Service Catalog

IX.Cloud
SWISS FINANCIAL CLOUD

Version: 2026-06-08 09:42

Table of Contents

- Backup
  - Service Architecture
  - Service Scope
  - Service Options
  - Recovery in Case of Disaster (RTO and RPO)
- Service Request Management
- Service Level Reporting
- Service Desk

- **ix.Cloud Edge**
  - Service Description: Public DNS Service ix.Cloud Edge
  - Service Architecture
  - Service Scope
  - Service Options
    - Initial Setup
    - Hidden Primary & Zone Distribution
    - GeoLoad Balancing
    - Security & Administration
  - Supported Record Types
- **Datacenter Services**
  - Cloud Connect
    - Service Architecture
    - Service Scope
    - Service Options
  - Rack Collocation
    - Service Architecture
    - Service Scope
    - Service Options
- **Network Services**
  - Firewall
    - Service Architecture
    - Service Scope
    - Service Options
  - Server Proxy
    - Service Architecture
    - Service Scope
    - Service Options
  - Load Balancer
    - Service Architecture
    - Service Scope
    - Service Options
  - Web Application Firewall

        - Service Architecture
        - Service Scope
        - Service Options
    - Private DNS
        - Service Architecture
        - Service Scope
        - Service Options
    - Secure Mail-Relay
        - Service Architecture
        - Service Scope
        - Service Options
    - Hosted Software-Appliance
        - Service Architecture
        - Service Scope
        - Service Options
- **Storage Services**
    - File Storage
        - Service Architecture
        - Service Scope
        - Service Options
    - Object Storage
        - Service Architecture
        - Service Scope
        - Service Options
- **Compute Services**
    - Virtual Machine
        - Service Architecture
        - Service Scope
        - IT Baseline Protection
        - Service Options
- **System Management Services**
    - Managed OS
        - Service Architecture
        - Service Scope
        - Service Options
    - Metrics Monitoring
        - Service Architecture
        - Service Scope
        - Service Options
    - Software Deployment
        - Service Architecture

        - Service Scope
        - Service Options
    - Software and Release Cycles
        - Linux Software
        - Windows Software
        - Dealing with 3rd Party Software
        - Operating System and Software Release Cycles
- **Database Services**
    - Managed xSQL-Instance
        - Service Architecture
        - Service Scope
        - Service Options
    - PostgreSQL HA Managed Service
        - Service Architecture
        - Service Scope
    - Managed noSQL-Instance
        - Service Architecture
        - Service Provisioning
        - Service Scope
        - Service Options
    - Managed Service Database Security Audit
        - Service Architecture
        - Service Scope
    - IT Baseline Protection Database Service
        - Patch Management
        - Logging
        - Malware Protection
- **Container Services**
    - IT Baseline Protection
        - Upgrade and Patching
        - Logging
    - Container Registry
        - Service Architecture
        - Service Scope
        - IT-Grundschutz
        - Service Options
    - Agile Factory
        - Service Architecture
        - Service Scope
        - IT Baseline Protection
        - Service Options

- AnyCloudK8s
  - Service Architecture
  - Service Scope
  - IT Baseline Protection
  - Service Options
- Container Namespace
  - Service Architecture
  - Service Scope
  - Service Options

- **Splunk Index**
  - Available Service Options
  - Available Meta-Information
- **Enterprise Streaming Service**
  - Inventx Scope of Services
    - Managed Streaming Cluster
    - Schema Registry
    - Access Control & Tenant Separation
    - Topic Management
    - Monitoring & Alerting
    - Security
    - Topic Monitoring & Logs
  - Customer Responsibility
    - Integration of Applications
    - Network Connectivity
    - Schema Registration
    - Monitoring Usage
  - Service Level Parameters (SLP)
- **Character Legend**

Introduction

This service catalog describes the services provided by Inventx as standardized cloud services within the ix.Cloud product line.

The services of ix.Cloud are produced on Inventx's own Community Cloud and exclusively in Swiss data centers. The architecture of the Community Cloud follows a shared infrastructure approach, in which the tenants share the hypervisor and the underlying hardware (network, storage, and compute).

Inventx defines a tenant as the highest organizational instance, which represents a self-contained unit in terms of data and organization. This ensures the necessary isolation between customers. Each tenant is always assigned to an explicit customer and serves not only for isolation but also for the use of resources and services.

The following graphic illustrates the available service models (IaaS, PaaS, SaaS) and service categories in the middle section, and the digital customer interfaces (Digital Customer Interfaces) for service management in the upper section. Complementary to this, the general Inventx standards in the areas of information security and compliance are listed in the lower section.

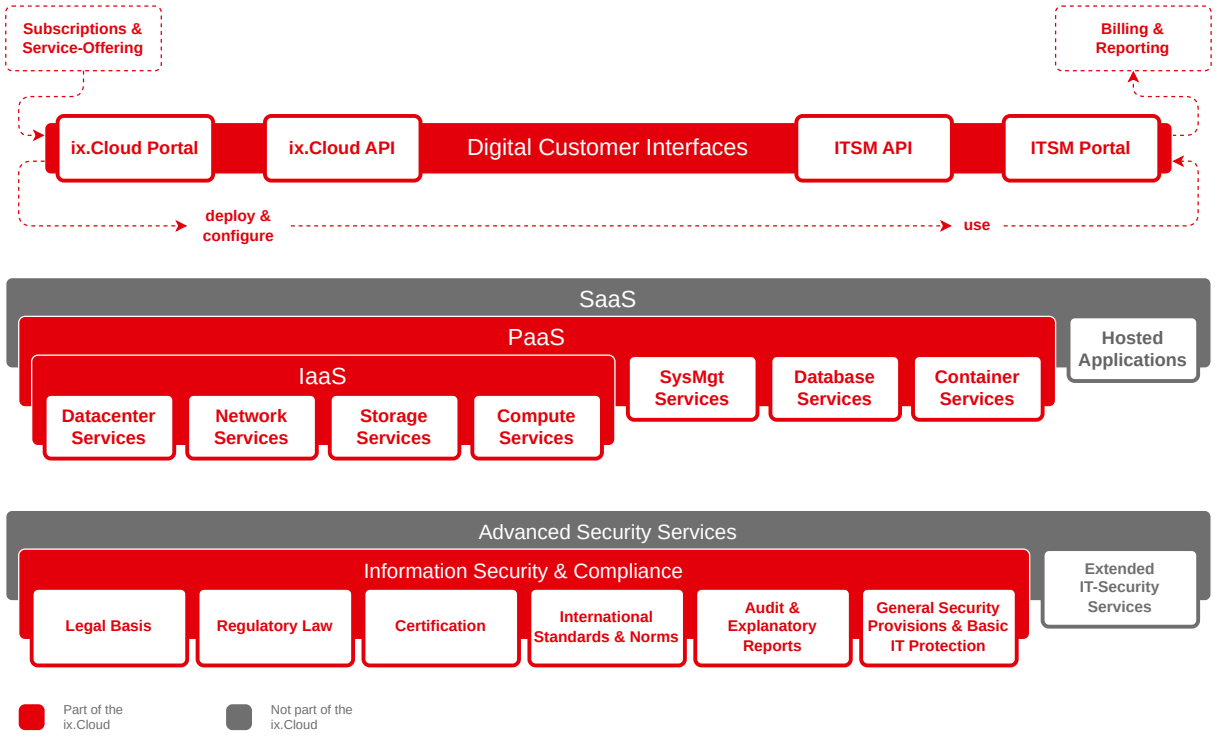

Image: ix.Cloud Map

Service Models and Service Categories

The table below shows the service models and service categories available in ix.Cloud.

Table: Service Models and Service Categories

Model	Category	Description
Infrastructure as a Service (IaaS)	Datacenter Services	Infrastructure in Inventx's own data centers (DCs). Thanks to secured access, fire protection, secure power supply, and cooling, applications remain available at all times and data is secured.
	Network Services	Offers highest service quality by networking cloud and on-premises infrastructures.
	Storage Services	Provides secure, scalable cloud storage for data, apps, and workloads.
	Compute Services	Pre-configured Availability Sets and computing power from the cloud on demand, billed on a usage basis.
Platform as a Service (PaaS)	System Management Services	Operations-optimized services for efficient provisioning and operation of business applications on virtual machines.
	Database Services	Fully managed database services enable barrier-free and highly scalable data management.
	Container Services	Continuous Delivery with simple and reliable tools for even faster development - innovation at its core.

Depending on the service model, the supply obligations of the customer and Inventx differ. This means that both parties must assume corresponding responsibilities so that a target application can be provided with the desired depth of production and the required security features.

The following graphic illustrates the responsibilities per service model.

Image: Responsibility Matrix

:::tip Managed Services

Additional Managed Services can be agreed upon between the customer and Inventx for all service models and service categories. These are individual services that are not or only partially covered by this service catalog.

:::

Digital Customer Interfaces

Inventx offers customers highly available digital interfaces according to SLA Rhodium, through which ix.Cloud services and resources can be obtained and managed around the clock.

Authorization for the digital customer interfaces is defined and implemented via a separate onboarding project.

ix.Cloud Portal

With the ix.Cloud Portal, authorized customer users have access to a web portal where the corresponding services can be ordered and managed around the clock. Every customer has at least one user for the ix.Cloud Portal after successful onboarding. The customer-specific user management is

controlled via a customer's identity provider (e.g., Active Directory) and transmitted to the portal via an identity proxy (e.g., AD Federation Services), which ensures authorization.

The attributes to be transmitted are:

Table: User Management ix.Cloud Portal

Criterion	Description
Username	Transmission of the User Principal Name
Group name	Transmission of the Group Name

ix.Cloud API

The ix.Cloud API forms the basis for standardized and automated management of ix.Cloud services and simultaneously the prerequisite for integrating IaC (infrastructure as code). IaC is the process of managing and deploying resources via code instead of physical configurations or using configuration tools.

ITSM Portal

The Inventx ITSM Portal is available to defined customer users for processing shared IT Service Management processes within the scope of Service Requests and Incidents.

ITSM API

The ITSM API forms the basis for linking and automating ITSM processes between the customer and Inventx.

Resource Deletion Process

The deletion of a resource occurs automatically via an asynchronous job. Under normal conditions, the following steps are performed: 1. The corresponding resources are removed from the target system. 2. Associated entries in surrounding systems, such as DNS, monitoring, backup, and billing, are also deleted. 3. In the IT Service Management System (ITSM), the status of the record is set to "deleted".

In case of an error, the following procedures apply: 1. Job execution is automatically retried at specified intervals. 2. If the retries do not succeed, an incident is automatically created. 3. The processing of the incident is carried out by the responsible expert teams, who also ensure the completeness of the measures.

Subscriptions and Service Offering

The ix.Cloud is based on three fundamental elements - Tenant, Subscription, and Resources. A Tenant represents a company, which can be divided into several organizational units - Subscriptions. A Resource

is a manageable service element that can be ordered via the portal and consumed via a Subscription.

A Tenant must contain at least one Subscription. A Subscription serves to separate and map organizational structures through the following points:

- User and rights management
- Managing the service offering
- Managing and ordering resources
- Accounting for and allocating costs

Billing and Reporting

Billing and Reporting refers to the invoicing of consumed services described in this service catalog.

The services are invoiced based on the Consumption Report (Metered Services) and billing is done monthly in arrears.

Backup

The backup service is based on a highly available, scalable and performant platform in Inventx's data centers, which ensures the backup of data and, if required, its restoration for a service or a complete VM.

Service Architecture

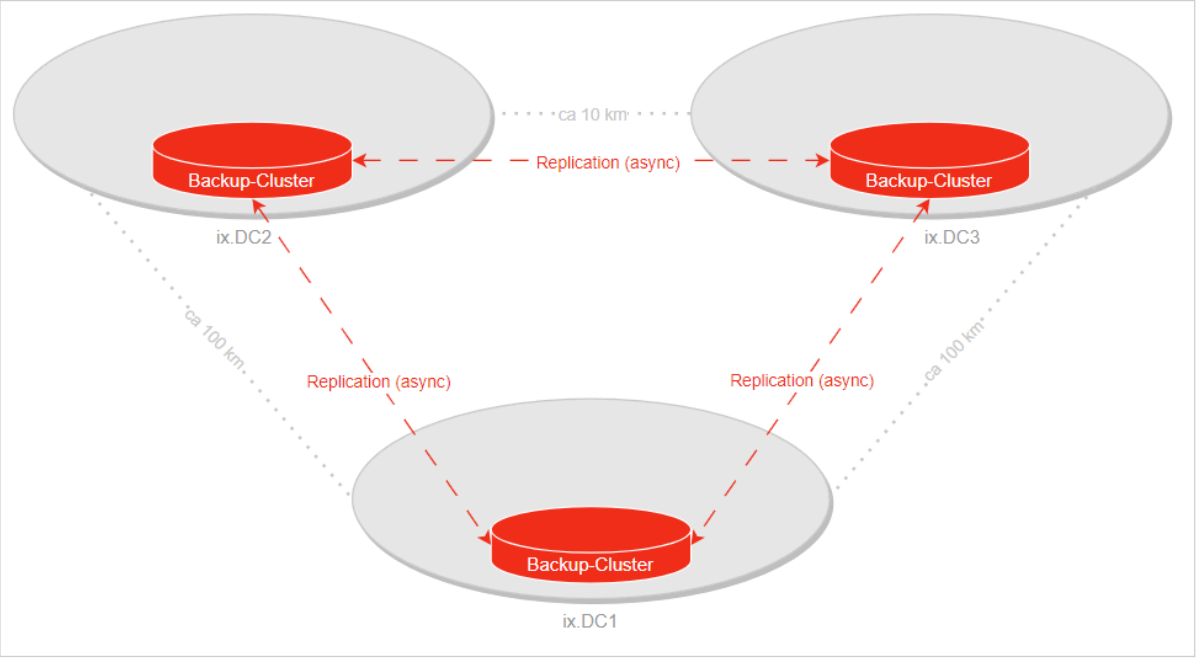

Figure: Calculating the recovery time

Service Scope

Table: Backup Service Scope

Features	
Access Control	■
Data Encryption	■
Immutable Backups	■
Backup Profiles	■
Monitoring & Logging	■
Data Integrity	■
Data Recovery	■
Disaster Recovery (RTO and RPO)	■

Service Options

Access Control

Access to the backup infrastructure is role-based (RBAC/ixPAM) and uses multi-factor authentication (MFA).

Data Encryption

Backup data is encrypted both during transfer and throughout its retention period (AES-256, FIPS 140-2 compliant encryption of data in flight and at rest).

Immutable Backups

Backup data cannot be changed or deleted throughout its retention period (Immutable Backup & DataLock/WORM). This provides effective protection against ransomware or malicious deletion of backup data.

Backup Profiles

The following backup profiles can be selected:

Table: Backup Profiles

Service Level	Bronze	Silver	Gold	Platinum*	Rhodium
Backup Location	Local	Remote		Remote & DR-Datacenter	
Backup Interval	Daily				

Backup Retention	No Backup, 14, 40, 100, 200 or 400 Days

* the Platinum service level is not available for the Community ix.Cloud

Backup Location

Primary data is always backed up to the backup infrastructure in the remote data center. For Platinum*/Rhodium service levels, an additional replication (copy) of the data is made to the DR data center.

:::info As an option, additional backup copies within or outside the Inventx data centers can be ordered via "Generic Request". :::

Backup Interval

The interval defines the time period at which data backups are created. An incremental backup is performed at least once daily; for databases, logs (archive, transaction logs, etc.) are backed up several times per hour.

Backup Retention

Retention determines the retention period of data backups in days. After the selected period expires, the data will be irrecoverably deleted.

:::caution For the "No Backup" profile, the customer explicitly waives data recovery! This applies to primary data at the OS, database, application level, etc., and also includes cloning via backup restore to a VM/DB for which the "No Backup" profile was selected! :::

Monitoring & Logging

Inventx monitors the backup process and ensures that it is carried out at the planned regularity. All relevant accesses, changes to configuration or system parameters are recorded in an audit log.

Data Integrity

The backup solution used ensures that data is encrypted during transmission and, after the copy is stored, is protected against accidental or malicious changes and premature deletion using WORM technology.

Data Recovery

Monolithic recovery of a complete VM is performed in self-service via the ix.Cloud portal. File- or database-based recoveries generally need to be requested via "Generic Request".

Disaster Recovery (RTO and RPO)

RTO and RPO define, in the event of a disaster, the maximum duration of recovery (RTO) of an application, system and/ or process and the maximum data loss (RPO).

Table 1: SLA - Disaster Recovery (RTO and RPO)

Service Level	Bronze	Silver	Gold	Platinum	Rhodium
Recovery Time Objective (RTO)	-	Best Effort	48h	2h	2h
Recovery Point Objective (RPO)	Best Effort	24h	15min	0min	0min

Recovery Time Objective (RTO)

The circumstances of a disaster can vary greatly and influence this service level. The value strongly depends on the number of simultaneous recoveries; i.e., with multiple simultaneous recoveries, the value per recovery may be lower. For a single recovery, a guideline of 200-400 MB/s can be assumed.

Recovery Point Objective (RPO)

Damage events caused by manipulated or corrupted data are exclusively covered by the backup-relevant quality elements in the SLA. This means that if corrupted data exists in the live system, it can only be corrected from a backup, and the specified RPO does not apply.

{/* Do not remove this is only used in the PDF the current date */} export const Datum = ({}) => (

Version: {new Date().toISOString().replace(/T/, ' ').replace(/.. /, '').slice(0, -8)}

);

This document is machine translated. German language is contractual binding.

Service Level Framework

:::info

The contractually binding service levels are governed in the framework agreement between the customer and Inventx. The values listed in this chapter are not contractually binding, but serve exclusively to provide context for the services described in the Service Catalog.

:::

Service Delivery Point

The Service Delivery Point (SDP) defines the technical interface at which the service delivery by Inventx ends and responsibility transfers to the customer. Depending on the service model, the SDP is defined at different levels, which results in separate Service Level Targets for each level.

Service Delivery Point Infrastructure

The SDP Infrastructure applies to all services whose service delivery ends at the infrastructure level (IaaS) – that is, where Inventx provides the underlying physical and virtualized resources. The following Service Level Targets refer to availability at this level and apply to all services that build on this service delivery point.

Table: SDP Infrastructure

SLP	Bronze	Silver	Gold	Platinum	Rhodium
Availability	n.a.	99.40%	99.40%	99.90%	99.90%
Outage Frequency p.m.	n.a.	2	2	1	1
Outage Frequency p.a.	n.a.	6	6	4	4

Service Delivery Point Platform

The SDP Platform applies to all services whose service delivery ends at the platform level (PaaS) – that is, where Inventx additionally provides and is responsible for the operational platform in addition to the infrastructure. The following Service Level Targets refer to availability at this level and apply to all services that build on this service delivery point.

Table: SDP Platform

SLP	Bronze	Silver	Gold	Platinum	Rhodium
Availability	Best Effort	99.20%	99.60%	99.80%	99.95%
Outage Frequency p.m.	Best Effort	1	1	1	1
Outage Frequency p.a.	Best Effort	9	9	6	4
Max. Net Downtime per Incident	Best Effort	Best Effort	Best Effort	Best Effort	Best Effort

Service Time

For each Service Level Target listed in the preceding sections, the following Service Time applies – regardless of the service delivery point:

Table: Service Time per Service Level

Tier	Service Time
Bronze, Silver	Standard
Gold, Platinum, Rhodium	7 x 24

Service Levels

The following tables show the service levels that Inventx guarantees to customers (explicitly marked as SLP (Service Level Parameter)), respectively indicative key performance indicators (KPI) that Inventx aims for. All information refers to the Swiss time zone.

Operating and Service Times

The following times apply as operating and service times, taking into account the defined maintenance windows and holidays.

Inventx ensures monitoring of services during operating time. During service time, qualified personnel are available to customers for support requests. The service time forms the basis for calculating the agreed service levels.

Table: SLA - Operating and Service Times

Service Level	Bronze	Silver	Gold	Platinum	Rhodium
Operating Time (excl. Maintenance Windows and incl. Holidays)	7 x 24h Monday - Sunday 00:00 - 24:00 hours				

Service Time (excl. Maintenance Windows and incl. Holidays)	5 x 11h Monday - Friday 07:00 - 18:00 hours	7 x 24h Monday - Sunday 00:00 - 24:00 hours
Monitoring	7 x 24h	
Customer Maintenance Window	on-demand	
Service Maintenance Window	Service patching in the selected patch week/day (between 19:00-23:00)	
Provider Maintenance Window	4h per month	
Emergency Maintenance Window	on-demand	
Holidays	11 holidays	

Operating Time

Operating time is the time during which all IT components relevant to overall service delivery (systems, applications, networks) at all levels of system architecture are in operation. Maintenance windows are the exception.

Service Time

Service time (or support time) is the portion of operating time during which the availability of the support organization is guaranteed. Within the service time, incidents are received and processed. Outside of service time, incidents are received but not processed.

Monitoring

During operating time, Inventx continuously monitors services automatically. Any events are logged, forwarded to the support organization, and processed during service times.

Maintenance Windows

Maintenance windows define the agreed times for system maintenance. The maintenance window is agreed as closely as possible after service time, taking technical feasibility into account. In approved exceptions, other or longer times may be defined.

Customer Maintenance Window

Customer-specific maintenance windows serve the customer's maintenance needs and are individually agreed and planned. If such maintenance windows affect Inventx's smooth service delivery, they must be reported early via "Generic Request".

Provider Maintenance Window

Provider Maintenance Windows are used to reserve periods during which Inventx can perform necessary maintenance work on the central infrastructure. Inventx strives to perform maintenance work as efficiently as possible and to keep any service disruption as brief as possible.

Table: SLA - Provider Maintenance Window

	Minor Changes	**Major Changes**
Interval	Monthly	2 times per year
Day	on Thursday after the 2nd Monday of the month	on the 2nd Sunday in February and June
Time	19:00 - 23:00 hours	09:00 - 20:00 hours
What	minor adjustments to ix.Cloud management environments	major adjustments to ix.Shared/ix.Cloud infrastructures
Business Impact	YES - short interruptions (15') with business impact possible	YES - downtimes or reduced redundancy
Risk Probability	medium	high
Risk Impact	medium	high

Emergency Maintenance Window

During ongoing operations, situations arise that require urgent intervention, typically the closure of a critical security vulnerability. This means that maintenance work - possibly with downtime - must be performed unplanned. Inventx aims to inform customers as early as possible before execution.

Holidays

The following days are considered holidays for Inventx's support organization:

Table: SLA - Holidays

Holiday	**Duration**

New Year (January 1)	All day
Berchtold's Day (January 2)	All day
Good Friday	All day
Easter Monday	All day
Labour Day (May 1)	All day
Ascension Day	All day
Whit Monday	All day
National Holiday (August 1)	All day
Christmas (December 25)	All day
St. Stephen's Day (December 26)	All day

Availability

Availability is defined and reported via the following service levels.

Table: SLA - Availability

Service Level	Bronze	Silver	Gold	Platinum	Rhodium
Availability in % per year	Best Effort	98.32	99.7	99.9	99.9
Maximum downtime per year	-	48h	24h	8h	8h
Maximum downtime per incident	-	8h	4h	2h	2h
Maximum number of failures per year	-	12	6	4	4

Availability in % per year

The availability in % per year of a service indicates the percentage of the agreed service time during which this service is available. This value is calculated as follows:

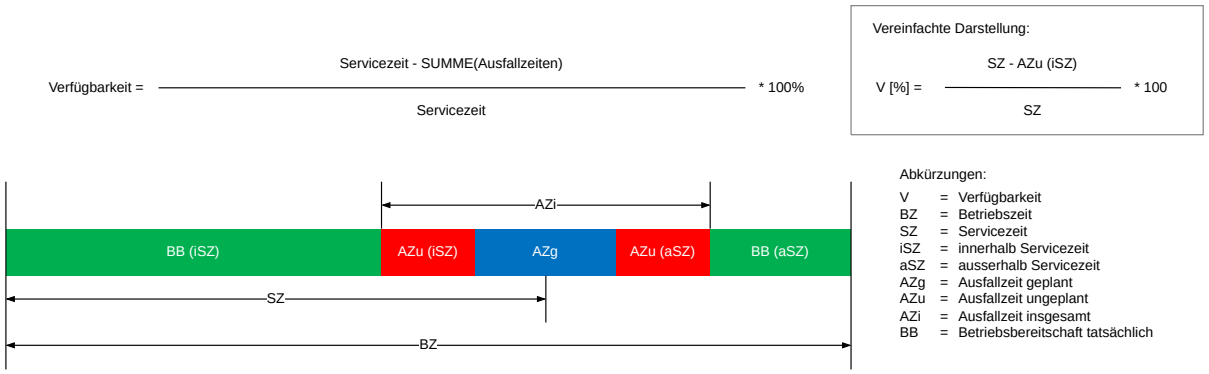

Figure: Calculation of availability

Maximum downtime per year

The Service Level "Maximum downtime per year" is the maximum downtime measured in hours per service and year.

Maximum downtime per incident

The Service Level "Maximum downtime per incident" specifies the maximum downtime per incident of the respective service within the agreed service time. The start time of the downtime is the earliest of the following times:

- Report by the customer to Inventx or
- Report by a system management tool, network management tool, end-to-end reporting tool

The end time of the downtime is the time when the incident is resolved and the service functions as contractually agreed.

Maximum number of failures per year

The Service Level "Maximum number of failures per year" defines the maximum number of failures of a service per year.

Incident Management

Within the scope of Incident Management, qualified support requests from qualified Inventx employees are processed during the agreed support time. Different reaction times and recovery times apply for on site time and on call time:

Table: SLA - Onsite and on call Times

Service Level	Bronze	Silver	Gold	Platinum	Rhodium
on site time (excl. Holidays)	Monday - Friday, 07:00 - 18:00 hours				

on call time (incl. Holidays)	Remaining times

on site time

During on site time, incident reports of Incident Priority Classes P1 to P2 are received and processed.

