

PoINT Archival Gateway


Unified Object Storage
with Disk and Tape



Contents

- 5 — **Introduction**
- 6 — **Product Overview**
 - Key Features
- 7 — **Use Cases**
 - Backup Target
 - Broadcast
 - Research
 - S3 Archiving
 - Tiering/ILM for On-Prem Object Storage
 - Cloud/Object Storage Backup
- 9 — **Architecture and Design**
 - Interface Nodes
 - Database Nodes
- 11 — **Functionality**
 - S3 REST API
 - S3 Compatible Storage Classes
 - Single Namespace
 - Flexible Configuration of Storage Classes
 - Automatic Disk/Tape Replication
 - Offline Media Management
 - Lifecycle Policies
 - Direct Access to Tape
 - Erasure Coding
 - Administration and Logging
- 14 — **Installation Options**
 - Enterprise Edition
 - Compact Edition
- 15 — **Supported Storage Systems**
 - Disk Systems
 - Tape Systems

PoINT Archival Gateway



Introduction

The growth of unstructured data is one of the greatest challenges for the IT infrastructure of companies. Recent studies by various research institutes confirm this. Studies also show that a large part of the unstructured data is inactive. While this data is rarely used, it must be retained for business and legal reasons. However, the exponential growth of cold data drives up costs and energy consumption, and increases floor space requirements and the carbon footprint.

Today, inactive data is stored predominantly on hard disk-based storage systems in local data centers or with a cloud service provider. This is an expensive and very energy-intensive storage method.

The solution for this problem are software-defined object stores with standardized S3 REST API that integrate multiple storage classes such as disk and tape.

Object storage – like file and block storage – is a method of storing data.

Unlike file and block storage, however, object storage was developed to meet the highest demands for scalability, cost efficiency, reliability and high availability. Data storage is independent of the underlying storage technology, and data can be stored on different storage media in an object storage system.

A software-based object store that flexibly unifies disk and tape media in one system takes advantage of different storage technologies for internal tiering and automatic replication. In addition, removable media such as tape create an “air gap” to protect data.

PoINT Archival Gateway is a unified object storage in which several storage technologies with different quality characteristics – especially hard disk and tape – are integrated in a homogeneous architecture. Thus, the specific properties of the different technologies can be used optimally.

Storage Class	Purpose
Disk	Warm, frequently accessed data (low access times, within msec)
Tape	Cold, long-lived, archive data (direct access, within sec to minutes)

Product Overview

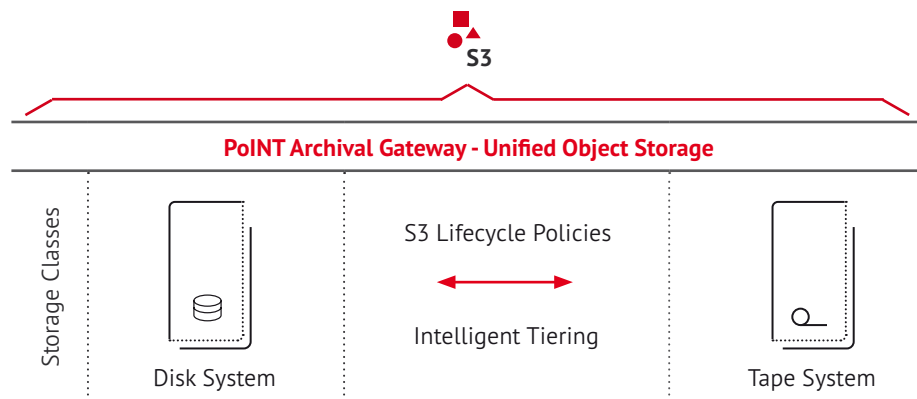
PoINT Archival Gateway is a software defined scalable object storage system that integrates the storage classes disk and tape. It is designed to store large amounts of data with high performance and provides services and functions necessary to meet enterprise archiving requirements.

PoINT Archival Gateway uses standardized interfaces and protocols, such as 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.

