


PoINT Archival Gateway Tape-based Object Storage



Your data is
the key, so don't
lose it.

Contents

4	Introduction	16	Functionality
6	Product Overview		S3 REST API
	Key Features		Performance
8	Use Cases		Erasure Coding
	Backup of Object Storage Systems		Offline Media Management
	Backup of Cloud Data		Redundancy
	Tiering / ILM for On-Prem Object Storage		Data and Management Security
	S3 Archiving to Tape		Encryption
10	Architecture and Terms		Retention Management
	Object Repository (Bucket)		Lifecycle Management
	Archival Storage Partition		Audit Logs
	Archival Storage Device	21	Supported Tape Storage Systems
11	Design and Concept		
	Interface Nodes		
	Database Nodes		
	Networks		
14	Workflow		
	Writing Objects		
	Reading Objects		
15	Installation Options		
	Enterprise Edition		
	Compact Edition		

Introduction

Growth of unstructured data is one of the biggest IT data center infrastructure challenge for enterprises, according to recent surveys of research firms. This challenge cannot be solved economically and technically meaningful with file and block storage (or NAS and SAN) based on hard disk or flash technology.

Additionally, studies show that the majority of unstructured data is inactive. This data is not frequently used, but still needs to be preserved for business or compliance reasons. It does not make sense to store huge amounts of inactive data on hard disk or flash based storage systems.

Software-defined object storage offering a compatible S3 REST API using standardized tape technology as storage medium, like

LTO and 3592, provides a solution to this challenge.

Like file and block storage, object storage is a way to store data. The difference between file, block and object storage is that object storage has been designed for scalability, cost-effectiveness, reliability and high availability. In principle the way to store data is independent from the underlying storage media. E.g., in the same way as files can be stored on hard disk and tape in a file system structure, also objects can be stored on different storage media.

Considering the massive growth of unstructured data within the next years in combination with the fact that majority of this data becomes inactive very fast, tape is the only storage media which provides

sufficient data capacity with an acceptable price ratio. In combination with its WORM capabilities tape technology also fulfills archiving requirements.


Additionally, removable media like tape create an “air gap” against malware. Thus, tape technology is perfectly suited for data protection and backup.

For these reasons, an approach which combines object storage and tape technology as storage medium provides an innovative and economic solution for the data growth problem. Additionally, compared to all the storage media

available, tape technology has the potential for essential capacity improvements.

PoINT Archival Gateway is a software-defined scale-out object storage system designed to store and manage massive amounts of data directly on tape libraries with outstanding performance. Hard disk buffers for caching are not necessary. PoINT Archival Gateway provides a compatible S3 REST API and supports a wide range of tape libraries.

In this technical white paper, a detailed technical description of PoINT Archival Gateway is provided.



Archival Gateway



PoINT Archival Gateway offers standardized interfaces and protocols like the S3 REST API which may be used by the client applications for storing and reading data on the supported archival storage devices, drives and media. Essential services provided by PoINT Archival Gateway include user, data and storage management as well as access control, logging and monitoring.

— KEY FEATURES

Product Overview

PoINT Archival Gateway is a high-performance, scale-out, software-based object storage solution designed to manage massive amounts of data.

PoINT Archival Gateway is building a bridge between client applications or systems, respectively, and archival storage systems like tape library systems by providing services and functions which are mandatory for data archival solutions in enterprise data centers.

- High performance in terms of data and object rates
- Direct writing to tape without hard disk buffers
- High availability and reliability
- High scalability (incl. load-balancing, redundancy and failover)
- Industry standards (S3 REST API, LTO and 3592 tape drives)
- Object versioning
- Data protection (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 provides a solution for a wide range of requirements. With its standardized S3 REST API PoINT Archival Gateway can be used by many S3-capable applications.