The following reaction times apply during on site time for the individual Incident Priority Classes:

Table: SLA - Reaction times on site time

Incident Priority Class	Bronze	Silver	Gold	Platinum	Rhodium
P1	-	2h	2h	2h	2h
P2	4h	2h	2h	2h	2h

on call time

During on call time, only incident reports of Incident Priority Classes P1 are received and processed.

:::note

P1 incidents must be reported by the customer to the Service Desk by telephone.

:::

:::tip

The customer can submit P2-P4 incident reports to Inventx during on call time. However, these will only be processed on the following business day and in accordance with the SLA of on site time.

:::

The following reaction times apply during on call time for the individual Incident Priority Classes:

Table: SLA – Reaction times on call time

Incident Priority Class	Bronze	Silver	Gold	Platinum	Rhodium
P1	-	2h	2h	2h	2h
P2	Next Business Day	Next Business Day	Next Business Day	Next Business Day	Next Business Day

Incident Priority Classes

The determination of Incident Priority Classes (P1-4) is made when the incident is recorded with the customer's incident reporter. The following definition applies:

Table: SLA - Incident Priority Classes

Incident Priority Class	Severity	Description
P1	Outage or high impact	A central service or base service is not available. One or more sites are affected. It is a global issue.
P2	Medium impact	An important functionality cannot be executed. There is no workaround. Many users are affected.
P3	Minor impact	Other cases, in particular a functionality cannot be executed. However, there is a workaround. Only one user is affected.
P4	No impact	General information, inquiries, etc.

Priority changes may be made during processing if the incident's impact requires a different priority. Downgrades of incidents are only made after consultation with the customer.

Reaction Time

Reaction time is the time span within the service time between the recording of the incident in the ITSM system by the support organization, the user, or a monitoring tool (the earliest applies) and the start of the support organization's troubleshooting effort.

Recovery Time

Recovery time is the actually measured time in hours - within the service time - to resolve an outage or provide a workaround.

In case of involvement of third parties or delays caused by the customer (e.g., missing approval for a proposed measure), the measurement of recovery time is interrupted as follows:

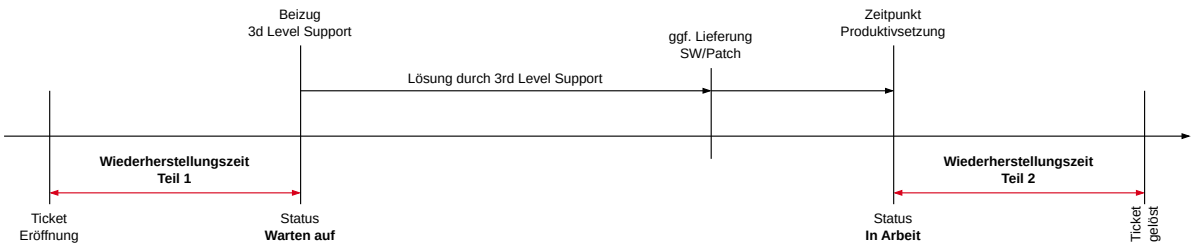

Figure: Calculation of recovery time

Backup

The Backup Service is based on a highly available, scalable, and performant platform in Inventx's data centers, which ensures the backup of data and, if needed, the recovery of a service or a complete VM.

Service Architecture

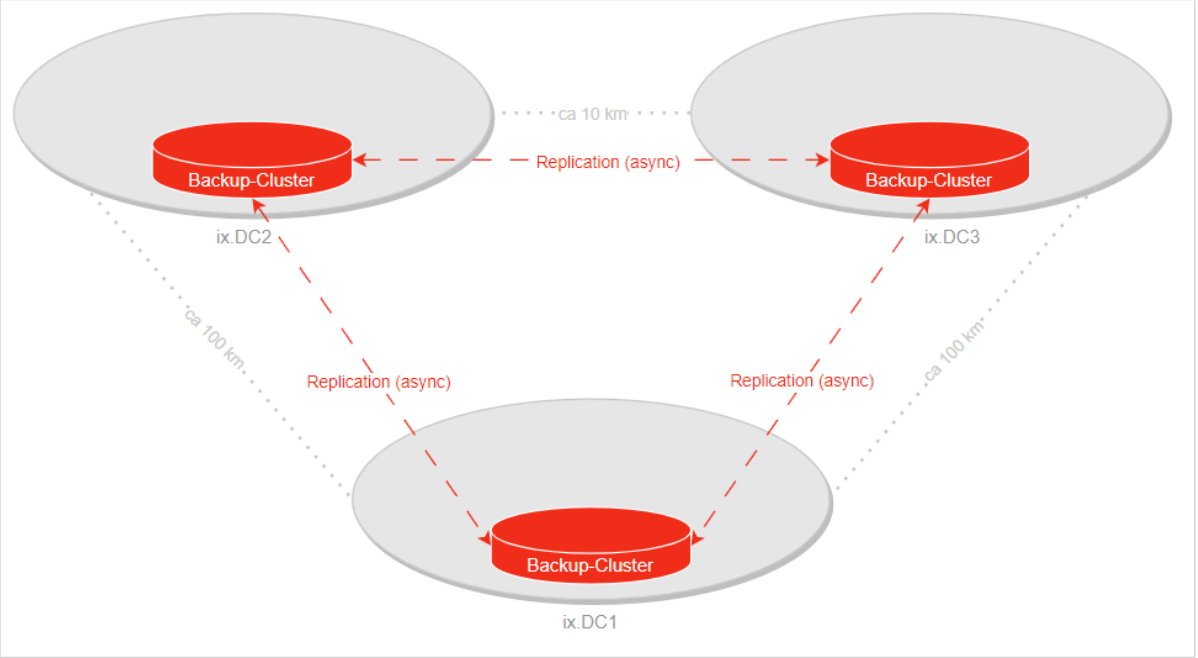

Figure: Calculation of recovery time

Service Scope

Table: Backup Service Scope

Features	
Access Control	■
Data Encryption	■
Immutable Backups	■
Backup Profile	■
Monitoring & Logging	■
Data Integrity	■

Data Recovery	■
Disaster Recovery (RTO and RPO)	■

Service Options

Access Control

Access to the backup infrastructure is role-based (RBAC/ixPAM) and with multi-factor authentication (MFA).

Data Encryption

Backup data is encrypted both during transfer and throughout the entire retention period (AES-256, FIPS 140-2 compliant encryption of data in flight and at rest).

Immutable Backups

Backup data cannot be modified or deleted during the entire retention period (Immutable Backup & DataLock/WORM). This provides effective protection against ransomware or deliberate deletion of backup data.

Backup Profile

The following backup profiles can be selected:

Table: Backup Profile

Service Level	Bronze	Silver	Gold	Platinum*	Rhodium
Backup Location	Local	Remote		Remote & DR Datacenter	
Backup Interval	Daily				
Backup Retention	No Backup, 14, 40, 100, 200, or 400 days				

* the Service Level Platinum is not available for Community ix.Cloud

Backup Location

Primary data is generally backed up to the backup infrastructure in the remote data center. For Service Level Platinum*/Rhodium, an additional replication (copy) of the data to the DR data center is performed.

:::info As an option, additional backup copies within or outside the Inventx data centers can be ordered via "Generic Request". :::

Backup Interval

The interval defines the time interval at which a data backup is created. At least one incremental backup is performed daily; for databases, logs (archive logs, transaction logs, etc.) are backed up multiple times hourly.

Backup Retention

Retention determines the retention period of backups in days. After the selected period expires, the data is permanently deleted.

:::caution With the backup profile "No Backup", the customer waives data recovery according to explicit request! This applies both to primary data at OS, database, application level, etc., and also includes cloning via backup-restore to a VM/DB for which the backup profile "No Backup" was selected! :::

Monitoring & Logging

Inventx monitors the backup process and ensures that it is performed according to the planned schedule. All relevant access, changes to configuration or system parameters are recorded in an audit log.

Data Integrity

The backup solution used ensures that data is encrypted during transfer and, after the copy is stored, is protected against accidental or malicious changes as well as premature deletion through WORM technology.

Data Recovery

Monolithic recovery of a complete VM is performed via self-service through the ix.Cloud Portal. File or database-based recovery typically requires ordering via "Generic Request".

Recovery in Case of Disaster (RTO and RPO)

In the case of a disaster, RTO and RPO define the maximum duration for recovery (RTO) of an application, system, and/or process, and the maximum data loss (RPO).

Table 1: SLA - Recovery in Case of Disaster (RTO and RPO)

Service Level	Bronze	Silver	Gold	Platinum	Rhodium
Recovery Time Objective (RTO)	-	Best Effort	48h	2h	2h
Recovery Point Objective (RPO)	Best Effort	24h	15min	0min	0min

Recovery Time Objective (RTO)

The situation in a disaster can vary greatly and has an impact on this service level. The value depends very much on the number of simultaneous recoveries, i.e., with multiple simultaneous recoveries, the value per recovery can be lower. For a recovery, a reference value in the range of 200-400 MB/s can be assumed.

Recovery Point Objective (RPO)

Damage events caused by manipulated or corrupted data are covered exclusively by the backup-relevant quality elements in the SLA. This means that if corrupted data exists in the live system, it can only be corrected from a backup, and the specified RPO thus does not apply.

Service Request Management

Service Request Management allows the customer to order services defined in the service catalog. The following target values apply:

Table: SLA - Service Request Management

Service Level	Entry Point	KPI
Automatic Service Request	ix.Cloud Portal	Max. 2h after order (7 x 24h)
Standard Service Request (non Automatic)	ITSM Portal	Best Effort or as specified in ITSM Portal
Generic Request (non Automatic)	ITSM Portal	Best Effort

Service Level Reporting

The Service Level Report is electronically sent to the customer monthly as proof of the agreed service levels of services provided by Inventx. This report is sent to the customer no later than the 5th business day of the following month.

Service Desk

The Inventx Service Desk is available to the customer's 1st level support for all qualified 2nd level requests regarding services provided on the basis of this service catalog during defined service times excl. holidays. Essentially, the following services are provided:

Table: SLA - Service Desk

Feature	Description

Incident reception	Reception of incidents and requests during service times via customer portal, telephone, ITSM interface, or email in German and English
Incident analysis and triage	Conducting incident analysis and assigning the incident to the responsible resolution group
Ticket management	Recording and maintenance of information in Inventx's IT Service Management Tool
Deadline monitoring	Monitoring of incident processing and escalation in case of deadline overruns
Information provision	Information on open incidents

ix.Cloud Edge

Service Description: Public DNS Service ix.Cloud Edge

The Public DNS Service is part of the Internet Perimeter ix.Cloud Edge and enables authoritative name resolution of public zones. Zone distribution takes place globally via an Anycast network. The service meets regulatory requirements (revDSG / FINMA) through localized data storage and audit-proof processes.

The service is available exclusively under the SLA Rhodium and must be ordered via "Generic Request".

Service Architecture

Zone data management is handled centrally on the Inventx infrastructure via a server that is not accessible from the Internet (Hidden Primary). DNS queries from around the world are answered via an upstream Anycast network.

This distributed architecture eliminates single points of failure and reduces latencies.

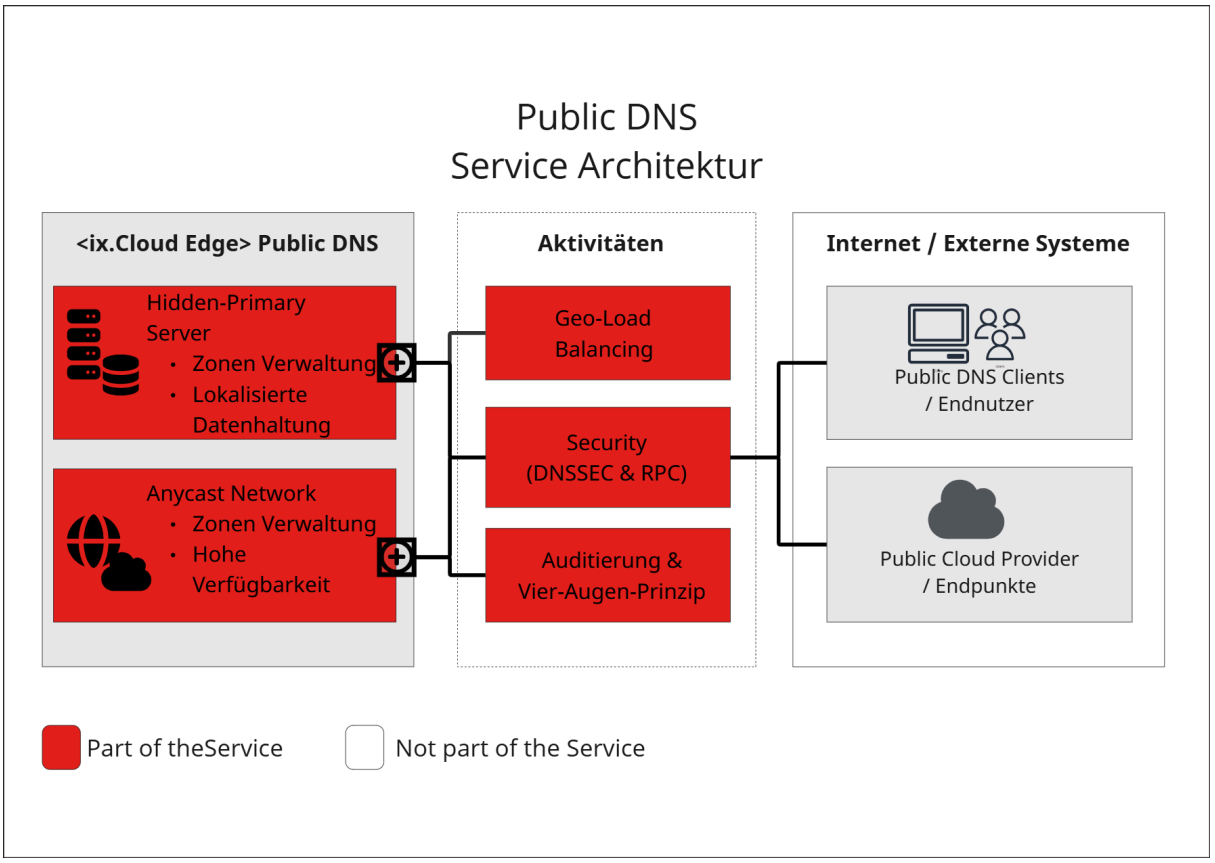

Image: Public DNS Service Architecture

Service Scope

Table: Public DNS Service Scope

Performance Feature	SLA Rhodium
Initial Setup	■
Hidden Primary Architecture	■
Anycast Zone Distribution	■
GeoLoad Balancing	■
Security (DNSSEC & RPC Listen)	■
Audit Security & Compliance	■

(■ = Included in standard service, □ = Project-based / one-time service)

Service Options

The following performance elements define the service and its operation:

Initial Setup

The initial specification, configuration, and migration of existing public zones are carried out as part of an initial setup in collaboration with the customer.

Hidden Primary & Zone Distribution

Zone management takes place on an internal primary server (Hidden Primary) to reduce the attack surface. Zone information is distributed to clients exclusively via the Anycast network. Sovereign fallback runs via the Hidden Primary.

GeoLoad Balancing

DNS queries can be dynamically delegated to Internet endpoints across different locations. This includes ix.Cloud locations as well as connected public cloud providers for distributed load management.

Security & Administration

- **DNSSEC:** Cryptographic protection of DNS responses against manipulation.
- **RPC Listen:** Implementation of Response Policy Zones for active filtering and control of name resolutions.
- **Administration:** Zone management strictly according to the four-eyes principle including audit-proof auditing.

Supported Record Types

The following DNS records are supported for public zones (forward mapping):

Table: Supported Record Types

Record Type	Forward Mapping	Reverse Mapping	Purpose / Description
A Record	■		Resolution of a hostname to an IPv4 address.
AAAA Record	■		Resolution of a hostname to an IPv6 address.
CNAME Record	■		Alias entry that points one hostname to another.
MX Record	■		Definition of the responsible mail servers for email receipt for the domain.

NS Record	■		Definition of the responsible name servers for a zone or subzone (delegation).
PTR Record		■	Resolution of an IP address to a hostname (reverse mapping), often used to verify mail servers for spam prevention.
SRV Record	■		Definition of the availability of specific services (including port and protocol).
TXT Record	■		Storage of text information, often used for security and verification purposes (e.g., SPF, DKIM, DMARC).
CAA Record	■		Definition of which certificate authorities (CAs) are authorized to issue TLS/SSL certificates for the domain.

Datacenter Services

Inventx's Datacenter Services encompass the entire lifecycle from planning (Plan) through construction (Build) to operation (Run) of the central Datacenter infrastructure in the Inventx datacenters ix.DC1 (Chur), ix.DC2 (St. Gallen), and ix.DC3 (Gais). The following table provides an overview of the service features per service model. This forms the basis for all services described in this service catalog.

Table: Datacenter Service Scope

Service Feature	IaaS	PaaS
Datacenter locations on Swiss territory	✓	✓
Data center complex distributed across two geographically separate site areas with route-redundant backbone connectivity	✓	✓
Autonomous alarm system for all critical infrastructure components	✓	✓
Access control systems against unauthorized entry additionally with isolation systems and personnel airlocks	✓	✓
Intervention and escape routes in case of emergency	✓	✓
All accesses designed with break-in resistance	✓	✓

Intrusion detection system	✓	✓
Video surveillance	✓	✓
Two independent, separate power supplies	✓	✓
Redundant uninterruptible power supply (UPS)	✓	✓
Power generators with network backup system (Diesel)	✓	✓
Overvoltage protection and lightning protection	✓	✓
Redundant power supply within the rack	✓	✓
Climate monitoring	✓	✓
Redundant cooling of racks	✓	✓
Early fire detection	✓	✓
Hand-held fire extinguishers	✓	✓
Water sensors	✓	✓

The following Datacenter Services are available:

Table: Datacenter Services

Service Name	Service Description
Cloud Connect	Connecting on-premise IT with ix.Cloud
Rack Collocation	Rental of rack space in the Inventx data center

Cloud Connect

This service includes the operation and administration of the communication infrastructure in Inventx's data centers as a link between the system hardware components and their communication interfaces for the entry and exit of the data centers.

Service Architecture

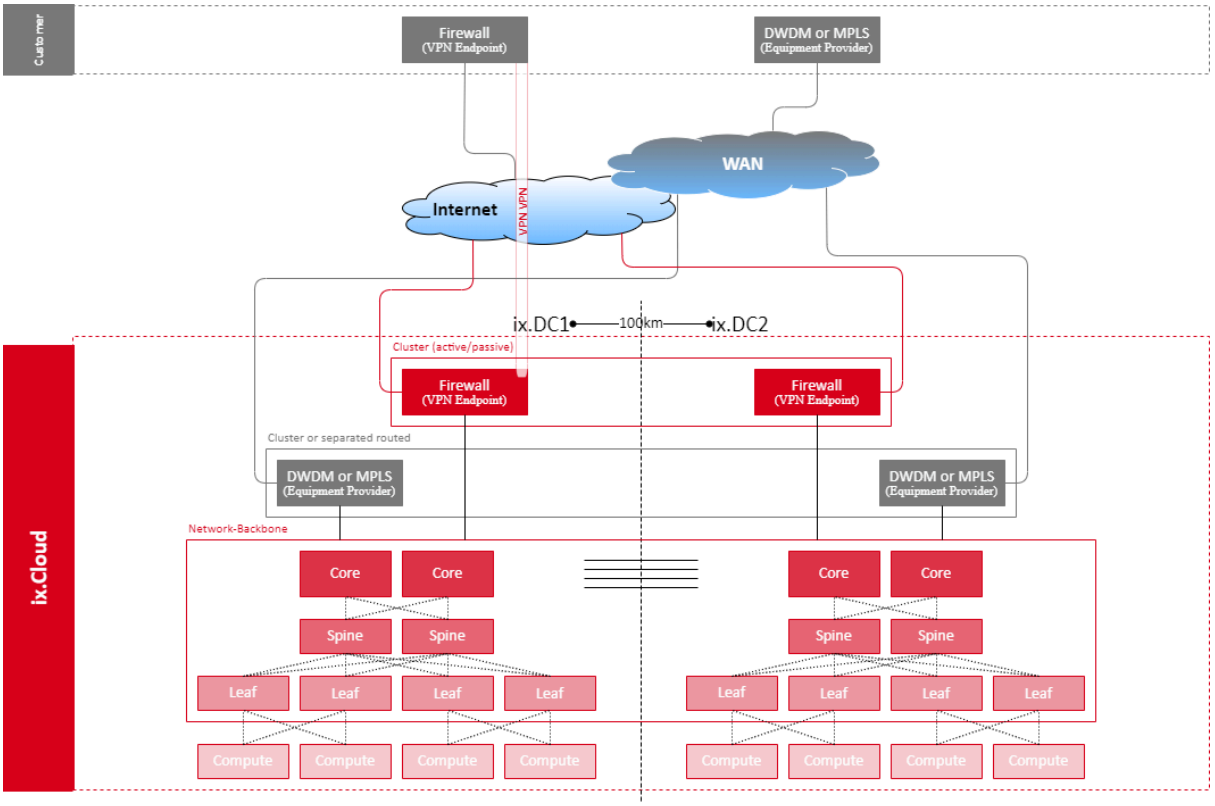

Figure: Cloud Connect Service Architecture

Service Scope

Table: Cloud Connect Service Scope

Service Feature	IaaS	PaaS
Datacenter LAN Infrastructure	■	■
Datacenter Interconnect	■	■
Private Connectivity Provider Equipment	□	□
Shared Internet Access	□	□

Service Options

With the "Cloud Connect" service, the customer can establish a private connection to their services in ix.Cloud, which can be customized according to requirements.

Datacenter LAN Infrastructure

The services of the IaaS and PaaS service models are operated using the datacenter LAN infrastructure managed by Inventx based on the following service features:

Table: Cloud Connect Datacenter LAN Infrastructure

Service Feature	IaaS	PaaS
Strict zone concept at infrastructure and security level	■	■
Zone separation through firewalls	■	■
No sharing of infrastructure components across zones	■	■
Dedicated customer zones	■	■
Dedicated and shared service zones	■	■
Scalable and redundant security components	■	■

Datacenter Interconnect

The Inventx data centers are securely and efficiently connected through the Datacenter Interconnect. This WAN connection is operated by Inventx as follows and as an integral part of the corresponding IaaS and PaaS services:

Table: Cloud Connect Datacenter Interconnect

Service Feature	IaaS	PaaS
Route-redundant, private connection between Inventx data centers	■	■
Encrypted communication	■	■

Private Connectivity Provider Equipment

The customer can establish a private network connection customized according to their requirements to the Inventx data centers via the WAN infrastructure of their connectivity provider. The customer provides a connection via MPLS/DWDM or Dark Fiber, which is then terminated in the Inventx data center and thus provides access to the services operated in ix.Cloud. The necessary infrastructure of the customer's connectivity provider is operated on the basis of the Rack Collocation service in the Inventx data center. The customer must ensure compliance with IT security standards commonly used in the financial industry, in particular Distributed Denial of Service (DDoS) and 1st firewall level.

Shared Internet Access

Shared Internet Access is based on Inventx's own IP range including BGP peering to two different providers and terminates with one connection each to the global firewall instance in the geo-redundant Inventx data centers. This global, transparent firewall instance from Inventx is a mandatory connection element to the customer's first firewall level, which serves as a VPN endpoint. The following firewall policies are implemented on this global, transparent firewall instance:

Table: Cloud Connect Shared Internet Access Firewall Policies

Service Feature	IaaS	PaaS
Distributed Denial of Service (DDoS)	■	■
Botnet Control Services	■	■
Application Control Analytics	■	■

The IP addressing between the global, transparent firewall instance and the customer's first firewall level is provided by Inventx. These chargeable IP addresses can be obtained in block form in the following quantities: 2/4/8/12/16/20/30/40/50 addresses. The number of IP addresses can be changed on the 1st of the following month via "Generic Request" while observing a notice period of 3 business days.

This service is offered exclusively in the Platinum SLA. It should be noted that Inventx does not assume responsibility for the customer's on-premise VPN endpoint and only assumes responsibility for the VPN endpoint in ix.Cloud if it terminates on a component operated by Inventx.

The desired bandwidth can be reserved for the customer and can be changed on the 1st of the following month via "Generic Request" while observing a notice period of 3 business days. The following bandwidths are available:

Table: Cloud Connect Shared Internet Access Bandwidths

Service Feature	IaaS	PaaS
20 Mbps	■	■
50 Mbps	■	■
100 Mbps	■	■
200 Mbps	■	■

Rack Collocation

With the Rack Collocation service, the customer rents a complete, dedicated rack or the desired number of rack units in a shared rack in Inventx's data centers via "Generic Request". The systems are managed by the customer themselves in this case. Power consumption is billed individually based on actual consumption. The electricity price is adjusted annually according to the price level of electricity suppliers.

This service is not available as a standalone service, but only in combination with other services from this service catalog. Customers benefit from this service in cases where, in addition to the cloud service, a

solution for hosting applications, systems, or appliances that are not cloud-capable is also needed.

Service Architecture

n/a

Service Scope

Table: Rack Collocation Service Scope

Service Feature	Shared	Dedicated
Dedicated rack	-	■
Individual rack units (RU)	■	-
Network connections to the appropriately defined customer network zones and outward (e.g., Internet)	■	■
Power on Demand	■	■
Remote Hands and Eyes	□	□
Customer Access	□	□

Service Options

As a supplement to the Rack Collocation service, the customer can obtain additional services by arrangement as follows:

Remote Hands and Eyes

Required hands and eyes services must be registered via "Generic Request". These are billed according to the agreed hourly rate on a time and materials basis.

Customer Access

Customer access is exclusively possible upon advance registration via "Generic Request" and only in strictly defined exceptional cases that have been approved via a defined process. The following guidelines must be observed:

- Customers are accompanied in at least 1:1 supervision by Inventx employees.
- The registered persons of the customer must identify themselves with an official identification document before accessing the data center.
- A log of accesses is maintained.

- Persons are not permitted to bring mobile phones or smartwatches with photo functions or cameras into the data center rooms.
- All work performed by Inventx is billed on a time and materials basis according to the master agreement.