— KEY FEATURES

- High performance for data and object transfer rates
- High availability through redundant server nodes
- High scalability including load balancing
- Single namespace across multiple storage classes (disk and tape)
- S3 and S3 Glacier compatibility including lifecycle policies
- LTO and 3592 tape drives
- Object Versioning
- Data protection through erasure coding, object locking, authentication and encryption
- Self-monitoring, reporting and alerting
- User management based on domain services (AD, LDAP)



Use Cases

PoINT Archival Gateway meets the requirements of different use cases. With its standardized S3 REST API PoINT Archival Gateway can be used for all S3-enabled applications. Due to the multi-tier function with disk and tape PoINT Archival Gateway is especially suitable as secondary and archive storage to store and archive all kinds of “cold” data.

In the following, some exemplary use cases are listed which, in addition to the tape storage class, can especially use the disk storage class for data access with low access time:

— BACKUP TARGET

Backup applications e.g. from Commvault, Veeam, Rubrik or Cohesity can use PoINT Archival Gateway as backup storage target and thus include tape media into the backup process.

— BROADCAST

In the media and entertainment sector, many applications realize data storage via an S3 connection. PoINT Archival Gateway is suitable for these solutions as storage target. Active data can be stored on the disk storage class at first. Later, when they

are no longer needed productively, they can be moved automatically to the tape storage class.

— RESEARCH

In the research area, e.g. in DNA sequencing, very large amounts of data are generated. The different phases of data processing can be ideally covered by PoINT Archival Gateway. In the analysis phase, the research data remain accessible on the fast disk storage class and are automatically moved to the tape storage class for archiving after completion of the analysis.

— S3 ARCHIVING

With the help of WORM functionality and integrated retention management PoINT Archival Gateway fulfills archiving and compliance requirements. Therefore, the data is not only protected against unintentional deletion, but also against intentional changes (e.g. by ransomware attacks). Based on tape media, PoINT Archival Gateway enables long-term data management to fulfill legal and corporate archiving requirements. Thus, appropriate retention rules can be activated on the level of an object repository. These rules define how and when existing objects may be modified or deleted.

The following use cases are particularly suitable for the so-called tape-only configuration (i.e. without disk storage class). Further information can be found in the document “PoINT Archival Gateway – Tape-based Object Storage”.

— TIERING / ILM FOR ON-PREM OBJECT STORAGE

Public cloud storage providers, such as AWS and Microsoft Azure, provide different classes of storage with different performance characteristics. On-premises object storage used as a private cloud, on the other hand, offers only one, usually disk-based, storage class. This is inefficient because both active and inactive data are backed up on the same storage technology. The combination of disk-based object storage with PoINT Archival Gateway allows tiering of inactive data from disk to tape. Many object storage systems already offer integrated ILM functions.

— CLOUD / OBJECT STORAGE BACKUP

Cloud and object data must be protected by a backup. For example, cloud storage can also be affected by malware attacks. This also applies to data on on-prem object storage. Therefore, a backup of cloud and object data is indispensable. PoINT Archival Gateway offers the possibility to create backups of cloud and object data in native S3 format on tape. Thereby, object data is stored in the same structure as the original data. In case of emergency, the data can be accessed directly via the S3 interface of PoINT Archival Gateway. A time-consuming restore process is not necessary.

Architecture and Design

The software of PoINT Archival Gateway consists of two software packages, each of which can be installed on different server systems.

— INTERFACE NODES

An Interface Node (IFN) is primarily the communication partner for client applications or systems. By means of dedicated interface modules, it provides the S3 REST API for client applications and systems for storing and reading objects and transfers data between client applications or systems and the storage classes (disk and tape). The S3 REST API is to be understood as a web service that in addition to the HTTP protocol also supports S3-specific protocol elements.

