



Almato Bardioc → Semantic data platform

Content

01 Introduction

02 Overview

2.1	Challenges	7
2.2	Operating System for Data	7
2.3	Components	8
2.4	Functional Architecture	8
2.5	Digital Sovereignty	9

03 Almato Bardioc

3.1	Access Manager	11
3.1.1	Data Protection and Data Security	11
3.1.2	Identity and Access Management	12
3.1.3	Access Token Concept	12
3.2	Data Manager	13
3.2.1	Scheduling	13
3.2.2	Memory Management	13
3.3	Message Bus	14
3.4	Knowledge Core	15
3.4.1	Storage Technologies	15
3.4.2	Scalability	16
3.4.3	Historisation	16
3.5	Knowledge APIs	17
3.5.1	APIs for Specialised Function	17
3.5.2	APIs for Real-time and Near-time Applications	17
3.5.3	Connector SDK	18
3.5.4	Functional Data Queries	18
3.6	Desktop	19
3.7	Graph Applications	21

04 Ontology

4.1	Concepts and Properties of the OGIT	23
4.2	OGIT Ontology Extensions	24



05	System requirements and operation	5.1	Software as a Service / Platform as a Service	26
		5.2	Hyperscaler	26
		5.3	On-Premises	26
		5.4	System and Integration Requirements	26
		5.5	Open Source	26
		5.6	Certiability	26

06	Integration and extension options	6.1	Knowledge Automation	28
		6.2	Large Language Models / Retrieval-Augmented Generation	28

07	Summary	7.1	Comparison with Graph Databases	31
		7.2	Competitive Advantages of Almato Bardioc	32

Turn data into competitive advantage



Almato Bardioc is a semantic data integration and processing platform that has been researched, designed and developed over a period of 20 years. It integrates numerous open-source technologies and combines them with our proprietary technology to create a complete operating system (OS) for data. The platform is designed to integrate, process and manage very large data sets from different sources. Complex relationships about meaning and context of the data are stored to retain its full meaning within the platform regardless of its source or primary use.

Introduction

Bardioc provides all necessary components to develop scalable and networked applications for data integration, data processing, data analysis and data visualisation based on a data pool. The existing applications for semantic data analysis in real time, combined with various AI technologies, enable companies and organisations to make informed decisions and harvest the full value of their data.

Bardioc offers unique functions for efficient and economical scaling for very large data sets (several hundred petabytes) with high query and processing velocity. The efficient execution of distributed algorithms and the merging of results are fully controlled by the platform. Sophisticated and tested security mechanisms enable granular data security and data protection concepts to be implemented while meeting the highest compliance requirements.

Bardioc is suitable as an ontology-based OS for companies and organisations from industry and government to make highly relevant data- and AI-driven decisions, raise operational efficiency to a new level or achieve hard-to-replicate competitive advantages. Mapping of knowledge and information within a defined context also increases the agility and innovative capacity of companies and organisations.

Key → Takeaways

-
- Bardioc is the central OS for managing and storing very large amounts of data with contextual information. The foundation of this OS is a highly scalable semantic graph.

<TIMESERIES>

✚ Bardioc

02

Overview

CONTAINS

RELATES

<ORGANIZATION>


LEADS

CONNECTS

02 Overview


2.1 Challenges

Many companies and organisations work in complex ecosystems where data is generated and stored across numerous platforms, departments or even external partners. Bardioc enables centralised integration of different data sources and therefore breaks down the conventional purpose-specific data silos. This is the only viable approach to create a comprehensive understanding of an organization's data landscape, which in turn forms the foundation for cross-functional collaboration, enabling companies and organisations to make holistic and well-informed decisions and overcome critical challenges.