Network Services

The Network Services in ix.Cloud include the following options to network services within ix.Cloud with each other and to connect external services with services in ix.Cloud:

Table: Network Services

Service Name	Service Description
Firewall	Elementary security between different subnets
Server Proxy	Indirect and restrictive communication for ix.Cloud servers
Load Balancer	Distribution of incoming connections to applications or service endpoints
Web Application Firewall	Provides security for online services against malicious internet traffic and filters threats such as OWASP TOP 10 which negatively impact online applications
Private DNS	Resolution of IP addresses to DNS names within ix.Cloud
Secure Mail Relay	Secure email delivery from all your systems within ix.Cloud
Hosted Software Appliance	Virtual servers for software appliances, without support for Microsoft Hyper-V

Firewall

The Firewall Service is a cloud-based network security service managed by Inventx that protects virtual network resources within ix.Cloud. This service is exclusively available in the Platinum SLA.

Policies for application and network connectivity are created and logged centrally across all ix.Cloud subscriptions and virtual networks. Networks and IP addresses can be enabled between different sources and destinations. Connections are restricted to defined services (particularly TCP/UDP ports).

Service Architecture

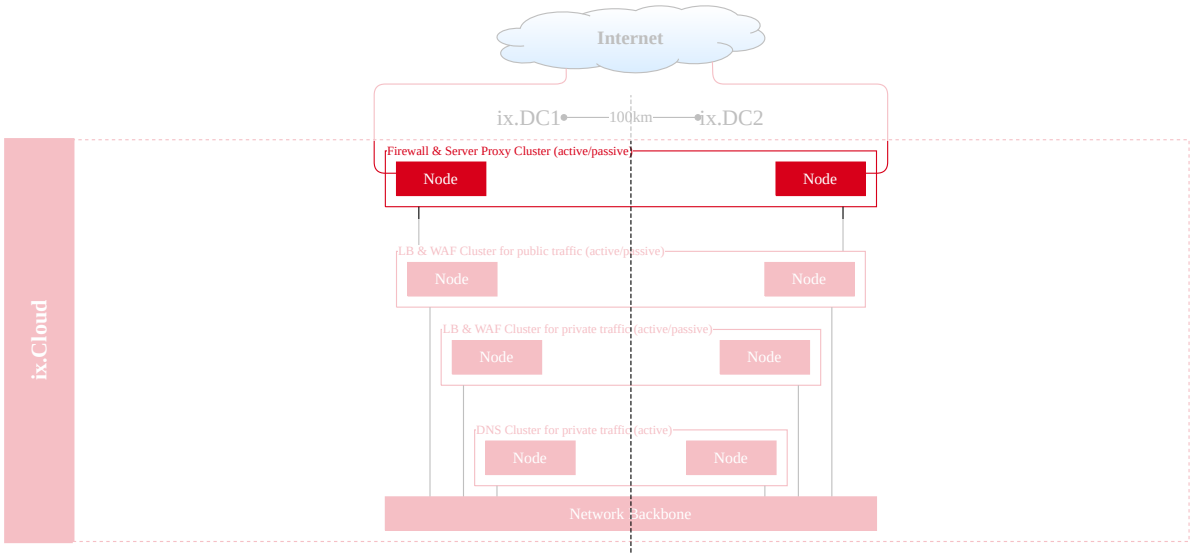

Image: Firewall Service Architecture

Service Scope

Table: Firewall Service Scope

Feature	Platinum
Initial Setup	☐
IT Basic Protection	■
FQDN Application Filter Rules	■
Filter Rules for Network Traffic	■
SNAT Support	■
Logging	■
DNAT Support	☐
UTM Features	☐

IT Basic Protection

Upgrade/Patching

The firewall infrastructure is updated twice yearly, unless critical security vulnerabilities occur that require immediate upgrades.

Logging

Logging is performed via a central instance to which all logs are sent for analysis and evaluation. Logs are retained live for 14 days and then archived for 800 days.

Malware Protection

Malware protection is based on hardened appliances that are checked for integrity during boot via secure boot. Secure Boot ensures that only trusted and signed firmware and software are loaded onto the hardware. The boot process starts with immutable code embedded in the hardware and verifies subsequent components. Devices regularly check the integrity of installed firmware through cryptographic signatures. Firewall ASICs are specially developed hardware security processors that accelerate many security functions such as encryption and Deep Packet Inspection (DPI) at the hardware level. System partitions that are responsible for executing the operating system and configuration data are isolated from user data and applications.

Service Options

With the Firewall Service, the customer can obtain additional services by arrangement:

Initial Setup

Implementation of the initial configuration of the Firewall Service is charged as part of a project. During this process, the customer-specific ruleset is specified in collaboration with the customer and then implemented. All changes must be requested via a "Generic Request".

FQDN Application Filter Rules

Customers can restrict outgoing HTTP/S traffic to a specified list of fully qualified domain names (FQDN), with wildcard entries needing to be implemented via the UTM Features option. The FQDN feature does not require SSL termination.

Filter Rules for Network Traffic

Customer-specific network filter rules for allowing or denying based on source and destination IP address, port, and protocol are maintained centrally by Inventx. The firewall is stateful, allowing distinction between legitimate packets for different types of connections. Rules are enforced and logged across all ix.Cloud subscriptions and virtual networks.

SNAT Support

All IP addresses for outgoing traffic from the virtual network in ix.Cloud are translated into the firewall's public IP address (Source Network Address Translation). You can identify and allow traffic from your virtual network to remote destinations on the Internet. SNAT functionality can optionally also be implemented for internal traffic within ix.Cloud.

Logging

All connections terminating at Inventx and for which Inventx is responsible are logged. This means that all incoming and outgoing connections from external and internal sources are recorded.

The retention period is 2 years. Log reporting to the customer can be requested as needed via "Generic Request".

DNAT Support

Incoming network traffic to the firewall's public IP address in ix.Cloud is translated into private IP addresses in the customer's virtual networks (Destination Network Address Translation) and filtered. DNAT functionality can optionally also be implemented for internal traffic within ix.Cloud.

UTM Features

Optional UTM Features (Unified Threat Management) are specified together with the customer and then operated by Inventx.

Server Proxy

If servers in the ix.Cloud should not communicate directly with the Internet to increase IT security, the Server Proxy enables server systems in the ix.Cloud to call only certain addresses on the Internet via defined rules. Access to this explicit proxy is controlled and restricted through various technologies such as web filters, virus filters, categories, and application control. All accesses (exceptions possible) are checked using Deep-Inspection (breaking down traffic), with Inventx adhering to legal data protection requirements.

The Server Proxy is available as a Shared Service exclusively in the Platinum SLA and must be ordered and managed via the "Generic Request". The Server Proxy can only be used by server-based systems and is not available for clients.

Service Architecture

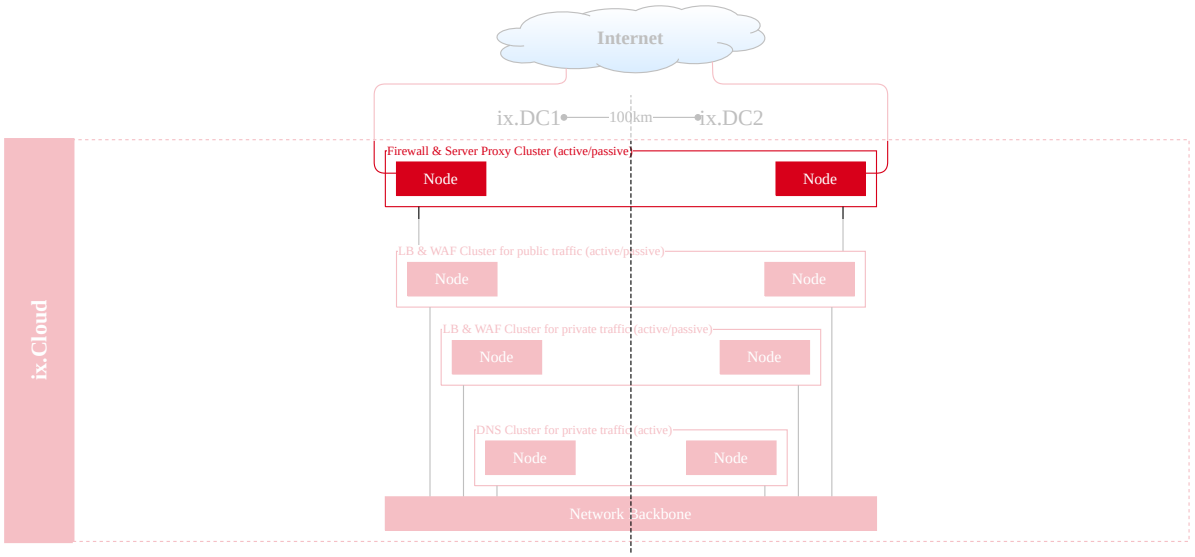

Image: Server Proxy Service Architecture

Service Scope

Table: Server Proxy Service Scope

Feature	Platinum
Initial Setup	□
IT Basic Protection	■
Default Policy for Protocols and Ports	■
Default Policy for Web Filter	■
Default Policy for Virus Filter	■
Default Policy for Application Control	■
Default Policy for Deep-Inspection	■

IT Basic Protection

See description in section Firewall IT Basic Protection.

Service Options

Currently, customers cannot obtain optional services with the Server Proxy Service.

Initial Setup

The Server Proxy as a Shared Service cannot be individually customized to customer requirements and only global filter configurations are available.

If individual configuration of the Server Proxy Service is desired, this will be developed within the scope of a project and charged accordingly. During the project, the customer-specific ruleset is specified in cooperation with the customer based on the Inventx Standard Ruleset and implemented on a private Server Proxy.

Default Policy for Protocols and Ports

Inventx maintains a standard policy for all server-based systems in the ix.Cloud. The Inventx standard permits the following services and ports:

Table: Server Proxy Default Policy for Protocols and Ports

Protocol	Proxy Port(s)	Socks Port(s)
HTTP	80	-
HTTPS	443	-
SSH	-	22

Default Policy for Web Filter

The web filter categorizes all Internet pages based on predefined algorithms (manufacturer specification), which are either allowed or blocked. The global configuration is based on Inventx standards, with the following categories being permitted:

Table: Server Proxy Default Policy for Web Filter

Web Category	Allowed
Business	■
Finance and Banking	■
Information Technology	■
Information and Computer Security	■

Default Policy for Virus Filter

Inventx's global standard policy determines which incoming and outgoing content is scanned for viruses, thus preventing the introduction of malicious software, with all HTTP and FTP traffic being analyzed.

Default Policy for Application Control

Using Application Control, unwanted features of websites are disabled. For example, streaming of audio and video files, chatting, and uploading and downloading of files can be prevented. Inventx's global ruleset is defined as follows:

Table: Server Proxy Default Policy for Application Control

Web Category	Blocked
Webmail (e.g. Gmail or GMX)	■
Game	■
Mobile	■
P2P	■
Remote Access	■
Social Media	■
Video/Audio	■
VOIP	■
Unknown Applications	■

Default Policy for Deep-Inspection

The Server Proxy inspects HTTPS packets, with the categories "Health and Wellness" and "Finance and Banking" not being analyzed for data protection reasons. The HTTPS traffic is decrypted, inspected, re-encrypted, and forwarded to the destination.

Load Balancer

A load balancer distributes the data traffic of a specific service endpoint across multiple targets. It detects faulty targets and forwards data traffic only to intact targets. This allows the availability and performance of an application to be optimized.

For HTTP/HTTPS applications, a Layer 7 load balancer is recommended, while for applications using TCP/UDP protocols, a Layer 4 load balancer is recommended.

This service is only available in the Platinum SLA. Layer 4 load balancers can be ordered and managed via the portal in self-service. The remaining load balancers must be ordered and managed via the "Load Balancer" service request.

Service Architecture

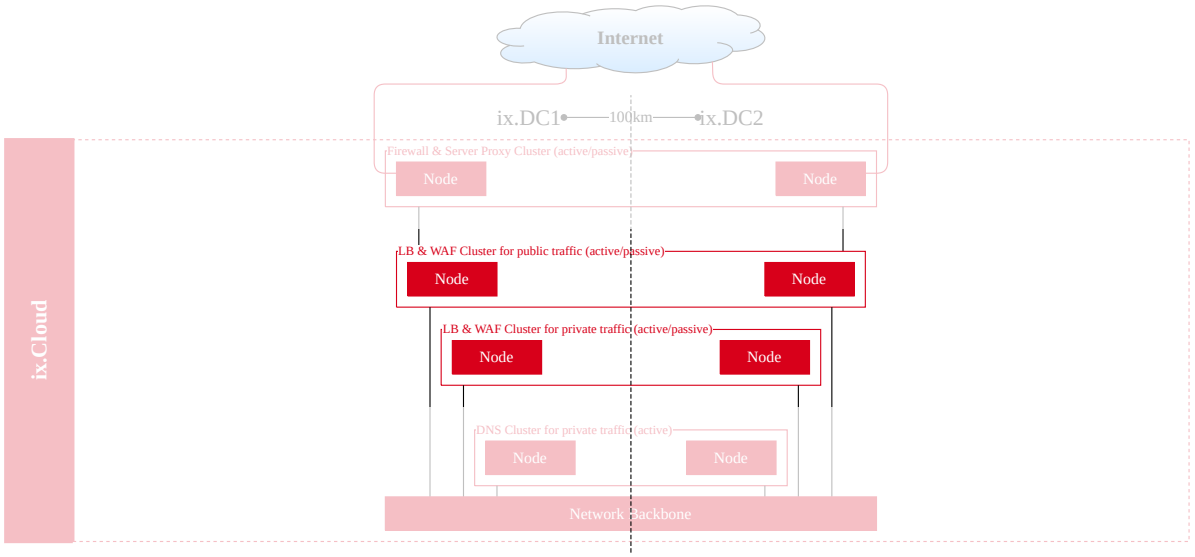

Image: Load Balancer Service Architecture

Service Scope

Table: Load Balancer Service Scope

Feature	Layer 4	Layer 7
Initial Setup	☐	☐
IT Basic Protection	■	■
Service Management	■	■
Bandwidth (5 MBit/s)	■	■
Service IP Address	■	■
Service FQDN	-	■
Protocols/Ports	■	■
Load Balancing	■	■
Persistence	■	■
X-Forwarded-For	■	■
Default DDoS Profile	■	■

Health Monitoring	■	■
Error Page	-	■
Maintenance Page	-	■
SSL Offloading/Bridging	-	■
Host Header Forwarding/Rewriting/Redirecting	-	■

IT Basic Protection

Patch Management

The load balancer infrastructure is updated at least twice yearly, unless critical security vulnerabilities occur that require immediate upgrades.

Malware Protection

Malware protection is based on hardened appliances that are checked for integrity during boot. The load balancer infrastructure uses Secure Boot and Image Signing to ensure that only signed and trusted software components are executed. The load balancers support RASP functionality (Runtime Application Self-Protection) that monitors and protects applications during runtime.

Service Options

Through the options listed in this section, a load balancer can be configured in different ways.

Initial Setup

The effort required to set up a load balancer is heavily dependent on the customer's desired individual requirements. Therefore, the initial setup of a load balancer is charged on a time-and-materials basis.

Service Management

Service management includes updating the software components and security patterns used, resource management, and backup of the infrastructure. Certificate Lifecycle Management (creation/request, integration, replacement/renewal of certificates) is charged separately on a time-and-materials basis.

Bandwidth

Bandwidth (data throughput) is individually configurable per service. A bandwidth of 5 megabits per second (MBit/s) is included in the base price. In steps of 5 MBit/s, the service can be scaled up to a maximum of 40 MBit/s according to a separate price list and ordered according to requirements (see table below). Billing is based on the number of 5 Mbit/s units ordered.

If more data is transported via the load balancer than the ordered bandwidth allows, packet losses (packet drops) are generated. If packet losses are detected, a bandwidth increase can be ordered. For a

Layer 4 load balancer, the bandwidth increase can be made directly in the portal. For a Layer 7 load balancer, the bandwidth increase must be requested via a Load Balancer service request.

Table: Load Balancer Bandwidth

Bandwidth	Layer 4	Layer 7
5 MBit/s	■	■
10 MBit/s	□	□
15 MBit/s	□	□
20 MBit/s	□	□
25 MBit/s	□	□
30 MBit/s	□	□
35 MBit/s	□	□
40 MBit/s	□	□

Service IP Address

One IP address can be assigned per service. If it is a private IP address, this is free of charge. A public IP address incurs additional costs according to the price list. Private IP addresses are not routed on the internet and can only be used within a local network. Public IP addresses are routed on the internet.

Service FQDN

One or more URLs can be pointed to (DNS entry) a Layer 7 service IP address (VIP). A DNS entry for the respective URL is a prerequisite for end-to-end (client-server) communication.

Protocols/Ports

Unless otherwise specified at the time of order, Layer 4 will use standard ports 80/443 and Layer 7 will use the HTTPS protocol for service setup.

The following table shows the possible protocols and ports per service and layer.

Table: Load Balancer Protocols/Ports

Protocol/Port	Layer 4	Layer 7
TCP (all possible ports)	■	□
HTTP	-	□

HTTPS	-	■

Load Balancing

By default, None (Round-Robin) is enabled. The load balancing option serves to distribute the load with the goal of equally loading the end systems.

Table: Load Balancer Load Balancing Options

Load Balancing Option	Layer 4	Layer 7
None (Round Robin) Each new request is sent to a server in the pool, then top-down from the beginning again.	■	■
Least Connection Connections are sent to the server that currently has the fewest open connections.	□	□
Least Load Connections are sent to the server that is currently least loaded.	□	□
Fastest Response Connections are sent to the server that responds the fastest.	□	□
Fewest Servers An algorithm calculates how many servers are needed to handle the request. Requests are only sent to the first server in the pool; once it reaches its capacity limit, traffic is passed top-down to the next server in line.	□	□

Persistence

By default, no "Persistence" is configured. By using the Persistence option, the session is bound to a specific end system. This ensures that requests during a session are always processed by the same end system.

Table: Load Balancer Persistence Options

Persistence Option	Layer 4	Layer 7
Client IP The client IP is used as an identifier and assigned to the server.	☐	☐
TLS The information is embedded in the client's SSL/TLS ticket ID.	-	☐
APP Cookie Reads existing server cookies or embedded URI data such as JSessionID.	-	☐
HTTP Cookie Inserts a cookie into the HTTP response(s).	-	☐
Custom HTTP Header The customer can create custom specifications for mapping header values to specific servers.	-	☐

X-Forwarded-For

With X-Forwarded-For, it is possible to transmit client IP addresses (original IP) to the target system via the header. The target system can use this information to, for example, show where the request originates from or to enable server-side black/white lists. This option can only be used with Layer 7 load balancer.

Default DDoS Profile

By default, a DDoS profile (built-in) is enabled, which detects and prevents network attacks on Layer 3, 4, and Layer 7.

Table: Load Balancer Default DDoS Profile

Default DDoS Profile	Layer 4	Layer 7
Layer 3 SMURF, ICMP Flood, Unknown Protocol, Tear Drop, IP Fragmentation	■	■

Layer 4 SYN Flood, LAND, Port Scan, X-mas Tree, Bad RST Flood, Fake Session, Bad Sequence Number, Malformed/Unexpected Flood, Zero/Small Window, Rate Limiting CPS per IP, SSL Errors, SSL Renegotiation	■	■
Layer 7 Request Idle Timeout (10,000ms), SlowPost (30,000ms), SlowLoris (30,000ms), Invalid Requests	-	■

Health Monitoring

With Health Monitoring, the load balancer sends requests to the target system at intervals and expects a response within a set time window for each request.

If the respective requests to a target system are not answered, the target system is marked as unreachable. Consequently, client-server requests are no longer forwarded to that target system.

Table: Load Balancer Health Monitoring Options

Health Monitoring Option	Layer 4	Layer 7
TCP (custom-client-request/custom-server-response) Waits for a complete TCP connection on a specifically requested port.	□	□
ICMP Sends a ping and expects a response from the "pinged" server.	-	□
DNS (request/response) Checks whether the "name server" can correctly resolve a name to a specified entry.	-	□
HTTP/S (custom-client-request-header/-body, custom-server-response) Checks the specified "response code" for correctness.	-	□
External	-	□

Customer-specific health checks can be performed via script command. (wget, netcat, curl, dig, mysql-client, snmpget)		

Error Page

By default, a "Default Error Page" is displayed for a Layer 7 service, which informs the client about the connection error. If the layout or content of the page does not meet requirements, a "Custom Error Page" can be created and provided to Inventx for integration.

Maintenance Page

If a maintenance page should be displayed for maintenance mode, this can be arranged via a request to Inventx or by the customer themselves. In the latter case, a script-based solution approach must be used; please inform us of your specific use case.

SSL Offloading/Bridging

When a Layer 7 service is ordered, SSL offloading is enabled by default. This option enables the load balancer to decrypt encrypted traffic to, for example, detect network attacks and prevent them based on WAF policies.

With SSL offloading, traffic is decrypted:

- Client to load balancer = encrypted
- Load balancer to target = unencrypted

With SSL bridging, traffic is decrypted and then re-encrypted:

- Client to load balancer = encrypted
- Load balancer to target = encrypted

For certificate issuance/integration, an existing PKI infrastructure in the customer environment is required. If this is not available, the customer must provide the required certificates to Inventx.

Certificate Lifecycle Management is not part of this service and must be ensured by the customer.

Host Header Forwarding/Rewriting/Redirecting

If a Layer 7 service is ordered, it is possible to perform forwarding, rewrites, and redirects based on host header information. Additionally, HTTP to HTTPS redirection can be performed upon request.

Web Application Firewall

The Web Application Firewall (WAF) operated by Inventx supports service connectivity via the load balancer by checking incoming HTTP traffic for security vulnerabilities or unauthorized data transmission

before it reaches the application server. Thus, the WAF service serves as an enforcement point for security policies that take place between the web application and the client application.

The WAF intercepts all HTTP requests and checks them using the previously defined rule set (security model) to identify whether it is unwanted data traffic (cross-site scripting, SQL injection, etc.). This approach prevents L7-DDoS attacks, which attempt to exploit security vulnerabilities in web-based applications or negatively impact the service.

This service is only available in the Platinum SLA and must be ordered and managed via the standard service request "Web Application Firewall".

Service Architecture

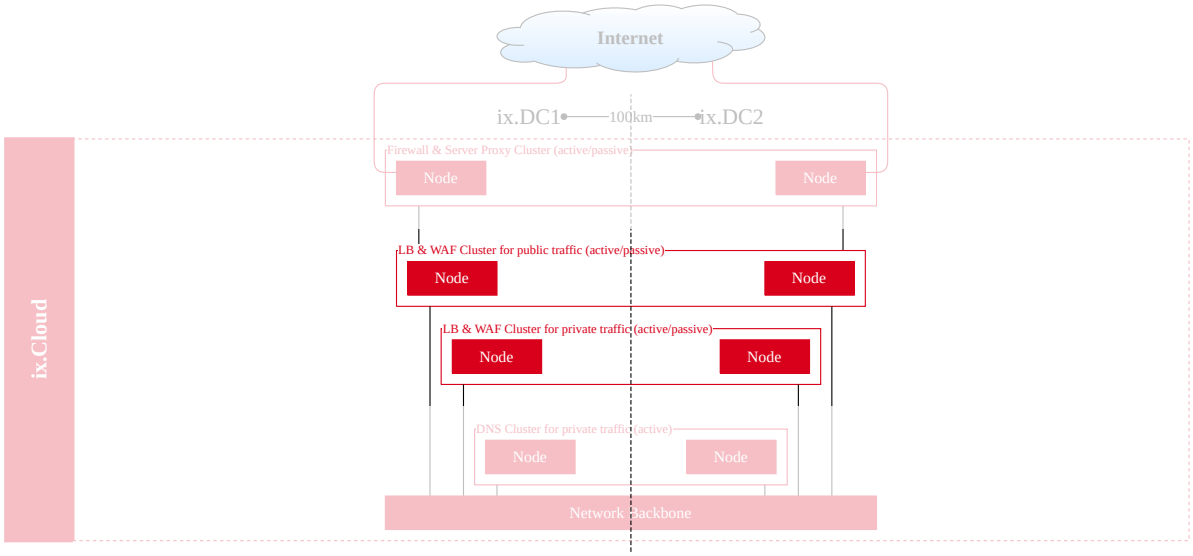

Figure: Web Application Firewall Service Architecture

Service Scope

Table: Web Application Firewall Service Scope

Service Feature	Platinum
Initial Setup	□
IT Basic Protection	■
Load Balancer Layer 7	■
Service Management	■
OWASP TOP 10 Rule-Set	■

Customer-Specific Rule-Set	□
Operating Mode Enforcement	■

IT Basic Protection

See description in chapter Load Balancer IT Basic Protection.

Service Options

With the WAF service, customers can obtain additional services by arrangement.

Initial Setup

The effort required to set up a WAF service is highly dependent on the desired individual requirements of the customer, especially regarding the rule sets to be defined. Therefore, the initial setup of a WAF is charged based on actual effort.

Load Balancer Layer 7

The basic configuration for the WAF consists of a Load Balancer Layer 7. Corresponding service options are described in the Load Balancer section in the Layer 7 variant.

Service Management

Service Management includes, among other things, the updating of the software components and security patterns used, resource management, and infrastructure backup.

WAF lifecycle management (analysis/adaptation of rule sets) is charged separately based on actual effort.

OWASP TOP 10 Rule-Set

The standard rule set consists of the OWASP Top 10 vulnerabilities. The Open Web Application Security Project (OWASP) is an international non-profit organization dedicated to the security of web applications. Its most well-known project is called OWASP Top 10. This is a report that covers the 10 most critical risks.

Customer-Specific Rule-Set

Certain applications require individual configuration of the rule set for proper operation. For this purpose, for example, exceptions are created for the OWASP rule set for unwanted detections. Modifications to the rule set must be requested via a "Generic Request".

Operating Mode Enforcement

The WAF service operates in Enforcement mode. The configured rule set is applied in production and web applications are protected accordingly, whether the WAF service is used for the test or production environment.

Private DNS

The Private DNS (Domain Name System) Service enables server and client systems to perform authoritative/reverse resolution of IP addresses to DNS names and vice versa. A customer's access is mapped in a private view (zone) of the global DNS system, which is maintained and operated by Inventx.