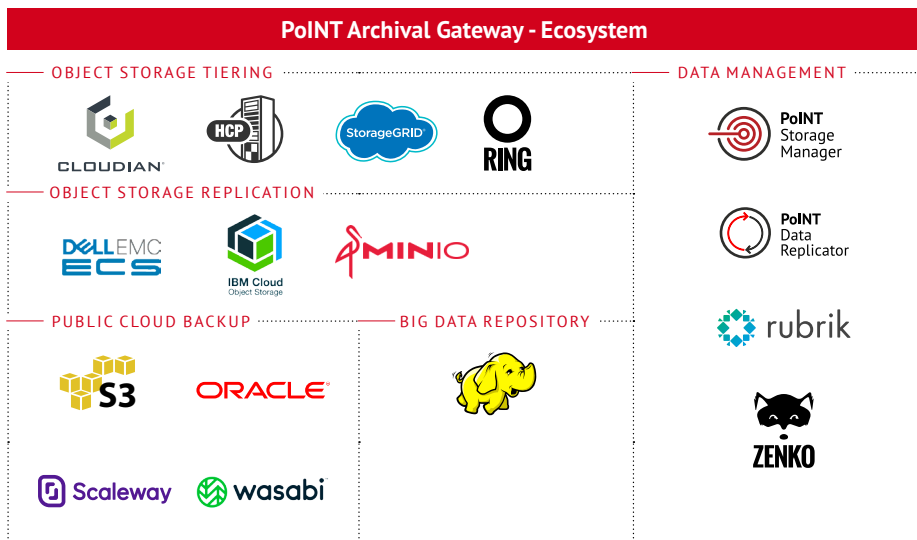
The following figure gives some examples of applications and environments which can be integrated with and connected to PoINT Archival Gateway.

— BACKUP OF OBJECT STORAGE SYSTEMS

Backup of object data becomes more and more important. Even if data in object storage is technically well protected by redundancies, risks such as software errors,

ransomware or malicious deletion remain. For this reason, backup of object data is essential.

PoINT Archival Gateway allows you to create backups of all object data in an S3-compatible object storage system including the object-specific metadata on tape. The data is stored in the same structure, i.e. under the same object keys, as your 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 OF CLOUD DATA

Data in the cloud must also be protected by a backup. It has already happened that cloud providers have discontinued their services and the stored data was no longer available. Clouds can also be affected by malware attacks. Therefore, a backup of cloud data is essential.

As with the backup of object data, PoINT Archival Gateway offers the possibility to create backups of S3 buckets stored in the cloud on tape. In case of emergency, local access to the backed up data is possible.

— TIERING / ILM FOR ON-PREM OBJECT STORAGE

In many cases public cloud storage providers offer different S3 storage classes with different performance features to their customers. On premises object store products which are used as private cloud offer only one storage class which is typically hard disk based. This is inefficient because active and inactive data are stored on the same storage technology.

The combination of hard disk-based object store with PoINT Archival Gateway allows to tier inactive data from hard disk to tape. For this purpose, many object store products provide already integrated ILM functions.

— S3 ARCHIVING TO TAPE

PoINT Archival Gateway fulfills archiving and compliance requirements thanks to WORM functionality and integrated retention management. As a result, saved data is protected not just against unintended deletion, but also against manipulation (e.g. from ransomware attacks).

PoINT Archival Gateway provides persistent data management for meeting legal and business data archival requirements. Appropriate retention rules can be enabled and specified on object repository level. The rules define how and when existing objects may be modified or deleted.

Architecture and Terms

The figure below provides an overview of the architecture and terms used in the scope of a PoINT Archival Gateway installation.

— OBJECT REPOSITORY (BUCKET)

PoINT Archival Gateway allows to create an arbitrary number of Object Repositories (Buckets), each of which can have different settings and restrictions (e.g. size limitations, access rights). Client applications and systems can store and read objects in the Object Repositories by using the S3 REST API of PoINT Archival Gateway. Object Repositories of the same Archival Storage Partition share the same Protected Volume Arrays.

