



# Delivering Performance, Scalability, and Availability

ON THE CELONIS EXECUTION MANAGEMENT SYSTEM



**Overview**

---

**Global Regions**

**Multi-tenant Architecture**

**Scalability**

**Performance**

**Availability**

**Backup and Recovery**

**Appendix and Glossary**

**Disclaimer**

# Overview

The Celonis Execution Management System (EMS) provides organizations with the capability to uncover process gaps and prescribe intelligent action to improve business execution. Our customers rely upon access to IT systems, business data, and Cloud services for their continued operation and success.

Celonis works to ensure that our backend technical infrastructure supports the front-end experience that our customers know and love. Our team is passionate about making infrastructure investments that support our customers and we want to provide you with the opportunity to understand our technical architecture.

This whitepaper provides an overview of Celonis' High Availability architecture—a key element in delivering a true enterprise-grade cloud, the Execution Management System. Our multi-instance architecture meets and exceeds stringent requirements surrounding data security, availability, and performance.

Overview

Global Regions

Multi-tenant Architecture

Scalability

Performance

Availability

Backup and Recovery

Appendix and Glossary

Disclaimer

# Global Regions

To support our global customer base, Celonis is available in data centers around the globe. Customers can select to host their Celonis tenant in the United States, Germany, UK, Netherlands, or Japan. The EMS is only hosted on best-in-class data centers. We partner with Microsoft Azure, Amazon Web Services (AWS) and IBM Cloud, the market leaders in cloud computing services. They provide best-in-class enterprise security, scalability and availability. Additional information on these topics can be found in the Appendix of this paper.



Overview

Global Regions

Multi-tenant Architecture

Scalability

Performance

Availability

Backup and Recovery

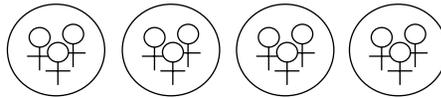
Appendix and Glossary

Disclaimer

# Multi-tenant Architecture

The EMS is a cloud-native platform built as a highly available, performant system. A multi-tenant architecture was chosen to achieve this goal. Multi-tenant applications share a single stack of resources to meet the needs of multiple tenants. There are various benefits of multi-tenancy including simple and smooth upgrades for both customer and vendor, flexibility, scalability, high performance, and low operational costs.

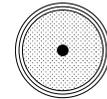
## Multiple Teams (Tenants)



## Single shared stack of software and hardware



Polymorphic Application



Database

One instance of the EMS is hosted in each data center location, which we refer to as an EMS realm. Within each EMS realm, there are multiple customer tenants and those tenants are separated following a metadata-driven approach. Each tenant is a private customer space in which a customer's application and analytical data is separated from other tenants. To find out more about our platform security, please review our [platform security whitepaper](#).

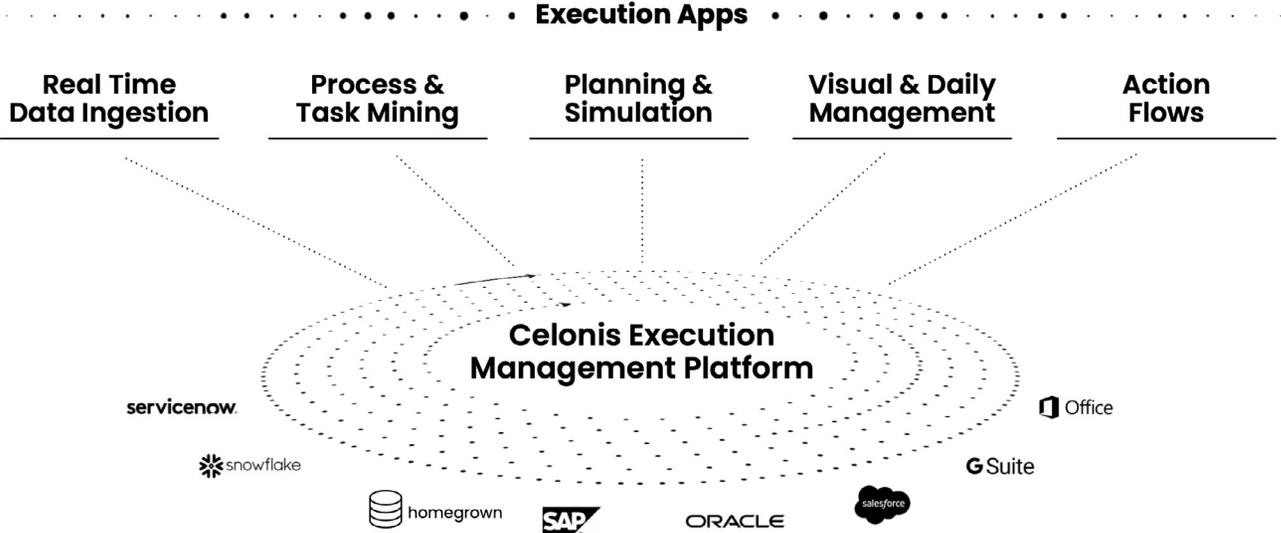
Our multi-tenant architecture also allows the Celonis Engineering team to collect and aggregate telemetry data across the full customer base, which facilitates frequent, incremental improvements that benefit the entire user community.

