DEVELOPMENT OF INFRASTRUCTURE FOR BM@N SOFTWARE COMPLEX

A. Chebotov[a,1], *K. Gertsenberger*[a], *I. Romanov*[a]

[a] Joint Institute for Nuclear Research, Dubna, Moscow region, 141980, Russia,

The BM@N experiment, as part of the NICA complex, produces a substantial quantity of physics data, necessitating the implementation of a sophisticated infrastructure for the software systems providing efficient storage, processing, and management of the data. In order to address these challenges, a comprehensive set of information systems has been developed. The complex includes an information system representing the Unified Condition Database (UniConDa) that stores and provides necessary parameters of experiment systems; the Geometry Database for operating with information on geometric models of detectors; the COnfiguration Online Platform (COOP) for managing hardware settings and a sequence of software tasks to be used in online; the Event Metadata System (EMS) for indexing and searching physics events for a particular analysis; and the Electronic Logbook (e-Log platform) to record and share information on experiment runs during the sessions. In addition, the BM@N uses various collaboration services, which have been already deployed, such as the official website, collaboration forum, document server (Wiki). Security policy is ensured through the Keycloak, authentication and authorization system, which centralizes access control to BM@N software systems. The paper covers the description of the deployed infrastructure on a cluster platform managed by the Proxmox system, which oversees virtualization and containerization of the components. The integrated single reverse proxy ensures centralized secure access to all the software systems of the experiment. As a result, the developed infrastructure with the information systems and software services ensures the management of information being necessary for physics analysis of experiment data within the BM@N experiment.

PACS: 07.05.Bx; 07.05.Kf

Introduction

Modern experiments in high-energy physics necessitate the implementation of a sophisticated infrastructure for the efficient data processing and analysis. The BM@N (Baryonic Matter at Nuclotron) experiment, which is part of the NICA (Nuclotron-based Ion Collider fAcility) project [1] at the Joint Institute for Nuclear Research, is conducted to investigate the interaction of heavy ion beams with a fixed target. This research offers a distinctive opportunity to gain insights into fundamental questions related to the state of matter at extreme densities and temperatures, as well as to study nuclear properties.

In order to achieve the objectives, a comprehensive software infrastructure has been established. The primary challenges associated with maintaining a set of BM@N information systems are ensuring the reliable operation of their

[1] E-mail: chebotov@jinr.ru

interconnected components and implementing unified security policy, including centralized access, authentication and authorization. This necessitates utilization of contemporary technologies.

In response to these requirements, the developed infrastructure is to provide a reliable, secure platform for a comprehensive ecosystem of the following information systems [2]. The Electronic Logbook provides collaboration members with interfaces for recording and sharing information on current parameters and operation modes of experiment subsystems, current events, encountered problems and taken actions during BM@N runs. The information system on the Condition Database is designed to ensure storage, unified access, search and management of parametric information on the experiment systems to be further used for processing of experimental and simulation data. The Configuration Online Platform is used to store and provide data on the configuration of hardware systems and a sequence of software tasks of the experiment to be run during online data acquisition. The Geometry Information System is intended to store and manage data on geometric models of detectors, and to provide a centralized repository for detector geometries, which are used to process and analyze simulated and experimental data. The Event Metadata System is based on the Event Catalogue, which contains summary information on obtained particle collision events and allows for search and selection, using the metadata, of a set of only those events that are needed for a particular physics analysis.

In addition, the software ecosystem of the BM@N experiment also contains various collaboration services, such as the official website of the experiment, discussion forum, document management server (Wiki), and project management system to guarantee effective communication and collaboration. The next sections present the description of the developed infrastructure for the BM@N software systems and services, with an emphasis on architectural solutions that facilitate the deployment and maintenance of the systems, which are of great importance to obtain timely, high-quality physics results in the BM@N experiment.