01

Fundamental architecture of operating systems



2.2 Data Operating System

Bardioc's DNA is operating system design, an OS for data. The core is a set of functions that allows programmes to access data, perform calculations and communicate with other programmes. The equivalent of a file system is the knowledge core, consisting of a graph database, a BLOB database and a time series database. The equivalent of an OS's processing core is indexing, queries and the ability to execute code in the form of programmes inside the platform. The equivalent of an OS's communication layer is the message bus, which enables all other components and programmes to communicate with each other and across the entire distributed system.

Similar to any conventional OS, access protection as well as programme and data integrity are guaranteed by a horizontally integrated and granular authorisation and identity system. This security system also ensures the encryption of data at all times during transmission and at rest as well as the management and recognition of authorised applications and users.

Features

- OS for data
- In-Memory processing for high performance
- Cost-optimised operation through targeted memory allocation

02 Overview

2.3 Components

For architectural reasons which ensure security, scalability and maintainability Bardioc has the following components:


- **Access Manager** as a central access control and security layer
- **Data Manager** handles scheduling tasks, caching, and pre-emption tasks to ensure cost-efficient storage management
- **Message Bus** integrates storage and communication technologies to ensure high data throughput
- **Knowledge Core** as redundant and distributed data storage
- **Knowledge APIs** as a central connectivity framework for accessing data
- **Desktop** framework for front-end development
- **Graph Applications** are individual programmes and applications based on the knowledge graph
- **OGIT Ontology** forms the semantic context definition of all data
- **OGIT Ontology Extensions** to allow customised extensions, and domain specific language adaption to the ontology

2.4 Functional Architecture

In accordance with the fundamental architecture for operating systems outlined, Bardioc is designed as an operating system for data and has the following functional architecture:

02

Functional architecture of Bardioc




02 Overview

2.5 Digital Sovereignty

Bardioc relies extensively on open source software to support digital sovereignty in the context of government, as well as in independent industry. The architectural approach designed by Almato offers maximum independence and flexibility. Thanks to the modular architecture and the focus on well-established standards, the development of programmes for the platform and the operation of the platform are optimally supported by the expertise of the communities.

At the same time, Bardioc offers a commercial solution that is unique on the market in terms of its scalability and innovative functionality. It enables companies and organisations to gradually build a seamless data integration and processing platform that offers extensive options for expansion and customisation.


Almato Bardioc



Access Manager

3.1 Access Manager

The Access Manager is responsible for data security and data protection within Bardioc and ensures compliance requirements are met with integrated Identity and Access Management (IAM) and extensive encryption.



01

Access Manager

3.1.1 Data Protection and Data Security

Data sources may contain personal or other sensitive information in need of special protection. The implementation of data protection and security measures in accordance with legal and regulatory requirements is essential to ensure confidentiality and integrity of data in the Knowledge Graph. This is particularly relevant for sensitive project data, e.g., classified data, data subject to secrecy, or data to be processed according to the need-to-know principle.

Data protection:

- Software made in Germany
- SaaS in German data centre
- Technology available for on-premises operation

- **SSL/TLS authentication** for encrypted data transfers
- **Encryption at Rest** (hard disk encryption) based on AES256 to protect data at rest
- **Audit Logs** with digital signatures to support regulatory requirements (permanent logging of all actions, cannot be deactivated, cannot be deleted, cannot be manipulated)

03 Almato Bardioc

3.1.2 Identity and Access Management


Identities are assigned to users and applications. The combination determines access authorisation to all data and resources within the platform. The identity assigned by the Access Manager to a person can either be an OAuth-enabled account on the platform or connect to existing Single-Sign-On solutions such as OneLogin or internal Active Directories or any X.509 source.

- **Identity and Access Management (IAM)** for providing accounts (Knowledge Core Accounts) based on OAuth and for managing and controlling roles and rights when accessing data and applications within Bardioc
- **Federation** with other, independent IAM systems or X.509-capable services for user authentication with trusted identity providers or for access to external resources and services via single sign-on using SAML
- **RBAC (Role-based Access Control)** and implementation of a need-to-know principle for protected data access at application level
- Only the combination of a **user** and an **application token** will allow access to data

3.1.3 Access Token Concept

The Access Manager implements a sophisticated access token concept. All users (knowledge and core accounts as well as accounts connected via federation) and all graph applications are provided with software tokens. Access to data in the knowledge graph is only possible through a combination of user and application tokens. Unauthorised user access to the data pool is therefore not possible.

