The XXVIII International Scientific Conference of Young Scientists and Specialists (AYSS-2024)

Status of the Event Metadata System for the BM@N experiment

O. Nemova, P. Klimai, K. Gertsenberger, I. Dunaev

October, 2024

Plan

- BM@N (NICA)
- EMS for NICA
  - Architecture
  - Database as the Event Catalogue
  - REST API and Web interface
  - Monitoring
- EMS for NICA. Current status (an update)
  - Database indexing
  - High availability
  - Collection and display of statistics
  - Automated deployment
  - Integration with KeyCloak-based authorization system
- EMS for NICA. Planning

NICA – **N**uclotron based **I**on **C**ollider f**A**cility

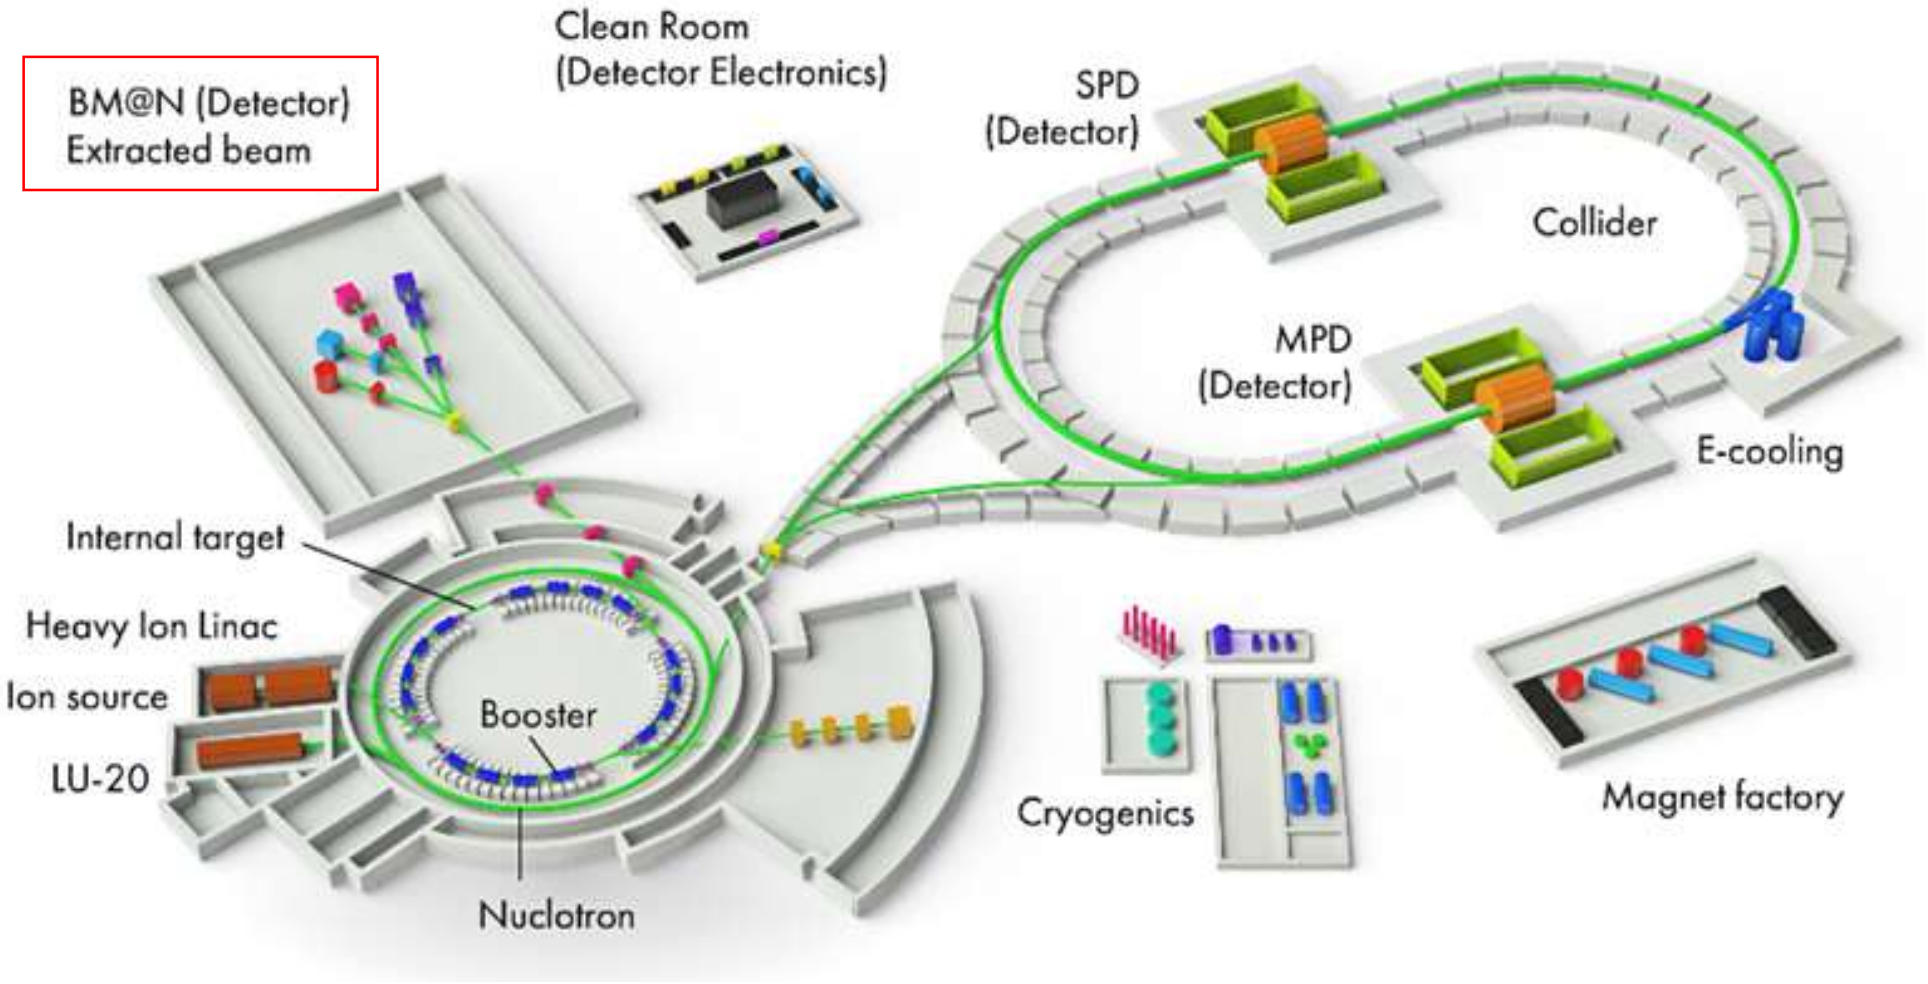

EMS (**E**vent **M**etadata **S**ystem) for NICA

- Indexing of all reconstructed events stored in DST ROOT files
- Storing necessary event metadata, such as:
  - Number of primary and all reconstructed tracks
  - Track number of +/- charged particles
  - Primary and secondary particles found
  - Number of hits by detectors
  - Software version
  - and etc.
- Flexibly tune per experiment (BM@N, BM@N SRC, MPD, SPD)
- Convenient access to the metadata (several APIs)
- Search for a required set of events
- Provide statistics and check the quality of the catalogue of physics events

