Development of Contemporary **Log Management Solution** for the Information Infrastructure of the BM@N Experiment

I. Romanov, A. Chebotov, K. Gertsenberger

Laboratory of High Energy Physics

13th Collaboration Meeting of the BM@N Experiment at NICA

Joint Institute for Nuclear Research

2024

Infrastructure Overview | **Past State**

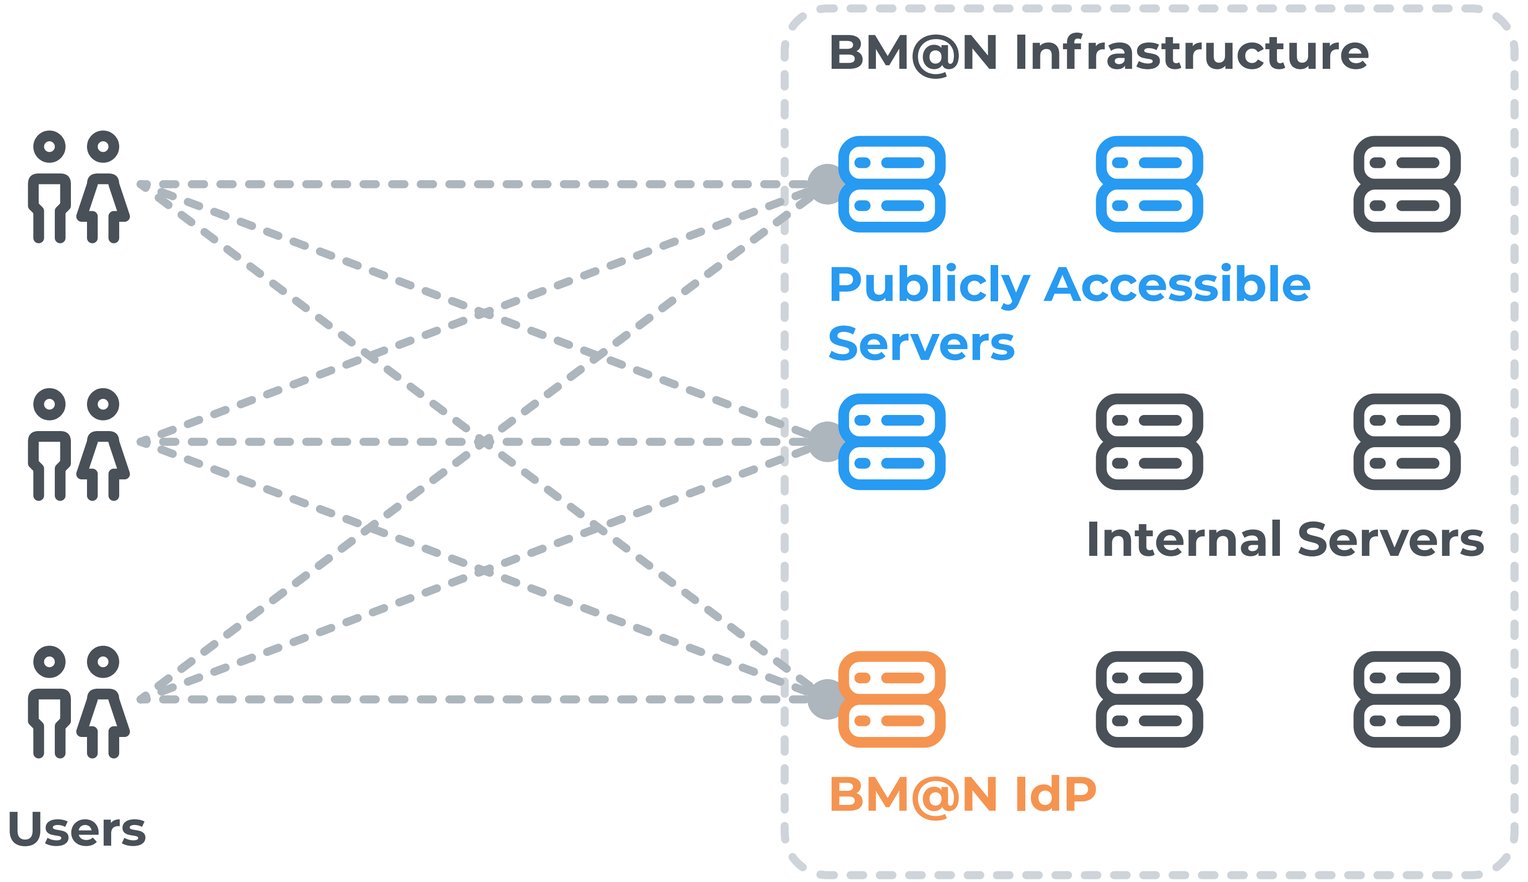

Infrastructure Overview | **Current State**

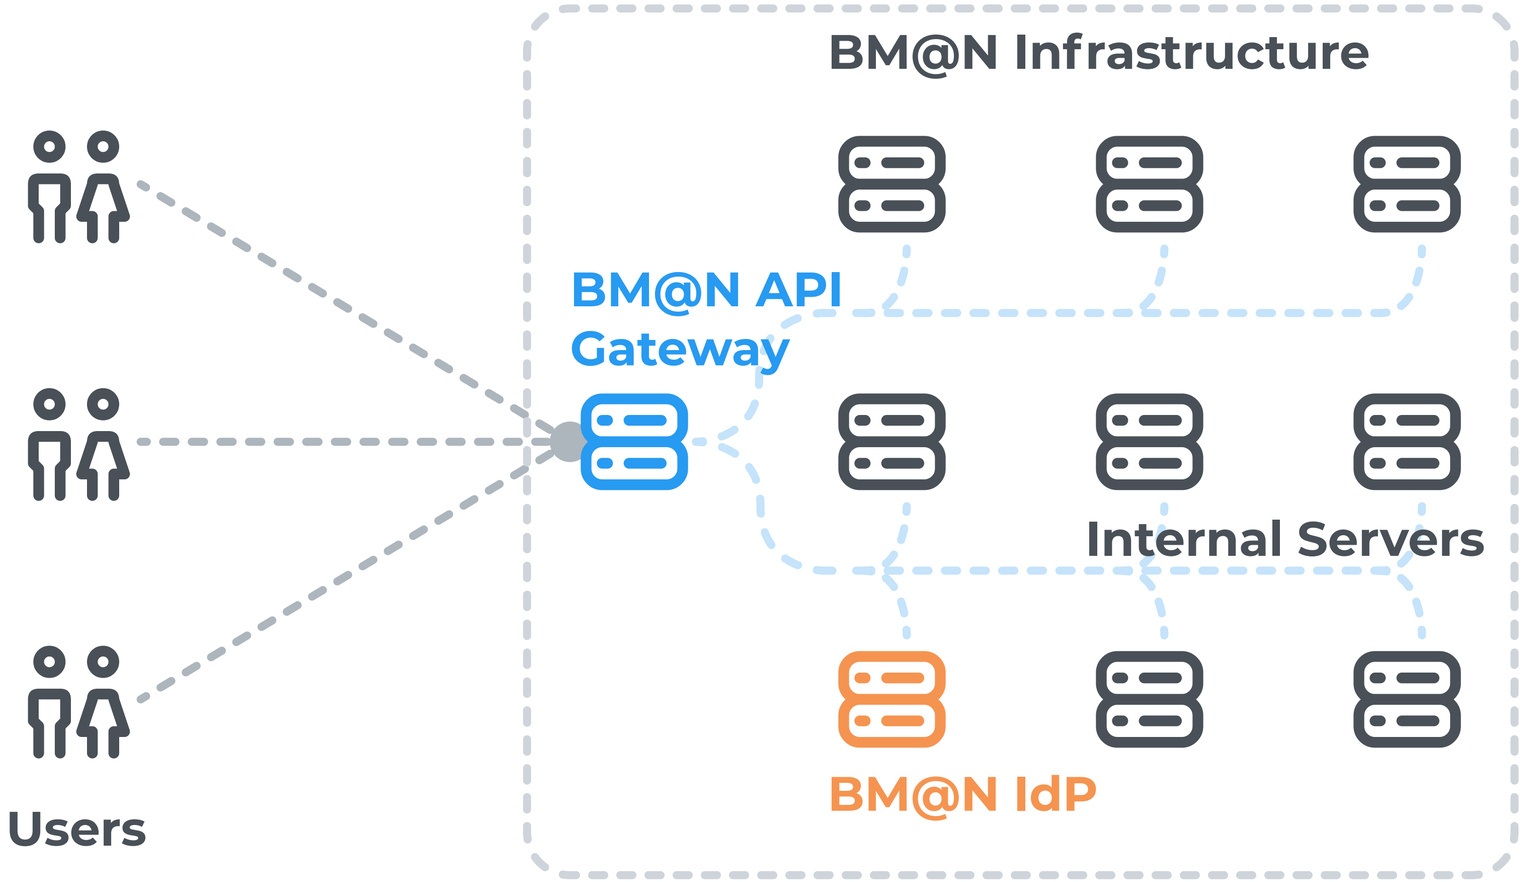

Infrastructure Overview | Log Management System

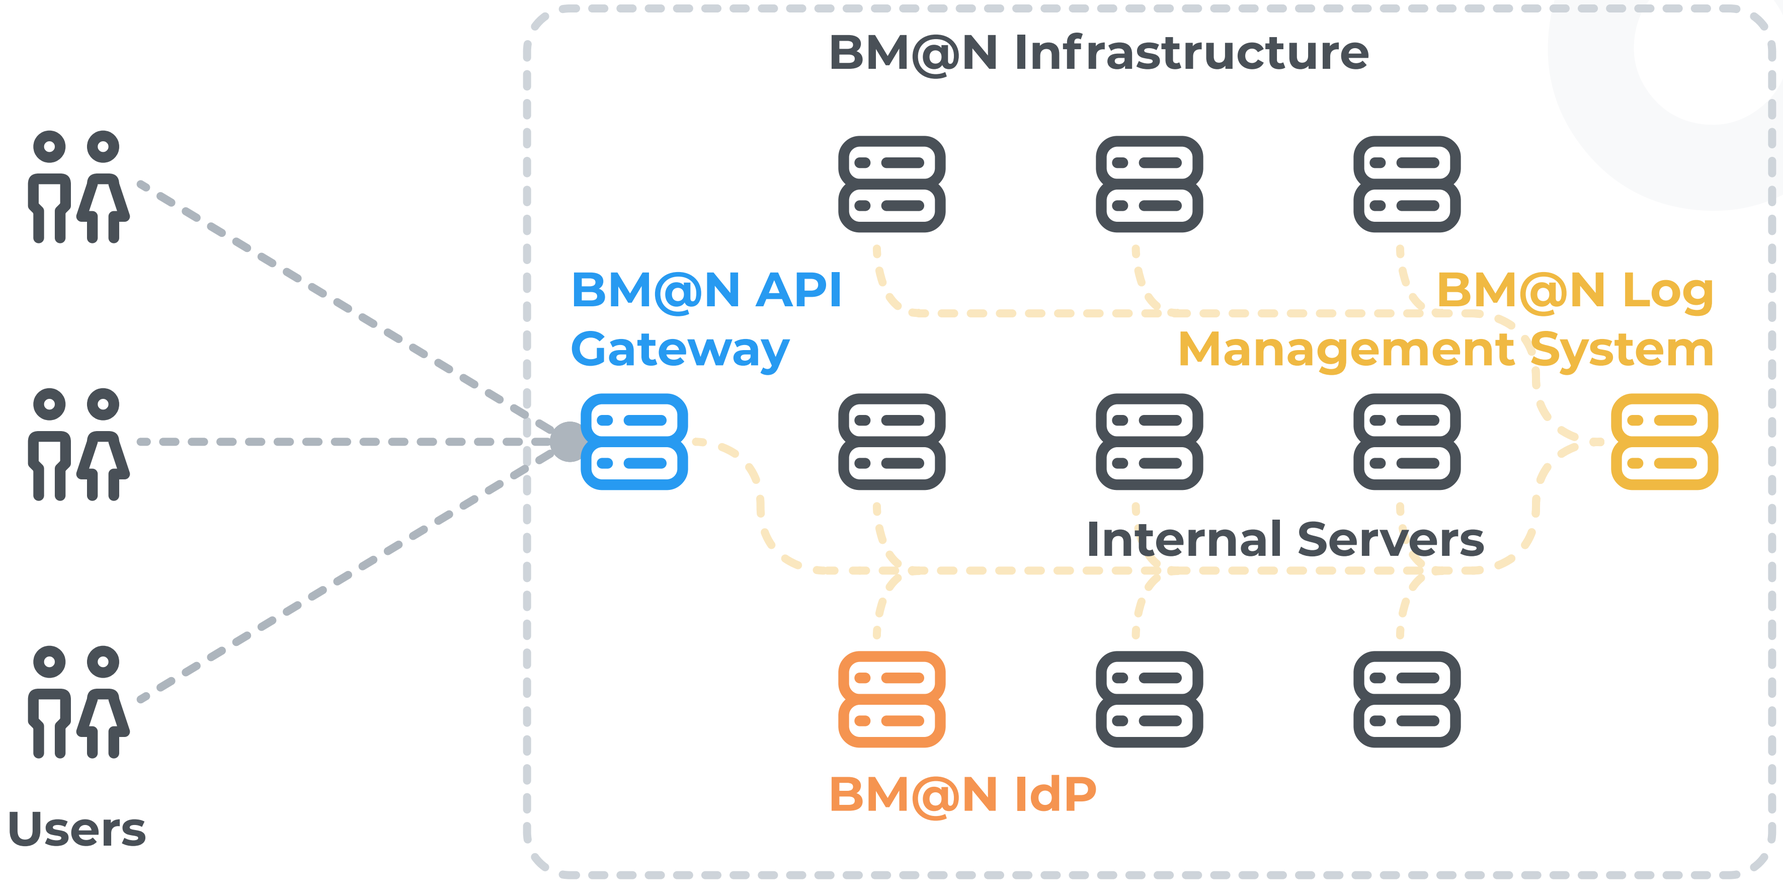

Motivation

The main issue is the large number of systems (**30+*** docker containers) distributed across different nodes.