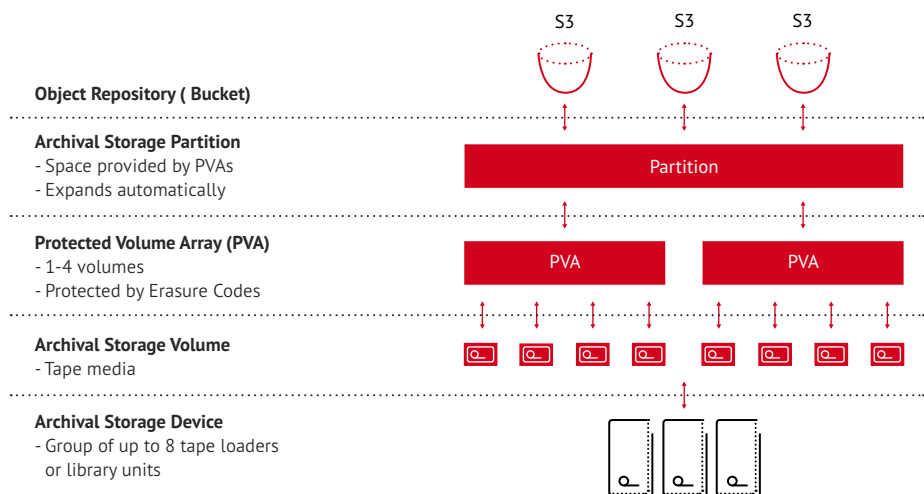
— ARCHIVAL STORAGE PARTITION

PoINT Archival Gateway uses a logical construct called Archival Storage Partition

as management units for the entire archival storage space available on all Archival Storage Volumes which are available in the Archival Storage Devices. Multiple Archival Storage Partitions can be created, each of which having different settings and restrictions (e.g. size limitations, access rights, levels of protection, redundancy or availability) and using separate Archival Storage Volumes. By default, PoINT Archival Gateway automatically extends the size of Archival Storage Partitions by allocating additional Protected Volume Arrays (see below) when the lastly allocated array has completely been filled with data.

— ARCHIVAL STORAGE DEVICE

Archival Storage Devices are media changer devices (tape libraries or tape loaders) with a specific number of integrated tape drives.



Design and Concept

PoINT Archival Gateway leverages the object storage architecture and manages data as objects. Each object includes a unique identifier of the object, the data itself as well as a variable amount of custom metadata. This architecture basically allows scalable retention of massive amounts of unstructured data on removable media like tape media in libraries.

PoINT Archival Gateway explicitly separates data and metadata. Unique identifiers and metadata are stored in common databases, while data are stored on tape media. Copies of the metadata and unique identifiers are stored on storage media, too, to support disaster recovery and verification processes.

Object storage architectures typically use databases to store metadata and unique identifiers. In contrast to other architectures (e.g. file systems and their hierarchical structures), databases do not impose any kind of limit on character sequences which can be used in unique identifiers and custom metadata or on the number

of object records. In addition, databases support arbitrary extension of existing records as well as indexes and highly performant queries covering all existing records. These features and functions are essential for supporting various kinds and versions of object storage interfaces like the S3 REST API.

This separation of data and metadata is reflected by the design of PoINT Archival Gateway which requires assigning the services to separate and dedicated server systems. Consequently, the PoINT Archival Gateway software consists of two software packages, each of which may be installed on a dedicated server system. A dedicated server system including its basic operating system software and the software package is called a gateway node. Defined by the installed software package, a gateway node serves exactly one of two possible purposes.

— INTERFACE NODES

A gateway interface node (IFN) primarily is the communication partner for client applications or systems. By means of dedicated interface modules, it provides an S3 REST API for storing and reading objects to client applications and systems and transfers data between client applications or systems and the Archival Storage Volume (i.e. tape media). The S3 REST API should be understood as a web service that supports beyond the HTTP protocol also S3 specific protocol elements.

In detail, the 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 metadata exchange with database node

— DATABASE NODES

A gateway database node (DBN) primarily provides central database services to the IFNs. The database holds the index of the objects which have been stored, e.g. object keys and metadata, storage locations of object data on the Archival Storage Volume (tape medium), as well as configuration and maintenance data of PoINT Archival Gateway.