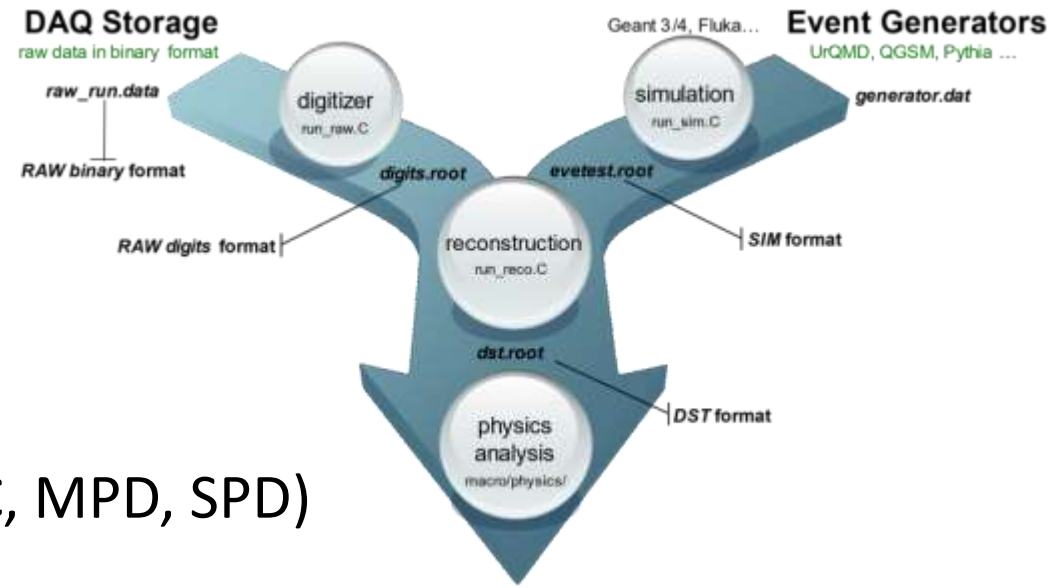

EMS for NICA. Architecture

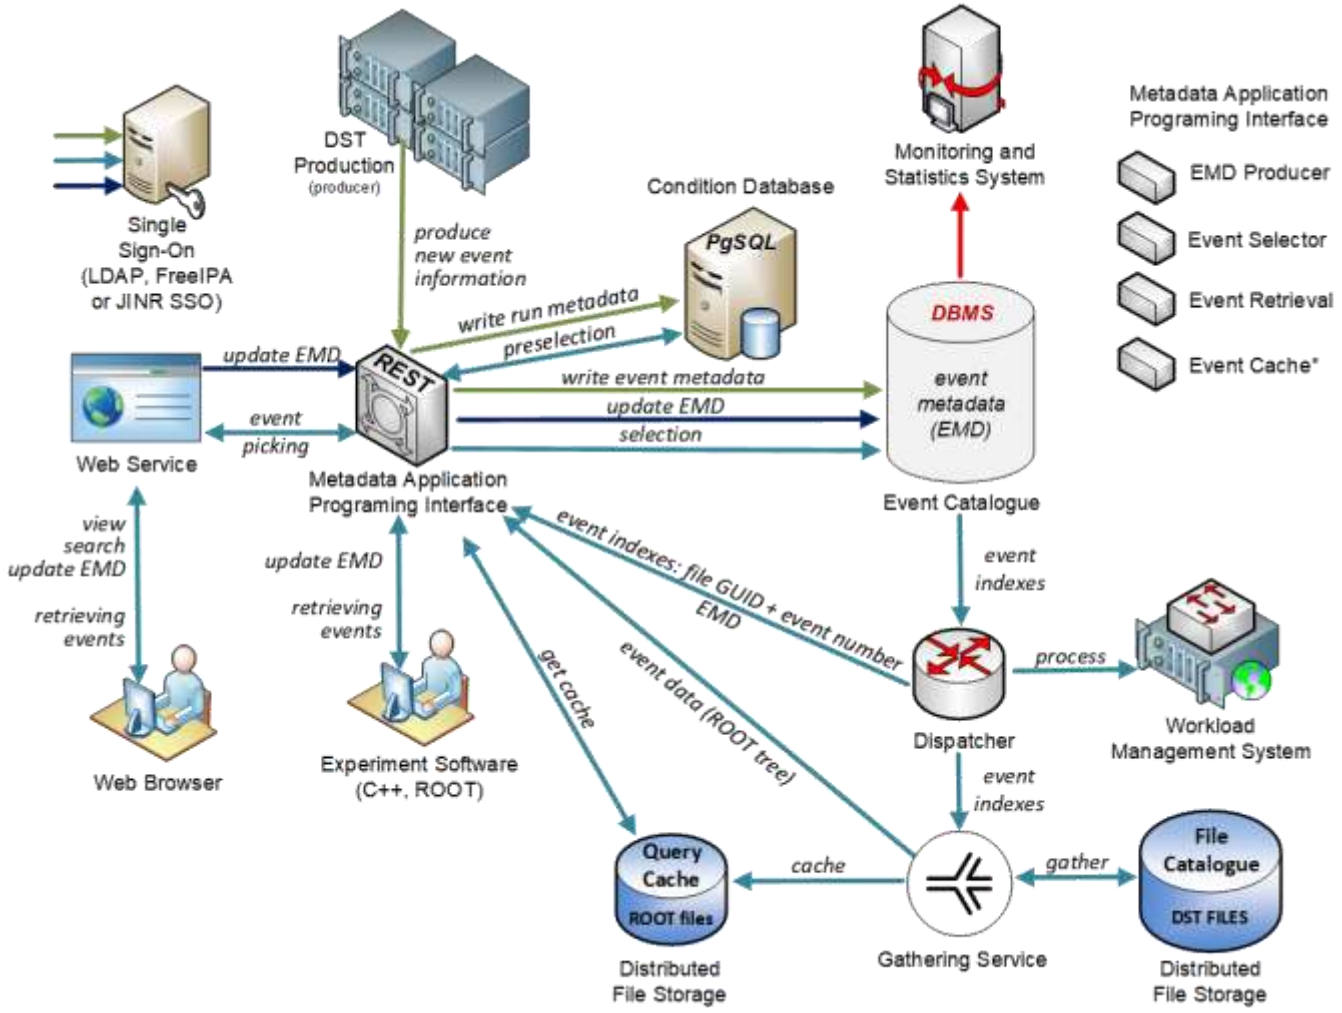

EMS for NICA. REST API and Web Interface

- Using Kotlin programming language
  - JVM runtime
  - ktor framework for Web UI and API
  - JDBC for database connectivity