The Private DNS Service is available exclusively in the Platinum SLA and must be ordered via "Generic Request".

Service Architecture

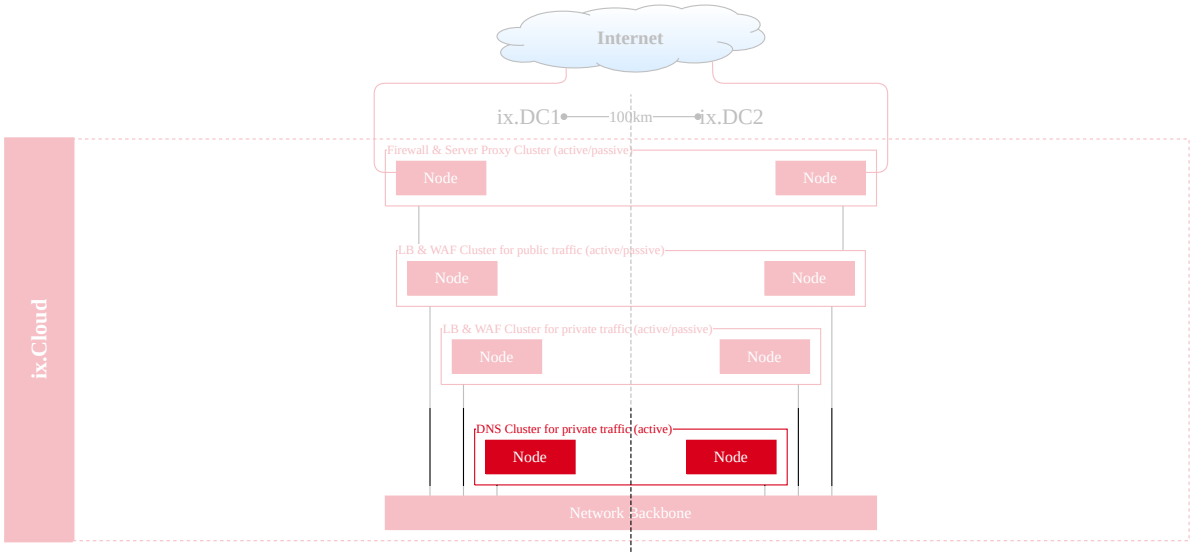

Figure: Private DNS Service Architecture

Service Scope

Table: Private DNS Service Scope

Feature	Platinum
Initial Setup	□
Authoritative Zone	■
Forwarding	■

Service Options

The DNS Service can be operated in different ways through the service elements listed in this chapter.

Initial Setup

The implementation of the initial configuration of the Private DNS Service is billed as part of a project. The configuration is specified in collaboration with the customer and then implemented.

Authoritative Zone

An authoritative zone is a zone for which the local (primary or secondary) DNS server references its own data when responding to queries. The local DNS server is responsible for the data in this zone and responds to queries without referring to another server. There are two zone types:

- **Forward-Mapping**: A forward-mapping zone is an area of the domain name space for which one or more nameservers are responsible for responding to name-to-IP-address queries.
- **Reverse-Mapping**: A reverse-mapping zone is an area of the network space for which one or more nameservers are responsible for responding to IP-address-to-name queries.

The following record types are possible per zone type:

Table: Private DNS Record Types

Record Type	forward-mapping	reverse-mapping
Host Record	■	-
A Record	■	-
CNAME Record	■	-
Alias Record	■	■
MX Record	■	-
NS Record	■	-
PTR Record	■	■
SRV Record	■	-
TXT Record	■	-

Forwarding

In a hybrid architecture, DNS forwarding can logically connect the ix.Cloud with a customer's on-premises environment. Through this option, customers can continue to use their existing local DNS servers as

authoritative.

Secure Mail-Relay

With the Secure Mail-Relay Service operated by Inventx, customers can securely send email (SMTP syntax) from ix.Cloud to the internet, with an Inventx address listed as the sender of the message.

Service Architecture

N/A

Service Scope

Table: Secure Mail-Relay Service Scope

Feature	Platinum
Initial Setup	□
Source and Destination	■
Malware Protection	■
Content Filtering	■
Session Handling	■
Addressing	■
Shipping over Internet	■

Service Options

The Secure Mail-Relay Service of ix.Cloud has the following features:

Initial Setup

The Mail-Relay Service is commissioned as part of a project. The service is specified in collaboration with the customer and then integrated. Orders and all changes must be requested via a "Generic Request".

Source and Destination

The Mail-Relay Service is only accessible within ix.Cloud, as all messages are received based on IP range and sender address. Any email address can be specified as the recipient.

Malware Protection

After receiving the message, a malware scan is performed. If a positive finding is detected, the message is rejected. Additionally, an antivirus outbreak filter with a 20-minute window is in place to enable timely malware identification.

Content Filtering

For security reasons, all messages are filtered. Files of type video, audio, archive as well as executables, scripts and encrypted files are filtered and replaced with an error text file. To prevent unauthorized data leakage, the following rules apply:

- Number of attachments per message: Maximum 5
- Message size: Maximum 10 MB
- Compression level of attachment: Maximum 12

Session Handling

A maximum of 1,200 messages per 30 minutes is possible, and a maximum of three parallel EHLO commands per SMTP connection. If these values are exceeded, throttling is automatically applied.

Addressing

Before sending to the destination, the sender is rewritten with a generic Inventx address (noreply@ixcloud.ch).

Shipping over Internet

Messages are always sent via the internet. Encrypted shipping (TLS) is strongly recommended (preferred), but not enforced.

Hosted Software-Appliance

For operating software appliances, virtual servers based on VMware ESX virtualization can be procured in cases where the software vendor does not offer support for Microsoft Hyper-V. Such VMs are provided exclusively in the SLA Rhodium and cannot be managed via the ix.Cloud Portal or the ix.Cloud API.

Service Architecture

See Virtual Machine representation SLA Rhodium.

Service Scope

Virtual servers based on VMware can be ordered in the Standard Hardware Profiles according to Virtual Machine. Such VMs are delivered without an operating system (OS). The customer is responsible for licensing, operation and maintenance of the OS (see "Customer Owned OS" under Virtual Machine). The Off functionality is not available.

Such VMs must be ordered via "Standard Service Request" - mutations and decommissioning via "Generic Request". Image import is performed according to the description for "Customer Owned OS" under Virtual Machine.

Service Options

No options available.

Storage Services

The ix.Cloud "Storage Services" provide the ability to store data in Inventx's highly available data centers or expand a local data center with additional storage capacity.

Table: Storage Services

Service Name	Service Description
File Storage	Data storage for office documents, desktop profiles, and WORM archives.
Object Storage	Scalable object storage for large amounts of unstructured data

File Storage

The File Storage Service provides managed file shares that can be accessed via industry-standard protocols (NFS or CIFS/SMB). Data on File Storage is permanently synchronized between Inventx data centers in Switzerland.

:::info

Orders and all changes to the File Storage Service must be requested via a "Generic Request".

:::

Service Architecture

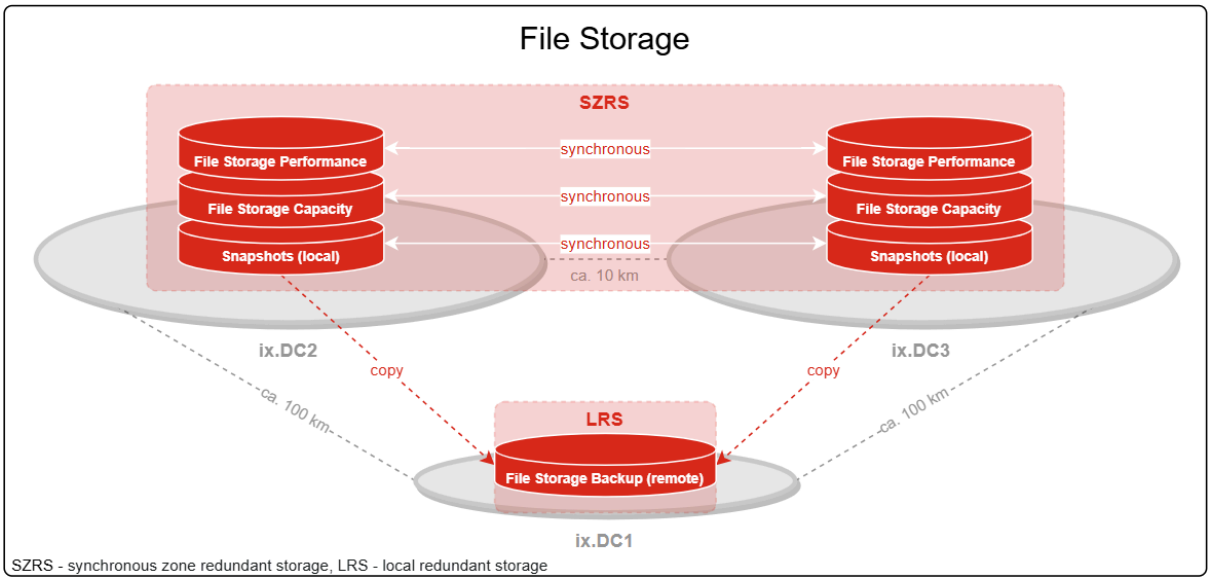

Image: File Storage Service Architecture

Service Scope

Table: File Storage Service Scope

Performance Features	
Redundancy and Replication	■
XTS-AES 256-Bit Encryption	■
AutoGrow & AutoShrink	■
Access Protocols NFS / CIFS (SMB)	■
Data Backup	■
Data Recovery	■
Customizable Tiering	■
Antivirus Protection	□
Ransomware Protection	□
WORM (write-once-read-many)	□

Service Options

The File Storage Service has the following options explained below.

Redundancy and Replication

To achieve the highest possible data availability, primary data and backups (snapshots) are permanently replicated synchronously across two data centers (zone redundancy). Additionally, for disaster recovery purposes, a backup copy is created to a third data center.

XTS-AES 256-Bit Encryption

Data on File Storage is encrypted with XTS-AES 256-bit encryption (Encryption@Rest). This encryption algorithm is one of the most commonly used and simultaneously most secure methods for encrypting data on storage.

AutoGrow & AutoShrink

Volume expansion and reduction occurs dynamically, with billing based on daily measured data consumption.

:::caution

Increases of more than 20% of total capacity according to the Consumption Report must be announced two months in advance.

:::

Access Protocols NFS / CIFS (SMB)

Access to the folder share or data occurs via NFS or CIFS protocol.

With the NFS protocol, access is restricted based on the client IP address or DNS name, and with the CIFS protocol, access is granted and managed via the customer's Active Directory.

:::caution

Multi-protocol access (NFS and CIFS) to a folder share is not supported.

:::

Data Backup

Primary data is backed up via snapshot technology at regular intervals (hourly and daily) and replicated synchronously across two data centers. The daily snapshot is additionally copied to a third data center.

Data backup can be configured according to needs using the following four backup profiles:

- 40d Backup
- 200d Backup (Standard)
- 400d Backup

- No Backup

:::caution

With the "No Backup" backup profile, no backup of primary data is created upon explicit request. The customer thus waives the possibility of data recovery.

:::

The retention period for the backup profiles is set up according to the specifications below. If the "40d Backup", "200d Backup", or "400d Backup" option is selected during backup, immutable snapshots of the primary data are created regularly. This results in conditional ransomware protection, as previous backups can be accessed.

Table: File Storage Retention Period "40d Backup"

Retention Period "40d Backup"		**Location**	
		local (snapshot)	**remote**
Interval	**hourly**	2 days	-
	daily	20 days	40 days (with WORM 20 days)

Table: File Storage Retention Period "200d Backup"

Retention Period "200d Backup"		**Location**	
		local (snapshot)	**remote**
Interval	**hourly**	10 days	-
	daily	40 days	200 days (with WORM 40 days)

Table: File Storage Retention Period "400d Backup"

Retention Period "400d Backup"		**Location**	
		local (snapshot)	**remote**
Interval	**hourly**	20 days	-
	daily	80 days	400 days (with WORM 80 days)

Data Recovery

Recovery of primary data is performed in self-service via the "Previous Versions" function in Windows File Explorer. Recovery from a remote backup must be requested via "Generic Request".

:::caution

If the "No Backup" option is selected in the Backup Profile (see Data Backup), recovery of primary data is not possible.

:::

Customizable Tiering

Customizable tiering makes it possible to move inactive data to a more cost-effective storage tier. The File Storage Service provides the following two storage tiers:

Table: File Storage Tiers

Storage Tier	Throughput / Volume	IOPs / Volume
Performance	max. 200 MB/s	max. 5'000
Capacity	max. 50 MB/s	max. 1'000

:::info

The KPIs "Throughput" and "IOPs" mentioned above are to be considered as guideline values and are fundamentally dependent on file size and the protocol used. For the Performance Tier, response times of an average of <5ms are to be expected (measured over 4 hours on the storage controller).

:::

All data initially remains on the Performance Tier and can then be moved to the Capacity Tier using "Auto Tiering" based on rules and automatically. The following profiles are available:

- No Tiering (data remains on the Performance Tier)
- Auto Tiering (inactive data is moved to the Capacity Tier after 40 days)

Antivirus Protection

CIFS (SMB) file shares can optionally be scanned with an antivirus scanner. Known file extensions of ransomware are additionally blocked.

:::note

For the Antivirus Protection option, two VMs with antivirus software are installed in the customer's network. Depending on the load on the file shares, more than two VMs may be required.

:::

Ransomware Protection

As an option, AI-based anomaly detection and pattern recognition with emergency snapshot creation can be selected. Thanks to this option, early detection of ransomware can be ensured. See attached graphic for 2-layer defense.

The scope of functions also includes the possibility of configurable user lockout, the possibility of differential recovery, and other analysis options.

Additionally, individual patterns and response policies can be defined. The configuration of honeypots is also available.

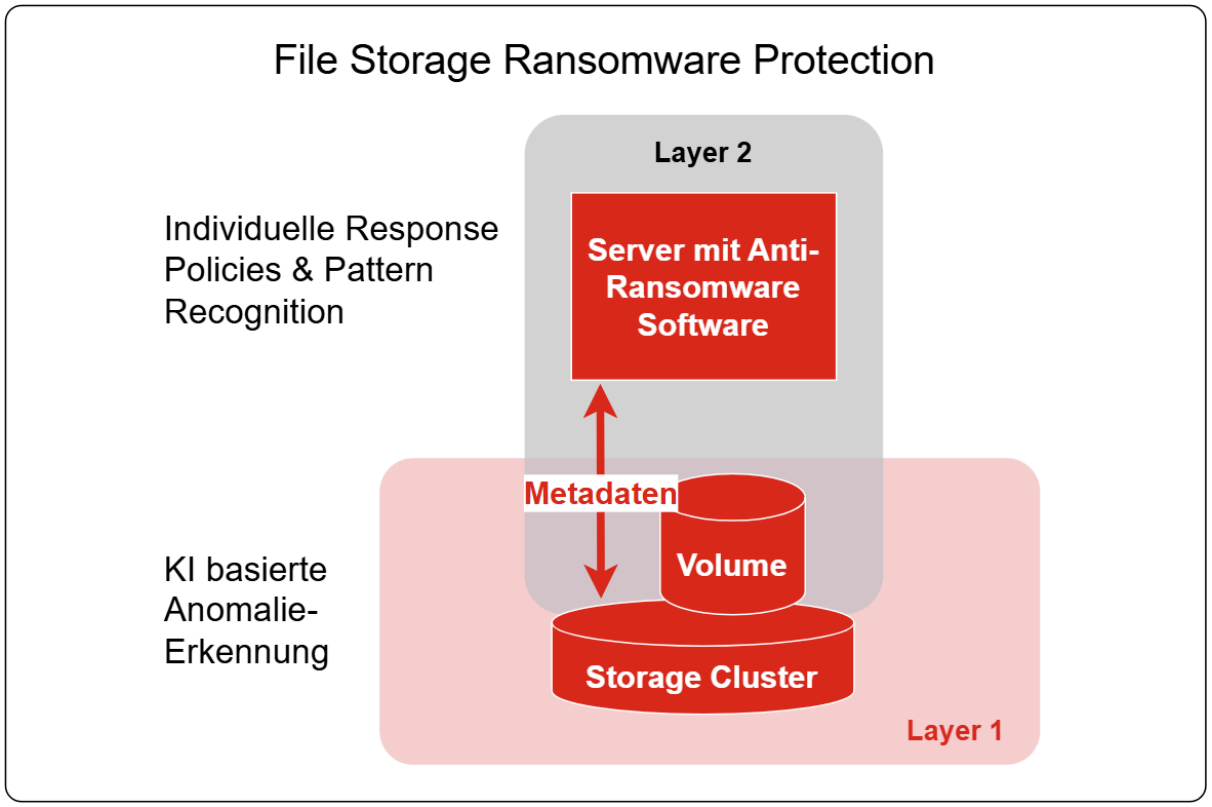

Image: File Storage Ransomware Protection

:::note For the Ransomware Protection option, a VM with the appropriate software is installed in the customer's network. Depending on the load on the file shares, more than one VM may be required. The Ransomware Protection option is implemented for the customer in collaboration with ix.CRC. Technically also possible with WORM volumes. Since WORM volumes are usually used in connection with archive applications, a more detailed analysis of the actions and their effects must be clarified. This can be implemented together with the customer upon request. :::

:::caution

This option can only be used with non-WORM volumes, as WORM data is already stored immutably.

:::

WORM (write-once-read-many)

To prevent files from being deleted, modified, or renamed, the WORM option can be selected if needed.

:::caution

Before initial setup, the dependencies and requirements of an archival solution must be checked.

:::

Object Storage

The Object Storage of ix.Cloud is an S3-compatible, scalable, and geo-redundant data store. Data is grouped in so-called vaults. This allows storage objects to be retrieved efficiently without knowing the physical location of an object - complex directory structures are eliminated. During data upload, data is automatically broken down into individual pieces using the Erasure Code method, extended with redundant information, and stored in physically different locations in the storage system. This allows corrupted or lost data to be reconstructed using information still available elsewhere.

Object Storage is ideal for storing large amounts of unstructured data, making it versatile in its applications: e.g., file storage, media, web content, data archiving, backup and restore.

Service Architecture

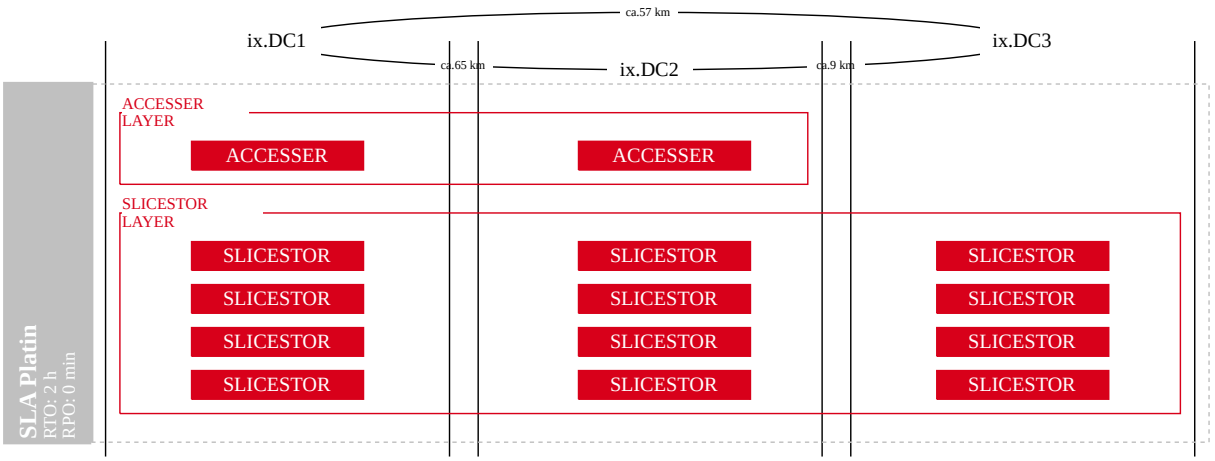

Image: Object Storage Service Architecture

Service Scope

Table: Object Storage Service Scope

Performance Feature	
Initial Setup	□
Access and Authentication	■
Encryption	■
Storage Management	■
Malware Protection	□

Service Options

Orders and all changes must be requested via a "Generic Request".

Access and Authentication

Logical separation and administration integrity is managed via an individual vault. After initial setup, access is provided via S3-API over HTTPS using username (Access Key ID) and password (Secret Access Key).

Encryption

Data transmission (upload and download) is encrypted with TLS (formerly SSL). All data on the storage system is stored using the Secure-Slice algorithm (256 bit), with the storage application implementing data encryption during the storage process.

Storage Management

Storage management of ix.Cloud Object Storage is based on AutoGrow and AutoShrink respectively. Consequently, no explicit storage size is reserved per customer. However, optionally per customer during initial setup or subsequently via "Generic Request", a maximum storage size per vault can be configured, with capacity management of such vaults being the customer's own responsibility. Billing is done monthly based on daily measured data consumption.

Compute Services

With "Compute Services," customers can obtain the required compute performance on demand based on different compute technologies and service models in Inventx's highly available data centers and align them according to the required SLA requirements.

Compute Services are used to deploy, host, and manage workloads. This section describes the different services and their options in relation to compute services in ix.Cloud.

Table: Compute Services

Service Name	Service Description
Virtual Machine	Virtual machine (VM) hosted in ix.Cloud, optionally also with management of the guest operating system via "Plus Services"

Virtual Machine

Through the ix.Cloud "Virtual Machine" service, a virtual computer (VM) can be provisioned within minutes in the geographic region that is suitable for the required workload.

Service Architecture

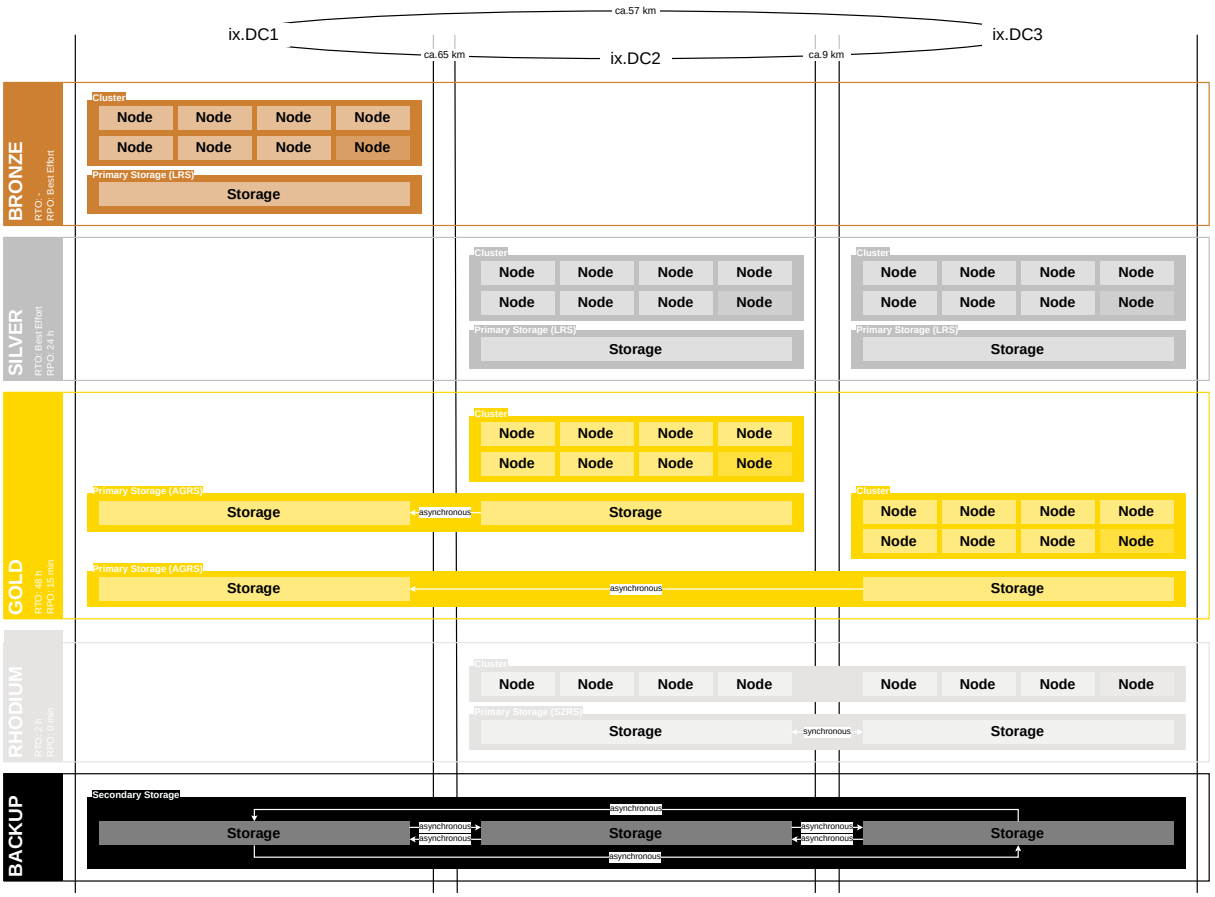

Figure: Virtual Machine Service Architecture

Service Scope

Table: Virtual Machine Service Scope

Feature	Feature Description
Definition, configuration, and provisioning of OS instances (VM with OS image)	• Automatic provisioning of a server from image (see Service Options) • Decommissioning of a server • Automated deployment of storage including choice of storage class • Change of SLA and change of location via "Standard Service Request"
Management and monitoring of computing systems and virtualization components	• Monitoring and reporting via the web portal (monitoring includes a configuration and capacity overview per VM as well as a list of attached resources) • Administrative rights at VM level (Start/Stop/Reconfigure) • Administrative rights in guest OS with local accounts • Authentication to customer AD via NTLM or Kerberos • Authentication via SAML (e.g., customer AD) • Access to the remote console of the VM, and depending on the selected operating system, access via Remote Desktop or SSH protocol • OS operation by customer
Management and monitoring of storage subsystems and storage networks	• Change of storage class • Adding or removing additional data disks
Management and monitoring of backup systems	• Backup of the virtual server monolithically at VM level (retention 14/40/90 days, daily) • "No Backup" option • Restoration of VM (monolithically at VM level) based on SLA • Encryption of backups with separate key per customer
Configuration and provisioning of virtual networks	• Configuration and use of virtual networks (SDN) • Provisioning and operation of virtual networks (VLAN). Limited to three VLANs per subscription • Additional VLANs are billed according to effort

Additional Services	• Assurance of manufacturer support for all infrastructure components • Ensuring capacity management of all infrastructure components (e.g., is there sufficient server hardware, storage hardware, or backup hardware?) • OS licenses and necessary licenses for infrastructure components included (e.g., backup, OS monitoring) • Automated reporting • Automated billing • Independent management and viewing via online portal • ON / OFF capability available

IT Baseline Protection

This section describes the general definitions of IT baseline protection for the compute service (at host level).