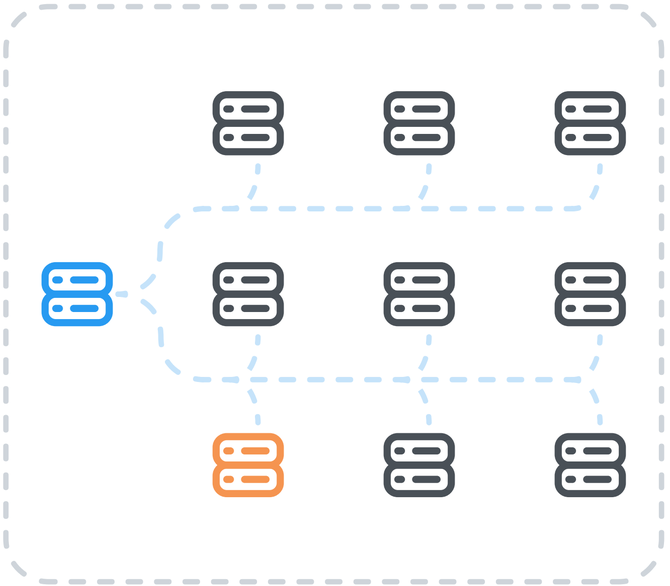

Consequences

- Incident investigation time is increased because logs must be manually searched at each node.
- Failure to detect security threats in a timely manner increases the risk of successful attacks and breaches.
- The inability to effectively monitor events and errors in real time makes it difficult to troubleshoot service issues.
- The lack of centralized log management limits process optimization and informed decision making.

*The calculation didn't consider the containers used by the log management system.

Requirements

- Open Source
- High Performance
- Cost-Effective Resource Allocation
- External Endorsement Reliability
- Scalability for Growth

Selecting a Storage System | Elasticsearch

Elasticsearch is a robust, distributed search and analytics engine designed to handle large data volumes.

Pros

- Powerful search.
- **Elastic Stack**.
- It has a large and active community of users.

Cons

- Not open source.
- High memory and disk space requirements.
- Configuration and management are complex.
- Response time may increase with large volumes of data.

Selecting a Storage System | Grafana Loki

Grafana Loki is a lightweight and rapid log aggregation and search tool constructed to operate in conjunction with **Grafana**.

Pros

- Open source.
- Easy integration with **Grafana**.
- Low resource requirements.

Cons

- Limited capabilities for complex analytical queries.
- Oriented to work in a **Kubernetes cluster**.

Selecting a Storage System | InfluxDB

InfluxDB is a high-performance time series database designed for the storage and analysis of real-time data.

Pros

- Open source.
- It is well-suited for the storage of temporary data and metrics.
- It offers high performance for writing and reading data.

Cons

- It places a greater emphasis on metrics than on logs.
- It has limited capabilities for complex data analysis.
- It may require a complex setup for large data volumes.

Selected Tech Stack | **ClickHouse**

ClickHouse is a column-oriented DBMS specifically designed for high-performance analytics and real-time work with large volumes of data.

Pros

- Open source.
- Efficient use of disk space and memory.
- It is optimized for processing large amounts of data and OLAP queries.

Cons

- Setting it up can be tricky.
- A scalable solution requires additional tools to manage the cluster.

ClickHouse was selected as the log storage system, as it most effectively fulfilled the specified requirements.

Selected Tech Stack | Vector

Vector is a high-performance tool designed for building observability pipelines consisting of collecting, transforming, and routing logs and metrics.

Pros

- Open source.
- It provides real-time data processing.
- It supports a wide range of data sources and sinks.
- It is easily scalable.

Cons

- Initial setup and configuration can be challenging for new users.
- No stable release at the moment.
- The tool has a modest community but its popularity is rising.

Deployment Topologies | **Distributed**

Pros

- This is the simplest solution.
- This is easily scalable.

Cons

- A high volume of requests may impact downstream services (in our case, **ClickHouse**).
- This may impact the performance of other applications on the same host.
- There is no guarantee of reliable data delivery.

Vector as a distribured agent

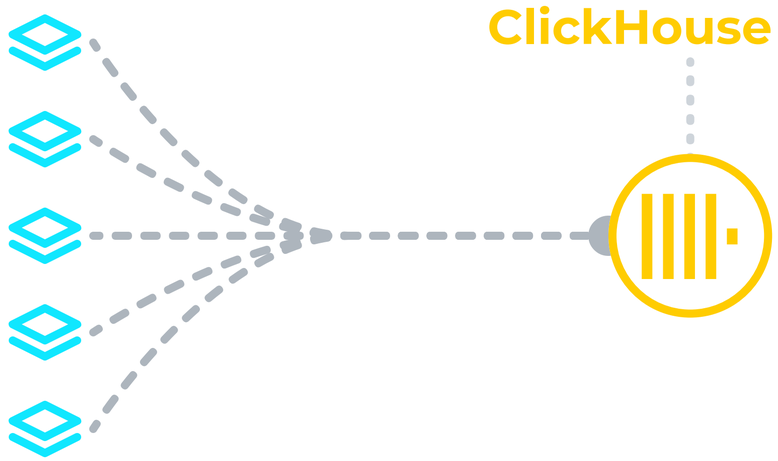

https://vector.dev/docs/setup/deployment/topologies/

Deployment Topologies | **Centralized**

Pros

- This provides a more reliable method of data delivery.
- This has low impact on other systems on the same host.
- This reduces the load on **ClickHouse** by sending in batches.

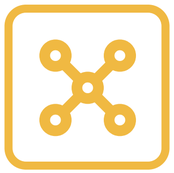

Cons

- This is a more complex approach.
- There is still no complete guarantee of reliable data delivery.