In detail, IFNs provide the following modules and services:

- HTTP Service Module (i.e. S3 REST API)
- Data Buffering Module
- Data Encoding Module (e.g. erasure coding, hashing, encryption)
- Driver module for tape drives
- Metadata Caching Module (e.g. caching of object metadata and configuration data)
- Communication Module for the exchange of 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 location of the object data on the storage classes and the configuration and maintenance data of PoINT Archival Gateway. Additionally, the database stores data of logging and monitoring processes and provides corresponding auditing services and log files. Further central services are the system configuration (Admin GUI), as well as management and control modules for the tape storage class (tape libraries), which are therefore also located on this node.

In order to work, PoINT Archival Gateway requires at least one running IFN and one running DBN. In the Compact Edition both software packages can also be installed on one server.

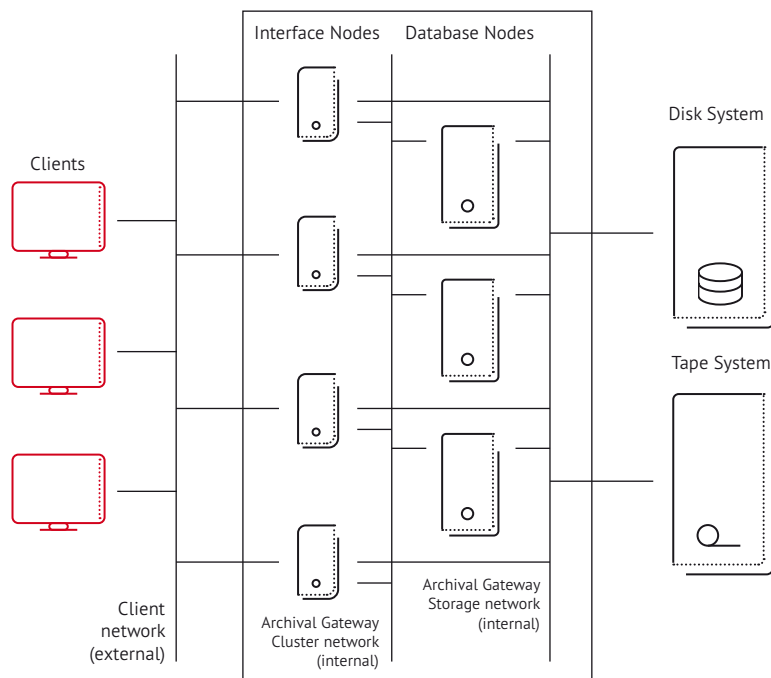
PoINT Archival Gateway meets highest requirements regarding performance, availability and scalability. The basis for this is the fully scalable and redundant design which provides scalable performance and redundancy levels for the service as well as for the data entities. Due to the server hardware and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

limited. Therefore, PoINT Archival Gateway in the Enterprise Edition supports the operation of multiple DBNs and IFNs within a single installation.

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

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

The following graphic illustrates the design of PoINT Archival Gateway in the Enterprise Edition and the network connections used by this solution on the basis of an exemplary installation.



Functionality

— S3 REST API

PoINT Archival Gateway offers a standardized S3 REST API. Thus, the software is suitable for the steadily growing number of those applications which support object-based storage by connecting S3 REST. Also supported are the S3 Glacier commands, which can be used especially to take into account the high latency of the tape storage class.

PoINT Archival Gateway features a highly scalable S3 REST web service which allows almost unlimited parallelization and very high data transfer rates.

— S3 COMPATIBLE STORAGE CLASSES

Disk and tape storage systems are integrated as S3-compatible storage classes and can be configured accordingly. Applications can access the various storage classes using the standardized S3 commands.

PoINT Storage Class	Corresponding AWS Storage Class
Disk	S3 Standard
Tape	S3 Glacier

— SINGLE NAMESPACE

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

— FLEXIBLE CONFIGURATION OF STORAGE CLASSES

PoINT Archival Gateway allows a very flexible configuration of storage classes. If the amount of data is still small, it is possible to start with the disk storage class. Later, this can be extended by a tape storage class to move inactive data from the disk class to tape media. Likewise, it is possible to start with a tape-only configuration, to which a disk storage class can optionally be added if, for example, read requirements with low latencies become important. Simultaneous use of both storage classes, e.g. for creating offline media, is also a valid use case.