DEPLOYMENT OF THE BM@N SOFTWARE INFRASTRUCTURE

The advancement of the BM@N software ecosystem for efficient processing of BM@N data has necessitated the implementation of modern methodologies that ensure high availability and security of all the systems. The developed infrastructure, as illustrated in Figure 1, is constructed on a multilevel basis, with each component performing a precisely defined function within the larger system. Previously, the infrastructure was distributed across multiple clusters, which introduced complications in system administration and monitoring. The majority of services provided direct network access to users, which introduced additional complexity to the configuration and security enforcement. A lack of containerization technologies on the local servers made it challenging to scale the projects and potentially led to conflicts in the system environment.

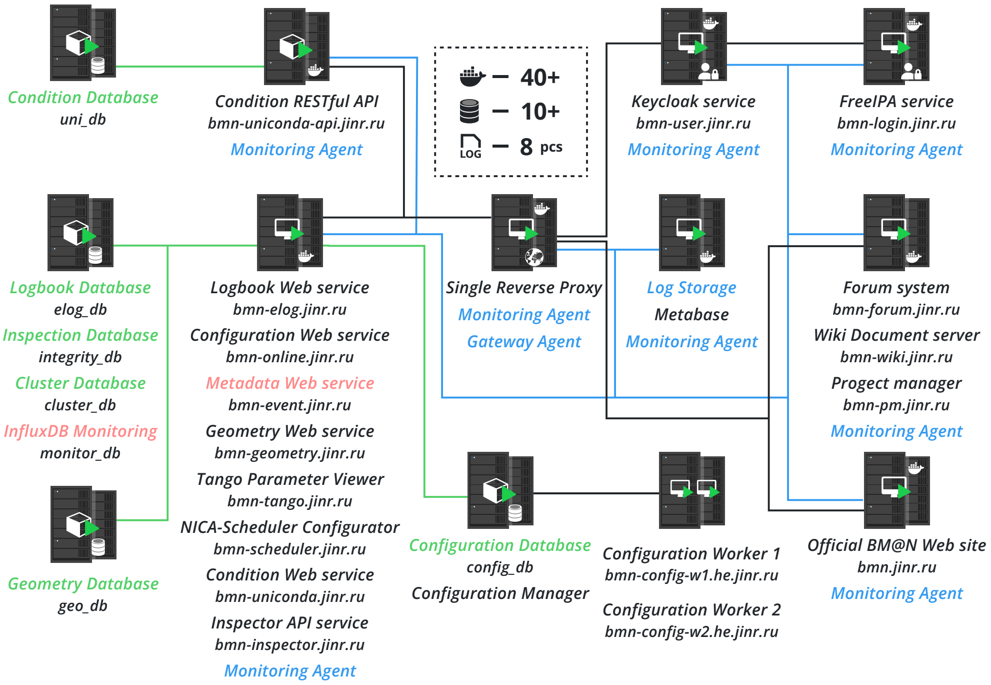

Fig. 1. The BM@N Software Systems and Services Infrastructure

A notable achievement has been the migration of the BM@N software systems to the C4 DAQ cluster, which is managed using the Proxmox [3] solution. This provides a stable virtualization environment with a high density of the software systems and services. The Proxmox platform facilitates the efficient administration of allocated resources, supports the live migration of virtual machines between cluster nodes, and provides integrated backup and recovery mechanisms. The solution consolidated the software components of the BM@N experiment into a single managed environment, significantly increasing the operational reliability of the entire ecosystem, especially during periods of intensive data collection.

The use of containerization via Docker has become a significant factor in increasing flexibility and simplifying deployment and maintenance of the systems. Docker enables the isolation of the software applications and ensures their stable operation within independent containers, reducing the likelihood of conflicts during updates and migration between different environments. Furthermore, containers enhance scalability of systems. This approach ensures stable operation and allows flexible scaling of computing resources in accordance with the current tasks of the BM@N experiment.

In addition, it is worth noting the Common Deployment System (CoDeS) [4] for the automation of the configurable deployment of the BM@N information systems on the implemented infrastructure. The deployment process is automated through the implementation of a set of templates and scripts that define the configuration, dependencies, and required resources associated with the information system. The solution automatically deploys the corresponding containers, configures network connections, installs necessary packages, and applies experiment-specific settings specified in the single configuration file. This not only unifies the deployment process but also significantly reduces the time and effort spent on it and the likelihood of errors during the installation process. The result of the configurable deployment of the information systems also includes a regular automatic update from the repository and backup of the stored data in case of software or hardware failures.