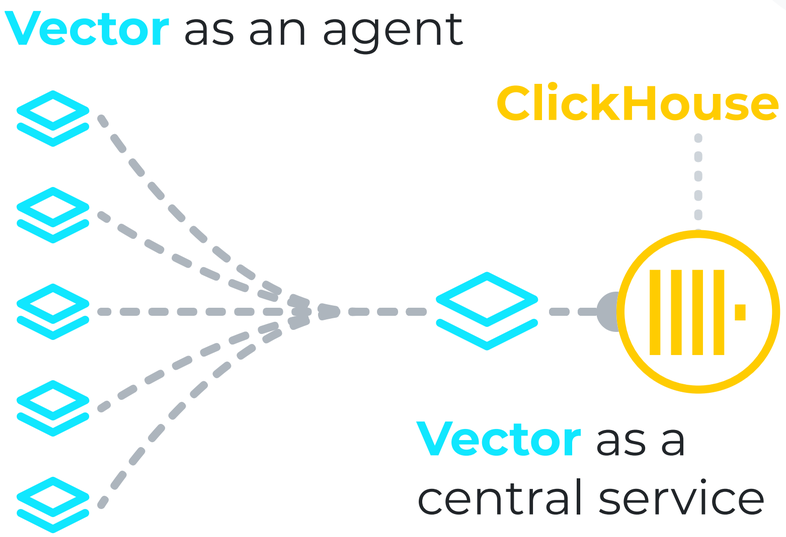

https://vector.dev/docs/setup/deployment/topologies/

Deployment Topologies | **Stream-based**

Pros

- This ensures the most reliable way to deliver data.
- This also has no significant impact on other systems on the same host.

Cons

- This is the most complex and expensive approach.

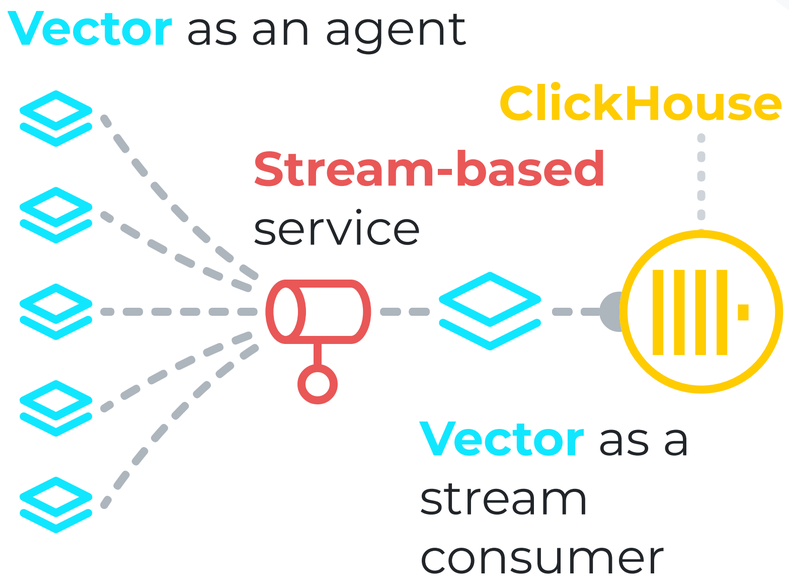

https://vector.dev/docs/setup/deployment/topologies/

Selected Tech Stack | **Metabase**

Metabase is a business intelligence platform for data analysis and visualization, enabling teams to gain insights and make data-driven decisions.

Pros

- Open source.
- Advanced analytics enables detailed analysis of log data beyond simple visualization.
- Intuitive interface makes complex log queries accessible to users without technical backgrounds.

Cons

- Performance may be negatively impacted by large log data sets or complex analytical queries.
- It is not suitable for real-time monitoring.