Additionally, the database stores data from logging and monitoring processes and provides corresponding auditing services and log files. Further central services are the system configuration (Admin GUI) and the management and control modules for Archival Storage Devices (tape libraries), which therefore are also located on this node.

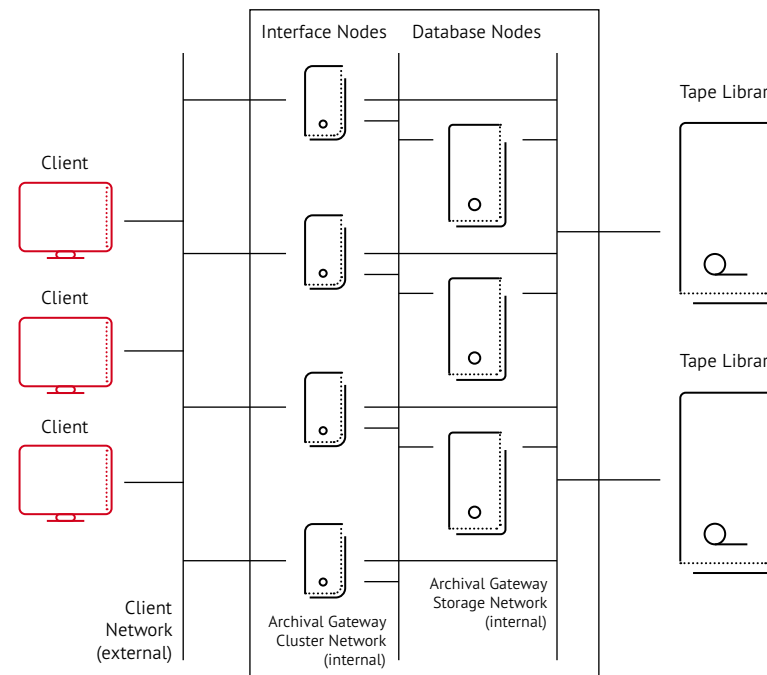
PoINT Archival Gateway requires at least one working IFN and one working DBN to be operable.

PoINT Archival Gateway fulfills the highest levels of performance, availability and scalability requirements by means of this fully scalable and redundant design, providing scalable performance and redundancy levels for both service and data entities. Because scalability and availability of a single node is restricted due to limitations of server hardware and operating systems, PoINT Archival Gateway supports installing multiple DBNs and IFNs in a single deployment in order to build clusters for increasing performance (i.e. load balancing) and availability (i.e. failover and redundancy).

PoINT Archival Gateway automatically performs all tasks which are required to reconstitute operability and consistency of a cluster node which had been affected by a failure before, of course provided that the system administrator has re-established the

operability of the server system, operating system and installed PoINT Archival Gateway software package before.

The following picture illustrates the design of PoINT Archival Gateway and the networks used by the solution by showing a sample deployment.



— NETWORKS

The storage network of the PoINT Archival Gateway can be an iSCSI or Fibre Channel network connecting Archival Storage Devices (e.g. tape libraries with tape drives) to the gateway nodes, while typically the other networks are Ethernet networks.

Workflow

Clients connect to the S3 REST API of the IFNs. Clients can send their requests to any IFN. Of course, distribution to all IFNs is also possible.

— WRITING OBJECTS

When a client desires to write a new object, it sends an appropriate PUT request including the object data. The addressed IFN receives the object data and stores it in data buffers in memory (q.v. data buffering module). In addition, the IFN computes a hash code of the data and adds it to the metadata of the object. Optionally, the IFN compares its hash code to a hash code which has been computed and transferred by the client in the request header, and sends a specific error response if the hash codes do not match. Then, it forwards the request and its associated data buffers to the driver module for the tape drive to write it to the tape media. In this module, the original data buffers are reorganized, interleaved with data buffers from other objects, protected by adding erasure codes and logical block protection codes, optionally

encrypted, and written to the Archival Storage Volume (tape media). Finally, after the data have successfully been written to the Archival Storage Volume, the HTTP service module stores the object metadata and the location of the object data on Archival Storage Volume in the database and sends a positive response to the clients.

