



Onum Technical Overview




Table of contents

Understanding the Challenge	2
Our Solution	3
Platform Capabilities	4
Data Reduction.....	5
Data Enrichment.....	5
Data Optimization.....	5
Data Routing.....	6
Architecture & Data Observability	6
Flexible Deployment without Limits.....	7
Management Layer.....	7
Communications.....	8
Multi-Tenant Management.....	9
API and SDK.....	9
Data Observation Layer.....	9
Data Distribution Stage.....	12
Data Processing Stage.....	14
Compliance and Audits	15
EU Data Residency.....	15
Third-Party Audits.....	16
Conclusion	16



Understanding the Challenge

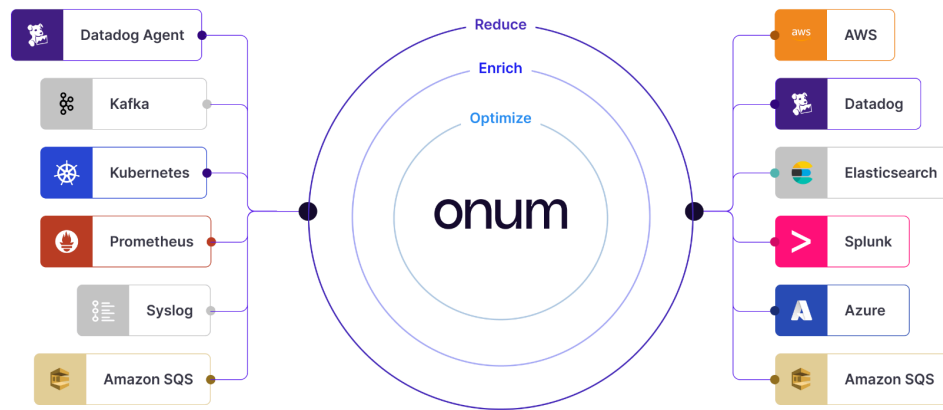
The volume of data generated globally is massive and rapidly increasing, and is estimated to reach 175 Zettabytes globally in 2025 (source [IDC Data Age 2025](#)). This data is diverse in typology and only a small part is business-critical and therefore requires immediate analysis. The large volume and variety of data makes it difficult to process and extract valuable information, which makes it challenging to generate informed decisions that are critical for the business. Data analysis platforms, while useful, do not help identify the value of each type of information before analysis, and can be expensive due to their data ingest pricing models and lack of capabilities to interact with other platforms in the ecosystem. Data observability tools offer some relief, but often introduce additional infrastructure costs and latency, and do not provide early detection or alerting when potential problems are discovered.

It is essential to maintain control over all company data and contain costs while processing that data from multiple sources to any destination — all in pure real time.

Let's see how Onum solves these challenges to facilitate making appropriate, instantaneous decisions for the business.



Our Solution



Onum is a cutting-edge solution for real-time observability, designed to enable organizations to harness the full potential of their data.

Onum's pioneering design and architecture for agnostic and instant data processing deliver unique, industry-leading capabilities that enable companies to:

- **Optimize data analysis costs:** Onum's intelligent data reduction and filtering capabilities ensure that only the most valuable data is sent directly to analytics platforms for immediate analysis, reducing costs and improving efficiency.
- **Regain control and traceability of data:** Onum can maintain a high-fidelity copy of the original data and provide complete traceability, avoiding vendor lock-in and streamlining regulatory compliance.
- **Orchestrate data agnostically:** Enables data processing from all sources to any destination, transparently, allowing companies to maximize their investment in existing infrastructure.




- **Get real-time alerts:** Onum provides immediate detection of anomalies, security threats, and problems with source systems, enabling proactive responses and minimizing downtime.
- **Unify hybrid networks:** Onum seamlessly integrates with on-premises, cloud, OT, and IoT environments, providing a unified view of the data landscape and eliminating information silos.
- **Deploy in any infrastructure:** Onum is scalable, highly efficient, and supports any type of installation, such as public cloud, private cloud, Onum's SaaS environment, on-premises, or any combination of these.
- **Enable any use case:** Onum's flexible architecture and open platform goes beyond data reduction to drive use cases such as cybersecurity, network performance and reliability, eCommerce point-of-sale transactions, and detection of Shadow IT.

Platform Capabilities

Onum is a modern, efficient, real-time observability platform that enables organizations to effectively handle large volumes of data from various sources while adding value to their data in transit.

The agnostic nature of Onum provides the flexibility to orchestrate data from any source to any destination.

To simultaneously serve a wide range of use cases and needs, Onum can make multiple copies of the organization's data, so that each team can have the context required to take decisive action. Each copy of the data is transformed into the format required by the receiving tool, thus avoiding dependence on any one analytics platform's schema and lack of interoperability ('vendor lock-in').