Patch Management

- Systems are patched on a 4-week cycle following our wave concept.
- Exclusions are generally not provided as these are critical security updates. An exclusion would only be considered in the case of a faulty patch and can be configured accordingly.
- In emergency patching, an identified critical vulnerability is patched immediately.

Logging

- Systems are logged by our monitoring, including event logs, logins, and administrator account records.
- System logs are retained for at least 9 months.
- Monitoring ensures that logs are continuously recorded. In case of failure, a ticket is automatically sent to operations.

Malware Protection

- An always up-to-date antivirus is the foundation of our malware protection.
- Updates are performed monthly together with patching. For critical vulnerabilities, an extraordinary update is performed.

Hardening

- We follow the CIS standard, which has been approved by the ATSB.
- All recommended security settings are implemented as long as they do not create functional restrictions.
- CIS recommendations are reviewed at regular intervals and considered in new software release cycles.

Configuration Management

- Systems (assets) are registered in our central CMDB
- Configurations are documented in the BHB and are implemented in the golden image of the deployment

Service Options

A VM can be adapted to current requirements via ix.Cloud Portal/API and through service requests by means of a variety of configurations and options.

Hardware Profiles

Hardware profiles are distinguished between two types (Standard and Highclock). The Standard hardware type has processors with a low clock frequency and is suitable for use with multithreading-capable applications. With the Highclock hardware type, on the other hand, processors with increased clock frequency are used; they are particularly suitable for non-multithreading-capable applications.

The following tables provide an overview of the available hardware profiles per hardware type. Additional hardware profiles can be requested via "Generic Request". Any implementation and pricing will be reviewed and approved by Inventx on a case-by-case basis.

Hardware profiles with 256 GB RAM can only be ordered via "Generic Request".

Table: Virtual Machine Hardware Profiles Standard

Standard		**RAM in GB**								
		4	**8**	**16**	**24**	**32**	**64**	**96**	**128**	**256**
Number of vCPU	**2**	■	■	■	■	■	-	-	-	-
	4	-	■	■	-	■	■	-	■	■
	6	-	-	-	-	-	■	-	■	■
	8	-	-	■	-	■	■	■	■	■
	12	-	-	-	-	-	-	-	■	-
	16	-	-	-	-	■	■	-	■	■
	24	-	-	-	-	-	-	-	■	■
System Disk		Windows: 100 GB / Linux: 80 GB								

Table: Virtual Machine Hardware Profiles Highclock

Highclock		RAM in GB								
		4	8	16	24	32	64	96	128	256
Number of vCPU	2	■	■	■	■	■	-	-	-	-
	4	-	■	■	-	■	■	-	■	■
	6	-	-	-	-	-	■	-	■	■
	8	-	-	■	-	■	■	■	■	■
	12	-	-	-	-	-	-	-	■	-
	16	-	-	-	-	■	■	-	■	■
	24	-	-	-	-	-	-	-	■	■
System Disk		Windows: 100 GB / Linux: 80 GB								

:::note

Highclock hardware profiles are offered exclusively in the "Rhodium" SLA.

:::

Storage Profiles

Storage profiles are composed of the storage type (redundancy) and storage class (speed), with the storage type depending on the selected SLA. For each virtual computer, only one storage type and one storage class is possible for all drives.

The table "Storage Classes" shows in the column "max IOPS / vDisk" the number of IOPS as limited by the hypervisor. This is the technically maximum possible number of IOPS, which however cannot be guaranteed in every case.

For the storage type "SZRS" in connection with the storage classes "High Performance" and "Ultra Performance", the specified throughput cannot be achieved if the block size is smaller than 32 KB. For the storage class "Ultra Performance", Inventx reserves the right to temporarily reduce capacity in case of performance bottlenecks.

Table: Virtual Machine Storage Types

Storage Type	Redundancy	Latency (Measurement duration 2h)	SLA

LRS	locally redundant	< 1 ms	Bronze & Silver
AGRS	asynchronously geo-redundant	< 1 ms	Gold
SZRS	synchronously zone-redundant	n/a	Rhodium

Table: Virtual Machine Storage Classes

Storage Class	Throughput (MB/s)	max IOPS / vDisk
Standard	40	5'120
Premium	150	19'200
High Performance	300	38'400
Ultra Performance	500	64'000

Operating Systems

As the operating system, either an image prepared by Inventx or a "Custom Owned OS" can be used.

The images provided by Inventx are patched at regular intervals after release by Inventx to the level of Wave 3. This is to ensure that new orders can be offered in Wave 3 or newer. An update to Wave 2 or Wave 1 is possible immediately after the ordering process. The update is implemented as follows:

- Windows: Once per quarter
- Linux: Upon each new minor release (e.g., RHEL V. 7.1 → 7.2)

After the staging process of a new VM, it is not patched directly to the desired wave. The first update process occurs after the "Customer Maintenance Window" defined by the customer, meaning after the selected wave and time choice. However, it is possible to patch a VM manually after deployment.

Among the images provided by Inventx are also license-free operating systems. For these operating systems, there is no manufacturer support. For this reason, Inventx can only guarantee the operation of these VMs on a "best effort" basis and reserves the right to suspend the recovery time according to SLA in the event of a malfunction (incident management) for errors beyond its control.

With customer-owned OS, the customer provides the image required for installation. The formats ISO and VHD are supported. In the case of ISO, the necessary settings for installation from the ISO file can be made directly in the self-service portal. If the VHD format is used, the settings for installation from the VHD file must be requested via "Generic Request". In both cases, manual work by Inventx is covered by the service request price.

Table: Virtual Machine Operating Systems

Operating System	Image	License	Plus Services
Windows Server 2022 Core	■	■	□
Windows Server 2022 Desktop Experience (DX)	■	■	□
Windows Server 2025 Core	■	■	□
Windows Server 2025 Desktop Experience (DX)	■	■	□
Red Hat Enterprise Linux 9	■	■	□
Red Hat Enterprise Linux 10	■	■	□
AlmaLinux 8	■	-	□
AlmaLinux 9	■	-	□
Customer Owned OS	□	-	-

Backup Profiles

The customer can use different backup profiles to implement data protection of the virtual server in different and customizable ways. Data protection is always created monolithically across the entire virtual server.

Table: Virtual Machine Backup Profiles

Feature	Bronze	Silver	Gold	Rhodium
Monolithic backup of the virtual server	■	■	■	■
Encryption with customer-specific key	■	■	■	■
Interval daily	■	■	■	■
Location				
• At the same location as the VM (local)	■	-	-	-
• At a remote location (remote)	-	■	■	-
• At two locations (local & remote)	-	-	-	■
Retention period				
• No Backup	■	■	■	■

• 14 Days	☐	☐	☐	☐
• 40 Days	☐	☐	☐	☐
• 90 Days	☐	☐	☐	☐
On-Demand Backup	■	■	■	■

Restore

If virtual servers are backed up (see Backup Profiles), they can be restored based on available backup copies. The customer can perform monolithic restoration of the virtual server via self-service.

Partial restoration (e.g., a specific file or folder) can be requested via "Standard Service Request".

Availability Set

To increase the availability of an application, it is recommended to set up two or more VMs with the same application and group them in an Availability Set. When multiple VMs are assigned to an Availability Set, the hypervisor will place these VMs on different hosts whenever possible. This configuration ensures that during a planned or unplanned outage of a host, at least one VM with the application remains available. Each Availability Set guarantees 3 hosts, with a maximum of one in maintenance mode at any time.

Availability Sets must be managed via a "Standard Service Request" to subsequently be able to assign VMs to an Availability Set via the ix.Cloud self-service portal.

System Management Services

For IT organizations to meet the demands of resilient infrastructure environments, from provisioning to operation, a series of security and monitoring activities are necessary in addition to the pure provisioning of a virtual machine (VM).

The System Management Services support customers in managing VMs and applications resiliently and scalably at the infrastructure level. The Application Owner can thus use a standard tool set and concentrate entirely on fulfilling their core elements, the management of their business applications.

The following table lists the individual services that support controlling and managing servers and workloads in ix.Cloud.

Table: System Management Services

Service Name	Service Short Description

Managed OS	Increases the security and availability of operating systems.
Metrics Monitoring	Monitoring of servers, applications, and services to optimize the performance and availability of IT services.
Software Deployment	Ensuring a homogeneous and resilient platform thanks to central software management.
Software and Release Cycles	Description of software repositories, handling of 3rd party software, and support and release cycles of Linux and Windows operating systems.

Managed OS

Managed OS is an optional add-on for [Virtual Machines](./compute-services/#virtual-machine) (VM) running an [Inventx Owned OS](./compute-services/#virtual-machine-operating-systems). If this add-on is activated on a VM, Inventx performs activities that contribute to increasing the security and availability of the operating system.

Service Architecture

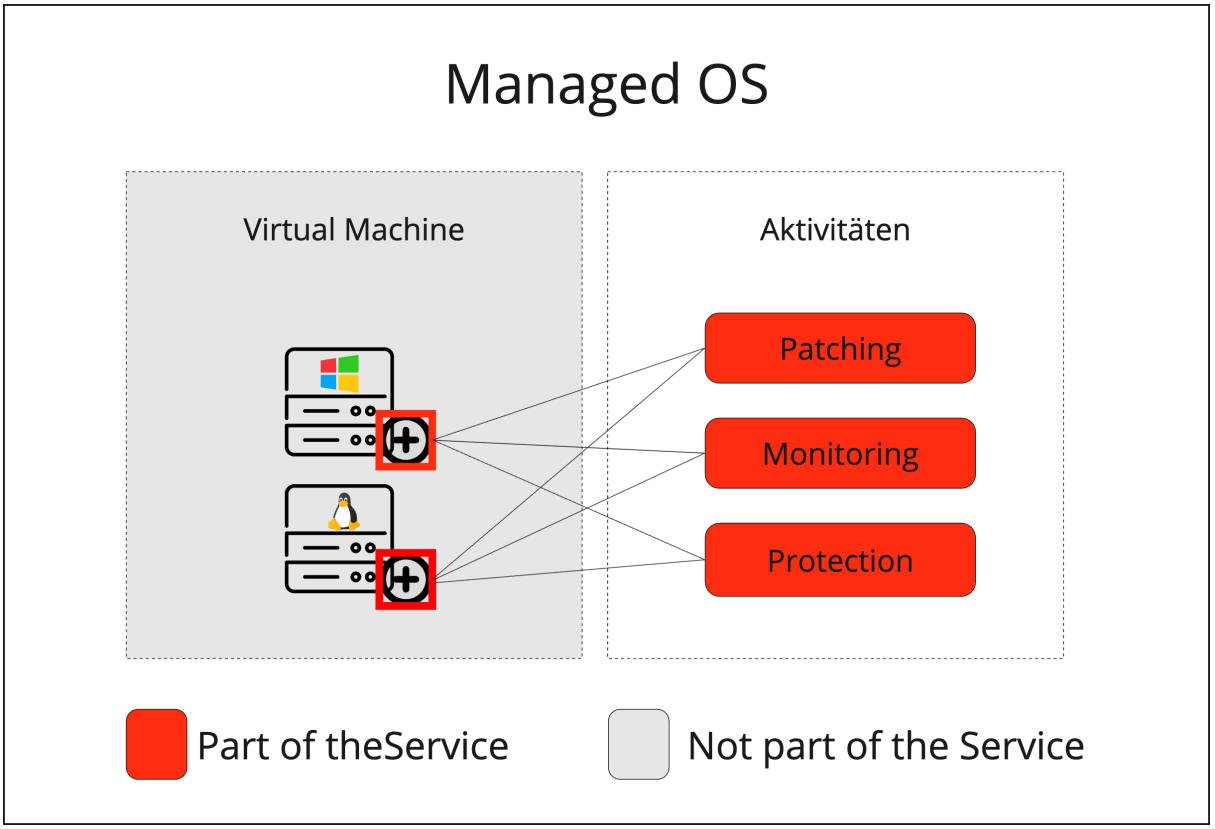

Figure: Managed OS Service Architecture

Service Scope

Table: Managed OS Service Scope

Features	Windows	Linux
Patching	■	■
Monitoring	■	■
Protection	■	■

Service Options

Patching Addon System Update Patching serves for the continuous improvement of stability, security, and currency of server operating systems.

The System Update addon includes an automatic update process that takes into account all software updates released by the manufacturer.

Table: Managed OS - Patching

Features	Windows	Linux
Update Types	Focus OS without software subsequently installed by the customer, i.e., with IE without frameworks. • Critical Updates • Security Updates • Service Pack • Update Rollup	Focus OS with software packages subsequently installed from RHEL-Repo.
Update Frequency	Monthly according to the defined Service Maintenance Window and the patch day configured on the VM. If automatic patching is not desired, there is the "No Automatic Patch" option.	
Update Cycle	The update process takes place once a month and can be configured flexibly: • No Automatic Patch ◦ The System Owner takes responsibility for installing the software updates. • Scheduling ◦ The System Owner selects the desired day and time window in which the automatic update process starts. The System Owner configures a response time (delay from the day of Inventx patch release until installation). ◦ The week of the second Tuesday of each month is exclusively available to Inventx.	

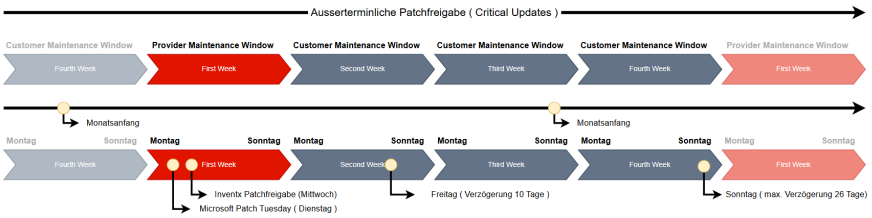

Figure: Out of Scheduled Release

One Time Update	Furthermore, the automatic update process can be initiated at any time – even outside regular maintenance windows – via the Cloud Portal using the One Time Update function. The defined time window for this must be at least 30 minutes in the future and have a minimum duration of four hours.	
Updated Products	• Windows Server 2016 Core • Windows Server 2016 Desktop Experience (DX) • Windows Server 2019 Core • Windows Server 2019 Desktop Experience (DX) • Windows Server 2022 Core • Windows Server 2022 Desktop Experience (DX) • Windows Server 2025 Core • Windows Server 2025 Desktop Experience (DX)	• Red Hat Enterprise Linux 8 • Red Hat Enterprise Linux 9 • Red Hat Enterprise Linux 10 • Alma Linux 8 • Alma Linux 9

:::info **Critical Updates**
System Management Services reserves the right to release Critical Updates even outside the scheduled patch release (second Tuesday of the month plus one day).
After an unscheduled patch release, Critical Updates are available to all systems:

- Systems that were already patched between the scheduled and unscheduled patch release can be brought up to date using a One Time Update.
- Systems patched after the unscheduled patch release will gain direct access to the critical updates.

:::

:::info To close security gaps more quickly, Edge Updates are released daily on the WSUS server.

After release, the update is available to the VM without a reboot.

- The update can be installed by the monthly update or a One-Time Update, which leads to a reboot of the VM.
- The update can be installed manually by the user in the OS.
- Alternatively, the standard Scheduled Task can be configured by the VM owner for installation.

:::

Monitoring

Monitoring is the surveillance of processes through systematic collection, measurement, and observation of an operation or process using technical aids. Based on the collected measurements, individual alerts can be set up and notified via a preferred communication channel.

Table: Managed OS - Monitoring

Monitoring	**Windows & Linux**
Virtual Machine	Active monitoring of performance behavior (CPU/RAM/IOPS)
Guest Operating System	Active monitoring and operation of the guest operating system
Usage and Performance Behavior	Monitoring and optimizing the usage and performance behavior of all infrastructure components to ensure SLA agreement and propose improvement possibilities

Protection

Endpoint Protection and Response (EDR)

Endpoint Detection and Response provides advanced threat detections that are near real-time and actionable. Security analysts can effectively prioritize alerts, gain insight into the full scope of a breach, and take response actions to remediate threats.

When a threat is detected, alerts are created in the system, which an analyst can investigate. Alerts associated with the same attack techniques or attacker are grouped into an entity called an incident. Aggregating alerts in this way makes it easier for analysts to collectively investigate and respond to threats.

Table: Managed OS – Protection EDR

Features	**Windows & Linux**
Cloud Protection	• Block Level ◦ High blocking level, aggressively blocking unknown items while optimizing device performance

	• Extended Timeout
	  ◦ This setting blocks a suspicious file for a certain period to perform an additional check in the cloud. The longer the block, the more time the cloud service has for an in-depth investigation.
	• Protection
	  ◦ Microsoft MAPS is the online community that helps you choose your reaction to potential threats.

Monitoring	Real-time behavior monitoring.
Scanning	• Archive files such as PLZ or CAB format • Downloaded files and attachments • Scripts • Removable media
Potentially Unwanted Application (PUA)	PUA protection is enabled. Potentially unwanted software will be blocked. Detected items are blocked. They will appear in the history along with other threats.
Quarantine	For the following threats • Severe Severity • Moderate High Severity • Moderate Severity • Moderate Low Severity
Exclusions	Exclusions are made in Self-Service via ix.Cloud Portal. Only for Windows servers • Auto Exclusions For Windows and Linux servers • Custom Exclusions • File Extensions and Folder Location Exclusions • Files opened by processes Exclusions • Contextual files and folder Exclusions
Operating Systems in Scope	• Windows Server 2016 (SQL / Core / DX) • Windows Server 2019 (SQL / Core / DX) • Windows Server 2022 (SQL / Core / DX) • Windows Server 2025 (SQL / Core / DX) • Red Hat Enterprise Linux 8

	• Red Hat Enterprise Linux 9 • Red Hat Enterprise Linux 10 • AlmaLinux 8 • AlmaLinux 9
Incident Management	Upon detection of a threat, the incident process is ensured by a defined security provider.
Reporting	A report is provided by the agreed security provider, identifying the monitored system and detected malware.

Prerequisites

For Inventx to properly deliver the services defined in this chapter, the following conditions must be met:

Prerequisite	Windows	Linux
The VM must be powered on	✔	✔
System components required for the service are exclusively configured by Inventx	Windows Update Agent	✔
The Azure Subscription required for the EDR service is created and managed by Inventx on the Azure customer tenant	✔	✔
Network targets required for the service are reachable from the VM	✔	✔
Inventx can access the VM over the network	WinRM and RDP	SSH
Inventx can access the VM via service accounts with required rights	Administrator rights	Root rights
The customer ensures that the disks on the system partition always have sufficient storage space and are not filled by application data and/or application logs	✔	✔
Additional software components that impair components for ensuring the service scope (e.g., proprietary antivirus or firewall software) must not be installed on the systems	✔	✔

:::caution The customer has administrative rights within the operating system and thus bears full responsibility for the operation of the virtual server if an SLA violation occurs due to incorrect customer

action (e.g., operating system update). :::

Metrics Monitoring

Metrics for business-critical applications collect and analyze data to improve the performance and availability of IT services. The use of metrics enables proactive monitoring, early detection of disruptions, and targeted alarming via defined contact points.

The "Metrics Monitoring" service is based on a highly available, scalable, and performant platform, thereby offering the necessary reliability required of a monitoring platform. Inventx ensures all necessary components related to Metrics Monitoring with this platform service. The customer can thus fully concentrate on monitoring their applications and services.

Billing is per subscription based on active series and the number of active users per month.

Service Architecture

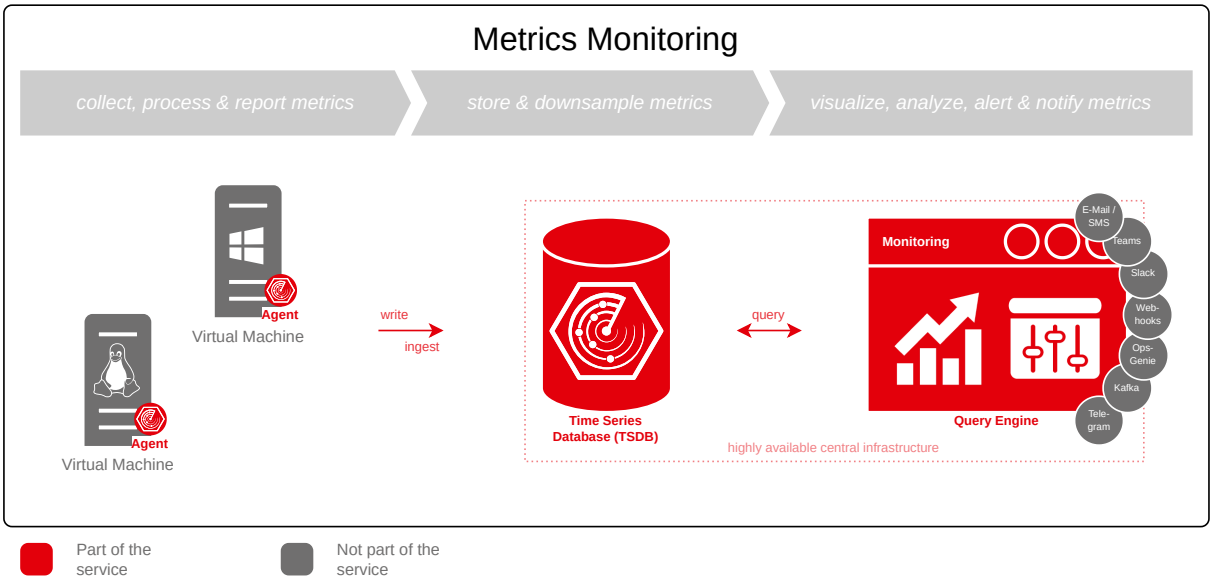

Figure: Metrics Monitoring Service Architecture

Service Scope

Table: Metrics Monitoring Service Scope

Features	
Monitoring Agent	■
Time Series Database	■

Query Engine	■
Default Metrics & Dashboard	■
Custom Metrics & Dashboards	■
Custom Alerts & Notifications	■
Interfaces to Notification Channels	■
Notification Channels	-

Service Options

The following chapters explain the individual options of this service in more detail.

Monitoring Agent

The "Monitoring Agent" is responsible for collecting, processing, and then forwarding data to the TSDB for storage. It is software used to monitor the respective system.

Inventx ensures that this component is installed on the defined systems, correctly configured at all times, and that the collection of metrics is guaranteed.

:::danger

If the installation and/or configuration of the monitoring agent is intentionally or unintentionally changed or damaged by third-party intervention, Inventx can no longer provide the services defined in the service.

:::

Time Series Database

The metrics collected by the Monitoring Agent are written to the Time Series Database (TSDB) and retained for 13 months. The TSDB is optimized for storing and retaining metrics and ensures performant data delivery.

:::info

For the Monitoring Agent to send the collected metrics to the TSDB, the IP address 10.94.12.36 and port 443 must be reachable.

:::

Query Engine

The Query Engine provides extensive options for visualizing, analyzing, alerting on, and notifying about metrics from the TSDB via various contact points.

:::info

The Query Engine is accessible via the URL https://monitoring.ixcloud.ch and follows the ix.Cloud authorization concept.

:::

Default Metrics & Dashboard

Upon activating the addon, the following user-optimized metrics are activated and written to the TSDB:

- CPU
- Memory
- Harddisk
- Network
- Services

Custom Metrics & Dashboards

In addition to the Default Metrics & Dashboard, custom metrics can be defined and configured. This enables writing customer-specific metrics from applications and services to the TSDB. Using the Query Engine, these metrics can be individually prepared and visualized as desired.

:::tip

A large number of different plugins are available on Github for agent configuration: https://github.com/influxdata/telegraf/tree/release-1.24/plugins

:::

Custom Alerts & Notifications

Based on the collected metrics, individual alerts can be set up using the Query Engine and notifications sent via a preferred communication channel.

Interfaces to Notification Channels

The Query Engine offers interfaces to the following common tools for notifications:

- E-mail / SMS
- Teams
- Slack
- Webhooks
- Ops-Genie

- Kafka
- Telegram

Notification Channels

Notification channels are not part of the service and must be provided by the customer.

Software Deployment

With the Software Deployment function, the provision and installation of software can be automated and managed from a central location via a portal. Thanks to the central control of software distribution processes, a homogeneous and resilient platform can be ensured.

The standardization of software on servers is a decisive step to ensure the security of the systems while optimizing effort and costs.

This addon is optional and can only be activated on Windows operating systems provided by Inventx. Subsequent deactivation of the addon is not possible.