The security mesh layer uses extremely efficient bit operations to verify authorisation for individual data access ensuring that these verifications are always executed for each data element without impacting access performance. As authorisation changes are not cached, they take effect immediately in the graph. This is particularly crucial for sensitive data and classified information.




01

Security Mesh Layer

Data Manager

3.2. Data Manager

The Data Manager is responsible for scheduling pre-emption, and caching within Bardioc.



3.2.1 Scheduling

The scheduler function analyses all accesses to data and routes all queries to the nodes, where data is physically stored. By analysing data access patterns and choosing suitable storage facilities according to these patterns, the Scheduler also helps to optimise operating costs of the platform. It also calculates probabilities for upcoming queries and uses idle resources to pre-compute them for improved access performance.

3.2.2 Memory Management

Bardioc's proactive memory management ensures that data is stored on the storage technologies suitable for its access patterns. Changes in access patterns are therefore reflected in the data storage. Data can be stored in memory, on fast block devices, slow block devices, or other storage media.

A sophisticated system of caches further optimizes performance by ensuring that processing takes place in working memory as far as possible. Copy-On-Write enables fast initialisation of applications, even if they load large amounts of data. This function also makes data only being

written upon changes. De-duplication to avoid unnecessary redundancies is an integral part of the system.

The data manager also provides functions to enable guaranteed real-time or near-time processing. Conventionally, messages and processing steps propagate through the system, resulting in a time delay (eventual consistency). This latency can be largely eliminated through scaling and prioritisation at query level.

Message Bus

3.3 Message Bus

A high-performance message streaming platform is used within Bardioc. The message bus integrates Bardioc's various storage and database technologies processing facilities, APIs and ensures high data throughput and low latency when processing data. Any event passing through the platform is handled here.

The Bardioc message bus provides interfaces and connectors (Producer, Consumer, Streams, Connect and Admin) that expose the following functions:




All communication within Bardioc is routed via the message bus according to the authorisations of the access manager. Additionally, the message bus offers advantages for stream-based data processing within the graph:

The use of a message backbone also supports dynamic scaling of Bardioc by parallelising and dynamically distributing load.

- System-level communication between the components of the data operating system (storage, I/O, and compute)
- Communication between Graph Applications and the Knowledge Core through the API layer
- Communication between the Graph Applications among each other (including Chain of Responsibility, which enables the operation of identical applications in different versions)
- Subscription to data streams by Graph Applications for near- or real-time processing

- **Dynamic data updating** helps to integrate data sources into the data stream, allowing the graph to be continuously and flexibly enriched with new data - even in real time.
- **Distribution of data changes** to the various Bardioc databases so that data changes are promptly reflected in the graph
- The **event-driven architecture** allows the graph to be designed to respond to specific events in the data streams, enabling dynamic and customisable data processing workflows
- The **orchestration** of data streams facilitates the integration and transformation of data from different source systems

Knowledge Core



3.4 Knowledge Core

The Knowledge Core is the actual Data Store of Bardioc. This is where graph data, time series data and BLOBs are stored, indexed, and cached in a redundant and distributed environment to facilitate high-performance yet cost-efficient access.