Data Reduction

- **Filtering:** The platform processes data in real time, intelligently based on user-defined data streams, and eliminates unnecessary noise by identifying and removing redundant, duplicate, or irrelevant information.
- **Sampling:** Not every data event is worth storing. Onum can avoid duplicates, redundant, and valueless data by applying techniques to select subsets of representative information.

Data Enrichment

- **Context aggregation:** Data from multiple sources can be combined and additional aspects can be added to improve the understanding and value of the data, such as threat intelligence, user details, status of other systems, and more.
- **Correlation:** Enables a more complete and contextualized view by looking across all data in the hybrid environment. To enrich the correlation, 'lookups' can be performed, which are searches on reference data sets containing relevant information.



Data Optimization

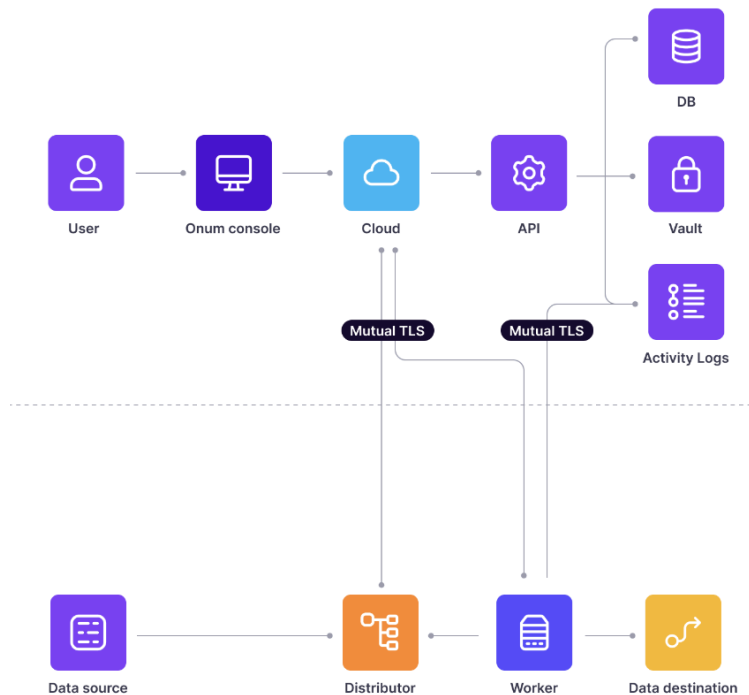
- **Compression and encoding:** The platform can apply compression and encoding techniques to reduce data size without loss of information, improving network and storage usage efficiency.
- **Transformation:** Data can be modified in real time to standardize formats, convert units, or extract specific information, making it more useful for further analysis.

Data Routing

- **Agnostic distribution:** The platform can route data from any source to any destination, such as analytics platforms, storage systems, or specific applications to enable any use case and meet business needs.
- **Real-time processing:** Thanks to Onum's pioneering design and architecture, data can be collected, processed, and delivered in milliseconds. Simultaneously, real-time alerts can be sent when the platform detects any anomalies, security threats, system issues, etc.



Architecture & Data Observability



Onum is purpose-built for technology-agnostic, secure data observation. It seamlessly interacts with any source or destination, providing unparalleled versatility for real-time processing and information routing.

Flexible Deployment without Limits

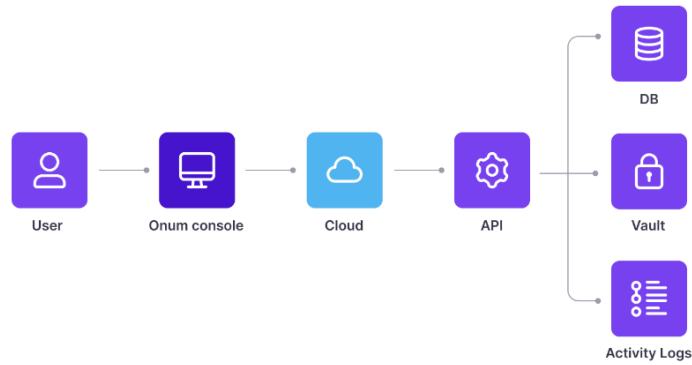
By default, Onum is offered as a SaaS platform. However, it can also be deployed as a hybrid model or with all elements located on-premises.

Additionally, the platform can operate in a completely isolated network (air-gapped), such as those typically found in operational technology settings, including industrial control systems (ICS) and supervisory control and data acquisition (SCADA) environments, or with classified information.

Regardless of deployment, Onum's infrastructure comprises data observation/processing and management layers.



Management Layer



Onum’s management layer comprises all necessary infrastructure elements required for centralized administration and operation of the platform.

