


PoINT Data Replicator

**Scalable File and
Object Backup &
Replication**



PoINT Data Replicator



Contents


4	Introduction
5	Product Overview
	Key Features
6	Concept & Architecture
	Multi-Threaded Processing
	Integrated Database for Job Optimization
	Background Operation
	Use of Messaging Services for Change Tracking
	Optimized for S3-to-Tape with PoINT Archival Gateway
11	Functionality
	S3-to-S3 Object Replication
	File-to-S3 Replication
	Filter-Controlled Replication
	Multi-Part Object Support
	Automatic Verification
	Database-Supported Synchronization
	Restore Process
13	Administration & Management
	Graphical User Interface (Web GUI)
	REST API
	Command Line Interface
	Logging & Event Reporting
14	Technical Information
	Operating Systems
	Supported Systems
15	Use Cases
	Protection of Cloud/Object Storage with Continuous Backup & Replication
	Cloud Repatriation
	Automatic Bucket Synchronization
	Migration from Legacy NAS to Object Storage
	Object/Cloud Backup & Replication to Tape with PoINT Archival Gateway

Introduction

The rising popularity of S3 object storage comes with challenges that are often overlooked or postponed for solution. Moving to the cloud can also have consequences that are rarely considered in advance. The same applies to traditional file systems that are to be migrated to S3 storage for reasons of capacity and scalability. Various use cases related to cloud and object storage are creating a need for a scalable software solution that can easily handle object storage replication, migration and backup.

PoINT Data Replicator is an easy to install, scalable and high performance software solution that performs replication processes from file or object storage systems to any S3-capable storage system. It is highly versatile, offering solutions for a wide range of use cases related to data protection and availability in the object storage sector.

The software uses standardized formats and APIs. It provides an extensive range of features that meet high standards for data



The key starting point for addressing these challenges is creating a 1:1 copy of the data, whether it is on an additional storage system or an actual remote site in the sense of geographic redundancy; or whether it is from a file-based source system to an object-based target system, while including all the original information the file system carried.

protection and backup. Additionally, PoINT Data Replicator allows for easy, seamless access to replicated data. This helps companies strengthen their data resilience and ensure business continuity.

This Technical White Paper describes the various use cases for PoINT Data Replicator and provides a detailed technical description.

Product Overview


PoINT Data Replicator is a software solution that provides data replication for a wide range of use cases. The use cases also comprise migration and backup applications. The product performs S3-to-S3 replication processes from any S3-capable source to any S3-capable target system. PoINT Data Replicator also offers File-to-S3 replication to migrate data from file systems to any target storage system with an S3 interface.

It uses standardized formats and APIs and supports any storage source and target system independent from the manufacturer.

The basic functionality of PoINT Data Replicator includes automatic verification, database supported synchronization and filter-controlled replication.

Key Features


- File-to-S3 and S3-to-S3 backup & replication
- Optimal throughput due to scaling and multi-threading
- Continuous backup & replication with S3 Bucket Notification
- Local database for fast retries without rescanning the source tree
- Copying and comparing of object metadata, access control list, tags and object locks
- Copying of objects filtered by time stamps
- Support for all mountable file systems (e.g. CIFS and NFS)
- Support for all S3-compatible cloud and object stores



Concept & Architecture

PoINT Data Replicator is a scalable solution consisting of several modules. The GUI, API, Scanner and Database Services run on the Coordinator Node, while the Copy Services are executed on the Data Mover Nodes.

The Coordinator Node manages the configured replication tasks, scans data sources and generates replication jobs. The replication jobs will be executed by one of the Data Mover Nodes. This allows scaling PoINT Data Replicator for higher performance needs just by adding more Data Mover Nodes.




The Copy Services can run on one or on multiple Data Mover Nodes. Increasing the number of Data Mover Nodes improves the performance of the replication process.

For smaller configurations, it is still possible to use a Single Node installation which combines Coordinator and Data Mover in one system service.

Multi-Threaded Processing

PoINT Data Replicator utilizes multi-threaded execution to perform multiple read and write operations in parallel. The number of threads per Data Mover determines how many objects or object parts are processed concurrently. To prevent bottlenecks during large-scale data movements, overall performance can be flexibly scaled: high parallelism is achieved through an architecture supporting multiple Data Mover Nodes and configurable worker processes. Additionally, adaptive memory limit controls ensure that RAM usage remains managed during parallel transfers.