Dashboard on User Access Statistics

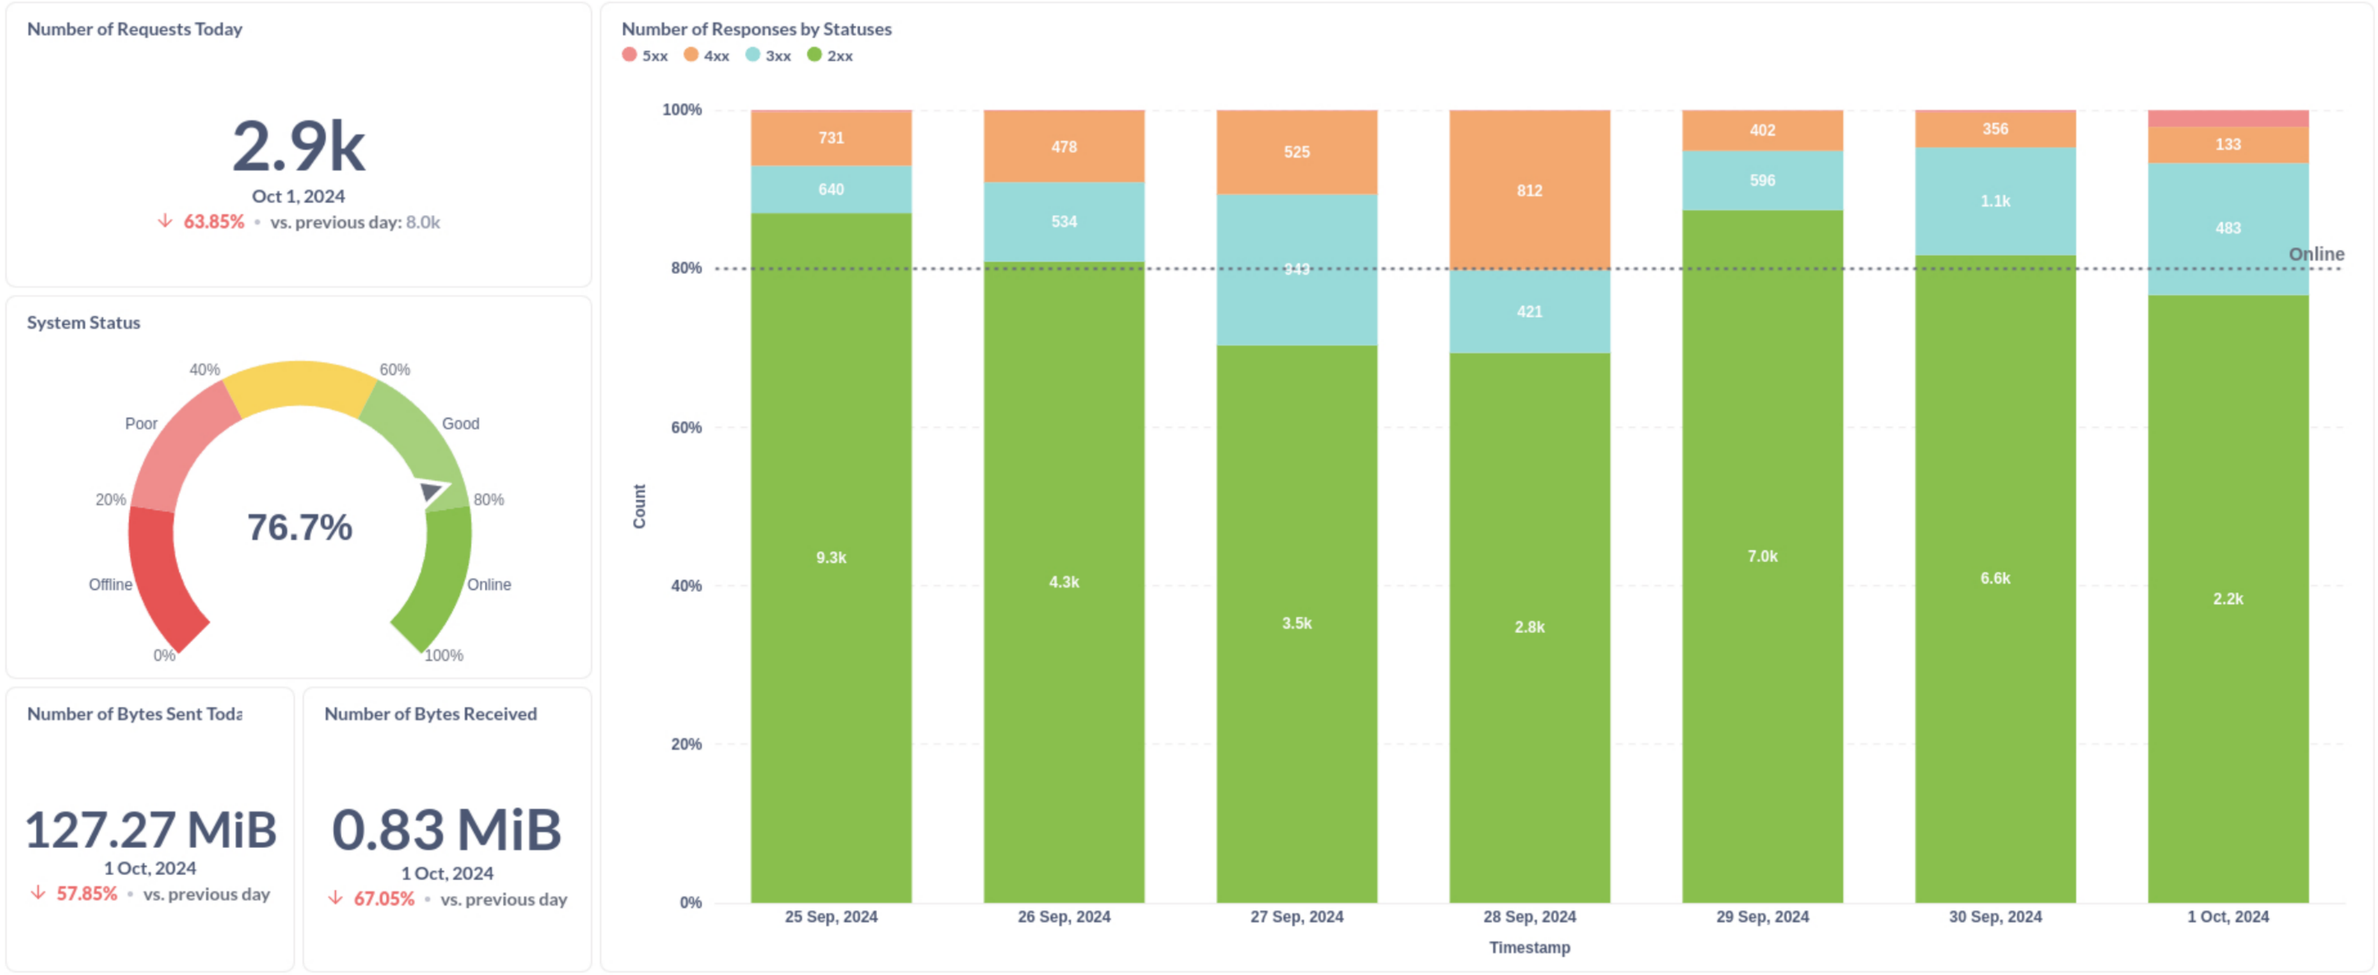

Geo Heatmap of Users Accessing BM@N Services

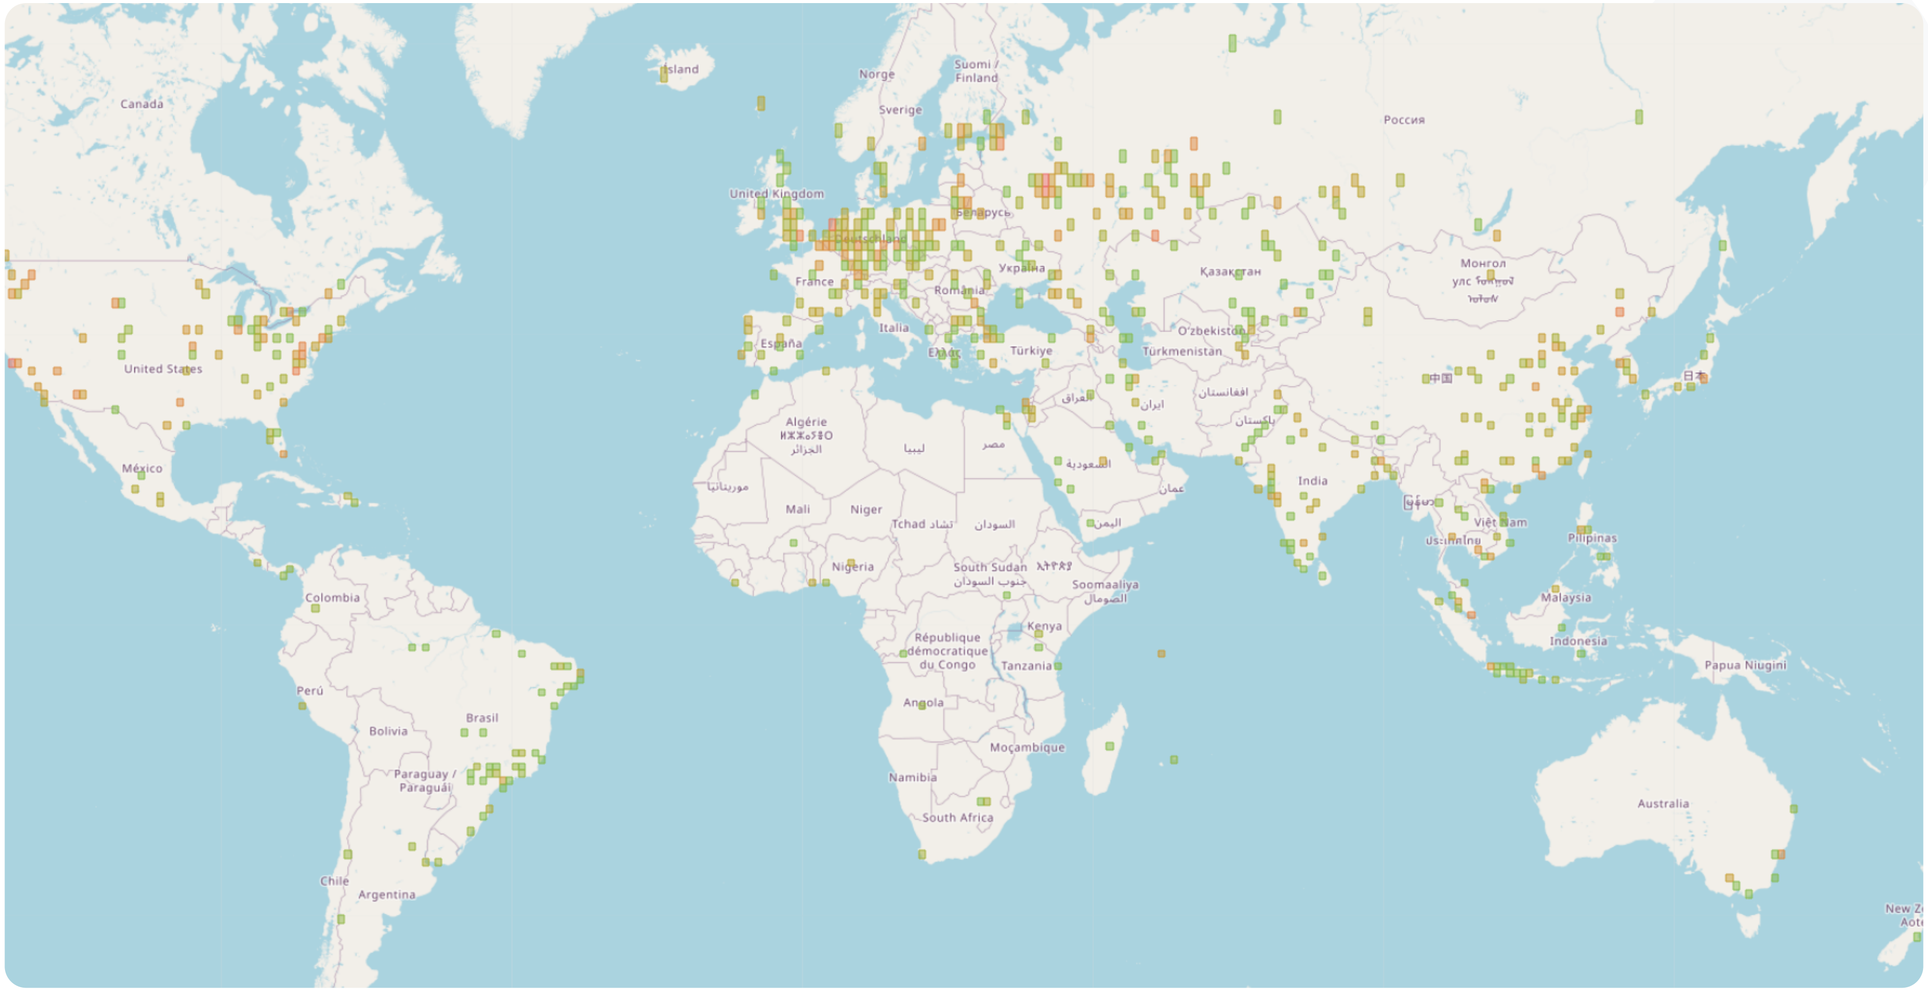

Implementation | Integration with API Gateway

The utilization of the **BM@N API Gateway** as a single entry point enables the collection of access logs for all systems in a single location.

Benefits

- Upon deployment of a new system, access log collection works out of the box.
- This enables the standardization of the log format across disparate systems.