3.4.1 Storage Technologies

The Knowledge Core uses and combines the following storage technologies to cater to the highest standards in scalability, performance and efficient resource utilisation:

- Scalable **Graph Database** for storing and querying semantic graphs with hundreds of billions of nodes and edges
- Scalable and highly available **NoSQL Database** for distributed data storage
- **Time-Series Database** (TSDB) specially optimised to support high query performance and integrated statistics functions
- Scalable **BLOB Database** (Binary Large Object) for storing unstructured data and objects (Photos, Videos, Containers etc.) in a high availability set up
- Scalable **Key-value Database** supporting high-performance queries in a high-availability setup
- Distributed and multi-client capable **indexing** with multi-tenant support

Log files are a special type of time series. The Knowledge Core supports multiple log entries at exactly the same time as well as a temporal normalisation of log entries (i.e. to defined time intervals) for statistical evaluations.

03 Almato Bardioc

3.4.2 Scalability

As the number of data objects and versions increases, so do the platform's resource requirements (memory, CPU and bandwidth). Efficient storage solutions and strategies are required to maintain scalability and data integrity in such a dynamic system. Bardioc achieves linear scalability for very large data sets under heavy usage (tested with several hundred petabytes) through:

- **Distributed Nodes** with distribution of data to several nodes to avoid single points of failure and multiple copies to maximise performance
- Support for **replication** and **distributed data centres** with fully implemented failover, disaster recovery and business continuity strategies
- **Linear scaling** of read and write throughput by self-scaling resources without downtime or other types of operational restrictions
- **Sharding and partitioning algorithms** for horizontal and vertical scaling and for optimized load distribution even for very large data sets
- Optimised **index mechanisms** for fast response times even when processing complex data queries
- Automated query and traversal acceleration by using **map-reduce** and other distributed computing algorithms on processing large and highly interconnected data sets

With Bardioc, extensive data analyses can be automatically distributed within the semantic graph. This prevents all queried data from first being transferred to the analysis programme.

3.4.3 Historisation

Any semantic data platform to manage historical data objects effectively requires a mechanism to enable queries of the graph in the context of a specific historical point in time of the graph. Data with different validity or temporal categorisation can be conveniently managed and queried on Bardioc. Data historisation is part of Bardioc's DNA offering management of a specific point in time:

- Support of a **temporal query capability in the overall context** of the graph through time and versioning concepts
- **Versioning** of objects (nodes, edges, properties)
- **Traceability** of changes to data objects over time to support regulatory requirements

Historisation is a basic prerequisite for compliance, audit compliance and explainability and therefore a prerequisite for the use of a data platform in any regulated environment.

Knowledge APIs

Knowledge APIs				
				
REST	WebSocket	Program	Streaming	Connector SDK

3.5 Knowledge APIs

Knowledge APIs provide a comprehensive connectivity framework for interacting with the Bardioc platform.

3.5.1 APIs for Specialised Functions

For optimum efficiency, API functions can be introduced via two different technical API concepts:

- **REST API:** The REST API provides a stateless interface to the Bardioc platform. Each individual call stands alone, is authorised separately, and can execute exactly one function call at a time.
- **WebSocket API:** The WebSocket API provides an interface that creates an open connection between an application and the platform. The connection is authorised at the beginning and is subject to regular renewal of encryption tokens. A sequence of different functional calls to the platform is possible per session.

3.5.2 APIs for Real-time and Near-time Applications

Bardioc has two access methods that have been specially designed for the development of real-time and data-intensive applications:

- **Streaming API:** The Streaming API establishes a permanent connection to the system in the same way as the WebSocket API. However, queries are registered in the system and the system only carries them out under certain conditions. For example, changes in the data can be automatically forwarded to an application or calculations can be carried out automatically under predefined conditions.
- **Program API:** With the Programme API, programmes that access data in the Knowledge Graph can be transferred to the system in binary form (e.g. as .jar). By analysing the data access of the programmes, the programmes can be transported to the platform resources on which

03 Almato Bardioc

the corresponding data reside. This allows calculations to be performed on local data with very short access latency, which are then automatically reassembled via the platform. This allows calculations to be carried out on very large databases at high speed.