Integrated Database for Job Optimization

The Coordinator Service maintains an integrated local database responsible for efficient job management and optimization. By storing the enumeration results of the source system in the database, execution routines (job re-runs) are significantly optimized. Failed transfers can be retried directly via Retry Jobs without requiring a full rescan of the source system. The database also tracks the status of copied objects, avoiding redundant checks or unnecessary re-transfers.

Background Operation


The components of PoINT Data Replicator run as independent system background services (managed via `systemctl` on Linux or as System Services on Windows). Beyond manual job execution, the system supports scheduled, automated workflows. It also offers automatic background bucket synchronization. Once configured, this feature handles the (delayed) deletion of objects on the target storage after they are removed from the source.



Use of Messaging Services for Change Tracking

To enable continuous change tracking, PoINT Data Replicator integrates with messaging services such as Apache Kafka (e.g. together with IBM COS, MinIO, or Scality RING) and Amazon SQS. Whenever objects are created

or modified in the source system, PoINT Data Replicator receives real-time notifications. This enables ongoing replication of changed data, eliminating the need for time-consuming, full scans of the source system.



Optimized for S3-to-Tape with PoINT Archival Gateway


PoINT Data Replicator is specifically optimized for seamless integration with PoINT Archival Gateway, enabling high-performance object replication directly onto tape storage. By leveraging the native S3 interface of PoINT Archival Gateway, PoINT Data Replicator allows organizations to automatically tier or replicate cold S3 objects to cost-effective, high-density tape media.

Key optimizations for S3-to-Tape with PoINT Archival Gateway include:

- **Tape-Aware Data Streaming:** When reading from tape storage managed by PoINT Archival Gateway, PoINT Data Replicator leverages the “Read in Tape Order” feature to optimize

restore performance. PoINT Data Replicator queries the physical tape positions of targeted objects directly from PoINT Archival Gateway and sequences read requests accordingly. This minimizes tape repositioning as well as seek times and avoids frequent tape drive start-stop operations (shoe-shining), thus substantially increasing overall restore throughput.


- **Direct S3 Access:** Data replicated to tape media via PoINT Archival Gateway remains natively accessible through standard S3 API calls. This ensures that data stored on tape can be directly queried and restored without proprietary lock-in.



Functionality


S3-to-S3 Object Replication

PoINT Data Replicator allows users to replicate S3 objects from any source to any target storage system, provided both systems are S3-compatible. They can be public clouds or on-prem object storage, for example. The S3-to-S3 function replicates objects, including their metadata, between S3 buckets, regardless of whether the source and target buckets are located on the same system or on physically different systems, providing different storage classes.



File-to-S3 Replication

The File-to-S3 replication function migrates file-based data from any file system to any storage system with an S3 interface. This function converts the original file path into an object key. Therefore, after replication, standard S3 browsers will display the same structure as the original file directory.



Filter-Controlled Replication

PoINT Data Replicator includes flexible, task-level filtering options to precisely control the replication process. Administrators can define criteria within the web-based UI to include or exclude specific data sets:

- **Time-Based Filtering:** Restricts replication using absolute or relative object modification timestamps.
- **Name & Path Filtering:** Uses prefixes, wildcards, or folder paths to target specific files or directories.
- **Size Filtering:** Defines minimum or maximum object size thresholds to skip irrelevant data.

Multi-Part Object Support

When replicating large datasets to S3-compatible target storage systems, handling large files efficiently and reliably is a crucial operational requirement. PoINT Data Replicator addresses this through robust support for Multi-Part Uploads (MPU).

Users can configure the MPU part size parameter on a per-task basis. Files or objects that exceed the configured threshold are automatically uploaded using the multipart transfer mode. Setting the part size parameter to 0 retains the original part sizes used on the source storage.

Automatic Verification

To ensure complete data integrity during replication, PoINT Data Replicator includes an automatic verification function using target-side checksum validation.

Once configured, PoINT Data Replicator calculates a checksum for the data using the selected algorithm or dynamically determines the algorithm based on the headers of the source object. Then, it injects this value into the data upload request headers.

Upon receipt, the target S3-compatible object storage independently calculates its own checksum of the incoming data stream and validates it against the provided checksum header before finalizing the write operation.

Database-Supported Synchronization

To ensure high performance and efficient data replication, PoINT Data Replicator utilizes an integrated database hosted on the Coordinator Node. Rather than relying on time-consuming full scans of the target storage system during routine synchronization or job re-runs, the software tracks object states and metadata in the integrated database.