# Scalability

The EMS is an integral part of our customers' daily work life, therefore a highly available, performant system is crucial. The EMS deployment architecture was built to optimize scalability and performance.

Each customer has access to the application landscape, shown in Diagram 1. The application landscape of each EMS realm includes 5 pillars of capabilities: Real Time Data, Process & Task Mining, Planning & Simulation, Visual & Daily Management, and Action Flows. These capabilities are customized and packaged into applications to meet the needs of our customers.

- Overview
- Global Regions
- Multi-tenant Architecture
- Scalability**
- Performance
- Availability
- Backup and Recovery
- Appendix and Glossary
- Disclaimer



The EMS is based on a distributed microservice architecture and its cloud-native applications are deployed in Kubernetes, an industry-standard container-based architecture used by Google, IBM, and others. A high level overview of these services can be seen in Diagram 2. These infrastructure choices allow us to add more nodes to all areas of the platform without downtime and serve our customers with the availability, scalability, and performance described in the following sections.

**Overview**

**Global Regions**

**Multi-tenant Architecture**

**Scalability**

**Performance**

**Availability**

**Backup and Recovery**

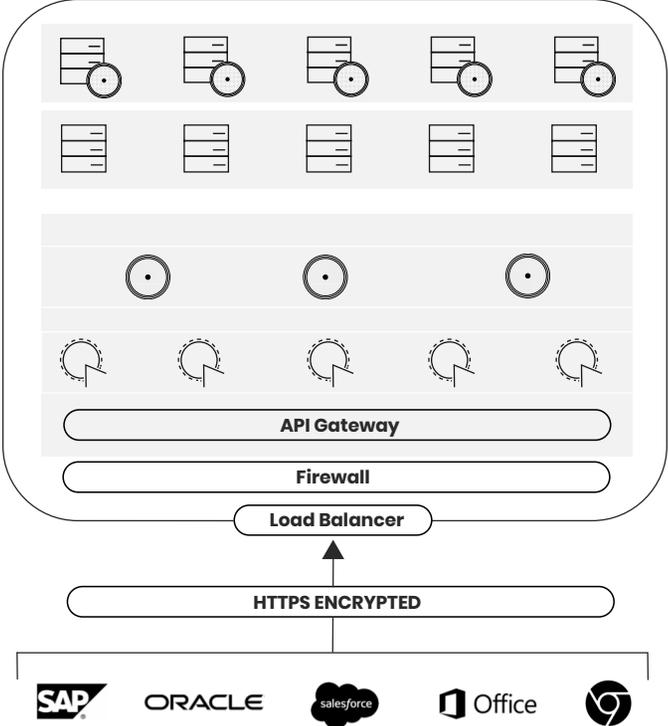
**Appendix and Glossary**

**Disclaimer**

**Data Lake**  
(analytics data)

**Process AI Engine**

**App Service Layer**  
App DB  
(metadata)  
EMS  
Applications



Our Kubernetes architecture allows for application replication and load balancing across the platform, enabling horizontal scaling across all EMS applications. Each application has its own database, further ensuring platform failure and reducing the risk of a single point of failure.

Overall, this approach allows us to handle load on the platform elastically for large-scale deployments and organizations with ease.

**Overview**

**Global Regions**

**Multi-tenant Architecture**

**Scalability**

**Performance**

---

**Availability**

**Backup and Recovery**

**Appendix and Glossary**

**Disclaimer**

# Performance

Performance is essential to deliver the best user experience. The EMS was designed to deliver best-in-class performance and to continuously deliver performance optimization improvements. Delivering a high performance system for our customers has several dimensions, including: data ingestion, data preprocessing, data analytics, and application performance.

Our extractor and data lake technologies handle data ingestion into the EMS and the preprocessing required for our operational and process mining use cases. Our extractors process large volumes of data by splitting a workload in micro-batches for high-speed data transfers. Micro-batches ensure that 1) the system load is kept at a minimum, and 2) processing steps, such as pseudonymization can be completed during the data transfer process. Operational applications require near real-time data extraction and processing. Our data lake, which is a cluster-based database, makes near real-time data preprocessing a reality as it is specifically designed to handle the workloads of big data processing.

