

TM


S3-to-Tape




S3 Object Storage on Tape



PoINT
Archival
Gateway



PoINT Archival Gateway



Contents

5	—	Introduction
7	—	Product Overview Key Features
9	—	Concept Interface Nodes Database Nodes
11	—	Functionality High Scalability Geo Distribution and Redundancy S3 Rest API S3 Compatible Storage Classes Single Namespace Tape-only Disk/Flash Systems S3 Lifecycle Policies Automatic Disk/Tape Replication Direct Tape Access Media Management Erasure Coding LTFS Support PoINT Data Replicator Logging and Monitoring Administration
18	—	Installation Options Compact Edition Enterprise Edition
21	—	Supported Tape Systems
22	—	Use Cases Object Storage Replication & Backup Backup Applications Tiering/ILM for On-Prem Object Storage HSM & Archiving S3 Capable Applications

A valuable good —


Introduction

The growth of unstructured data, high energy costs, data sovereignty, as well as cybercrime attacks are major challenges to enterprise IT infrastructure. These challenges cannot be solved either technically or economically by using storage systems based on hard disks or flash memory.

Most of the stored data is unstructured and inactive. Even though the data is rarely used, it still needs to be retained for

business reasons and because of regulatory requirements.

However, the exponential growth of cold data is driving up storage costs and energy consumption of the storage infrastructure. It also increases space requirements and carbon footprint. It makes no sense to store large amounts of inactive data on disk or flash-based systems.



The solution to this challenge is a software-based object storage that supports tape as a target storage system. Object storage systems meet the highest demands for scalability, cost-efficiency, reliability and high availability. As data volumes continue to grow, standardized tape technology is the recommended storage medium: Tape is the only storage medium that provides the necessary storage capacity at an acceptable cost for the data volumes predicted for the future.

In addition to its cost effectiveness, tape technology offers a number of other advantages. With its WORM properties, it meets archiving requirements for the immutability of archived data. In addition, removable media such as tape creates an “air gap” to protect against malware. This makes tape technology ideal for data protection and archiving.

The combination of object storage and tape provides an innovative and economical solution to the challenges of data growth. Compared to other storage media, tape also has the potential to meet the required capacity increases.

PoINT Archival Gateway is a software-defined, scalable object storage system. It is designed to store and manage large amounts of data with high performance directly on tape libraries. PoINT Archival Gateway provides a compatible S3 REST API and supports all common tape libraries. Temporary storage and formatting on expensive disks is not necessary. PoINT Archival Gateway thereby combines the convenience of S3 with the cost-effectiveness of tape.

PoINT Archival Gateway also offers the possibility to integrate an Amazon S3 compatible disk storage class. In this way, access times can be reduced, for example, if this is required by the specific use case. PoINT Archival Gateway integrates different storage technologies in a homogeneous architecture and uses their respective properties for internal tiering based on S3 lifecycle policies and automatic replication.

This Technical White Paper describes the use cases for PoINT Archival Gateway and provides a detailed technical description of the S3-to-Tape™ solution.

Product Overview

PoINT Archival Gateway is a high performance, scalable, software-based object storage solution that was developed for the high performance storage and management of large amounts of data. PoINT Archival Gateway connects S3 capable client applications or client systems with tape libraries as target storage.



PoINT Archival Gateway uses standardized interfaces and protocols like the S3 REST API. The basic functions of PoINT Archival Gateway include user, data and storage management as well as access control, logging and monitoring. PoINT Archival Gateway allows the optional integration of an additional disk/flash based storage class to meet the demands of use cases that require fast data access. Internal tiering using the standardized S3 Lifecycle Policies ensures optimized data and storage management.

— KEY FEATURES

- High data throughput thanks to parallelism and Erasure Coding
- High availability with redundant server nodes
- High scalability including load balancing
- Geo-distribution with automatic site failover and resynchronization
- Single namespace across multiple storage classes (disk/flash and tape)
- Compatibility with Amazon S3 Standard and S3 Glacier storage classes
- LTO and 3592 tape drive support
- Tape quality checking and monitoring, space reclamation and migration
- Object Versioning
- Data protection through Erasure Coding, Object Locking, authentication and encryption
- Self-monitoring, reporting, alerting and SNMP
- User management based on domain services (AD, LDAP)



Concept

The PoINT Archival Gateway software consists of two software services or packages, which can be installed either on different server systems as Enterprise Edition or on one server system as Compact Edition.