Management elements	
Management Server	The management server is composed of elements such as databases, a secret manager, or activity logs.
Admin Console	Manage every configuration and operation effortlessly with our intuitive, fully visual interface — no command line or coding required. Access it from any web browser for agentless, hassle-free administration. Monitor data



Management elements	
	flows and define operations with ease, ensuring optimal efficiency.

Communications

The management server and observation clusters communicate securely via mutual TLS encryption. Clusters initiate transmissions at configurable intervals, sending metrics and receiving updates on data flows.

North-south micro-segmentation isolates the management layer from production data processing. The management server receives only metadata and status information, preserving data privacy and ensuring control plane failures don't impact the data plane.

Multi-Tenant Management

The multi-tenant management server enables organizations to isolate information between different departments, services, or even countries to ensure clear separation of responsibilities and operations. For managed service providers, this multi-tenancy also guarantees secure data segregation between different clients on the Onum platform.

API and SDK

The management layer features an application programming interface (API) for seamless integration with external systems, enabling automated workflows and data exchange.



Additionally, our software development kit (SDK) empowers developers to programmatically create custom listeners, actions, and sinks for enhanced data flow control.

Data Observation Layer



Streamlined data collection, distribution, processing and routing is performed without the use of any agent.

Designed for edge computing, the observation layer can be installed on a Linux machine or Kubernetes cluster. Its minimal resource footprint enables deployment close to data sources, minimizing latency and maximizing efficiency.

Observation elements	
Data Source	These are the origin points for the company's data. Onum's agnostic nature allows it to seamlessly connect with various sources, including applications, storage systems, security devices, and more.

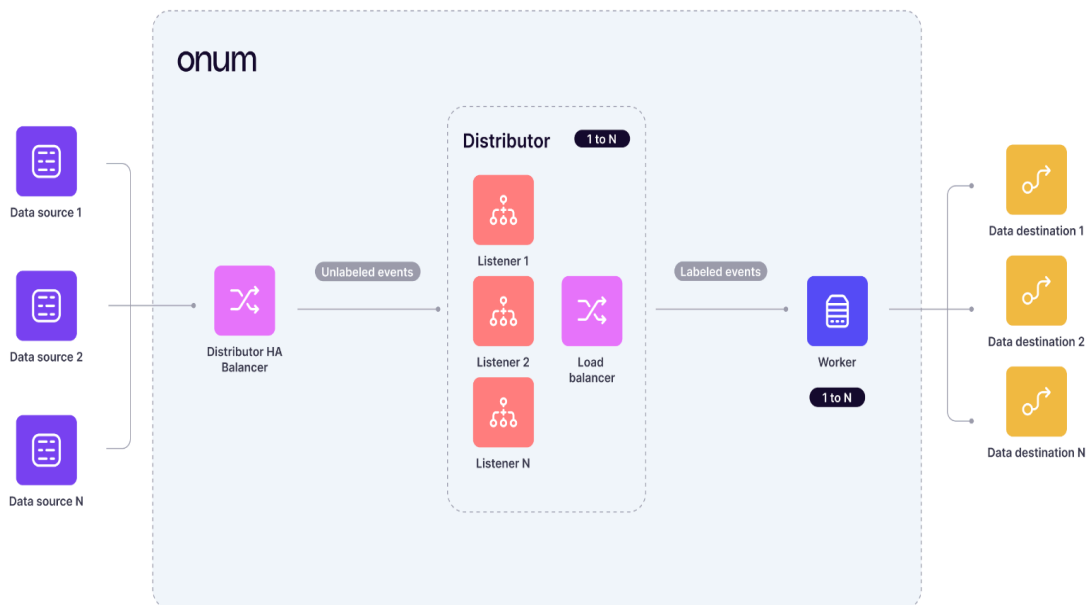


Observation elements	
Distributor	<p>A service that hosts listeners and load balancers. It efficiently distributes data to workers, enabling homogeneous workloads.</p>
Worker	<p>Workers are the pistons of Onum, executing the data flows defined in the console. They receive data from the distributor, apply the specified actions, and then deliver the processed data through the sinks to the data destinations.</p>
Cluster	<p>A Cluster is a group of distributors and workers that collaborate to distribute and process multiple instances of the same data flow using a round-robin load balancing approach. This ensures efficient and scalable data handling.</p> <p>Distributors and workers implement persistent queues of configurable size depending on available resources and processed data volume to ensure that there is no data loss in transit.</p> <p>A tenant can host as many clusters as necessary.</p>



Observation elements	
Data Destination	These are typically storage systems or analytics platforms, among others, where Onum delivers the transformed data for further analysis or treatment.

Data Distribution Stage



Data collection is precise and efficient, thanks to the intelligent and future-proofed design.