— READING OBJECTS

For reading an existing object, a client sends an appropriate GET request. The addressed IFN first reads the object metadata and the location of the object data from the database. Then, it forwards the request to the driver module for the drive units to read the object data from the Archival Storage Volume (tape medium). This module regenerates the object data from the data on the Archival Storage Volume, copies it to data buffers in memory, and immediately sends the data buffers to the client.

Installation Options

— ENTERPRISE EDITION


Typically, the software packages and thus the services of PoINT Archival Gateway should be installed on separate server systems (PoINT Archival Gateway - Enterprise Edition) because this kind of deployments provides the maximum levels of scalability, availability and performance.

— COMPACT EDITION

However, for systems which are not required to provide the maximum levels, a special software package (PoINT Archival Gateway - Compact Edition) allows installing one

database and one interface service on the same server system. In this case, the database and interface service functions are combined in a single and compact service module to optimize the common use of resources and to remove overhead caused by inter-service communication.

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



Functionality

— S3 REST API

The PoINT Archival Gateway offers a compatible S3 REST API. This makes the software suitable for use with the rapidly rising number of applications that support object-based storage using S3 REST.

As PoINT Archival Gateways writes data directly to the tape media and reads data directly from the media, no hard disk buffers for caching are necessary. This also eliminates the need for the clients to use the cumbersome S3 Glacier commands.

The PoINT Archival Gateway's design is independent of the storage technology or system used, meaning that users can choose and swap out such systems with confidence. The PoINT Archival Gateway's highly scalable S3 REST web service enables almost unlimited parallelization and very high data transfer rates.

— PERFORMANCE

The PoINT Archival Gateway is a software-based and high-performance object storage system with tape libraries as target storage systems. This combination makes it possible to store and archive hundreds of petabytes of data. The decisive factor that makes


this possible is the way the PoINT Archival Gateway quickly receives data and securely writes it to Archival Storage Volume (tape media) in a format that means this data can also be quickly read again afterwards. The PoINT Archival Gateway's high level of scalability means it can handle transfer rates of over 1 PB per day.

— ERASURE CODING

Data security is ensured by an erasure coding process. This saves data blocks redundantly on multiple media, meaning that data is not lost even if a medium fails.

Erasur code rates supported by PoINT Archival Gateway are 1/2, 1/3, 1/4, 2/3, 2/4 and 3/4. In combination with erasure coding, the level of data protection and redundancy can further be advanced by using two, three or four Archival Storage Devices (tape libraries) in parallel. A Protected Volume Array consisting of N tape media may span N libraries.

In the following two figures the write path and read path are illustrated in case an erasure coding rate of 3/4 has been selected.



— 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, the access fails and PoINT Archival Gateway sends a corresponding message to the application. An operator then has to re-import the medium into one of the operational libraries so that the access can be performed.

— REDUNDANCY

Redundant server nodes automatically managed by PoINT Archival Gateway mean that the server systems running the PoINT Archival Gateway remain stable.

— DATA AND MANAGEMENT SECURITY

The system’s approach to management security is made possible by a user and permission management concept. Access rights can be assigned for these guidelines. To this end, the administrative interface provides functionality for creating local security guidelines for users and groups. The system can also inherit external guidelines from an Active Directory domain.

Access rights for the objects in an Object Repository are granted by adding a local

principal of type user or an external security principal to the list of authorized object repository principals. While adding a principal, a security manager may define granular access rights by specifying if the principal shall be allowed to a list, read, write or delete objects in the object repository.

Principals created and used for providing access rights to object repositories may, but do not need to have management roles.

Access to the objects in object repositories is exclusively possible through the client interfaces and protocols, where the individual interface or protocol, respectively, defines how principals must authenticate themselves by providing their credentials.