- The consolidation of the log collection at the gateway has the effect of minimizing the impact on the performance of individual services.

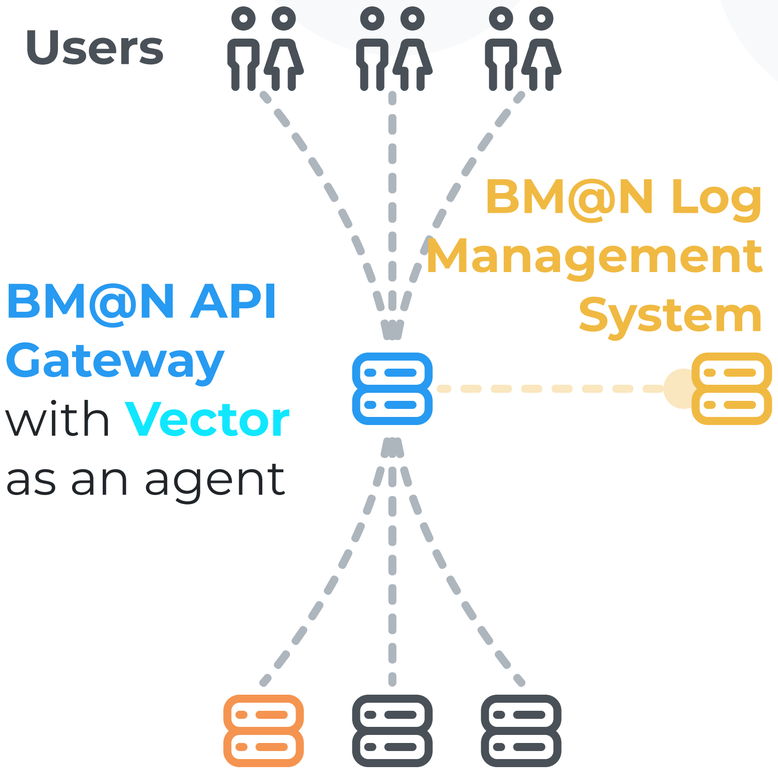

Implementation | **Log Collection from auth.log**

The **Server Monitoring Agent** is developed to monitor the operation of BM@N infrastructure servers. Currently, it collects data only from the *auth.log* file.

Benefits

- It permits the monitoring of authentication occurrences and the tracking of user activities.
- It provides the ability to notify administrators of any suspicious activity or potential security breaches.