3.5.3 Connector SDK

Integrated data from tabular data in relational or object-oriented data warehouses and the semantic graph is fundamental to the Bardioc platform. It requires a conversion and synchronisation between tables and graphs. This creates challenges in terms of query and write performance and requires specialised conflict management.

The Connector SDK of the Knowledge Platform provides the framework, documentation and ready-made skeleton code to enable such bi-directional connections between traditional data sources and the semantic graph.

The connector SDK significantly shortens development time for integrating with traditional data sources - even with complex data models - connected to enterprise data warehouses, data lakes, and EAI systems.

3.5.4 Functional Data Queries

In addition to the technical methods mentioned above, different functional accesses to data can be achieved via the Knowledge API layers:

-
- **Point queries via Open Graph Protocol (OGP):** OGP is the simplest protocol for querying the knowledge graph if exactly one node is to be searched for or queried. The edges of the node can be followed individually with separate API calls. OGP is suitable for less complex data queries.


 - **Graph-Traversals via Gremlin:** Gremlin is the most used graph traversal language supported for implementing real parallel algorithms, queries and searches on the graph. Gremlin enables both complex searches across paths in the graph and the parallelisation of queries using the platform's capabilities, as well as implementing specific graph traversals.

 - **Graph queries with GraphQL for tabular output:** GraphQL is available as a query language for graph data to convert semantic data into tables, matrices or vectors so that they can be processed statistically or with other evaluation tools optimised for linear data.

Desktop


3.6 Desktop

Desktop is Bardioc's runtime environment for applications with a graphical user interface. Bardioc Desktop includes the rights and roles concept of the Access Manager via tokens and controls communication with the graph.



Various standard applications are available, and the set can be extended with individual or shared applications to cover specific needs:

- **Selective queries via Open Graph Protocol (OGP):** OGP is the simplest protocol for querying the knowledge graph when exactly one node is to be searched or queried. The edges of the node can be followed individually with separate API calls. OGP is suitable for less complex data queries.
- **Profile Manager** for managing Bardioc accounts
- **Graph Explorer** for visualising graph data and displaying meta information
- **Ontology Manager** for managing and extending the OGIT ontology



03 Almato Bardioc

The Front-end SDK can be used to create customised applications that are also subject to the rights and roles concept of the Access Manager. The frontend SDK provides the following features for applications:


- | | |
|--|---|
| <hr/> <p>→ Authentication and authorisation with the Access Manager</p> <hr/> | <hr/> <p>→ Responsive Design Templates for the convenient creation of interfaces, including automatic adaptation of applications to the respective output device</p> <hr/> |
| <hr/> <p>→ Messaging System for communication with the graph and between the applications</p> <hr/> | <hr/> <p>→ Pipeline for deploying front-end applications on Bardioc</p> <hr/> |
| <hr/> <p>→ Error Handling for standardised error processing</p> <hr/> | |

Graph Applications

3.7 Graph Applications

Bardioc is an OS for data and, with Knowledge Core and Desktop, provides both the compute platform and the graphical user interface that can be expected from an OS. At the graph level, sophisticated algorithmic approaches and implementation of various concepts from graph theory are made available.

For example, you can train and use your own deep learning networks via the machine learning service. Basic applications are also available at the graphical level, which enable the management of profiles, platform components and the ontology.




An SDK can be used to create graph modules as well as front-end applications for data visualisation, data management and administration. Graph-Modules are programmes that run on the graph and are managed by the Data Manager. SDK libraries facilitate the use of functions such as identification, distribution, distributed computing, programme communication and security, thus accelerating the development of individual applications on your own data.