— ENCRYPTION

Data and metadata of objects may optionally be encrypted by PoINT Archival Gateway before they are written to the Archival Storage Volumes (tape media). Encryption leverages cipher blocks chaining algorithms, is applied on block level and can be enabled or disabled on Archival Storage Partition level. For encryption, either a common system key or a partition specific key may be used. The encryption algorithms supported by PoINT Archival Gateway use cryptographic keys of 128 (AES-1) and 256 (AES-2) bits length.

— RETENTION MANAGEMENT

PoINT Archival Gateway provides persistent data management for meeting legal and business data archival requirements. Appropriate retention rules can be enabled and specified on Object Repository level and apply to all objects in the repository. The rules define how and when existing objects in the object repository may be modified or deleted.

The retention rules basically allow specifying retention periods and to issue or retire legal holds and can be managed by means of the system configuration GUI.

Once enabled, retention rules cannot be disabled again, and affected object repositories cannot be deleted before the retention periods of all contained objects has expired. However, retention periods can be extended at any time and, in addition, administrative deletion of retained objects is possible, but PoINT Archival Gateway unconditionally logs administrative deletion of objects.

— LIFECYCLE MANAGEMENT

In addition to data retention management as described above, PoINT Archival Gateway provides a set of lifecycle management functions of AWS S3.

In particular, PoINT Archival Gateway supports policies defining expiration actions on objects and incomplete multipart uploads.

PoINT Archival Gateway removes expired objects and aborts incomplete multipart uploads one day after the specified date has been reached.

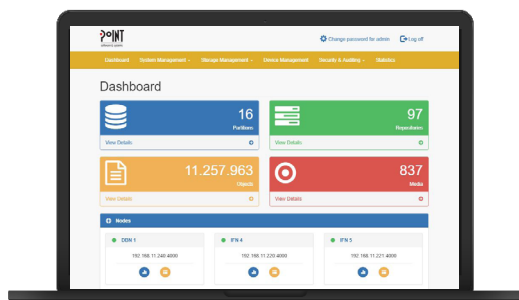
Expiration actions cannot delete objects of which the retention period has not yet elapsed. In such a case, the expiration action is performed at earliest as soon as the retention period has been elapsed.

— AUDIT LOGS

PoINT Archival Gateway supports data access audit logs and security audit logs.

Accesses to the data objects of an Object Repository are logged in associated access audit log files. Log records include identifier, time stamp and kind of access as well as an identifier of the principal who performed the access.

In addition, all logon actions of managers and all modifications applied by security managers or either way related to security settings are logged in a protected security audit log file. This log file includes information about the principals, including details about the modification.



— ADMIN GUI

The primary configuration interface is the Admin GUI provided by an HTTP service of PoINT Archival Gateway.

The service is hosted by Microsoft Internet Information Server instances running on the DBNs. The Admin GUI communicates with the configuration module via a local network connection. The configuration module forwards configuration data and corresponding changes to the affected nodes.

— ADMIN API

Besides the Admin GUI, PoINT Archival Gateway also provides a C/C++ Admin API. This API allows the integration of administrative functions into proprietary applications. Java and .NET wrappers are also available.

Supported Tape Storage Systems

The PoINT Archival Gateway supports a wide range of tape libraries and PoINT works closely with leading manufacturers. As a result, the list of supported products is being expanded all the time.

The following table provides an overview about the currently supported tape systems (loaders and libraries):

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	Diamondback
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 - TS1170) tape drives. Mixed drive configurations are also possible.

Please contact PoINT Software & Systems GmbH for support of tapes systems that are not yet listed.

PoINT Archival Gateway directly supports and integrates tape libraries. No additional drivers or software products are required.

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 of software solutions for storage and management of data using all available mass storage technologies like cloud and object storage, hard disk, magnetic tape and optical. Close collaboration with leading hardware manufacturers enables an early support of innovative storage technologies. Besides complete solutions PoINT also offers its know-how as Toolkits, which can be easily integrated in other applications by the programming interface. Furthermore we project entire storage solutions and provide consultancy with our long-term and versatile experience.

PoINT products are distributed in more than 25 countries world-wide and have been installed successfully in more than two million installations. Our customers range from end users expecting a compact and secure solution to large corporations, which 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