  - Jackson for (de-)serialization
  - Kotlinx.html
- Packed in Docker
- Configuration file provided in YAML
- Auto provisioning

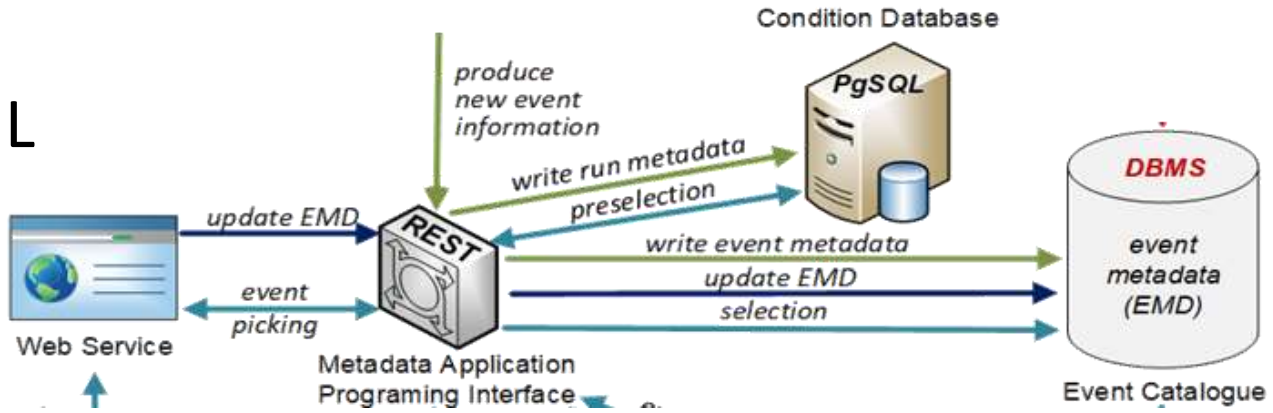

EMS for NICA. Database as the Event Catalogue

- DBMS choice – PostgreSQL (the best result in NFT (**N**on-**F**unctional **T**esting))

