Minerva™

Product Description

V1.2-2022

HENDRIKX ITC

Trademarks and Permissions

Minerva, 1OPTIC and other Hendrikx ITC trademarks are trademarks of Hendrikx ITC, Ltd. All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services, and features are stipulated by the contract made between Hendrikx ITC and the customer. All or part of the products, services, and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees, or representations of any kind, either express or implied. The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Hendrikx ITC, Ltd.

Address:	Hendrikx ITC Spoorlaan 21K Tilburg, Noord-Brabant The Netherlands
Website:	https://www.hendrikx-itc.nl
Email:	support@hendrikx-itc.nl

HENDRIKX ITC

Contents

1. Product Positioning 3
2. Product Characteristics 3
   2.1 Low Latency 3
   2.2 Robustness 3
   2.3 Scalability 3
   2.4 Vendor Independence 3
   2.5 Openness 3
   2.6 Security 3
3. Architecture 4
   3.1 Position 4
   3.1.1 Software Architecture 5
   3.1.2 Components 5
   3.3 Virtualization 5
4. Use Cases 6
   4.1 IoT 6
   4.2 Classic Networks 6
5. Configuration 6

HENDRIKX ITC

1. Product Positioning

Minerva is a high-performance ETL platform for standardized time series, attribute, and event type data. The standardization of these data types makes operations like aggregations and KPI calculations trivial and possible on both off- and online granularities. The combination of these standardized types of data enables advanced correlation capabilities.

2. Product Characteristics

2.1 Low Latency

Distributed, parallel processing allows for low latency of raw data availability. The advanced, standardized aggregation mechanisms allow even the aggregated data to be available for online analysis.

2.2 Robustness

The proven technology foundation of PostgreSQL enables very robust implementation of the system, with well-known characteristics in terms of availability, crash-recovery, backup, and restore.

2.3 Scalability

The platform is designed to maintain the online characteristics even for the largest of datasets.

2.4 Vendor Independence

The system is not tied to specific vendors of equipment as data sources, and the combining of data from multi-vendor networks is a key functionality.

2.5 Openness

The customer is always in control of its data by being able to access it in any way and with any tool that is required.

2.6 Security

Possibility of very fine-grained access control to data. User accounts can be configured to only be able to access parts of the data set that are required for their specific activities.

3. Architecture

Minerva consists of a core database, with distributed ingestion nodes and a set of worker services.

3.1 Position

The position of Minerva is at the foundation of a data-driven organization, connected to data delivering applications/devices and behind any number of data representation tools.