:::danger

For Inventx to properly deliver the services defined in the "Managed-OS" addon, the following software must not be distributed by the customer:

- Windows Updates (this includes Windows Security Patches, Windows Feature Updates and Windows Rollup Updates)
- .Net Updates
- Splunk Universal Forwarder
- McAfee Agent
- Zabbix Agent
- Snow Agent
- Telegraf Agent
- Microsoft Defender
- Azure Connected Machine Agent

:::

Service Architecture

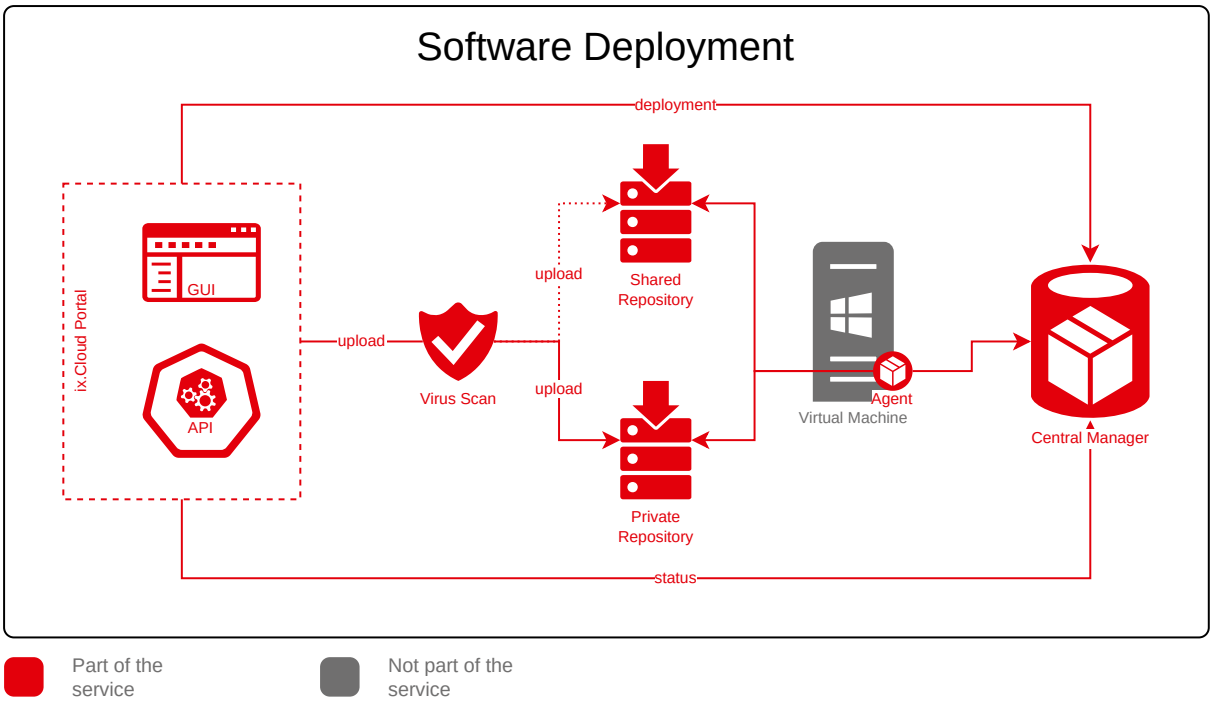

Figure: Software Deployment Service Architecture

Service Scope

Table: Software Deployment Service Scope

Features	
Shared Repository	■
Private Repository	■
Virus Scan	■
Automatic Update	■
Scheduled Deployment	■

Service Options

The following chapters describe the individual options of the Software Deployment addon.

Shared Repository

Through the Shared Repository, Inventx makes selected software packages available across ix.Cloud. The following software packages are made available to all customers via the Shared Repository:

- 7-Zip
- Adobe Reader
- Git
- Google Chrome
- Microsoft Edge
- Mozilla Firefox
- Notepad++
- Postman
- Visual Studio Code

:::info

The software packages in the Shared Repository have the Automatic Update option activated.

:::

Private Repository

The Private Repository serves as storage for customer-specific software packages. To store software packages in this repository, either a transfer from the manufacturer's Community Repository or an upload from the local computer can be performed.

When uploading from the local computer, the software packages are scanned for viruses before saving (see Virus Scan).

:::tip

For software packages from the manufacturer's Community Repository, the Automatic Update option can be activated.

:::

Virus Scan

As protection against malware, software packages are scanned for viruses during upload using a virus scan. If a virus is identified, the user is notified, and the upload is aborted.

Automatic Update

The Automatic Update option can only be activated for software packages originating from the manufacturer's Community Repository. This option cannot be activated for software packages uploaded from the local computer.

Software packages with this option activated are checked weekly on Sunday at 01:00 AM against the manufacturer's Community Repository for newer versions. If newer versions are available, they are

automatically downloaded and made available. This has the positive side effect that outdated installations are highlighted in the portal and can be updated with a few clicks.

Scheduled Deployment

A deployment can be scheduled over time. This way, the installation, update, or uninstallation of software can also be carried out at night.

Software and Release Cycles

Linux Software

On Linux systems, the software repositories listed below are essentially integrated via the ManagedOS Addon and considered in the update process. If the EDR Addon is enabled on the VM, Microsoft's Linux Software Repository is also included. Software can be installed from these software repositories on the target system at any time.

Table: Repos on Linux Systems

Linux Version	Repos
RHEL 8	• BaseOS, Appstream, CodeReady Linux Builder, EPEL* • Microsoft Software Repository (with EDR Addon enabled)
RHEL 9	• BaseOS, Appstream, CodeReady Linux Builder, EPEL* • Microsoft Software Repository (with EDR Addon enabled)
RHEL 10	• BaseOS, Appstream, CodeReady Linux Builder, EPEL* • Microsoft Software Repository (with EDR Addon enabled)
AlmaLinux 8	BaseOS, Appstream, EPEL*
AlmaLinux 9	BaseOS, Appstream, EPEL*

* The EPEL repository (Extra Packages for Enterprise Linux) is an additional package repository developed specifically for Enterprise Linux distributions such as Red Hat Enterprise Linux (RHEL), AlmaLinux, and Fedora. It offers a variety of additional open-source packages not included in the standard repositories of these distributions. The EPEL repo is a 3rd Party Software repo, for which the principles of the chapter "Dealing with 3rd Party Software" apply.

Windows Software

In addition to common installation procedures for software on Windows (e.g., with admin rights), the Software Deployment AddOn is available in the Self-Service to install software on a Windows system. The principles of the chapter "Dealing with 3rd Party Software" apply to this software.

Dealing with 3rd Party Software

The following principles apply to dealing with 3rd party software:

With administrator or root rights, it is always possible to install 3rd party software or packages or integrate your own software repositories. For this software, the responsibility, release management, and impact on operations lie entirely with the customer.

If the ManagedOS service is impaired by the use of 3rd party software, the corresponding SLA is no longer valid. In this case, Inventx cannot guarantee the functionality of the 3rd party software or stable ManagedOS operation. In extreme cases, this may lead to the complete affected VM having to be restored from backup by the customer himself or by Inventx on behalf of the customer. Additional expenses incurred by Inventx due to such incidents are not part of Inventx's business services and are to be reimbursed by the customer according to actual effort.

Operating System and Software Release Cycles

Windows and Linux major operating system release cycles are generally designed for 10 years, meaning that during this period, Systems Management Services, including software updates, are provided via the ManagedOS Addon, which the customer can configure for the respective VM via the portal. After these 10 years, the operating system is no longer supported, no new updates are available, and Systems Management Services are no longer developed for this operating system release. The customer is responsible for building a new VM with a newer major operating system release and migrating their application before the end of this 10-year period. In-place upgrades to a newer major operating system release on the same VM are not offered (e.g., from RHEL 9 to RHEL 10 or Windows Server 2022 to Windows Server 2025). If the customer performs an in-place upgrade themselves, they must ensure that all Systems Management Services continue to function properly on the new major operating system release. If the Systems Management Services are impaired by the customer's in-place upgrade to a higher major release, Inventx reserves the right to discontinue these services for the respective VM.

Support beyond these 10 years, e.g., through Extended Lifecycle Support (ELS) for RHEL or Extended Security Updates (ESU) for Windows, is generally not offered. In exceptional cases, this may still occur through special agreements with the customer. However, the conditions described by the manufacturer apply, and it cannot be guaranteed whether the Systems Management Services can still be provided with the same quality. This also entails any additional costs.

In addition to the 10-year major operating system release cycles, RHEL also has Appstream release cycles. This means that various applications in different major versions can be installed via the Appstream repository (e.g., PostgreSQL 13, 15, and 16 or .NET 6, 7, and 8, etc.). The responsibility for this major release management lies with the customer, as they can activate the corresponding channels on the system according to their needs. The update process of the ManagedOS Addon only considers upgrades within the activated major release and not to a higher major release. It should be noted here

that Appstream release cycles are often shorter than 10 years compared to the operating system release cycle. The exact details for all release cycles are published by the respective manufacturer.

Database Services

With "Database Services", customers can store and manage their business data in databases based on different database technologies and service models.

Table: Database Services

Service Name	Service Short Description
Managed xSQL-Instance	An MSSQL or MariaDB instance preconfigured and managed by Inventx with optional database operations.
Managed noSQL-Instance	A noSQL instance preconfigured and managed by Inventx.

Managed xSQL-Instance

The "Managed xSQL-Instance" service is based on the Virtual Machine service. In addition to the VM, Inventx installs, configures, and operates a SQL Server and SQL instance according to manufacturer specifications and Inventx best practices.

As an optional addition to this service, Inventx can take over database management.

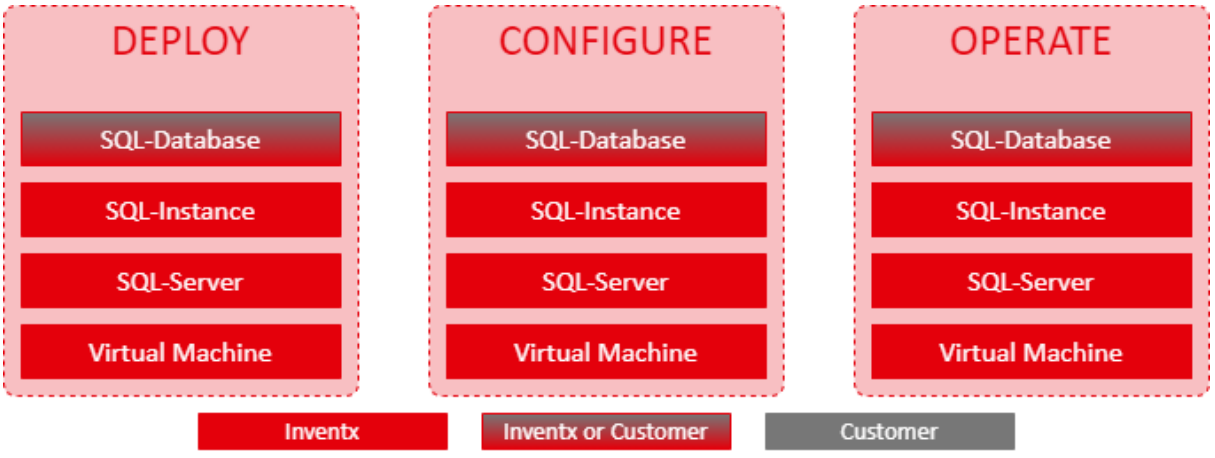

Figure: Managed xSQL-Instance responsibilities

Service Architecture

See Service Architecture from the Virtual Machine service.

Service Scope

Table: Managed xSQL-Instance Service Scope

Feature	Description
Licensing	See Licensing
Permissions	The customer receives no administrative rights on the instance. The customer is authorized as DBO (Database Owner) on individual databases.
Database Management	The customer can create databases within the database instance and must operate them themselves. Inventx does not provide maintenance services for such databases.
Database Backup (T-Log)	Database Backup
Database Restore	Database Restore
Database Clone	Database Clone
VM ON/OFF	For smooth operation of this service by Inventx, the customer must not turn the VM on and off.
Authentication	MS-SQL: Authentication is performed via Active Directory (NTLM or Kerberos). Authentication via SQL user must be requested as a Security Exception. PostgreSQL, Mariadb: Authentication is performed via SQL user.

:::info PaaS MSSQL Instance:
For security and operational reasons, the MSSQL Agent for any automations is not available. :::

Service Options

As part of the Managed xSQL-Instance service, customers have access to several technologies with specific additional features as follows:

Licensing

Inventx covers the licensing of the virtual server's operating system in this service (see Virtual Machine). For database server licensing, the following applies:

Table: Managed xSQL-Instance Licensing Responsibility

Licensing Responsibility	Inventx	Customer
Operating system virtual server	■	-
Microsoft SQL Server	-	■
MariaDB Server	■	-
PostgreSQL Server	-	-

:::note

- Optionally and by agreement between customer, Inventx, and Microsoft, "license mobility programs" are possible. These must be worked out individually between the parties.
- When using more than 16 CPUs or more than 128 GB RAM, Microsoft SQL Enterprise Edition is mandatory.

:::

Hardware Types and Hardware Profiles

All hardware profiles of the "Standard" hardware type according to the Virtual Machine service are available for selection.

The hardware profiles of the "Highclock" and "GPU" hardware types are not available.

Database Technologies

The following database technologies are available to the customer with this service:

Table: Managed xSQL-Instance Database Technologies

Database Technology	Community	Developer	Standard	Enterprise
Microsoft SQL Server 2019		■	■	■
Microsoft SQL Server 2022		■	■	■
MariaDB 10.6 as Managed Service	■			
MariaDB 11.8 as Managed Service	■			
PostgreSQL 15.0	■			■
PostgreSQL 16.0	■			■
PostgreSQL 17.0	■			■

Database Backup

With the data backup service for databases (Database Backup), Inventx backs up the customer's databases for the purpose of restoring databases in the event of IT disaster, data loss, or data corruption.

MS-SQL

Databases are backed up via the Service Agent of the Backup Service; this agent service is registered with gMSA. The gMSA is generated separately for each PaaS instance and has the necessary permissions on the MS-SQL instance according to the manufacturer.

PostgreSQL, MariaDB, MongoDB

Databases are backed up to an NFS share of the Backup Service. The Backup Service triggers the backup functions of the database instance remotely via SSH. The backup is stored in a dedicated directory on the NFS share.

Table: Managed xSQL-Instance Database Backup

Database Backup	MSSQL	MariaDB	PostgreSQL
Location	according to SLA	according to SLA	according to SLA
Interval			
• Full	daily	daily	daily
• Differential	-	-	-
• Transaction-Log	every 15 min	-	-
• Write-Ahead-Logging	-	every 15 min	every 15 min
Retention Period			
• No Backup	■	■	■
• 14 Days	■	■	■
• 40 Days	■	■	■
• 90 Days	■	■	■
On-Demand Backup	■	■	■

Database Restore

If databases are backed up (see Database Backup), they can be restored based on available backup copies via "Generic Request" as follows:

Table: Managed xSQL-Instance Database Restore

Database Restore	MSSQL	MariaDB	PostgreSQL
Database Restore	The customer can have individual databases restored from the last full backup via "Generic Request" by Inventx.		
Conditions	The database must be configured with transaction log and a valid backup retention period according to the "Database Backup Features" table so that point-in-time recovery can be implemented.		

Database Clone

A clone creates a copy of an existing database or a copy of individual database objects. Inventx provides the following variants, with clones being subject to a charge and must be ordered via "Standard Service Request".

Table: Managed xSQL-Instance Database Clone

Database Clone	Scope	Description
Full	Complete DB copy	1-to-1 copy of a database, where all database elements (schema and data) are copied.
Structure	DB copy without content	Essentially a 1-to-1 copy, but only with regard to the layout of a database. The database contents (data, jobs, procedures, etc.) are not copied in this procedure, only the tables.
Individual	Individual scope	Individual scope that is jointly specified: • Partial Clone: Copy of individual tables • Delta-Clone: Copy with subsequent mutations • Transaction-Realtime-Replication (TRR): Individual transactions in real time • Multiple Clone: Provision on multiple target DBs

Cloning requires two databases: a source DB and a target DB, which can have different names. The source DB is on an existing DB instance and must be backed up (active backup service). As a target DB, either a database on the DB instance of the source DB can be defined, or a database on a different DB instance, which must be operated in the same network zone as the source DB. During the cloning process, the target DB is not available.

Advanced Features

The following technology-specific additional functions are optionally available to the customer.

Table: Database Add-Ons Managed Service

Database Add-Ons Managed ServiceDatabase Add-Ons Managed Service	**MSSQL**	**MariaDB**	**PostgreSQL**
Always On / DB Clustering	☐	-	☐
Security Audit	☐	-	☐

Database Management

As an optional managed service, Inventx takes over database management based on the service described here. The following service agreement applies:

Table: Managed xSQL-Instance Database Management

Feature	**Description**
Order Management	New databases must be ordered via "Standard Service Request".
Permission	If Inventx takes over operational responsibility for the databases, Inventx removes the customer's permissions on the xSQL instance.
Database Deployment	Before deploying a new database, Inventx performs a check to determine whether it can be operated on the existing database instance or whether a new database instance should be created.

Database Operations	As part of regular operational responsibility, Inventx ensures the following services, whereby some services are charged separately (Change or Service Request):
	• Design and implementation of data security
	• Monitoring of database availability
	• Error analysis/resolution regarding database availability
	• Performing database backups
	• User management:
	  ◦ Management of personal DB users (Change Request)
	  ◦ Management of technical DB users
	• Performing database restores (Change Request)
	• Implementation of database optimizations (Change Request)
	• Cloning databases via backup (Service Request)
	• Migration of a database to a new instance (Change Request)

Extended Service Delivery	Additional work beyond database operations (e.g., performance analysis) can be performed by Inventx. However, the customer must order this individually via Change Request and it will be billed on a time and material basis.

PostgreSQL HA Managed Service

The PostgreSQL HA Cluster provides a PostgreSQL high availability solution as a Managed Service.

The PostgreSQL HA Cluster is deployed as 2 nodes (active, passive read-only) PostgreSQL Managed Service with virtual IP based on a licensed Failover Manager.

The cluster nodes are deployed and operated based on the PostgreSQL Managed Service.

The cluster must be ordered via Change Request.

Service Architecture

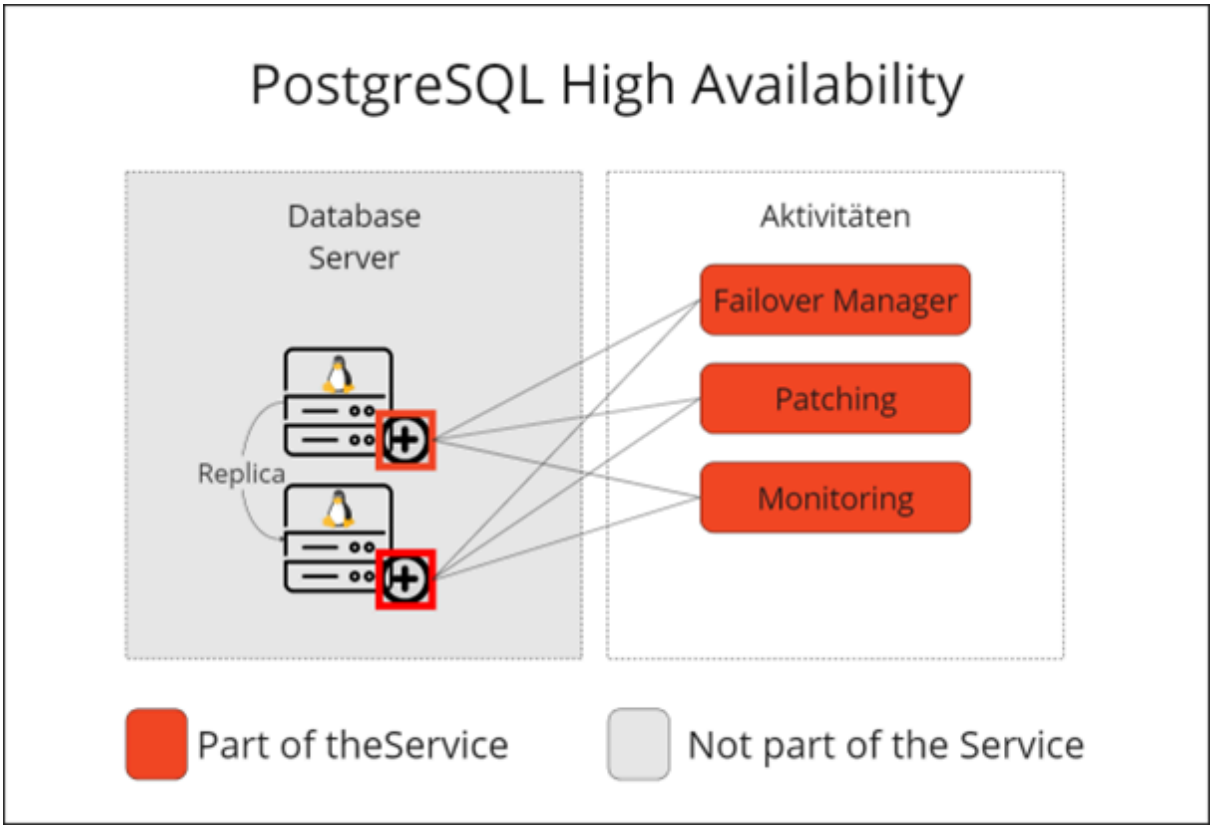

Image: PostgreSQL HA Managed Service Responsibilities

Service Scope

Table: PostgreSQL HA Managed Service Scope

PostgreSQL HA Cluster Node	Description
Primary	Active node for write operations and read operations. Virtual IP is assigned to this node.
Standby	Standby for failover case. Passive node for read operations

Managed noSQL-Instance

The "Managed noSQL-Instance" service is based on the Virtual Machine service. The server binaries and the instance are installed, configured, and operated on the VM according to the manufacturer's specifications and Inventx best practices.

:::note

With increasing data growth and the requirement to handle data flexibly and scalably, additional database management systems (DBMS) have emerged alongside traditional relational database management systems (RDBMS) that fundamentally differ from RDBMS systems. NoSQL DBMS are characterized by their horizontal and vertical scalability. As a rule, NoSQL systems are schema-free, which makes them suitable for big data environments and the development of geo-redundant, highly available DBMS clusters. NoSQL systems can typically not only handle SQL syntax; they are also often capable of using a variety of different and application-specific syntax.

:::

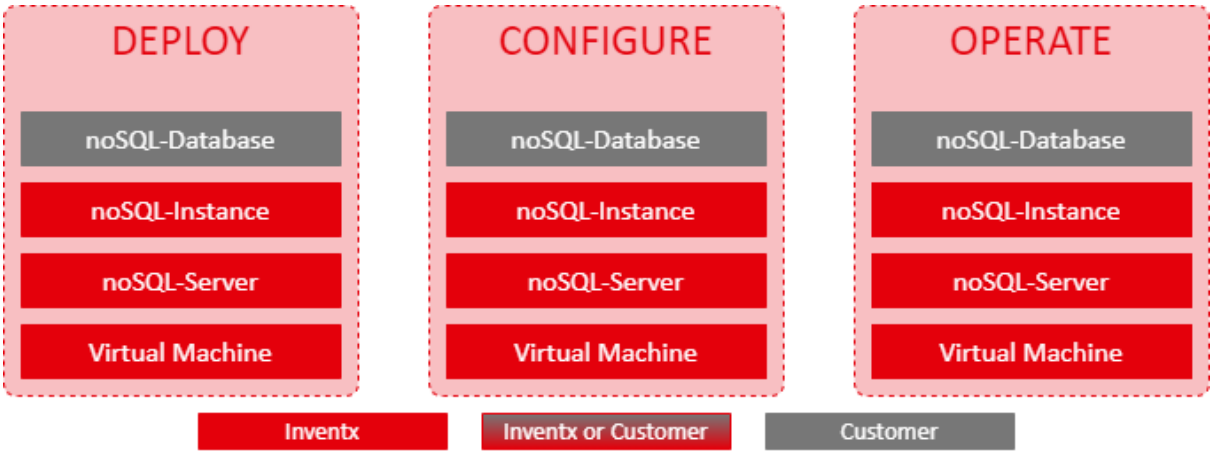

Image: Managed noSQL-Instance responsibilities

Service Architecture

See service architecture of the Virtual Machine service.

Service Provisioning

The host instance must be ordered as a **«Managed OS»** instance in the IX Portal.
The order for MongoDB Community must be requested via "Standard ServiceRequest".

Service Scope

Table: Managed noSQL-Instance service scope

Performance feature	Performance description
Licensing	See Licensing
Permissions	The customer receives no administrative rights on the instance. They will be granted DBO (Database Owner) permissions on individual databases
Database Management	The customer can create databases within the database instance and must operate them themselves. Inventx provides no maintenance services for such databases.
Database Backup	See Database Backup
Database Restore	See Database Restore
Database Clone	See Database Clone
ON/OFF of the VM	For smooth operation of this service by Inventx, the customer must not switch the VM on and off.
Authentication	Authentication is performed via SQL user.

Service Options

The following options are available as part of the Managed noSQL-Instance service.

Licensing

Inventx definitely covers the licensing of the operating system of the virtual server with this service (see Virtual Machine). For the licensing of the database server, the following applies:

Table: Managed noSQL-Instance licensing responsibility

Licensing responsibility	Inventx	Customer
Operating system of virtual server	■	-

MongoDB 7 as managed service	-	-
MongoDB 8.0 as managed service	■	■
MongoDB 8.2 as managed service	■	■

Hardware Types and Hardware Profiles

All hardware profiles of the "Standard" hardware type according to the Virtual Machine service are available for selection.

The hardware profiles of the "Highclock" and "GPU" hardware types are not available.

Database Technologies

The customer has access to the following database technologies with this service:

Table: Managed noSQL-Instance database technologies

Database technology	Community	Enterprise
MongoDB 7 as managed service	■	-
MongoDB 8.0 as managed service	■	■
MongoDB 8.2 as managed service	■	■

:::info For external end customers, only the Enterprise Edition (EE) is available. Due to a license change by the manufacturer (SSPL), we are no longer able to provide the Community Edition (CE) to external customers.

mongo/LICENSE-Community.txt at master · mongodb/mongo · GitHub :::

Database Backup

With the database backup service (Database Backup), Inventx backs up the customer's databases with the aim that the databases can be restored in the event of an IT disaster, data loss, or incorrect manipulation.

Table: Managed xSQL-Instance database backup

Database backup	MongoDB
Location	according to SLA
Interval	

• Full	daily
• Differential	-
• Transaction log	-
• Write-ahead log	-
Retention period	
• No backup	■
• 14 days	■
• 40 days	■
• 90 days	■
On-demand backup	-

Database Restore

If the databases are backed up (see Database Backup), they can be restored based on the available backup copies via "Generic Request" as follows:

Table: Managed noSQL-Instance database restore

Database restore	**MongoDB**
Database restore	The customer can have individual databases restored from the last full backup via "Generic Request" by Inventx.