ENSURING ACCESS SECURITY AND RELIABILITY

To guarantee the security of BM@N information systems and services, a comprehensive multi-layered protection strategy has been put in place. At the application level, a central element of the infrastructure is the NGINX-based single reverse proxy [5], which provides a single-entry point and control of all incoming traffic to the systems. The solution protects against unauthorized access, provides load balancing capabilities for traffic distribution, and guarantees security monitoring using centralized logging.

At the network level, the protection is implemented using iptables and Fail2Ban services, which prevent network attacks, unauthorized access attempts, and automatically block both incoming and outgoing malicious activity. These tools provide comprehensive protection against a wide range of threats, including brute-force attacks, port scanning, and vulnerability exploitation attempts.

The centralized user identity management is facilitated through the Keycloak [6] solution, which implements contemporary security mechanisms and provides Single Sign-On (SSO) capabilities across the BM@N software systems, incorporating support for OpenID Connect and OAuth 2.0 protocols. Keycloak centrally manages the lifecycle of access tokens and their scope, thereby reducing the complexity of security implementation at the individual service level. The BM@N systems need only verify token validity and user role matching the requested action, while all critical security logic is performed by Keycloak.

An important part of the developed infrastructure is an implemented monitoring and logging service, which provides comprehensive tracking of all user activities, requests to services, as well as the operation of servers and services, storing all logs in a central database. As a result, the comprehensive approach facilitates both fast detection and response to suspicious activities, and facilitates effective tracking of system performance and operational status.

CONCLUSIONS

As a result of the work performed, a comprehensive infrastructure for the BM@N software has been developed, ensuring efficient data management and processing. The successful migration to the DAQ C4 cluster with Proxmox has been completed, providing a stable virtualized environment, and the introduction of Docker-based containerization, which significantly improves the flexibility of the deployment of the systems and their maintenance. The implemented Common Deployment System automates the process of custom deployment of the information systems. In terms of security, multi-level protection has been implemented using the centralized single reverse proxy based on NGINX, iptables and Fail2Ban solutions. Furthermore. the Keycloak system has been integrated for unified authentication and authorization. The created infrastructure guarantees reliable operation of all the system components, providing the necessary level of security and scalability for efficient processing of physics data of the BM@N experiment.

Funding

This work was supported by ongoing institutional funding. No additional grants to carry out or direct this particular research were obtained.

Conflict of interest

The authors of this work declare that they have no conflicts of interest.

REFERENCES

1. *NICA Collaboration.* NICA White paper. Searching for a QCD mixed phase at the Nuclotron-based ion collider facility. — 2014.

2. *Gertsenberger K., Alexandrov I., Filozova I., Alexandrov E., Moshkin A., Chebotov A., Mineev M., Pryahina D., Shestakova G., Yakovlev A., Nozik A., Klimai P.* Development of Information Systems for Online and Offline Data Processing in the NICA Experiments // Phys. Part. Nucl. — 2021. — V. 52. — P. 801–807.

3. *Simon M., Huraj L.* VirtualBox and Proxmox VE in Network Management: A User-Centered Comparison for University Environments.

4. *Chebotov A., Gertsenberger K., Moshkin A., Slepov I.* Common Deployment Complex for the Information Systems of the BM@N Experiment // Phys. Part. Nucl. Lett. — 2023. — V. 20. — P. 1269–1271.

5. *Bukhari A.S., Iqbal M.* Public IP Efficiency and Data Center Security Enhancement with Reverse Proxy Implementation // Journal of Systems Engineering and Information Technology (JOSEIT). — 2024. — V. 3, no. 2. — P. 37–42.

6. *Chatterjee A., Prinz A.* Applying Spring Security Framework with KeyCloak-Based OAuth2 to Protect Microservice Architecture APIs: A Case Study // Sensors. — 2022. — V. 22, no. 5. — P. 1703.