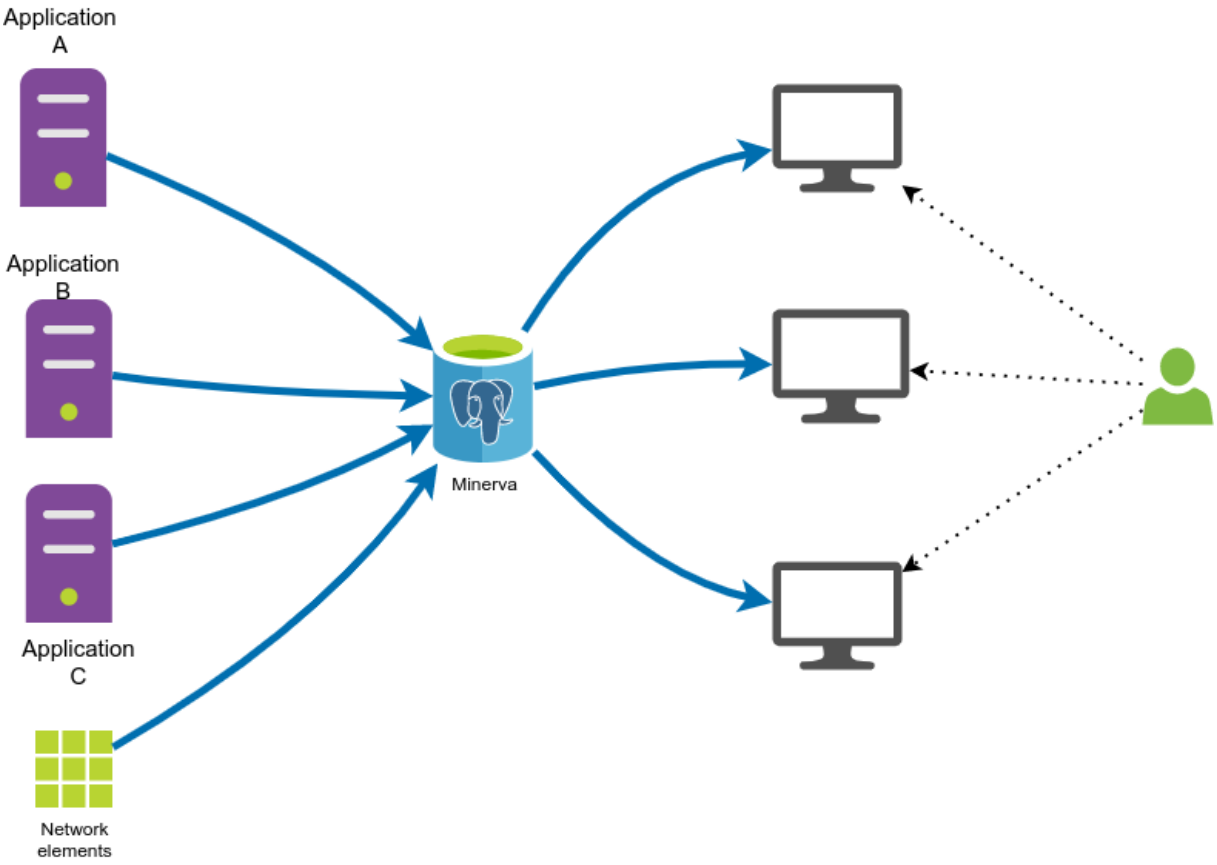

3.1.1 Software Architecture

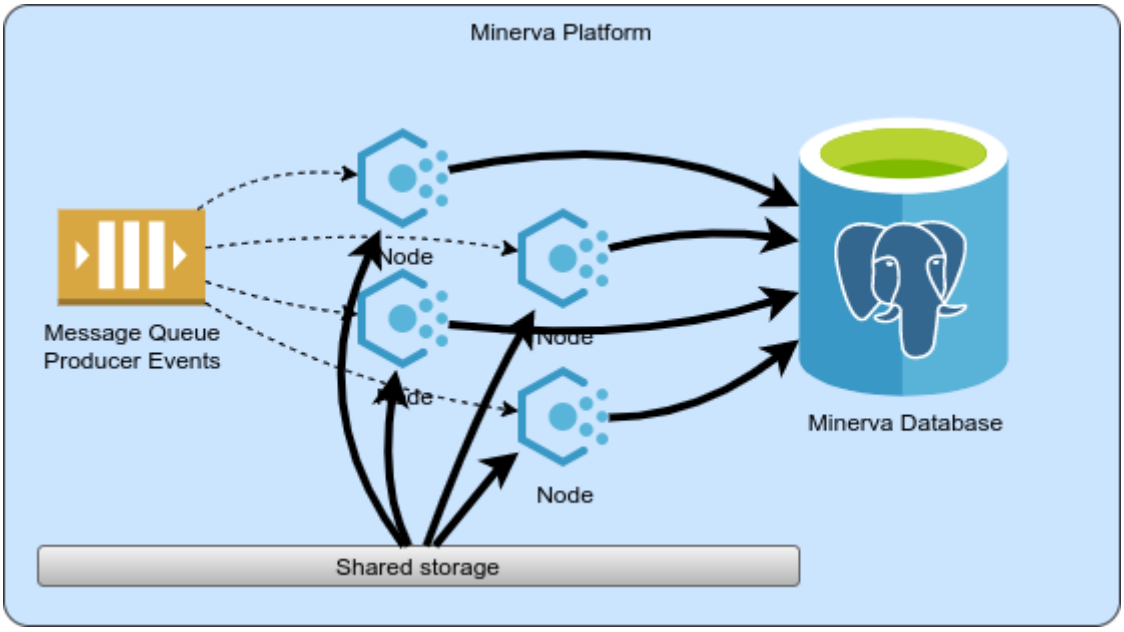

3.1.2 Components

1. Minerva Database - The core database that stores all data and executes all transformations on the data
2. Minerva Node - A parallelized ingestion node for loading data in various formats into the standardized data storage.
3. Job Queue - The queue that feeds the parallel running ingestion nodes.
4. Worker Services - A set of services that coordinate tasks like materialization (aggregation, KPI calculation), partition creation, and data retention enforcement.

3.3 Virtualization

All components of Minerva are designed to be able to run in a virtualized environment. Advantages of a Minerva deployment on a virtualized environment (non-exhaustive list):

1. Easy to scale up or down when the requirements for the system change.
2. Easy to create a strong logical separation of the components by setting them up on multiple virtual machines.
3. More options for dividing the load of the individual components.

4. Use Cases

Minerva is a generic, but opinionated ETL platform, and in this section, several specific use cases are provided.

4.1 IoT

The often heterogeneous nature of IoT requires a platform that supports multiple device types of multiple vendors. Due to the generalized and standardized data storage, this combining of data from lots of different IoT devices comes naturally.

4.2 Classic Networks

The classic networks with often high volumes of measurement data from the equipment are one of the original targets of the platform. Support for scheduled and gradual equipment upgrades of possible different vendors is what gives the customer freedom of choice without worrying about the data collection, processing, and insights.

5. Configuration

The configuration of Minerva is done using YAML-based configuration files that are human-readable and are designed to integrate well with configuration management systems (e.g. Puppet, Ansible, SaltStack). Administration tools are available for updating the database, based on the configuration files.

Component	**vCPUs**	**Memory**
Database (very dependent on data set size)	12	64GB
Worker Services (deployed on database host)	-	-
Minerva Node	4	4GB
Job Queue	4	4GB