Database Clone

A clone creates a copy of an existing database or a copy of individual database objects. Inventx provides the following variants, whereby clones must be ordered via "Standard Service Request" at a cost.

Table: Managed noSQL-Instance database clone

Database clone	**Scope**	**Description**
Full	Complete database copy	1-to-1 copy of a database where all elements of the database (schema and data) are copied.

Cloning requires two databases: a source database and a target database, whereby both can have different names. The source database is on an existing database instance and must be backed up (active backup service). The target database can either be defined on the database instance of the source database or on a different database instance, whereby it must be operated in the same network zone as the source database. During the cloning process, the target database is not available.

Managed Service Database Security Audit

The audits are generated through SQL internal functions and stored on the filesystem. These audits are forwarded to a Splunk load balancer in the customer's tenant.

Service Architecture

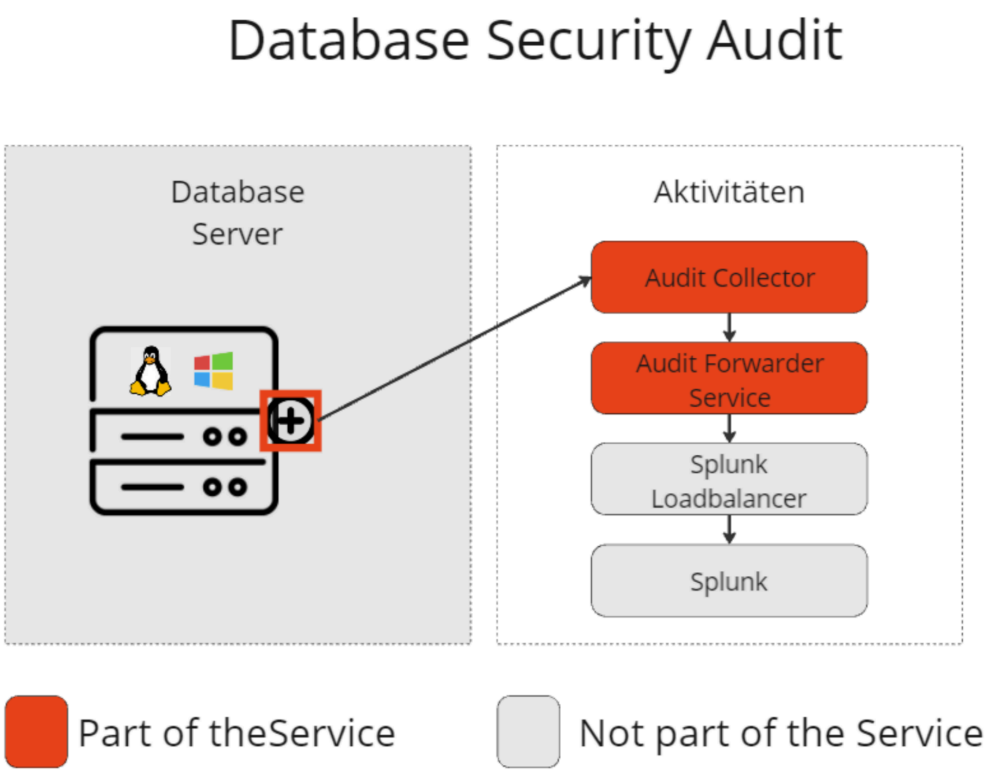

Image: Service Architecture

Service Scope

Mandatory Database Security Audit

Table: Mandatory Database Security Audit

Security Audit	Description

Audit Logins	- All user successful/failed logins - Logout

The mandatory security audit is installed by default on every instance and cannot be disabled.

Optional Database Security Audit

Table: Optional Database Security Audit

Security Audit	Description
Audit Privileges	- All create/delete/grant/revoke system privileges - All create/delete/grant/revoke database privileges
Audit System Settings	- All audit policy changes - All instance adjustments
Audit Activity High Privileged User	- All DDL actions by users with high privileges including database administrators The audit is performed at top level only on the metadata of the query - All DML actions by users with high privileges including database administrators The audit is performed at top level only on the metadata of the query

The optional security audit is installed by default on every instance and can be disabled if needed. This will be recorded in the audit upon deactivation and documented in the managed service configuration in the portal.

IT Baseline Protection Database Service

Patch Management

- Patches are deployed every 4 weeks in 3 waves
- Systems that cannot be patched automatically are patched manually on a monthly basis. The patching process and all relevant information are documented in Confluence.
- In case of emergency patching, an identified critical vulnerability is patched immediately.

Logging

- The systems are logged through our monitoring, including event logs, logins, and records of administrator accounts.
- Logs are stored online for 90 days; statistics are retained for 360 days.

- Monitoring ensures that logs are recorded continuously. In case of failure, a ticket is automatically sent to operations.

Malware Protection

- Malware protection is ensured through protection via the operating system.

Container Services

To enable IT organizations to meet dynamic requirements for implementing their business model, many companies are adopting modern IT platforms and organizational concepts such as DevOps. The following services support customers in optimally aligning application development and scalable IT operations with business dynamics.

With "Container Services," ix.Cloud offers a comprehensive tool set to deploy microservices in a fully automated manner and manage them efficiently. You focus entirely on your applications and processes, while Inventx operates the infrastructure for you and continuously develops it further.

Table: Container Services

Service Name	Service Short Description
IT Baseline Protection	Explanation of Container Service IT Baseline Protection
Container Registry	Store and manage your Docker container images securely in ix.Cloud
Agile Factory	Offers a customer-customizable DevOps environment based on Openshift Kubernetes
AnyCloudK8s	Is a flexible and agnostic container platform service
Container Namespace	Allows developers to group selected types together during programming and outsource frequently needed code into modules

IT Baseline Protection

This section describes general IT baseline protection for Container Services. If there are deviations for a specific service, these will be explicitly listed in the corresponding service scope.

Upgrade and Patching

- **Vendor-dependent requirements:** Patching and upgrades are performed according to the requirements of the respective vendors. The intervals depend on when new versions or patches are released.
- **No handling of exclusions:** Since this is a PaaS environment, no individual patches are distributed. All patches and upgrades are always performed at the platform level.
- **Emergency Patching:** When a critical vulnerability is identified, it is patched immediately or upgraded as soon as a corresponding solution is provided by the vendor for the platform level.

Logging

- **System logs:** System logs are stored in Splunk.
- **Retention period:** Logs are retained for a period of 3 months.

Container Registry

The Container Registry is a central location where container images can be deployed. During deployment, the image is automatically subjected to a vulnerability scan. This allows the execution of insecure images to be prevented.

By integrating the Container Registry into existing CI/CD structures, fully automated pipelines can be set up.

Service Architecture

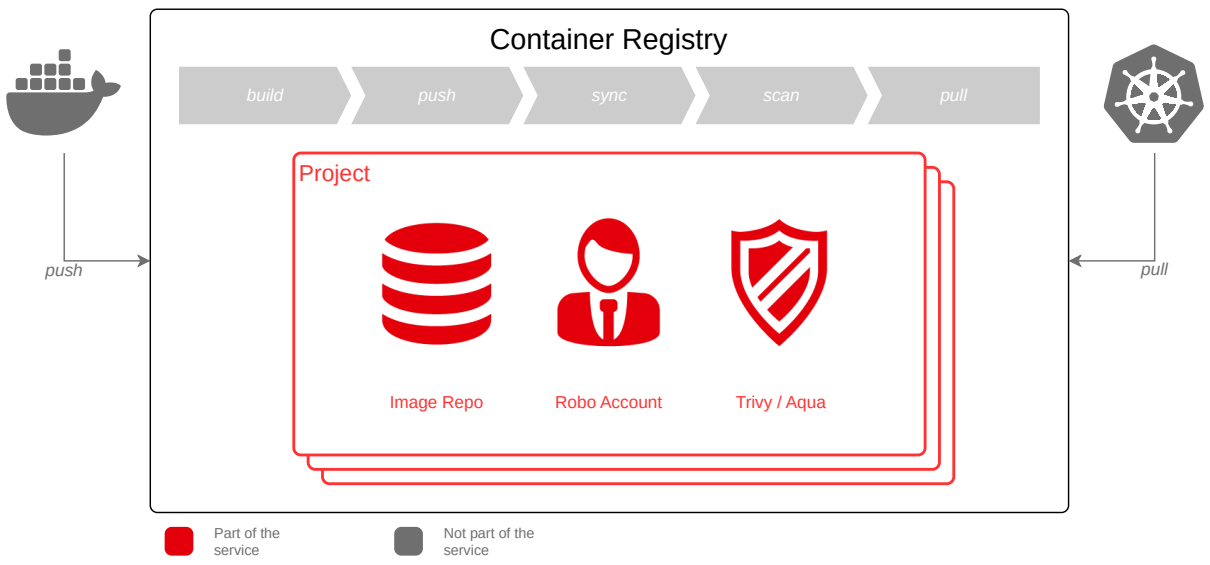

Image: Service Architecture Container Registry

Service Scope

Table: Container Registry Service Scope

Features	
IT-Grundschutz	■
Quota	■
Vulnerability Scan	■
Robot Account	■
Allow Public Access	■
Access to User Interface	■
CVE allowlist	■

IT-Grundschutz

Patching/Upgrade Interval

Table: Container Registry Patching/Upgrade

Day of Week	Time
Once per quarter, on the second Wednesday of the month	12:00-14:00

Hardening

- **External Review:** The service is reviewed by an external company every 2 years for security vulnerabilities and compliance.
- **Assurance through CI/CD:** The service is deployed via a pipeline (Continuous Integration/Continuous Deployment), which continuously ensures a secure and hardened setup.

Service Options

The following options are available in the "Container Registry" service and can be individually configured in self-service.

Quota

The quota defines the maximum storage capacity of the repository. Once this is reached, no further images can be uploaded.

The maximum storage capacity of the repository can be adjusted up or down at any time in the portal.

Vulnerability Scan

The Vulnerability Scan is enabled by default and cannot be disabled. All images are automatically checked daily by the Vulnerability Scanner. This ensures that a current vulnerability report is available for each image.

The following table lists and describes the available vulnerability scanners:

Table: Container Registry Vulnerability Scanner

Scanner	Description
Trivy	Simple scanning tool for applications that are not business-critical, or when working with less complex distributed architectures.
Aqua Enterprise	Advanced scanning tool for increased security. Particularly recommended for business-critical and complex cloud-native applications.

:::caution

The "Aqua Enterprise" Vulnerability Scanner generates additional costs.

:::

Robot Account

Robot Accounts are generally used for workflow, deployment, and test automation.

:::info

If the value "-1" is entered in the "Expiration time in days" field, the token has no expiration date.

:::

Allow Public Access

To authorize anonymous users with read access to a Container Registry, "Allow Public Access" must be enabled.

Access to User Interface

For advanced management of the Container Registry, access to the User Interface is provided. You can also view additional information via the UI, such as the Vulnerability Report of individual images.

CVE allowlist

An allow list can be created per repo/project, which applies to all images in this location. The validity period of the list can be defined with a date, or it can never expire. The default value is "Never expires".

Agile Factory

Agile Factory is a container platform service consisting of a variety of components that are interconnected to form an efficient and industry-optimized DevOps environment. All components and the connections between them are provided and operated by Inventx according to best practices, although individual components can also be provided by the customer.

With this service, cloud-native applications are created, managed, and scaled efficiently and simply. Inventx manages all components and provides relief for the customer. Developers can focus entirely on business logic and custom developments.

Agile Factory is offered in three standardized architectures/configurations (Basic, Top, and Premium) based on the Virtual Machine Service. The Basic configuration is recommended for non-production, the Top configuration for production, and the Premium configuration for business-critical applications.

The initial installation and configuration as well as any subsequent changes must be commissioned with a "Generic Request" or a "Service Request".

:::info

The Agile Factory service is based on the ix.Cloud standard Service Level, and the Worker Nodes of the respective configurations **Basic**, **Top**, and **Premium** define the applicable SLA.

For the restart of Agile Factory, an additional uplift of 2 hours must be expected in addition to the standard Service Level.

:::

Service Architecture

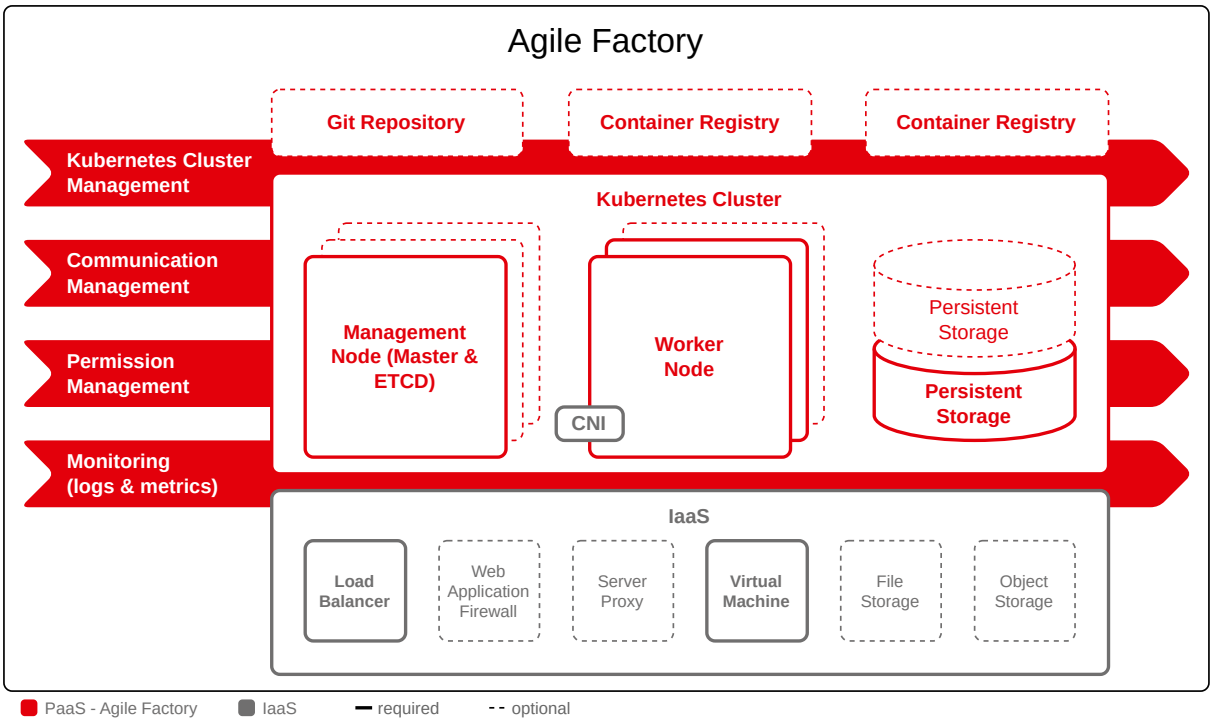

Image: Agile Factory Service Architecture for Rancher and Openshift

Service Scope

Table: Agile Factory Service Scope

Feature	Basic	Top	Premium
IT Baseline Protection	■	■	■
Initial Setup	□	□	□
Git Repository	□	□	□
Container Registry	□	□	□
Container Network	■	■	■
Management Node (Master & ETCD)	■	■	■
Worker Node	■	■	■
Persistent Storage	■	■	■
Load Balancer	■	■	■
Web Application Firewall	□	□	□

Server Proxy	□	□	□
Kubernetes Cluster Management	■	■	■
Communication Management	■	■	■
Permission Management	■	■	■
Monitoring (logs & metrics)	■	■	■
Cluster Network Interface	-	-	■

IT Baseline Protection

Patching/Upgrade Interval

Table: Agile Factory Patching/Upgrade Interval

Day of Week	Time
Friday after the second Monday of the month	02:00-06:00

Hardening

- The specifications were reviewed and approved in the ATSB
- These are implemented through automation
- This is ensured through GitOps

Malware Protection

All container images are scanned in the Container Registry for security vulnerabilities and weaknesses.

Configuration Management

Asset management is ensured through GitOps

Configuration Management

Asset management is ensured through GitOps

Service Options

The following describes the individual components of Agile Factory in detail. Some components are mandatory and some optional.

Initial Setup

The initial configuration of Agile Factory is implemented and billed as part of a project. Customer-specific configurations are specified in collaboration and then implemented.

Git Repository

For version control of deployment code, a Git Repository is mandatory in Agile Factory. This can be provided by the customer or by Inventx.

If the Git Repository is provided by Inventx, the customer receives an authorized user in Inventx's own Git, with which projects can be created, managed, and deleted. If the Git Repository is provided by the customer, the communication between the components is established and ensured in the Initial Setup.

Container Registry

Agile Factory uses a Container Registry for all deployments in the Kubernetes cluster. This component can be provided by the customer or by Inventx.

If Inventx provides the Container Registry, the Container Registry service is used. In this case, Inventx is responsible for trouble-free operation of this critical component.

If the Container Registry is provided by the customer, the communication between the two components Kubernetes Cluster and Container Registry is ensured in the Initial Setup. In this case, the customer is equally responsible for ensuring that communication between the components runs smoothly.

Container Network

Each Agile Factory requires a subnet. The size of the network determines how many Worker Nodes can be added to the cluster at most.

Table: Network Size

Product	Options
Openshift/Rancher	The network size is freely selectable

Management Node (Master & ETCD) for Openshift/Rancher

The Management Node is responsible for managing the Kubernetes cluster. Via the Management Node, the Kubernetes cluster and applications installed on it can be managed via CLI, GUI, or API. It serves as the control plane of the cluster and continuously manages the cluster's current state to the defined target state. The Management Node coordinates all tasks including scheduling and scaling of applications.

This cluster component is installed and configured according to Inventx best practices on dedicated Virtual Machine resources. To achieve fault tolerance, more than one Management Node is installed depending on the service configuration.

The initial configuration of a Management Node is fixed and defined as follows:

Table: Agile Factory Management Nodes for Openshift/Rancher

Feature	Basic	Top	Premium
Number of Servers	Rancher: 1, OpenShift: 3	Rancher: 3, OpenShift: 3	Rancher: 3, OpenShift: 3
Service Level	Rhodium	Rhodium	Rhodium
Hardware Profile	Rancher: P2/8, OpenShift: P4/32	Rancher: P2/8, OpenShift: P4/32	Rancher: P2/8, OpenShift: P4/32
Storage Class	High Performance	High Performance	High Performance
Backup Retention Time	14 days	14 days	14 days

Worker Node

The Worker Node is a Virtual Machine on which applications are executed and which is controlled by the Management Node.

Unlike the Management Node, hardware resources for the Worker Node can be selected and adjusted. Either additional Worker Nodes are added to the Kubernetes cluster (scale-out) or the hardware profile of existing Worker Nodes is changed (scale-up).

Resource adjustment of Worker Nodes can be commissioned with a "Service Request".

The initial configuration and expansion steps per service configuration for Openshift/Rancher are shown in the following table:

Table: Agile Factory Worker Nodes

Feature	Basic	Top	Premium
Number of Servers 1*	2 (initial) to 30	3 (initial) to 30	3 (initial) to 30
Service Level	Silver	Gold	Rhodium
Hardware Profile 2*	P4/16 (initial) P8/32 P8/64	P4/16 (initial) P8/32 P8/64	P4/16 (initial) P8/32 P8/64
Storage Class	Standard	Standard	Standard

Backup Retention Time	14 days	14 days	14 days

:::info

1* The more Worker Nodes the cluster has, the more spares (Worker Nodes) must be included.

- up to 16 Worker Nodes: 1 spare
- from 16 Worker Nodes: 2 spares

2* Further hardware profiles can be requested from Inventx.

:::

Persistent Storage

For persistent data, Agile Factory is equipped with persistent storage (SVM) in the Initial Setup. If needed, the Kubernetes cluster can be expanded with additional persistent storage.

Typically, persistent storage is obtained from the File Storage service. Optionally, the Object Storage service can also be used as external persistent storage.

All storage classes are geo-redundant. The following storage classes are provided and can be used through application-specific configuration.

Snapshots are triggered when corresponding Kubernetes resources are created by the application. This makes it possible to create a consistent snapshot of application data.

Table: Agile Factory Persistent Storage Classes

Feature	Description	Storage Classes	Reclaim Policy	Storage Classes (deprecated)
Block (Snapshot capable)	Creating snapshots is possible via kubeApi	block-std	Retain	netapp-block-std-ndr-$CUSTOMER-$ENV
Block Economy	No snapshot capabilities	block-eco	Retain	netapp-block-eco-ndr-$CUSTOMER-$ENV
File (Snapshot capable)	Creating snapshots is possible via kubeApi	file-std	Retain	netapp-file-std-ndr-$CUSTOMER-$ENV

File Economy	No snapshot capabilities	file-eco	Retain	netapp-file-eco-ndr-$CUSTOMER-$ENV

Recommendations for Using Storage Classes

In general, the eco storage classes should be preferred.

Table: Agile Factory Recommendations for Persistent Storage

Storage Class	Description
Block	Storage class for all SAN workloads that require their own data protection, triggered by scheduler
Block Economy	Storage class for all SAN workloads that do not require their own data protection
File	Storage class for all NAS workloads that require their own data protection, triggered by scheduler
File Economy	Storage class for all NAS workloads that do not require their own data protection. This storage class is marked as default in the cluster

:::note

Additional Storage Features

- When removing PVCs and PVs with the Reclaim Policy "Retain", the volume in the backend is deleted only after 14 days. Re-provisioning the volume must be ordered via Generic Request
- Autoscaling of PVCs using a File Storage Class is automatically expanded at a threshold of 79%. The expansion depends on the original size of the PVC. This usage check is performed every 5 minutes.

:::

Load Balancer

To build resilient applications in the cluster, the use of a Layer 7 Load Balancer or a Web Application Firewall is mandatory.

Web Application Firewall

To build resilient applications in the cluster, the use of a Web Application Firewall or a Layer 7 Load Balancer is mandatory.

:::note

Compared to a Load Balancer, a Web Application Firewall provides additional protection against unwanted requests to applications.

:::

Server Proxy

For applications to communicate to the Internet (outgoing traffic), the use of the Inventx Server Proxy is mandatory. See our Server Proxy service for details.

In the Initial Setup, it is defined whether a private or shared Server Proxy is used in the cluster.

Kubernetes Cluster Management

The Kubernetes clusters are centrally managed with one of the cluster management systems Openshift/Rancher or AnyCloudK8s. In the Initial Setup, the product to be used is defined with the customer.

Kubernetes Cluster Management essentially includes monthly patching and updating/upgrading of the entire platform as needed.

Communication Management

Inventx ensures that applications within the cluster cannot communicate with each other by default. Upon request, the customer can allow communication between applications with SSL encryption. SSL encryption is ensured through a self-signed certificate.

Permission Management

An individual RBAC concept can be implemented via the Authentication Authorization Infrastructure (AAI) and the LDAP Operator. For this purpose, customer-defined AD groups and Kubernetes cluster roles are linked together.

:::caution

For Inventx to perform all configurations during Initial Setup, the customer must provide the following information:

- desired AD groups
- an AD read-only user
- key material for LDAPS

:::

Monitoring (logs & metrics)

During operating hours, the Container Service is continuously monitored automatically by Inventx. Any events are logged, forwarded to the appropriate support organization, and processed during service

hours.

For all Agile Factory components for which Inventx is responsible, logs are collected centrally, evaluated, and appropriate measures are taken by Inventx when problems are identified.

:::note

Logs are retained for 90 days.

:::

AnyCloudK8s

AnyCloudK8s is a versatile, agnostic container platform service that consists of a variety of components. These components are optimally interconnected to provide an efficient and industry-specific DevOps environment. Inventx provides all components and their connections according to best practices and operates them. However, there is also the possibility that customers can provide individual components themselves.

With AnyCloudK8s, enterprises can efficiently create, manage, and scale cloud-native applications – simply and without complications. The platform offers maximum flexibility, regardless of which cloud provider you use.

Thanks to comprehensive management of all components by Inventx, customers are significantly relieved. Developers can focus entirely on business logic and custom developments without having to deal with infrastructure.

Additionally, AnyCloudK8s is available in the ix.Portal as a self-service – including easy provisioning and management of the environment directly via Portal.

Ordering can take some time, as network provisioning is not yet fully end-to-end automated. Once the network is available, the cluster will be automatically provisioned.

:::info

The AnyCloudK8s service is based on the ix.Cloud standard service level, and the Worker Nodes of the respective versions **Silver** and **Rhodium** determine the applicable SLA.

For the restart of AnyCloudK8s, an additional uplift of 1 hour must be expected beyond the standard service level.