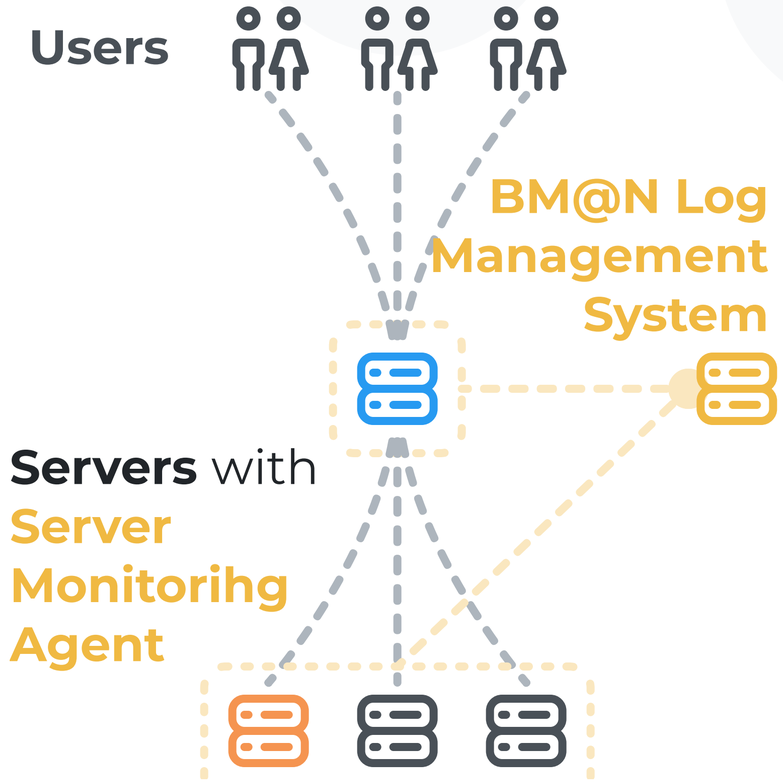

Implementation | Integration with Keycloak

The utilization of **Keycloak** as a Single Sign-On (SSO) **Identity Provider (IdP)** enables the collection of diverce authentication events.

Benefits

- This enables more efficient auditing of user access and facilitates the identification of potential security issues.
- The integration of **Identity Provider** events with **API Gateway** access logs in a centralized repository to more effectively respond to and remediate security threats.