— INTERFACE NODES

An Interface Node (IFN) is primarily the communication partner for client applications or systems. Using dedicated interface modules, it provides the S3 REST API for client applications and systems to store and read objects, and transfers data between client applications or systems and the storage classes (disk and tape). The S3 REST API is a web service that supports the HTTP protocol as well as S3 specific protocol elements.

Specifically, IFNs provide the following modules and services:

- HTTP Service Module (S3 REST API)
- Data buffering module
- Data encoding module (Erasure Coding, hashing, encryption, ...)
- Driver module for tape drives
- Metadata caching module (e.g. caching of object metadata and configuration data)
- Communication module for exchanging metadata with the database node

— DATABASE NODES

A Database Node (DBN) primarily provides central database services for the IFNs. The database contains the table of stored objects, e.g. the object keys and metadata, the storage location of the object data on the storage classes and the configuration and maintenance data of PoINT Archival Gateway. The database also stores data from logging and monitoring processes and provides corresponding auditing services and log files. Further central services are the system configuration (Admin GUI) as well as the management and control modules for the tape storage class (tape libraries), which are therefore also located on this node.


For operation PoINT Archival Gateway requires at least one running IFN and one running DBN. In the Compact Edition both software packages can be installed on one server.

PoINT Archival Gateway meets the highest requirements regarding performance, availability and scalability. The basis for this is the fully scalable and redundant design, which provides scalable performance and redundancy levels for the service as well as for the data entities. Due to the server hardware and operating system, the scalability and availability of a single node is limited. Therefore, PoINT

Archival Gateway Enterprise Edition supports the operation of multiple DBNs and IFNs within a single installation.

In this way, PoINT Archival Gateway builds clusters to increase performance (i.e. load balancing) and availability (i.e. failover and redundancy).

After an error situation, PoINT Archival Gateway automatically performs all necessary steps to restore the operability and consistency of a cluster node.




Functionality

— HIGH SCALABILITY

PoINT Archival Gateway is scalable by adding server nodes. The internal Interface Nodes cooperate with each other (e.g. for load balancing) and offer a highly scalable S3 REST Web Service, which allows almost unlimited parallelization with very high data transfer rates and S3 performance values.

- **50+ billion** objects per S3 bucket
- **10,000+** S3 buckets
- **100,000+** parallel HTTP connections
- **1,000+** parallel S3 applications



— GEO DISTRIBUTION AND REDUNDANCY

PoINT Archival Gateway provides redundancy at system and data level. Interface Nodes can replace each other and Database Nodes offer synchronous replication and failover. Flexibly selectable erasure coding methods ensure the protection of stored data.

Fail-safe operation is also possible across multiple locations. It is robust against high WAN latencies. In an installation spread across two locations, automatic replication and resynchronization after a site failure are supported.

— S3 REST API

PoINT Archival Gateway offers a standardized and Amazon compatible S3 REST API. This makes the software suitable for the growing number of applications that use object-based storage by connecting to S3 REST. In addition, the S3 Glacier commands are supported, which especially take into account the high latencies of the tape storage class.

— S3 COMPATIBLE STORAGE CLASSES

Disk/flash and tape storage systems are integrated as Amazon S3 compatible storage classes and can be configured accordingly. Applications can access the different storage classes using standardized S3 commands.

PoINT Storage Class	Corresponding Amazon Storage Class
Disk/Flash	S3 Standard
Tape	S3 Glacier Flexible Retrieval and S3 Glacier Deep Archive


PoINT Archival Gateway allows a very flexible configuration of storage classes. For example, if the data volumes are still small, you can start with the disk/flash storage class. Later you can add a tape storage class to move inactive data from the disk class to tape media. Similarly, you can start with a tape-only configuration and optionally add a disk storage class if, for example, low-latency read requirements become important. Using both storage classes simultaneously, for example to create offline media, is also a practical use case.

— SINGLE NAMESPACE

The disk/flash and tape storage classes are available together under one interface as a “single namespace”. This significantly simplifies the use of different storage classes for S3 applications.


— TAPE-ONLY

In the tape-only configuration, tape-specific applications can write to and read from the media directly. No disk caches are required. This allows for the most effective and high performance integration of tape.



— DISK/FLASH SYSTEMS

All NAS products can be configured in PoINT Archival Gateway as disk/flash systems in the disk storage class. Please contact PoINT regarding the support of disk-/flash-based object storage systems.




— S3 LIFECYCLE POLICIES

PoINT Archival Gateway is compatible with the Amazon S3 Lifecycle Policies. This allows to move data between storage classes based on individual policies. For example, it can be specified that data is first stored on the disk/flash storage class for a certain period of time and then automatically moved to the tape storage class.