:::

Service Architecture

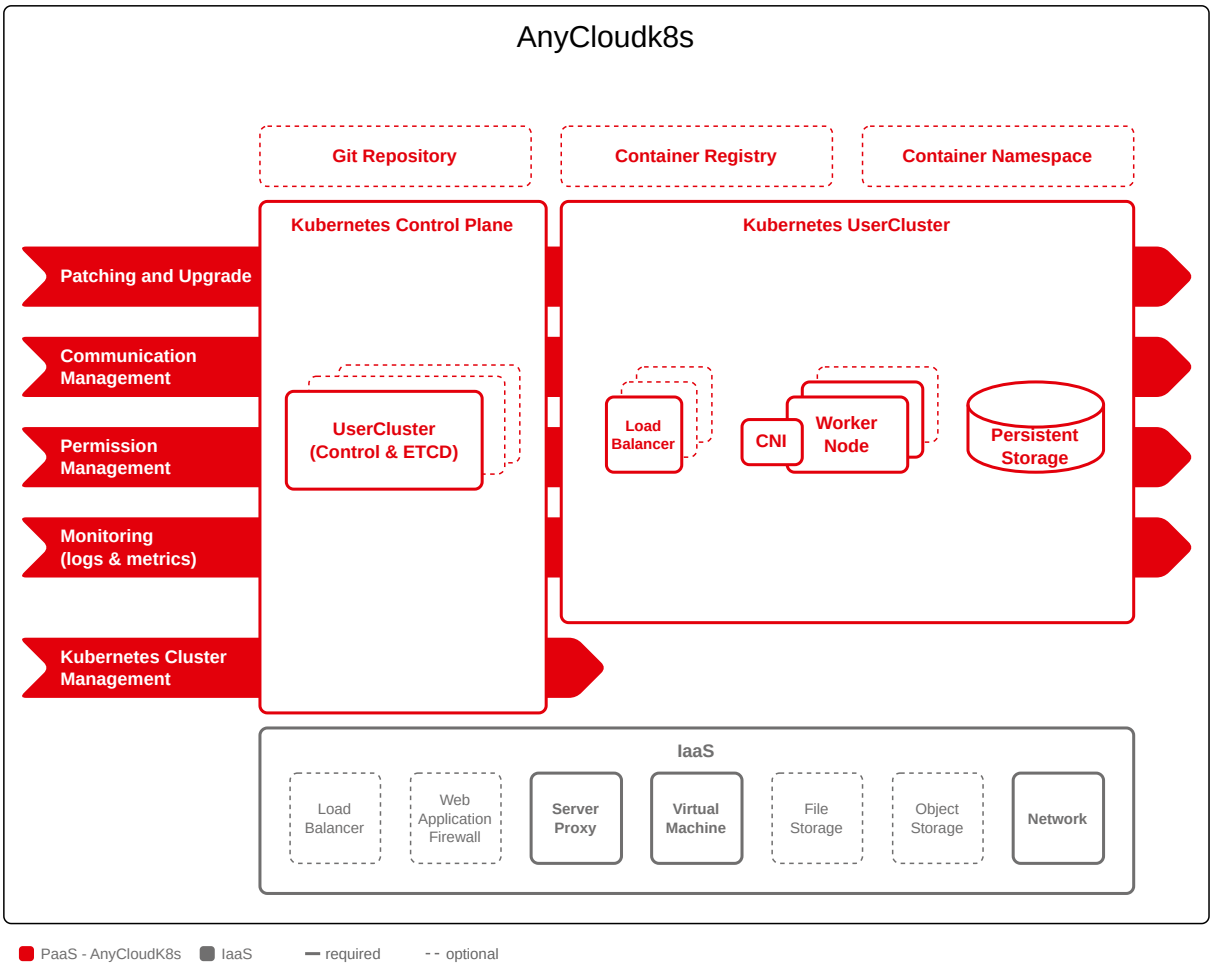

Image: Service Architecture for AnyCloudK8s

Service Scope

Table: AnyCloudK8s Service Scope

Feature	Basic	Premium
IT Baseline Protection	■	■
Initial Setup	■	■
Git Repository	□	□
Container Registry	□	□
Container Network	■	■
Cluster Control Plane	■	■
Worker Pool	■	■

Worker Node	■	■
Persistent Storage	■	■
Load Balancer	■	■
Web Application Firewall	□	□
Server Proxy	■	■
Kubernetes Cluster Management	■	■
Communication Management	■	■
Permission Management	■	■
Monitoring (logs & metrics)	■	■
Cluster Network Interface	■	■

IT Baseline Protection

Patching/Upgrade Interval

Table: AnyCloudK8s Patching/Upgrade Interval

Day of Week	Time
Friday after the second Monday of the month	02:00-06:00 hours

Hardening

- Requirements were reviewed and approved in the ATSB
- These are implemented through automation with GitOps

Malware Protection

All container images are scanned in the Container Registry for security vulnerabilities and weaknesses.

Configuration Management

Asset management is ensured through GitOps and CRs (ix.Cloud Operators)

Service Options

The individual components of AnyCloudK8s are described in detail below. Some components are mandatory and some are optional. This is evident in the table above AnyCloudK8s Service Scope.

Initial Setup

When ordering a new cluster, provisioning can take up to 5 days. The service is not yet fully automated, as the network is ordered separately from the responsible team as a managed service. Once the network is ready, the cluster is subsequently provisioned fully automatically.

After initial provisioning, all further service functions are fully automated, such as patching/updates and worker pool management (expanding, modifying, shrinking).

Git Repository

For versioning code of deployments, a Git repository is mandatory in AnyCloudK8s. This can be provided by the customer or by Inventx.

If the Git repository is provided by Inventx, the customer receives an authorized user in Inventx's own Git with which projects can be created, managed, and deleted. If the Git repository is provided by the customer, communication between the components is established and ensured during Initial Setup.

Container Registry

AnyCloudK8s uses a container registry for all deployments in the Kubernetes cluster. This component can be provided by the customer or by Inventx.

If Inventx provides the container registry, the Container Registry service is used. In this case, Inventx is responsible for trouble-free operation of this critical component.

If the container registry is provided by the customer, communication between the two components (Kubernetes cluster and container registry) is ensured during Initial Setup. In this case, the customer is equally responsible for ensuring that communication between the components runs without disruption.

Container Network

For the AnyCloudK8s environment, an additional subnet is needed. The size of the network determines how many worker nodes can be added to the cluster at most.

:::note

Please note that the entire IP range is not available for the worker nodes. Additional management and PaaS components are provisioned in this network, which are essential for service delivery. These components require their own IP resources, which restricts the usable area for the worker nodes accordingly.

:::

Network sizes are predefined in three T-shirt sizes: small(/27), middle(/26), and large(/25)

Table: Network Size

T-shirt Size	Maximum Worker Nodes
Small	10
Medium	26
Large	58

Additionally, a DHCP VM is deployed, which manages IP addresses in the container network. This service is automatically provisioned and fully configured when the network for AnyCloudK8s is set up. This ensures that all components can communicate seamlessly with each other and network resources are managed efficiently.

The **valid-lifetime = lease time** is fixed at **3600 seconds**.

Cluster Control Plane

The control plane is used to manage the Kubernetes cluster. Through it, the Kubernetes API can manage the cluster and applications installed on it. It serves as the control layer of the cluster and constantly manages the actual state of the cluster to the defined target state. The control plane coordinates all tasks, including scheduling and scaling of applications.

To ensure high fault tolerance and availability, the control plane is deployed on a seed Kubernetes cluster.

Required resources are dynamically ensured at the IaaS level.

Table: Management Nodes AnyCloudK8s

Feature	Basic	Premium
Seed Cluster	Local Redundancy	Geo-Redundancy
Service Level	Silver	Rhodium
Storage Class	High Performance	Standard
Backup Retention Time	14 Days	14 Days

Worker Pool

A cluster can consist of one or more worker pools. For each worker pool, you define a uniform worker node (e.g., CPU/RAM profile) and can thus address different requirements within the same cluster.

When multiple worker pools make sense

Multiple worker pools are used to combine different worker nodes – such as general-purpose nodes and those with large amounts of memory – in one cluster and distribute workloads accordingly. Using taints and tolerations, you ensure that certain workloads only run on designated nodes.

Targeted placement via pod specification: Control scheduling via nodeSelector (or node affinity if needed) so that pods are scheduled to the appropriate worker pools and worker nodes.

Advanced Worker Pool Management

Worker pools can not only be created and expanded in the portal, but also specifically adjusted. This allows different worker nodes, hardware profiles, and scheduling requirements to be cleanly separated within a cluster and controlled throughout their lifecycle.

ADVANCED MANAGEMENT FOR WORKER POOL ADJUSTMENTS

With advanced features, you can adapt worker pools to current requirements without rebuilding the entire cluster:

- Modify worker pools in detail Change pool-specific settings such as worker node/flavor, resource profile, metadata, and scheduling parameters. This allows you to tailor workloads to appropriate node profiles even after deployment.
- Scale down worker pools Reduce the number of worker nodes in a pool in a controlled manner. Nodes are orderly removed from the pool to reduce capacity and optimize costs (e.g., after peak loads or after project completion).
- Define taints/labels/annotations per worker pool Store rules and metadata directly on the worker pool:
    - Labels for targeted scheduling (e.g., nodeSelector / Affinity)
    - Taints to reserve pools for special workloads (e.g., GPU, high-memory) and only allow them via tolerations
    - Annotations for additional control and integration information (e.g., operations, monitoring, automation)

Worker Node

The worker node is a Virtual Machine on which applications are executed and managed by the Management Node.

Unlike the Management Node, hardware resources can be selected and adjusted on the worker node. You can either add additional worker nodes to the Kubernetes cluster (scale-out) or change the hardware profile on existing worker nodes (scale-up).

An expansion can be performed independently through the Worker Pool.

Table: Virtual Machine Hardware Profile Standard 1*

Standard		RAM in GB								
		4	8	16	24	32	64	96	128	256
Number of vCPU	2	-	-	-	-	-	-	-	-	-
	4	-	■	■	-	■	■	-	-	-
	6	-	-	-	-	-	-	-	-	-
	8	-	-	■	-	■	■	-	■	-
	12	-	-	-	-	-	-	-	-	-
	16	-	-	-	-	■	■	-	■	■
	24	-	-	-	-	-	-	-	-	-
System Disk		FlatCar: 80 GB								

:::info

The difference to AgileFactory is that in AnyCloudK8s a worker node is first added to a cluster, so no additional spare nodes need to be provisioned. The expansion is automated in the backend and can be provisioned faster.

1* Further hardware profiles can be requested from Inventx.

:::

Persistent Storage

For persistent data, AnyCloudK8s is equipped with persistent storage (SVM) during Initial Setup. Only one persistent storage can be provisioned per cluster in the subnet.

Typically, persistent storage is obtained from the File Storage service. Optionally, the Object Storage service can also be used as persistent storage from external sources.

With AnyCloudK8s, storage classes are automatically provisioned according to the cluster's SLA.

Table: AnyCloudK8s Persistent Storage

Feature	Basic	Premium
Cluster	Local Redundancy	Geo-Redundancy
Persistent Storage	Local Redundancy	Geo-Redundancy

The following storage classes are provided and can be used through application-specific configuration.

Table: AnyCloudK8s Persistent Storage Classes

Feature	Description	Storage Classes	Reclaim Policy
Block (Snapshot capable)	Creating snapshots is possible via kubeApi	block-std	Retain
Block Economy	No snapshot capabilities	block-eco	Retain
File (Snapshot capable)	Creating snapshots is possible via kubeApi	file-std	Retain
File Economy	No snapshot capabilities	file-eco	Retain

:::info

Snapshots are triggered by the application when creating corresponding Kubernetes resources. This allows creating a consistent snapshot of application data.

:::

Recommendations for Using Storage Classes

In general, the eco storage classes are preferable.

Table: AnyCloudK8s Recommendations for Persistent Storage

Storage Class	Description
Block	Storage class for all SAN workloads that require their own data backup, triggered by scheduler
Block Economy	Storage class for all SAN workloads that do not require their own data backup
File	Storage class for all NAS workloads that require their own data backup, triggered by scheduler
File Economy (default)	Storage class for all NAS workloads that do not require their own data backup. This storage class is marked as default in the cluster

:::note

Additional Storage Features

- When removing PVCs and PVs with the reclaim policy "Retain", the volume in the backend is deleted only after 14 days. Reprovisioning of the volume must be ordered via generic request
- Auto-scaling of PVCs using a file storage class is automatically expanded when a threshold of 79% is reached. The expansion depends on the original size of the PVC. This usage check is performed every 5 minutes.

:::

Load Balancer

To build resilient applications in the cluster, the use of a Layer 4/7 Load Balancer or a Web Application Firewall is mandatory.

When provisioning an AnyCloudK8s cluster, **no load balancer is automatically** installed or configured. If a load balancer is needed, the customer defines it **themselves in the cluster via manifest** (e.g., as a service of type LoadBalancer, according to the provided *template*).

```
apiVersion: v1
kind: Service
metadata:
  name: loadbalancer-sample
  annotations:
    lb.ixcloud.ch/enable: ""
    lb.ixcloud.ch/retain: ""
    lb.ixcloud.ch/fqdn: "only-test.[FQDN]"
    lb.ixcloud.ch/loadBalancingAlgorithm:
"LeastConnections"
    lb.ixcloud.ch/maxThroughput: "10"
    lb.ixcloud.ch/costCenter: "[Cost124]"
spec:
  type: LoadBalancer
  selector:
    app: my-app                              # must match
the pod labels
  ports:
  - name: http
    port: 80
```

```
    targetPort: 8080
    protocol: TCP
```

Once the load balancer is defined in the cluster, our automation ensures that **backend configuration is updated automatically upon changes** (e.g., when adjusting an IP address or patching/upgrading).

:::info

Features

- Automated adjustments: Changes to backend configurations occur without manual intervention.
- Scalability: Multiple load balancers can be managed simultaneously per cluster to support more complex requirements.

:::

This automation ensures efficient and error-free management of network components and simplifies scaling and expansion of cluster infrastructure.

Web Application Firewall

To build resilient applications in the cluster, the use of a Web Application Firewall or a Layer 7 Load Balancer is mandatory.

:::note

- Compared to a load balancer, a web application firewall provides additional protection against unwanted requests to applications.
- If an **L7 load balancer** or **WAF** is needed, first deploy an **L4 load balancer** via manifest in the cluster. Then it can be converted to L7 or WAF accordingly via a **Generic Request**.

:::

Server Proxy

For applications within the AnyCloudK8s cluster to communicate with the internet (outbound traffic), the use of the Inventx server proxy is mandatory. For more details, see the documentation of our Server Proxy service.

When provisioning an AnyCloudK8s cluster, only the shared server proxy is used. This ensures that all outbound connections meet Inventx's security and compliance requirements.

Kubernetes Cluster Management

Kubernetes clusters from AnyCloudK8s are centrally managed by Inventx. A key feature of this approach is the decoupling of the control plane from the cluster. This gives the customer admin rights on the

cluster, while Inventx ensures overall management.

Inventx regularly provides supported and current patches as well as new Kubernetes versions, which the customer can use for upgrades as needed. Announcements of new versions and deprecations of existing versions are made monthly via release notes according to the ixCloud process.

This management model ensures the security, stability, and currency of the Kubernetes environment, while the customer retains maximum flexibility for managing their cluster.

Communication Management

Inventx ensures that applications within the cluster cannot communicate with each other by default. Upon request, the customer can allow communication between applications with SSL encryption. SSL encryption is ensured through a self-signed certificate.

Permission Management

As described in Kubernetes Cluster Management, the customer receives admin rights for the provisioned cluster. These permissions enable the customer to have full control and flexibility in managing and using the cluster.

The corresponding kubeconfig can be easily downloaded via the portal. With this configuration file, admins can access the cluster directly via their preferred tools (e.g., kubectl).

Monitoring (logs & metrics)

During operating hours, Inventx continuously monitors the container service by machine. Any events are logged, forwarded to the appropriate support organization, and processed during service hours.

For all components of AnyCloudK8s for which Inventx is responsible, logs are centrally collected, evaluated, and appropriate measures are taken by Inventx when problems are detected.

:::note

Logs are retained for 90 days

:::

When provisioning a cluster, a Time Series Database is automatically set up in the portal under the same subscription. The cluster's metrics are transferred to this data source.

The data source serves to store and analyze metrics, enabling continuous monitoring and optimization of the cluster.

This integration ensures that all relevant metrics are centrally collected and can be used for monitoring or reporting purposes as needed.

Cluster Network Interface

Clusters can be provisioned by default with the container network interface (CNI) **Canal** or **Cilium**.

Canal combines the functionalities of Flannel (for network overlay) and Calico (for network security policies) to provide a robust and flexible networking setup for Kubernetes clusters.

Cilium is based on eBPF and enables performant, scalable, and security-oriented network implementation. In addition to pod-to-pod connectivity, Cilium supports granular network policies and advanced observability features.

Container Namespace

The Container Namespace Service simplifies the use of Agile Factory and AnyCloudK8s. All required Clusterbase permissions for namespaces are provided in this service.

Namespaces are central in Kubernetes and a key element. Namespaces allow developers to build their projects modularly and outsource specific aspects to separate files. This enables them to modernize their applications and reduce complexity, allowing developers to focus on what is essential.

Service Architecture

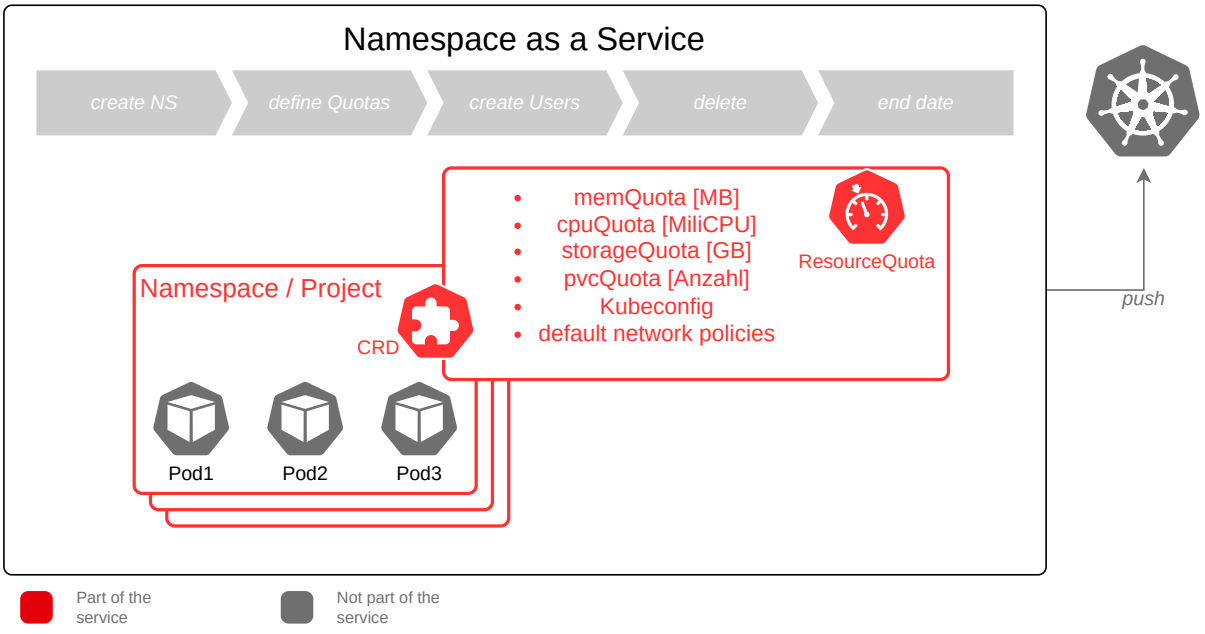

Figure: Namespace Service Architecture

Service Scope

Table: Namespace Service Scope

Feature	Rancher	OpenShift	AnyCloudK8s
Initial Setup	□	□	□
Target Destination	■	■	■
Network Policy	■	■	■
Quotas	■	■	■
KubeConfig	■	■	■
Annotations/Labels	■	■	■

:::info To ensure that system-critical namespaces remain protected, we have introduced a blacklist filter. This ensures that certain namespaces used by the cluster cannot be ordered. The following namespaces are affected:

openshift openshift- ; kube- ; trident ; argocd ; cert-manager ; patch-operator ; patch-automation ; conjur ; k8s-operators ; ix-ocpbackup-system ; ix-aai ; alertmanager-zabbix-webhook ; collectorforopenshift ; cerberus ; kyverno ; aqua :::

Service Options

The individual components of Container Namespace are described in detail below.

Initial Setup

The basis for using a Container Namespace is an existing Agile Factory or AnyCloudK8s. The subscription used for the order must be activated beforehand on the Target Destination.

For the release of available target destinations and all subsequent changes, a "Generic Request" must be submitted.

Target Destination

It is possible to define per subscription which target destinations can be used. This allows you to specifically define which users can deploy a Container Namespace or an application on which Agile Factory or AnyCloudK8s.

Network Policy

When creating a namespace on the target destination, predefined network policies are set. These must be adjusted when deploying the application according to the application's requirements.

List of network policies set by Inventx:

- Allow Namespace
- Allow DNS
- Default denyAll

Quotas

When creating a namespace, quotas must be defined. This creates the prerequisite that applications can be operated resiliently on the corresponding cluster. Quotas ensure that an application cannot exhaust all cluster resources and adversely affect other applications.

The following quotas must be specified when creating:

Table: Container Namespace Quotas

Feature	Description
Memory	Memory quota in MB on the namespace
CPU	CPU quota in millicpu on the namespace
Storage	Storage in GB on the namespace
PVC	Maximum number of PVCs in the namespace

All these quotas have a minimum value defined. The definitions can be seen in the table below:

:::note

There are applications that cannot work with the ResourceQuotas or Network Policy of the namespace. These can be disabled when ordering the namespace.

:::

Table: Container Namespace Quotas minimum value

Quota	Minimum Value
Memory	> 256MB
CPU	> 200m
Storage	>= 1
PVC	>0

The maximum value is not limited, so it is at the customer's discretion to define this according to available resources on the Agile Factory or AnyCloudK8s.

:::note

The memory quota must be chosen carefully, as it can cause applications to crash if it is chosen too low.

:::

KubeConfig

Each Container Namespace has a service account. This account can be ordered with different permissions. Either as an Admin Service Account or as a Viewer. In addition to the service account, a token is also created, which expires after the defined "Time to Live" period. The definition of whether the permission should be Admin or Viewer can only be set when creating the Container Namespace.

The KubeConfig can be displayed and copied. If the KubeConfig has expired (reached the Time to Live), it can be renewed. This ensures that the permission on the namespace does not always have the same credentials.

:::note

The Container Namespace can only be deleted on the target cluster if it was ordered with Admin rights.

:::

Annotations/Labels

This feature enables the setting of labels and annotations when creating and managing namespaces, which promotes automation and transparency in the cluster.

Labels enable clear and structured categorization of namespaces, for example by team affiliation, environment, or application type. Annotations provide the ability to store additional metadata such as descriptions, responsibilities, or operational information.

:::caution

The key of a label or annotation cannot be changed after creation. The associated value, however, can be adjusted at any time.

:::

Splunk Index

Splunk Index creation is a central location for managing Splunk indices to ensure efficient and structured management. The service provides the capability to create a Splunk Index that supports various service

options. Both technical and organizational information can be retrieved after the indices have been created.

Available Service Options

Table: Splunk Indices Service Options

Service Option	Service Option Description
Type	Splunk Index or Splunk Summary
Stage	Production or test environment
Online Retention	Retention time in days. Defines how long the logs remain in the index.
Description	Description of the index or summary. In general, "Index" should be selected. The purpose of "Summary" is to filter, aggregate, and store data from existing indices.

Available Meta-Information

Table: Splunk Indices Meta-Information

Meta-Information	Meta-Information Description
Subscription	The ix.Cloud Subscription in which the index or summary was created
Organizational Unit	The organizational unit to which the index or summary is assigned
Cost Center	The cost center for billing
Owner	The responsible person or team
Tags	Keywords or labels for categorization
Custom Properties	Additional custom properties

Enterprise Streaming Service

Enterprise Streaming Service – fully managed event streaming platform of ix.Cloud for real-time processing of events, transactions, and state changes.

:::info Status: Alpha The Enterprise Streaming Service is a community service of ix.Cloud. :::

It provides a fully managed event streaming platform that is shared by multiple customers (shared service). Tenant separation is ensured by the platform at the topic level.

The service is operated in Inventx data centers. The service meets regulatory requirements for data retention, tenant isolation, and auditability.

The following services are provided by Inventx and are included in the service.

Inventx Scope of Services

Managed Streaming Cluster

Inventx operates and monitors the entire streaming infrastructure. This includes broker management, cluster scaling, rolling upgrades, patching, and capacity planning. The customer uses the platform without having to worry about operations.

Schema Registry

Inventx provides a central schema registry. This enables producers and consumers to manage schemas centrally and ensure compatibility between applications.

Supported schema types:

- Avro
- Protobuf
- JSON Schema

Access Control & Tenant Separation

Each customer's topic begins with the customer identifier defined in the Inventx system. Within this prefix, additional topic groups can be created to make access control more granular. Permissions are granted per user and topic prefix.

Supported authentication methods:

- SASL/SCRAM
- mTLS
- OAuth

The registration and management of clients is performed by Inventx. The customer receives reports that allow them to view and verify permissions within their topics.

Topic Management

After assignment of a topic prefix and associated permissions, topics can be created and configured within this area:

- Directly via the Kafka Admin API (if appropriate permissions are available)
- Via Inventx Self-Service (planned)

Configurable topic settings:

Setting	Default	Maximum
Retention Time	7 days	30 days
Partition Count	12	96

Fixed settings:

Setting	Value
Replication Factor	3

Monitoring & Alerting

Inventx monitors cluster health, throughput, and partition distribution in real time. Proactive alerting occurs when anomalies and SLA-relevant thresholds are detected.

Security

Table: Security

Type	Description
Encryption in Transit	Always guaranteed when using mTLS authentication. With OAuth or SASL/SCRAM, the connection must be established via TLS.
Encryption at Rest	Ensured through the standard disk encryption of ix.Cloud.
Application-Level Encryption	Not covered by the service. It is the responsibility of the respective application.

Topic Monitoring & Logs

Monitoring and log information for the customer's own topics and clients are provided via reports.

Internal cluster and operational logs (e.g., broker/system logs) are not provided. Log forwarding to customer-owned systems is not supported.

Retention	Duration
Online	90 days
Offline	2 years

Customer Responsibility

The following items are the responsibility of the customer.

Integration of Applications

The customer is responsible for the development, configuration, and operation of their producer and consumer applications. The following conditions apply:

- Maximum message size: 1 MB
- Topics must use the assigned topic prefix (the application name cannot be freely chosen)
- Connectors that access shared resources within the cluster (e.g., MirrorMaker 2) cannot be used
- The Kafka version is specified by the Enterprise Streaming Service
- The replication factor is fixed and cannot be changed

Network Connectivity

The customer ensures network connectivity between their consumer/producer applications and the streaming platform (in collaboration with Inventx Network Services).

Schema Registration

Schemas must be registered in the schema registry before production.

Monitoring Usage

The customer actively uses the provided monitoring reports to monitor their topics and capacity.

Service Level Parameters (SLP)

The service level parameters are based on the performance handover point (PaaS) platform level Platinum.

Character Legend

The following rules apply to the characters per column or row in a table:

- Positions marked with "■" are included in the base price according to a separate price list.
- Positions marked with "□" are not included in the base price but can be ordered optionally. Billing is according to a separate price list.
- Positions marked with "-" are not available.