The following graph algorithms are integrated as standard:

Traversal Recipes	
Between Vertices	
Centrality	Degree Centrality Betweenness Centrality Closeness Centrality Eigenvector Centrality PageRank Centrality
Collections	
Connected Components	Small graph traversals Small graph scalability Large graphs
Cycle Detection	
Duplicate Edge Detection	
Duplicate Vertex Detection	
Moving an Edge	
Element Existence	
If-Then Based Grouping	
Looping	Conditional Looping with Max Depth Emitting Loop Depth Optional Loop Depth
Operating on Dropped Elements	
Pagination	
Recommendation	
Shortest Path	
Traversal Induced Values	
Tree	Lowest Common Ancestor Maximum Depth Time-based Indexing
OLAP traversals with Spark on YARN	Approach Prerequisites Running the job Explanation Additional configuration options
Anti-Patterns	
Long Traversals	
Unspecified Keys and Labels	
Unnecessary Steps	
Unspecified Label in Global Vertex lookup	
Steps Instead of Tokens	
has() and Traversal Arguments	

Ontology



From one-time integration to exponential returns.

With semantics as your foundation, each data integration unlocks exponentially more value over time. Once your data is semantically linked, it becomes reusable across projects, systems, and decisions, no duplication or new modeling needed. This turns every integration into an investment that repeatedly pays off.

Ontology

The central element of the Bardioc data model is the OGIT ontology, which was defined as part of the development of the platform and has been evolving ever since. The ontology maps the semantic context for structuring, integrating and analysing data. The process for further evolving the ontology and for integrating standard domain specific languages (DSL) are part of Bardioc's DNA.

4.1 Concepts and Properties of the OGIT Ontology

The OGIT ontology published by Almato defines and organises the relationships between different data entities by mapping an information technology-based view of the world based on business processes and people.


The Bardioc OGIT ontology is available as open source. In addition, individual extensions and evolution of the ontology are possible with private ontology extensions.

A defined process is available for the evolutionary development and maintenance of the ontology and taxonomy. This process is managed through an open-source approach and is publicly available at all times. This process documents who is responsible for which ontology extensions and how links between such sub-ontologies are decided and approved by the ontology board.

- | | |
|---|---|
| <ul style="list-style-type: none"> → Mapping of nodes (data entities, node, vertex) as basic building blocks of the ontology to represent real objects, concepts or events that are modelled in the data system. Examples could be customers, people, products, transactions, processes or locations. | <ul style="list-style-type: none"> → Mapping of bridge nodes to connect different semantic data networks with each other. |
| <ul style="list-style-type: none"> → Mapping of properties (label, attribute, labelling) to describe specific characteristics. For example, a customer entity could have attributes such as name, address, date of birth and customer number. | <ul style="list-style-type: none"> → Hierarchies can be used to organise entities in order to map complex data structures. Categorisation and classification can be used to organise products into categories and sub-categories, for example. |
| <ul style="list-style-type: none"> → Mapping of edges (connections, link, edge, relationship) to connect data entities and to represent the relationship between the entities. The edges form the basis for semantic analysis and enable navigation through the data network. An example is the relationship between customers and orders, which shows which customer has placed which order. | <ul style="list-style-type: none"> → Rules and constraints can be used to define fixed rules and restrictions in the ontology in order to ensure data integrity. For example, it can be specified that each order must have at least one customer or that a date of birth must always be in the past. |

04 Ontology

- **Semantic enrichment** allows data to be enriched with additional contextual information using the ontology. This increases the understanding of the data and improves the analysis options. External data sources can be integrated or AI technologies such as NLP (Natural Language Processing) can be used for this purpose.
- Since different data sources often use different terminologies and concepts to describe similar things, so-called semantic gaps arise. The OGIT ontology ensures **semantic heterogeneity** by defining a common vocabulary for relationships and their meaning within the graph.
- Thanks to the use of open standards, the OGIT Ontology is characterised by a high degree of **interoperability** and **integrability** with data from different sources.
- The OGIT ontology offers **dynamic adaptability** through evolutionary development. New requirements for the data model or new data sources can easily be integrated without losing existing context of the model.