AUTOMATIC DISK/TAPE REPLICATION

With automatic replication, data can be stored simultaneously on disk/flash and tape media. On the one hand, this enables media break; on the other hand, the “air gap” of the tape media provides additional data security.



DIRECT TAPE ACCESS

PoINT Archival Gateway supports direct writing to and reading from tape media. If applications are designed for high data throughput, they can write objects directly to the tape storage class. The objects do not have to be stored on the disk/flash storage class first. When reading, this feature avoids cumbersome and lengthy restore processes to disk. The data is delivered directly to the reading application.

MEDIA MANAGEMENT

PoINT Archival Gateway also manages tape media that have been exported from a library and are therefore “offline.” Offline media are listed in the Admin GUI, including the name of the library where the media was last online. If a client application accesses data on offline media, it receives an error code in response. At the same time, a request list is created and displayed in the Admin GUI. Corresponding SNMP notifications can also be configured. With this information, an operator can import the required media. After import, the request list is updated automatically. Furthermore, an operator can manage the list, e.g., delete entries that are to be ignored.

The Tape Quality Check performs automatic checks for data integrity and readability. In critical conditions, it performs migrations to new tapes.


The Space Reclamation function consolidates storage space by copying tapes that are fragmented with logically deleted objects.

Functions for migration to new tape formats, e.g., to a more recent LTO generation, are also integrated.

ERASURE CODING

Data protection on the tape media is provided by Erasure Coding. This process stores blocks of data redundantly on multiple media. This means that even if one medium fails, the data will not be lost. PoINT Archival Gateway supports the Erasure Code (EC) rates 1/2, 1/3, 1/4, 2/3, 2/4 and 3/4. In combination with Erasure Coding, data protection and redundancy can be further increased, e.g. by using two, three or four tape media in parallel in the tape storage class. Such a combination of

multiple media is called a Protected Volume Array. A Protected Volume Array consisting of N tape media can also extend over N tape libraries. The EC rates 1/2, 1/3, 1/4 indicate the automatic creation of copies. For the tape storage class, this means that multiple tape copies can be created (even in different libraries). Throughput rates can be significantly increased with EC rates that distribute data across multiple media (EC 2/3, 2/4, and 3/4).



LTFS SUPPORT

PoINT Archival Gateway offers support for the LTFS format. When writing data sets to tape, users can choose between PoINT’s optimized proprietary data format and the standardized LTFS format. With LTFS, data objects are visible and readable as individual files directly,

without any intermediary steps such as containers. Alternatively, PoINT’s proprietary tape format offers maximum performance and efficiency, as well as optional Erasure Coding. Both formats can also be used in combination for different data sets.


POINT DATA REPLICATOR

PoINT Archival Gateway is optimized for integration with the product PoINT Data Replicator. This integrated solution is designed for use cases like protection of cloud and object storage and cloud repatriation.

PoINT Data Replicator supports S3-to-S3 and File-to-S3 replication modes. This enables replication of public clouds, on-premises object storage, and CIFS/NFS-

based file servers to tape. The S3 replication includes metadata, tags, object lock, ACLs and multi-part objects. The solution also supports continuous S3-to-S3 backup & replication in combination with messaging services such as Kafka and SQS.

Data restore from tape is optimized between PoINT Data Replicator and PoINT Archival Gateway, ensuring that read requests to PoINT Archival Gateway are sorted according to their location on tape.



LOGGING AND MONITORING

PoINT Archival Gateway also supports Data Access Audit Logs and Security Audit Logs. Access to the data objects of an object repository is logged in assigned access audit log files. These log records contain the identifier, timestamp and access type as well as the identifier of the principal who performed the access.

All manager logon activity is also logged in a protected security audit log file, as are


any changes made by security managers or otherwise related to security settings. This log file contains information about the principals and details of the changes made.

Standard monitoring systems can be connected via the supported Simple Network Management Protocol (SNMP). This enables complete monitoring of all important functions and services. Warnings and error messages are issued proactively.

ADMINISTRATION

The administration is done via a web-based Admin GUI, which is provided by PoINT Archival Gateway via an HTTP service. In addition to the Admin GUI, a REST Admin

API with Swagger support is available. This API allows the integration of administrative functions into custom applications.




Installation Options

PoINT Archival Gateway supports Windows as well as Linux operating systems. The software can be installed on several servers in the Enterprise Edition or on one server in the Compact Edition.

COMPACT EDITION

PoINT Archival Gateway - Compact Edition allows the installation of only one database

instance and only one interface service on the same server system. In this case the database and the functionalities of the interface service are combined in a single compact service module. This optimizes overall resource consumption and reduces the overhead of inter-service communication.