Key technical advantages of this feature include:

- **Optimized Job Execution & Fast Retries:** By referencing the local database, PoINT Data Replicator can execute job re-runs and quickly retry failed objects without rescanning the entire source file tree or target bucket.
- **Automated Source/Target Synchronization:** The software continuously tracks object states, enabling automatic background synchronization.
- **Managed Deletion Handling:** Changes on the source system, including deletions, are reflected in the database. Deletions on the target storage can be executed immediately or with a configurable delay, preventing accidental data loss.
- **Compare & Resync Capabilities:** To account for out-of-band changes (e.g., data added bypassing PoINT Data Replicator), an explicit Compare Job with a database update option can be executed to reconcile target storage states with the local database.

Restore Process

PoINT Data Replicator provides flexible restore capabilities designed to restore data from S3 object storage back to a target environment. PoINT Data Replicator supports

two distinct restore functions depending on the recovery scenario: restoring the latest production state or recovering from a ransomware attack or accidental deletion.

- **Current Status Restore:** The Current Status Restore workflow recovers data to its most recent state. PoINT Data Replicator evaluates the current state and restores the latest active version of each object. Data can be restored to an S3-compatible object storage system or directly to a file system (CIFS/NFS). When restoring to a file system, PoINT Data Replicator preserves the original directory paths and file structures by using the object keys. When restoring to S3 object storage, custom metadata, Access Control Lists (ACLs), tags, and object locks can be restored along with the object itself.
- **Point-in-Time Restore:** The Point-in-Time Restore feature provides flexibility and security when restoring data. It allows restoring data based on a selectable date. Thus, for example, a targeted recovery can be performed after a ransomware attack, enabling a fast return to an uninfected data set. This process relies on S3 versioning enabled on the backup target (source of the restore). By evaluating historical object versions, PoINT Data Replicator retrieves the exact version of each object that was current at the specified snapshot timestamp.


Administration & Management

PoINT Data Replicator offers a flexible and comprehensive administration layer. System administration can be performed through an intuitive graphical web interface, a REST API, or an independent command-line tool (CLI).

Graphical User Interface (Web GUI)

The primary interface for system configuration, monitoring, and task administration is the Web GUI. Key capabilities of the Web GUI include:

- **Task Management & Workflows:** Allows administrators to configure, clone, sort, and manage replication tasks. Tasks support replication modes including S3-to-S3, File-to-S3, and point-in-time Restore tasks.
- **Task Filters & Search:** Features built-in filters to display tasks based on status (e.g., Running Jobs, Recently Run, Next Scheduled, and Jobs With Failures) alongside a keyword search bar for target/source paths and aliases.
- **Data Mover Control:** Provides a real-time overview of connected Data Mover nodes, current processing throughput, worker assignment, queued partial jobs, and node approval/rejection management.
- **Storage Configuration Helpers:** Enables inline browsing of file system paths, S3 buckets, and prefixes, as well as copying and pasting S3 configurations across tasks.
- **Global Settings & Scheduling:** Controls job scheduling parameters, SMTP email notifications, global log verbosity, certificate options, and security key management.



REST API

For integration into standard IT management tools, custom scripts, and orchestration workflows, PoINT Data Replicator provides a RESTful API. The REST API endpoint is hosted directly by the PoINT Data Replicator Coordinator Service. It covers full product functionality, enabling programmatic creation, execution, and monitoring of replication jobs, querying of task statuses, and configuration updates.

Command Line Interface

In addition to central service-managed tasks, PoINT Data Replicator includes an independent command-line program. The S3 Copy Tool operates independently of the Coordinator Service and does not require pre-configured GUI tasks. It enables swift copy or restore operations directly from the command line. Parameters can be specified via command-line arguments or saved into profile sections within a user

configuration file, allowing pre-defined defaults to be overridden on demand. The S3 Copy Tool supports file list input files, glob pattern include/exclude filters, multipart upload chunk size selection, concurrency tuning, and point-in-time sync rules.

Logging & Event Reporting

PoINT Data Replicator incorporates comprehensive logging and alert mechanisms to provide operational visibility, facilitate auditing, and streamline troubleshooting. A notification icon on the Web GUI toolbar alerts administrators in real time to job executions and completions. Each task maintains a historical event log tracking start, stop, and failure events. Additionally automated email notifications can be configured to send alerts for job completions, failures, or change notification processing errors.

Technical Information

Operating Systems

- Windows
- Linux

Refer to the document “PoINT Data Replicator - System Requirements” for detailed information.

Supported Systems

Sources

- Support for all mountable file systems (e.g. CIFS and NFS)
- Support for all S3-compatible cloud and object stores

Targets

- Storage systems supporting S3 for ingest


Use Cases

Protection of Cloud/Object Storage with Continuous Backup & Replication

A common use case for PoINT Data Replicator is the protection of cloud or object storage systems. Data stored at public cloud providers or on local object storage is typically not protected by a backup. Replicating this data to a physically separated S3 bucket protects it from human error, ransomware, malware, or technology failure. In principle, any S3-capable system can be used by PoINT Data Replicator as target system. This can be an on-premises object storage or an additional cloud storage.