01


Meta Model of the OGIT Ontology

4.2 OGIT Ontology Extensions

The OGIT ontology can be extended with domain-specific extensions. Existing ontologies from the user's own environment or standardised and publicly available ontologies can also be integrated. OGIT Ontology Extensions ensure that these domain-specific languages can be integrated without adaptation, can evolve independently and are still fully integrated into the semantic context of the OGIT Ontology.

Thus, ontologies or other semantic structures that have already been developed and recognised as standard can be included in the Bardioc data pool without increasing the maintenance effort. At the same time, newly added data can be linked to existing data leaving limitless semantic interpretation or data queries.

System Requirements and Operation



Bardioc is designed for extreme scaling, high security and performance. The requirements for infrastructure and operation are correspondingly demanding. The platform can be used in various operating models.

05 System Requirements and Operation

5.1 Software as a Service / Platform as a Service

Almato operates Bardioc in DATAGROUP's own and geo-redundant data centres in Frankfurt am Main. The data centre operation is certified according to the security standards of ISO 27001 and the quality standards of ISO 20000. Continuous monitoring of all IT components, services and capacities as well as a high degree of automation enable high service levels and cost-effective operation.

The functions of Bardioc can be used as Software as a Service (SaaS) on the tenanted, centralised platform. Alternatively, a dedicated platform can be provided. This option is also available with a platform as a service approach.

5.2 Hyperscaler

Bardioc is compatible with all major cloud providers, in particular with the platforms of the major hyperscalers Amazon Web Services, Microsoft Azure and all cloud platforms based on OpenStack. Appropriate reference architectures and platform specific definitions and automations are available for deploying the platform in such environments.

5.3 On-Premises

For customers with high confidentiality and security requirements or an explicit make strategy, Bardioc's technology is also available for on-premises operation in their own data centre. Due to the infrastructure and system resources required, an on-premises strategy should only be considered by customers with very high rollout requirements.

5.4 System and Integration Requirements

The system and integration requirements for on-premises or cloud installations are determined in customer-specific requirements workshops.

5.5 Open Source

Bardioc uses a range of proven open source components that are carefully selected and used specifically according to their strengths. A list of the components with the corresponding licences can be found at: <https://bitbucket.org/almatoag/workspace/projects/KG>

5.6 Certifiability

Bardioc has been successfully certified to ISO 27001 and ISO 25051 in customised environments. Certifications can be carried out for specific customised installations if required.


Production ready:

-
- Platform as a Service from Almato in German data center

 - All Hyperscalers (AWS, Azure, OpenStack)

 - On-premises operation in your own data center

Integration and Expansion Options



Bardioc can be expanded in many ways to meet specific customized requirements.

06 Integration and Expansion Options

6.1 Knowledge Automation

The Knowledge Automation Platform (see separate product description) combines an automation engine based on problem solving artificial intelligence (reasoning engine) with a knowledge management engine based on deep learning and thus offers a framework for the formalisation of operational knowledge and the sustainable automation of highly individual tasks or processes.

Thanks to this unique combination of technologies, the Knowledge Automation Platform can solve tasks comparable to a human expert and is not dependent on the mandatory adherence to well-defined procedures, runbooks or scripts for automation.

With the help of knowledge acquisition tools, complex tasks can be broken down into atomic units of information, so-called knowledge items. The Knowledge Automation Platform finds the right solution for given tasks based on the context of the problem and the available knowledge items.


6.2 Large Language Models / Retrieval-Augmented Generation

In combination with a Large Language Model (LLM), Bardioc can be expanded into a RAG system (Retrieval Augmented Generation). The LLM can be used to submit queries to the knowledge graph in natural language, whereby the answer is also generated in a user-friendly and comprehensible form. The RAG system provides a concrete answer based on the relevant data or structured knowledge stored in the knowledge graph. This means that not only isolated facts, but also their semantic relationships and temporal dependencies can be considered when answering.