The Compact Edition can also be installed as a failover cluster, i.e. on two server systems in two different locations.

ENTERPRISE EDITION

In general, the software packages and thus the services of PoINT Archival Gateway should be installed on separate server systems (Interface Nodes and Database Nodes). The distribution on different server systems ensures maximum scalability, availability and performance.

This installation option is offered by PoINT Archival Gateway - Enterprise Edition.

The following figure shows the architecture of PoINT Archival Gateway - Enterprise Edition and the network connections used by this solution based on an exemplary installation.





Supported Tape Systems

The design of PoINT Archival Gateway is independent of the storage technology and storage systems used. This means that the user can select and replace such systems at will.

PoINT Archival Gateway supports a wide range of tape libraries. PoINT Archival Gateway integrates tape libraries directly, no additional drivers or software products are required.

PoINT Archival Gateway supports LTO and IBM 3592 tape drives. Mixed drive configurations

are also supported. A PoINT Archival Gateway installation supports up to 12 tape libraries with a maximum of 384 tape drives.

The following table lists the currently supported tape systems (loaders and libraries).


Please contact PoINT Software & Systems if you need support for tape systems not currently listed.

Vendor	Product
actidata	actiLib LTO Family
ADIC	Scalar Series
BDT	FlexStor, MultiStak, MultiStor, Orion MC6
Cristie	GigaStream T Series
COMBACK	TL Series
Fujitsu	Eternus Series
HP/HPE	MSL Series
IBM	TS Series, Diamondback
Overland Tandberg	NEO Series
Qualstar	Q Series
Quantum	Scalar Series
Spectra Logic	T Series, Cube

Use Cases

PoINT Archival Gateway is already certified and validated by many technology partners. All applications that support either S3 Standard or S3 Glacier can use PoINT Archival Gateway to store data on tape

systems and benefit from the associated advantages. The following graphic shows an overview of possible use cases including certified product solutions.



— OBJECT STORAGE REPLICATION & BACKUP

Cloud and object data must also be backed up. For example, cloud storage can be vulnerable to malware attacks. This also applies to data on on-premises object storage. Therefore, a backup of cloud and object data is essential. PoINT Archival Gateway offers the possibility to backup cloud and object data in native S3 format to tape. The object data is stored in the

same structure as the original data. In case of an emergency, the data can be accessed directly via the S3 interface of PoINT Archival Gateway. A time-consuming restore process is not necessary.

— BACKUP APPLICATIONS

Backup applications, e.g. from Commvault, Veeam, Rubrik or Cohesity, can use PoINT Archival Gateway as backup storage target

and thus include tape media in the backup process. Within the technology partnership with Veeam, PoINT Archival Gateway is validated for Veeam.

— TIERING/ILM FOR ON-PREM OBJECT STORAGE

Public cloud storage providers, such as AWS and Microsoft Azure, offer different storage classes with different performance characteristics. On-prem object storage used as a private cloud, on the other hand, offers only one class of storage, usually hard disk-based. This is inefficient because both active and inactive data is stored on the same storage technology. The combination of disk-based object storage with PoINT Archival Gateway enables the tiering of inactive data from disk to tape. Many object storage systems like Cloudian HyperStore and NetApp StorageGRID already offer integrated ILM functionality.

— HSM & ARCHIVING


With the help of WORM functionality and integrated retention management, PoINT Archival Gateway fulfills archiving and compliance requirements. The stored data is not only protected against accidental deletion, but also against intentional changes (e.g. ransomware attacks). PoINT Archival Gateway enables long-term data

management to fulfill legal and corporate archiving requirements based on tape media. Retention policies can be activated at the level of an object repository. These policies define how and when existing objects may be modified or deleted. Especially HSM and archiving applications can use PoINT Archival Gateway as target storage system for secure and compliant archiving.

— S3 CAPABLE APPLICATIONS

In principle, all S3 capable applications can use PoINT Archival Gateway as target system for tape integration. For example, many applications in the media and entertainment sector implement data storage via an S3 interface. PoINT Archival Gateway can be used as storage target for these solutions. Active data can initially be stored on the disk storage class. Later, when the data is no longer needed, it can be automatically moved to the tape storage class. Research applications, such as DNA sequencing, generate very large amounts of data. The different phases of data processing can be ideally covered with PoINT Archival Gateway. During the analysis phase, the research data remains accessible on the fast disk storage class and is automatically moved to the tape storage class for archiving after the analysis is completed.

PoINT Software & Systems



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-pag-s3-to-tape_e_20260925).



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

www.point.de