— AUTOMATIC DISK/TAPE REPLICATION

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

— OFFLINE MEDIA MANAGEMENT

PoINT Archival Gateway also manages tape media which have been exported from a library, i.e. which are “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, PoINT Archival Gateway sends a corresponding message to the application. An operator then has to re-import the medium into one of the operating libraries.

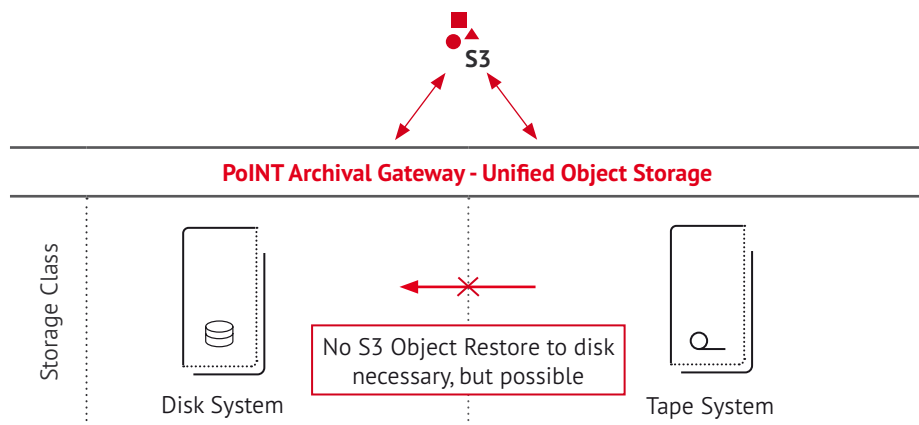
— LIFECYCLE POLICIES

PoINT Archival Gateway is compatible with AWS S3 Lifecycle Policies. This allows data to be moved between storage classes based on individual policies. For example, you

can specify that data is stored on the disk storage class for a certain period of time and then automatically moved to the tape storage class.

— DIRECT ACCESS TO TAPE

PoINT Archival Gateway supports direct writing and reading to and 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 cached on the disk storage class first. When reading, this feature avoids cumbersome and lengthy restore processes to disk. The data is delivered directly to the reading application.

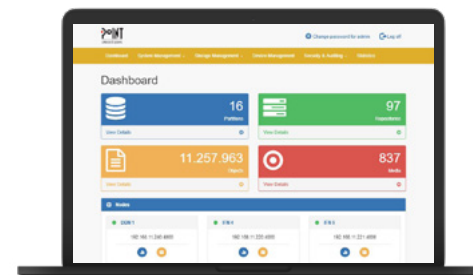


— ERASURE CODING

Data security is ensured by erasure coding. The process stores data blocks redundantly on several media. Thus, the data is not lost even if one medium fails.

PoINT Archival Gateway supports Erasure Code (EC) rates of 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 composite of several media is referred to as a “Protected Volume Array”. A Protected Volume Array consisting of *N* tape media can also cover *N* tape libraries.

The EC rates 1/2, 1/3, 1/4 mean the automatic creation of copies. For the tape storage class, this means that multiple tape copies (even in different libraries) can be created. With EC rates that distribute the data over multiple media (EC 2/3, 2/4 and 3/4), throughput rates can be increased significantly.



— ADMINISTRATION AND LOGGING

Administration is done via a web-based Admin GUI which is provided by PoINT Archival Gateway via HTTP service.

In addition to the Admin GUI, a C/C++ Admin API is available. This API allows the integration of administrative functions into custom applications. Java and .NET wrappers are also available.

PoINT Archival Gateway also supports Data Access Audit Logs and Security Audit Logs. Accesses to data objects of an object repository are logged in assigned Access Audit Log Files. Log records contain identifier, time stamp and access type as well as the identifier of the principal who performed the access.