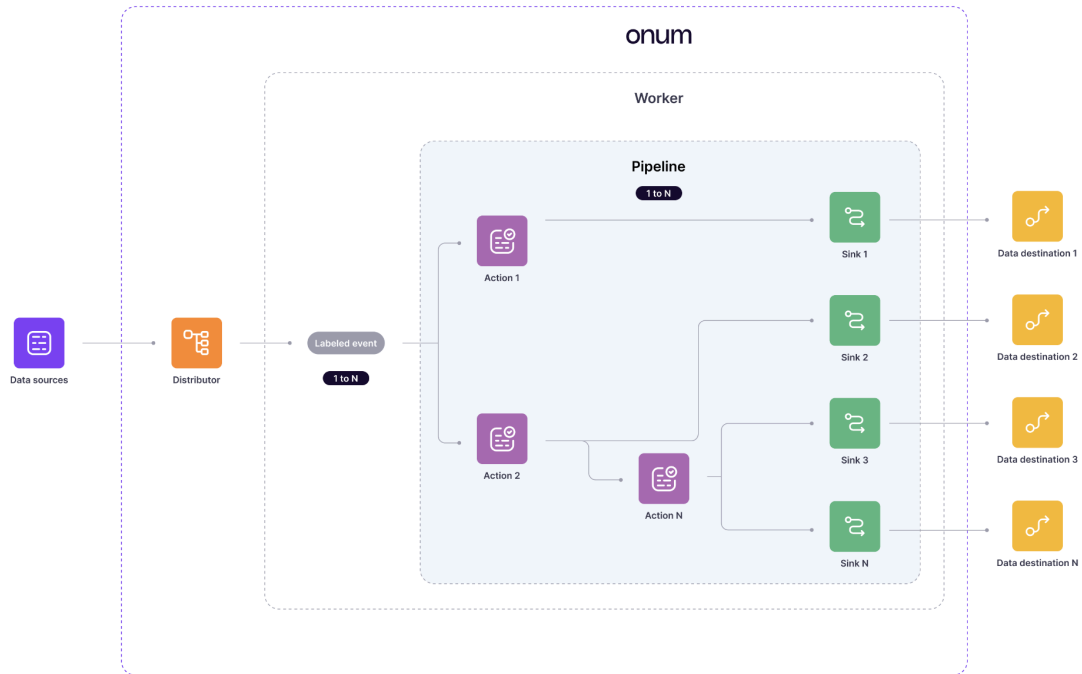
Listeners capture the raw input from data source(s) while the HA balancer ensures optimal assignment among distributor(s) that categorize and label data. Only then is the data ready to enter the processing stage, handled by the capable worker(s).



Distribution elements	
Distributor HA Balancer	This component ensures high availability and fault tolerance for the data flows and intelligently distributes incoming data across multiple distributors for optimal performance and resilience.
Listener	A system process residing in the distributor, responsible for receiving data via various protocols such as TCP, Syslog, NetFlow, or HTTP.
Load Balancer	This component evenly distributes incoming data across multiple workers within a cluster, maximizing throughput and preventing bottlenecks.
Labels	Enable data categorization based on configurable filters, providing a powerful way to manage data flows and organize them visually in the console.



Data Processing Stage



This stage orchestrates the data's transformation. A versatile pipeline ingests data from multiple sources and applies actions — singular, multiple, or even nested — before seamlessly delivering the refined output through the sink(s) to its final destinations.

Cloud-native to its core, it leverages container orchestration to streamline the data processing pipeline. Deployment, management, and scaling are automated, freeing the administrator to focus on extracting insights, rather than managing infrastructure.



Processing elements	
Actions	<p>These are the operations applied to data during processing. Actions include filtering, enrichment, transformation, formatting, aggregation, and more, enabling data to be tailored to handle specific needs.</p> <p>It is possible to nest several actions and a single action can lead to multiple sinks.</p>
Sinks	<p>Represent the output channel in the pipeline towards the final data destinations.</p>

Compliance and Audits

EU Data Residency

Onum's infrastructure runs on AWS cloud, specifically in the Ireland region (eu-west-1).

Ireland is part of the EU and adheres to the General Data Protection Regulation (GDPR), one of the strictest data privacy laws in the world.



This offers customers peace of mind regarding data residency in Europe and allows customers to comply with GDPR, allowing them to leverage the platform's real-time data observation capabilities while ensuring the protection of personal data.

Third-Party Audits

Onum has carried out, in conjunction with external collaborators, the process to audit its security controls and ensure that they are implemented adequately.

As a result, Onum has received SOC 2 Type I compliance.

Simultaneously, Onum is preparing to adapt its controls and measures to receive SOC 2 Type II compliance by the end of 2024.



Conclusion

Onum is not just another data observability platform. It is a comprehensive, real-time observability solution that allows organizations to optimize costs, gain deep insights in real time, and drive innovation through their data.

Onum enables organizations to transform their data from a challenge to a strategic asset.