The automated S3 Bucket Notification enables the continuous backup of very large data sets. Only new objects are identified and replicated, which makes the replication process more efficient and avoids long scan times on the source system.


The S3 Bucket Notification of PoINT Data Replicator works in combination with queuing services like Kafka and SQS.




Cloud Repatriation

Cloud repatriation refers to the process of moving data, applications, and workloads from a public cloud back to your on-prem infrastructure. The primary reasons for shifting away from the cloud include concerns about data governance, data security, and especially financial considerations.

PoINT Data Replicator is designed for such migration projects. With the S3-to-S3 replication feature, data is seamlessly migrated from the cloud to on-premises storage systems, preserving the native format and metadata. Migrated data can be accessed directly via S3 on object storage.




Data retrieval from cloud-based S3 Glacier buckets is supported by PoINT Data Replicator in conjunction with PoINT Archival Gateway.



Automatic Bucket Synchronization

PoINT Data Replicator performs automatic bucket synchronization. This process takes place in the background. Optionally, objects that have been deleted from the source can also be deleted from the destination storage (immediately or later).

With automatic bucket synchronization, companies can increase redundancy, reliability, and resilience in the event of data loss as the replicated data is available for rapid disaster recovery.



Migration from Legacy NAS to Object Storage


As organizations face exponential data growth, traditional Network Attached Storage (NAS) systems often reach their limits in terms of scalability and flexibility. Migrating data from legacy NAS architectures to cloud-native object storage can help achieve scalability, reduce storage costs, and ensure long-term data preservation. However, such migrations present the challenge of transitioning from file-based protocols (NFS/SMB) to an object-based API (S3).

migration processes, facilitating data migration from legacy NAS devices to modern object storage systems.

Key features:

- Multi-threaded replication ensures optimized throughput.
- Metadata preservation accurately maps file attributes, access control lists (ACLs), and folder structures to the target object storage metadata framework.
- Continuous verification guarantees data integrity post-migration.

PoINT Data Replicator executes File-to-S3




Object/Cloud Backup & Replication with PoINT Archival Gateway

Tape-based S3 object storage as a target system is the optimal solution for protecting cloud/object storage. Air-gapped tape provides malware protection. Using a standardized S3 REST API, large data sets can be stored and restored quickly. PoINT Archival Gateway is a scalable and high-performance S3-to-Tape solution to store and archive large amounts of data on tape media. PoINT Data Replicator and PoINT Archival Gateway complement each other and work together perfectly.

Since PoINT Archival Gateway provides an object storage for tape, applications and users can directly access the replicated data. Long-term retrieval to the original cloud/object storage is not necessary.

PoINT Archival Gateway supports a wide range of tape products. A list of supported tape libraries is available on the PoINT web page at www.point.de.



PoINT Software & Systems is specialized in the development and distribution of software products for storage, management and long-term archiving of data using all available mass storage technologies like hard disks/flash, magnetic tapes, optical media, object store and cloud storage. We work jointly together with leading manufacturers of storage systems. Thus, we can offer an early support of innovative storage technologies. Furthermore, we plan entire storage solutions and provide consultancy with our long-term and versatile expertise.

PoINT products allow efficient usage of storage systems and help to reduce costs and issues caused by data growth. The software solutions fulfil compliance and archiving requirements and provide independence from storage technologies and vendors. PoINT products are distributed by our partners world-wide and have been proven in more than two million installations. Our customers include many well-known companies from different industries, who comply with our solutions their complex demands by providing the necessary reliability and perfection.

PoINT Software & Systems GmbH believes the information in this publication is accurate as of its publication date. Such information is subject to change without notice. PoINT Software & Systems GmbH is not responsible for any inadvertent errors. The PoINT logo and the brand name "S3-to-Tape" are registered trademarks of PoINT Software & Systems GmbH. All other trademarks belong to their respective owners. Software and documentation are available in English. © 2026 PoINT Software & Systems GmbH All rights reserved. No portions of this document may be reproduced without prior written consent of PoINT Software & Systems GmbH. Printed in Germany September 2026 (wp-pdr-scalable-file-and-object-replication_20260923).



POINT Software & Systems GmbH
Eiserfelder Straße 316
57080 Siegen, Germany
+49 271 3841-0
info@point.de

www.point.de