roll back the entire file system structure to a point before the loss occurred. Users don't need to know the precise time of loss to find their data.



Do it Yourself

The easy to use GUI puts users in the driver's seat without the help of an IT administrator. Anyone can simply connect with their standard credentials to access a single, global namespace and find their data.

Powerful and simple wildcard searches find files in milliseconds, no matter when and where they existed – even when there are billions of files or more.

Press a button to recover your data – immediately accessible with the storage locations completely transparent – through standard file shares.

Mediaflux Point in Time provides:

- Security and resilience at scale.
- Every important file change is recorded throughout all time, and then accessible.
- Easy navigation to any point in time, and the ability to see files at all points in time.
- A view of files from the past, present, and future, simply by moving a slider. Apply that point in time to your file shares at the press of a button.

Experience the Point in Time difference.

Mediaflux Point in Time redefines resilient enterprise storage – the convergence of storage and data resilience in the one system.

There are flexible deployment options. It can run on dedicated hardware, virtual machines, containers, cloud infrastructure or any combination of these. Mediaflux Point in Time supports multi-node scaling to match your performance requirements. It can span hybrid on-premises and cloud deployments.

Virtualize storage

Unify disparate storage (disk, object, tape, cloud) under a single namespace and a single mount-point or share – all with Mediaflux Point in Time. Dynamically add or remove storage to match changing needs. Mix technologies of different types from different vendors at any time to maximize cost and resilience objectives.

Only With Mediaflux

Mediaflux Point in Time is a core component of every Mediaflux system. It adds temporal resilience to any form of data including structured and unstructured, file and object. Achieve an RPO and RTO without compromise – at scale.

Interested?

Talk to our team at: <http://www.arcitecta.com/conversation/>