A protected Security Audit Log File also logs all logon activity by managers, as well as any modifications made by security managers or otherwise related to security settings.

This log file contains information about the principals and details about the modifications made.

Installation Options

PoINT Archival Gateway supports both Windows as well as Linux operating systems.

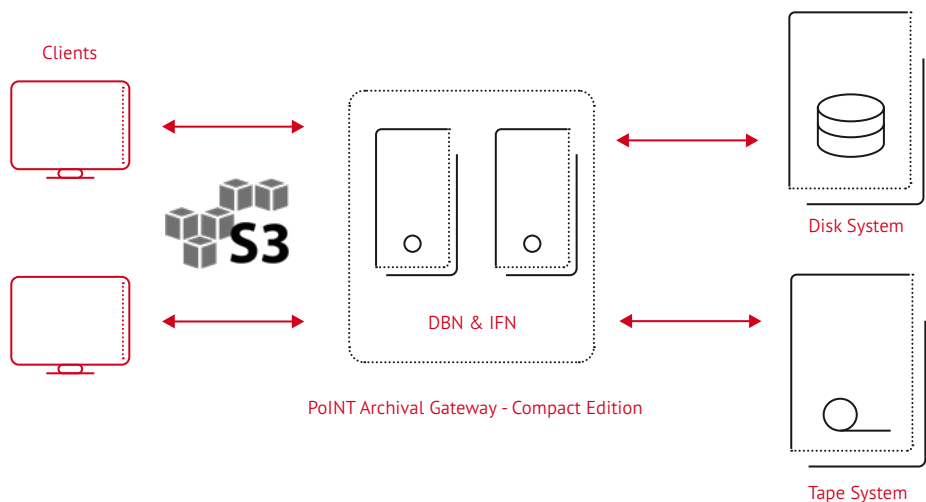
— ENTERPRISE EDITION

As a rule, the software packages and thus the services of PoINT Archival Gateway should be installed on separate server systems (Interface Nodes and Database Nodes). The separation on different server systems ensures maximum scalability, availability and performance. This installation option is offered by PoINT Archival Gateway – Enterprise Edition.

— COMPACT EDITION

For systems which do not require these performance levels, a special software package (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 Interface Service functionalities 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.



Supported Storage Systems

The design of PoINT Archival Gateway is independent of the storage technology and storage systems used. This means that users can select and exchange such systems as desired.

— DISK SYSTEMS

Any NAS and JBOD systems as well as supported disk-based object storage systems can be configured as disk systems in the disk storage class in PoINT Archival Gateway.

— TAPE SYSTEMS

PoINT Archival Gateway supports a wide range of tape libraries in the tape storage class. PoINT Archival Gateway integrates tape libraries directly; no additional drivers or software products are required. The following table provides an overview of the currently supported tape systems (loaders and libraries). Please contact PoINT Software & Systems if you need support for currently not listed tape systems.

Vendor	Product		
actidata	actilib Library 2U		
BDT	FlexStor II		
Cristie	GigaStreamT8	GigaStream T24	GigaStreamT48
HPE	MSL2024	MSL8096	StoreEver MSL6480
IBM	TS3100 Tape Library	TS3200 Tape Library	TS3500 Tape Library
	TS4300 Tape Library	TS4500 Tape Library	
Overland-Tandberg	NEO Series		
Qualstar	RLS-8560	RLS-85120	XLS-832700
Quantum	Scalar i3	Scalar i6	Scalar i6000
Spectra Logic	T Series		

PoINT Archival Gateway supports LTO (5 to 9) and IBM 3592 (TS1150, TS1155, TS1160) tape drives. Mixed drive configurations are also possible.

A PoINT Archival Gateway installation supports up to 8 tape libraries with a maximum number of 256 tape drives.

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
Eiserfelder Straße 316
57080 Siegen, Germany

P +49 271 3841-0
M info@point.de
W www.point.de