Knowledge Graph-based RAG systems offer the following advantages:

- **Improved accuracy** of the generated answers through the use of structured knowledge
- **Extended contextuality** through the linking of entities and their relationships
- **Greater explainability of the answer** thanks to the known origin of the information, which increases transparency and traceability, especially in critical applications

06 Integration and Expansion Options



01

Architecture for Retrieval-Augmentation Generation based on Bardioc

<TIMESERIES>

✚ Bardioc

07

Summary

CONTAINS

RELATES

<ORGANIZATION>

LEADS

CONNECTS

Summary

7.1 Comparison with Graph Databases

Compared to graph databases, Almato Bardioc offers the following advantages:

Attribute	Graph Databases	Almato Bardioc
Development and operation of the solution in Germany	-	+
Fully integrated semantic data platform (operating system for data)	-	+
Integrated ontology with an information technology characterised view based on processes and persons	-	+
Linear scalability up to several hundred petabytes of data volume	-	+
High-performance handling of super nodes (highly networked nodes)	-	+
High speed for complex queries or transactions in the graph	-	+
Efficient memory management thanks to sophisticated memory management	-	+
De-duplication to avoid unnecessary data redundancies	-	+
Inherent system security through token-based authorisation concept for users and applications	-	+
Effectiveness of authorisation changes in real	-	+
SDKs for the development of Graph Applications, frontend applications and the integration of APIs	-	+
Historization for temporal query capability in the overall context	-	+
Audit logs for the highest compliance requirements	-	+

Summary

7.2 Competitive Advantages of Almato Bardioc

- Bardioc offers **security by design**. Encrypted data storage and transmission, granular authorisations at data record level and token-based access mechanisms are inherent design criteria of the platform and not optional add-ons.
- Bardioc offers **linear scaling options** for very large data volumes. Many data projects start as pilot projects with small amounts of data. If the data volumes scale quickly and requirements for distributed systems, independent and efficient resource allocation without manual intervention and parallelisation of tasks are added, simple graph databases very quickly reach their limits.
- Bardioc's **historization function** creates the basic prerequisites for compliance, audit security and explainability - prerequisites for using a data platform in any regulated environment. At the same time, it ensures a high level of efficiency by automatically allocating the most effective storage medium.
- With the OGIT ontology, Bardioc has its **own ontology** that supports the integration of other domain-specific languages (DSL) as well as the integration of ontology extensions. The defined and transparent process for evolving the ontology creates the conditions for an integrated semantic data pool.
- The stringent design of Bardioc as an **operating system for data** and the availability of extensive APIs and SDKs means that programmes and applications can be fed to the data in the graph instead of having to extract the data from the graph and feed it to the programmes.

Start your revolution.

Almato

About Almato

Almato AG is Europe's leading specialist in AI-driven semantic data platforms. With Bardioc, Almato provides a sovereign, high-performance, AI-powered semantic data platform for the intelligent integration, analysis, and utilization of data. Based on an open ontology, the platform operates securely and transparently in line with European standards and can be deployed in various operating models, including as a Platform as a Service (PaaS) offering. It enables organizations to utilize their data flexibly and on a large scale across various operating environments.

Around 200 Almato employees develop and maintain the platform, supporting organizations in industry, the public sector, and defense in using their data efficiently and securely, from integration and semantic modeling to AI-driven analysis and automation.

Almato collaborates with leading partners specializing in implementation and consulting, particularly in the fields of data science, security, and solution development. Together, they help customers build comprehensive, data-driven strategies.

In addition to its headquarters in Stuttgart, Almato has offices in Bonn, Reutlingen, Neu-Isenburg and Barcelona. In 2025, Bardioc was named 'AI-Powered Semantic Data Platform of the Year in Europe' by CIO Review magazine.

Company

Almato AG
Theodor-Heuss-Str. 9
70174 Stuttgart
Phone: +49 711 3406-7810
E-mail: info@almato.com

Offices

Stuttgart
Bonn
Neu-Isenburg
Reutlingen
Barcelona

Release 7 - 10.2025