All applications built upon the EMS platform perform data analytics using the Process AI Engine. The Process AI Engine is a proprietary in-memory database which serves many customers in parallel, each using data models with hundreds of millions of events and in some cases even more than one billion events. The Process AI Engine is a core component of the EMS platform, allowing our customers and partners to easily build applications without needing to worry about scaling compute capacity.

Customer requirements and data volumes constantly grow, therefore our platform needs to keep up. We continuously work to improve our platform performance by rolling out improvements to our entire customer base.

Overview

Global Regions

Multi-tenant Architecture

Scalability

Performance

Availability

Backup and Recovery

Appendix and Glossary

Disclaimer

# Availability

We aim to ensure that the EMS is always operational for our customers. No vacation, no extended upgrade or maintenance windows, no single points of failure. We focus on near-perfect availability and have built redundancies into every layer of our cloud platform. This means that even standard EMS platform upgrades do not require downtime. When deploying a new version of a service we gradually add more and more replicas of the new version while slowly reducing the number of replicas of the old version. During termination, open connections to the service are allowed to complete. Our cloud-native architecture allows us to scale services on-demand and upsize our infrastructure without downtime to ensure that even in peak-load scenarios, the platform is still available for our customers.

To ensure smooth operations, Celonis has a dedicated operations team with on-call duties to coordinate quick issue resolution and mitigation. The Celonis Service desk can be reached 24/7 and following every incident, the operations team ensures that each incident is logged to identify its root cause and prevent it from reoccurring in the future.

The live availability status of all EMS realms can be checked on the **Celonis Status Page**:

REALM	LINK	STATUS
EU-1	<a href="https://status.eu-1.celonis.cloud">https://status.eu-1.celonis.cloud</a>	Operational 
EU-2	<a href="https://status.eu-2.celonis.cloud">https://status.eu-2.celonis.cloud</a>	Operational 
EU-3	<a href="https://status.eu-3.celonis.cloud">https://status.eu-3.celonis.cloud</a>	Operational 
JP-1	<a href="https://status.jp-1.celonis.cloud">https://status.jp-1.celonis.cloud</a>	Operational 
US-1	<a href="https://status.us-1.celonis.cloud">https://status.us-1.celonis.cloud</a>	Operational 
US-2	<a href="https://status.us-2.celonis.cloud">https://status.us-2.celonis.cloud</a>	Operational 

**Overview**

**Global Regions**

**Multi-tenant Architecture**

**Scalability**

**Performance**

**Availability**

**Backup and Recovery**

**Appendix and Glossary**

**Disclaimer**

# Backup and Recovery

While High Availability is the primary means to restore service in the case of disruption, the EMS also leverages backups of its application and analytics data. Celonis protects data from accidental or intentional destruction due to system failures, or malicious attacks. The automated backup system is configured to perform daily incremental data backups of production databases, which are retained for 30 days.

Nevertheless, Celonis should not be considered as a backup solution for customer data - in the unlikely event of data loss on Celonis side, the customer still possesses all the data and can re-upload to Celonis. We do not support recovery of data that is lost due to customers deleting the data. The backup and recovery ensure data restoration in case of problems on our side as the service provider.

Regular disaster recovery exercises are performed to ensure that each EMS realm can be recovered as quickly as possible in case of disaster events, such as hardware malfunction or other large infrastructure outages.

Additional information and updates on Celonis' platform infrastructure and security can be found at [www.celonis.com/trust-center](http://www.celonis.com/trust-center).

# Appendix

———— **AWS Infrastructure Security**

———— **Microsoft Azure Infrastructure Security**

# Glossary

———— **Horizontal Scaling** provides the ability to scale out in order to accommodate larger workloads. Data is logically distributed (sharded) across multiple servers.

———— **Vertical Scaling** provides the ability to scale up in order to accommodate larger workloads on the same machine. Vertical scaling is limited by the maximum server size.

———— **Kubernetes** is an open-source system for automating deployment, scaling and management of containerized applications.

Overview

Global Regions

Multi-tenant Architecture

Scalability

Performance

Availability

Backup and Recovery

Appendix and Glossary

Disclaimer

**Overview**

**Global Regions**

**Multi-tenant Architecture**

**Scalability**

**Performance**

**Availability**

**Backup and Recovery**

**Appendix and Glossary**

**Disclaimer**

---

# Disclaimer

This document is protected by copyright laws and contains material proprietary to Celonis SE, its affiliates (jointly “Celonis”) and its licensors. The receipt or possession of this document does not convey any rights to reproduce, disclose, or distribute its contents, or to manufacture, use, or sell anything that it may describe, in whole or in part.

This document is provided for informational purposes only. It represents Celonis’ current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of Celonis’ products or services, each of which is provided “as is” without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances. The responsibilities and liabilities of Celonis to its customers are controlled by Celonis agreements, and this document is not part of, nor does it modify, any agreement between Celonis and its customers.