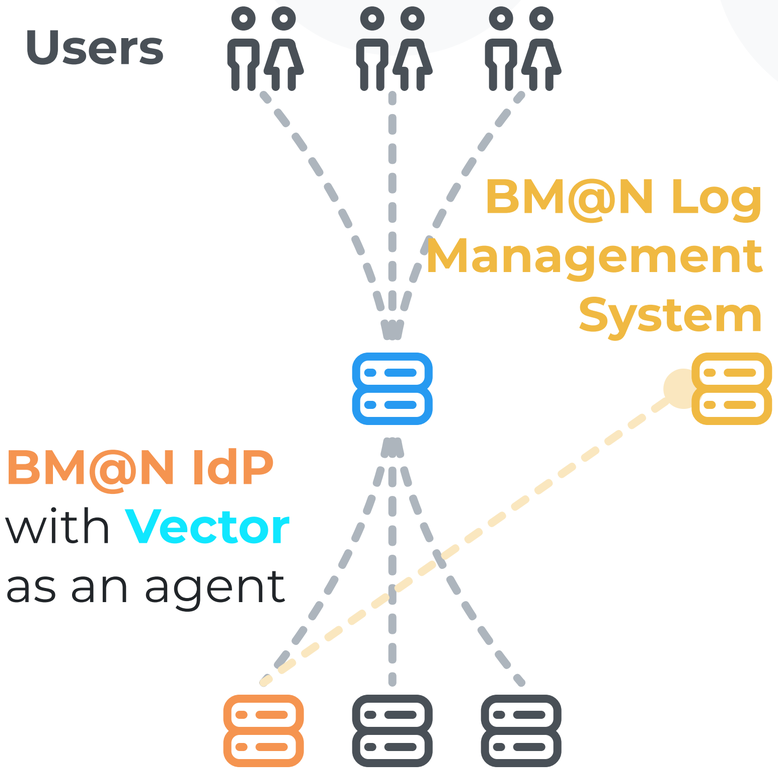

Final Infrastructure Design

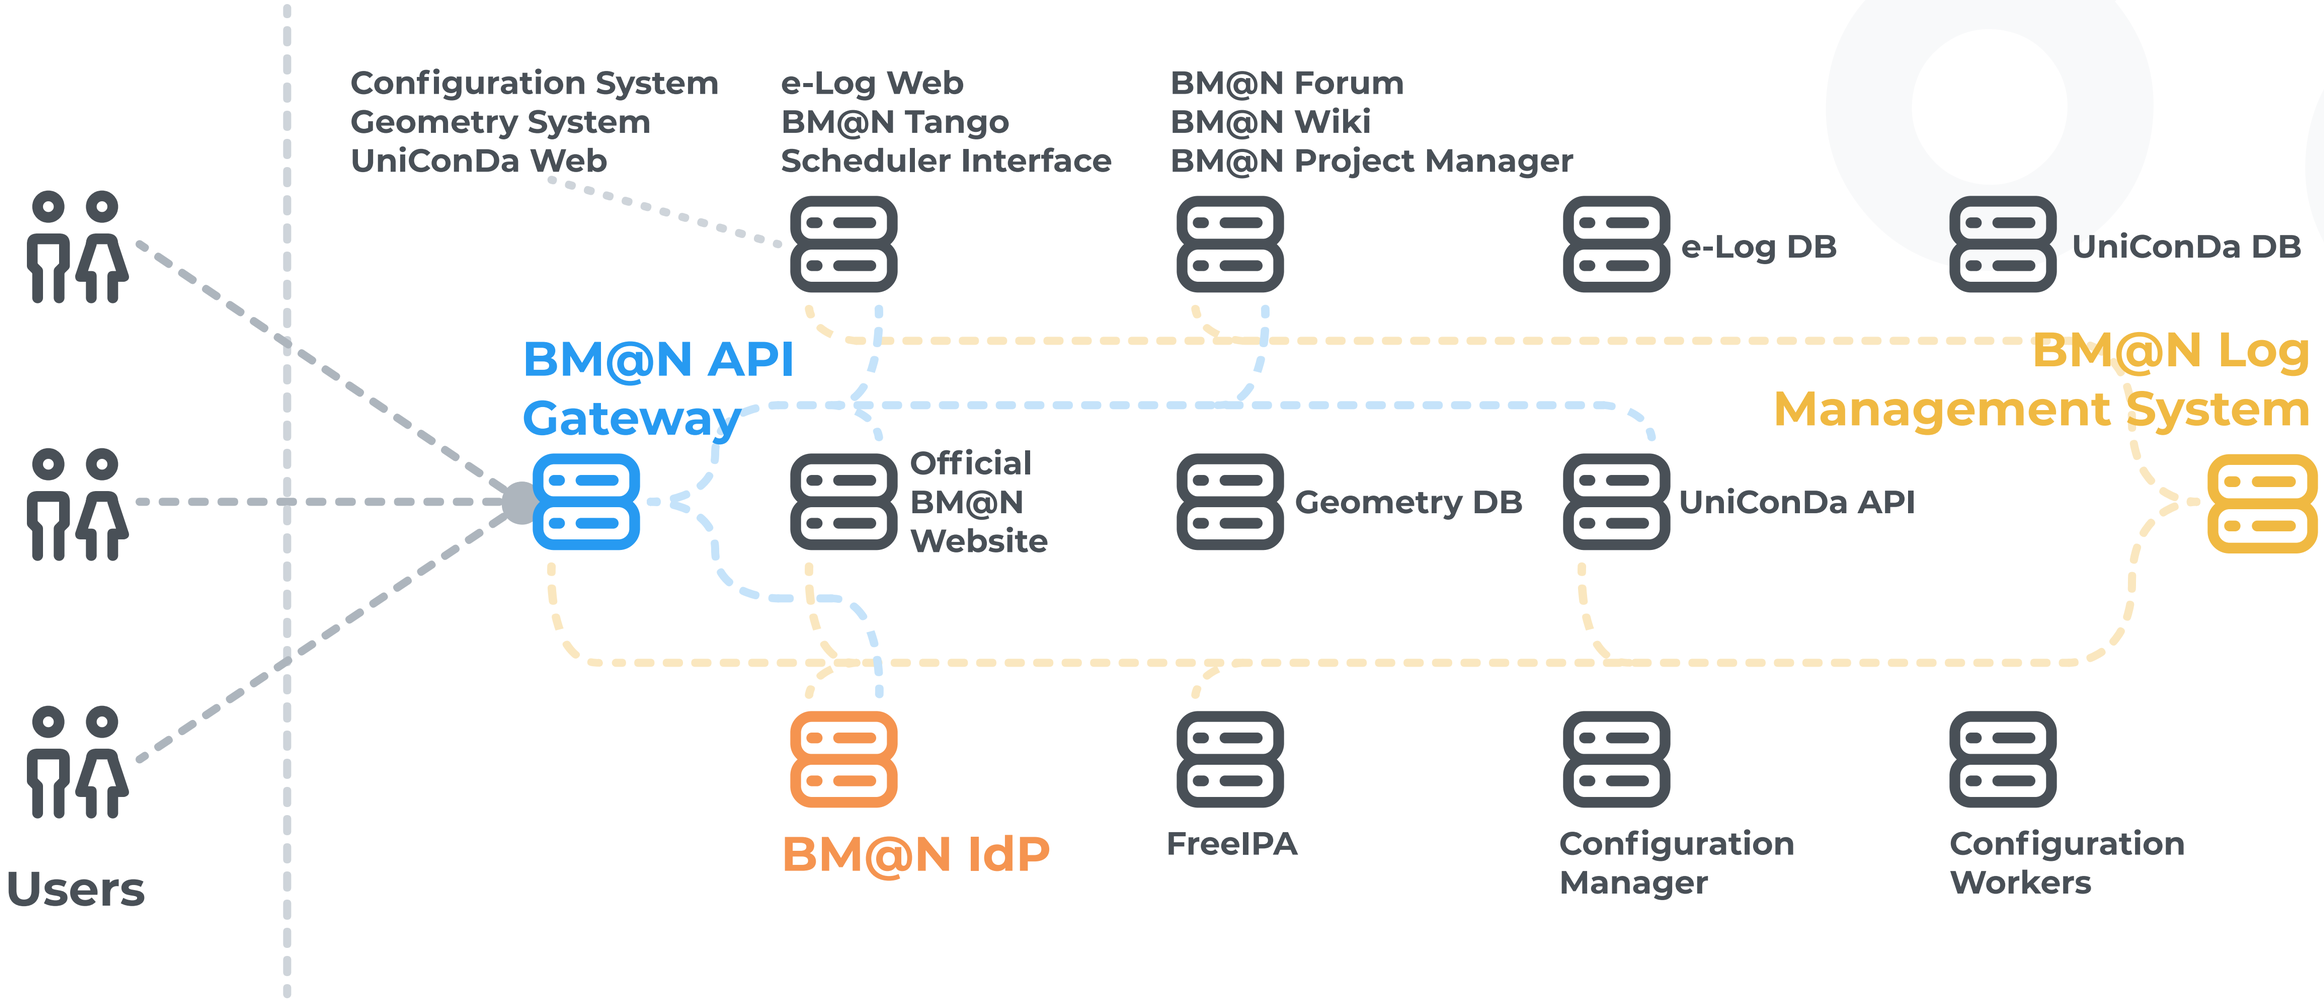

Conclusions

- The **Log Management System** helps to quickly identify and fix security issues, making the BM@N infrastructure more resilient.
- The system's architectural design is intended to facilitate scalability, enabling the seamless integration of supplementary data sources and analytics tools upon the introduction of new systems into the experiment's ecosystem.
- Implementing analytics tools like **Metabase** provides deep insight into application performance and user behavior, enabling more informed data-driven decisions.

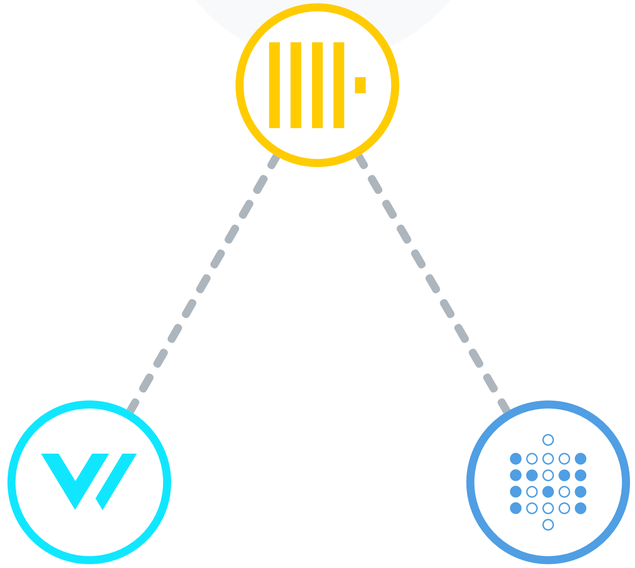

Future Plans

- Deploy the **Server Monitoring Agent** to all servers in the infrastructure and extend its capabilities.
- Full integration with all existing systems of the BM@N experiment (Official BM@N Website, BM@N Forum, BM@N Wiki, Configuration System, Geometry System, UniConDa, e-Log Platform).
- Add support for real-time monitoring tools such as **Grafana** and related data sources as a full implementation of the system's capabilities.
- Establish a robust notification system that uses multiple communication channels to effectively disseminate information.

Thank You for Your Attention!

13th Collaboration Meeting of the BM@N Experiment at NICA

Joint Institute for Nuclear Research

2024