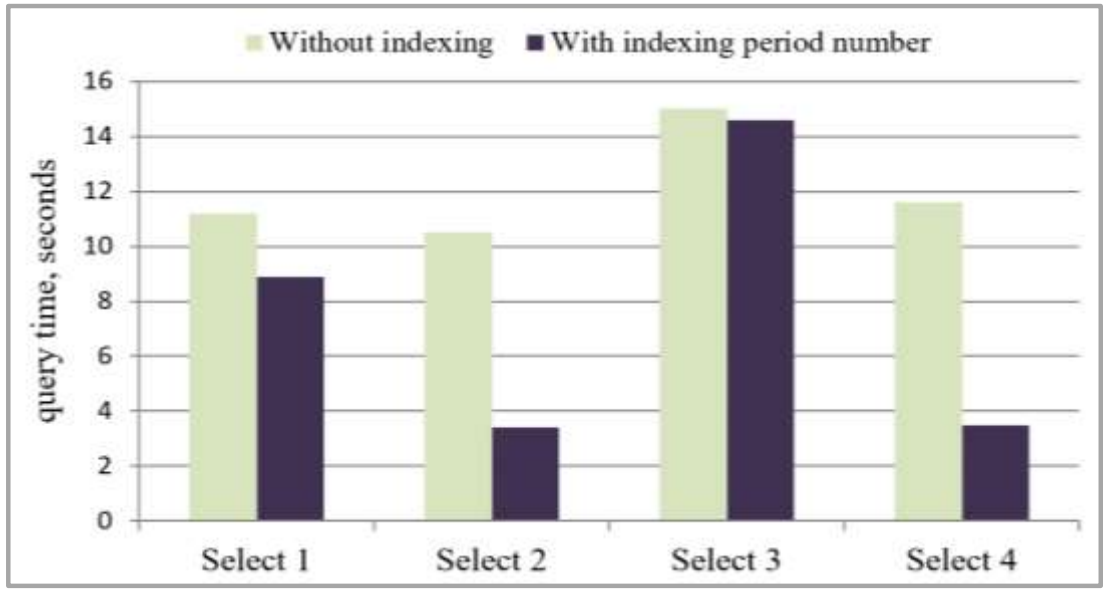

HW Config 2:
- Intel Core i9-10900F 2.80GHz (20 CPU cores),
- 64 GB RAM,
- 1TB NVMe SSD disk,
- CentOS Linux 8.2,
- PostgreSQL 12.5

<- Example test result for 500M events

SELECT 1: *PERIOD_NUMBER*=6 AND *SOFTWARE_ID*=0 AND *PRIMARY_TRACKS* > 5;
SELECT 2: *PERIOD_NUMBER*=6 AND *SOFTWARE_ID*=2 AND *PRIMARY_VERTEX* = FALSE AND *DETECTOR_HIT*[BITS:0-5];
SELECT 3: *PERIOD_NUMBER*=5 AND *RUN_NUMBER* > 100 AND *PRIMARY_VERTEX* = TRUE AND *PARTICLES*[BITS:5-9] > 4;
SELECT 4: *PERIOD_NUMBER*=4 AND *SOFTWARE_ID*=1 AND *PRIMARY_TRACKS*>6 AND *ALL_TRACKS*>40 AND *PARTICLES*[BITS:10-14].

NEW

EMS for NICA. Current status (an update)

EMS for NICA. PostgreSQL Indexing

- Indexing methods such as:
  - **BRIN (Block Range Index) –**
    - **the most effective**
  - BTREE (B-Tree) –
    - the best target result
    - ineffective
    - too heavy

- Indexing based on **BRIN implemented** for columns:
  - period_number,
  - run_number,
  - software_id.

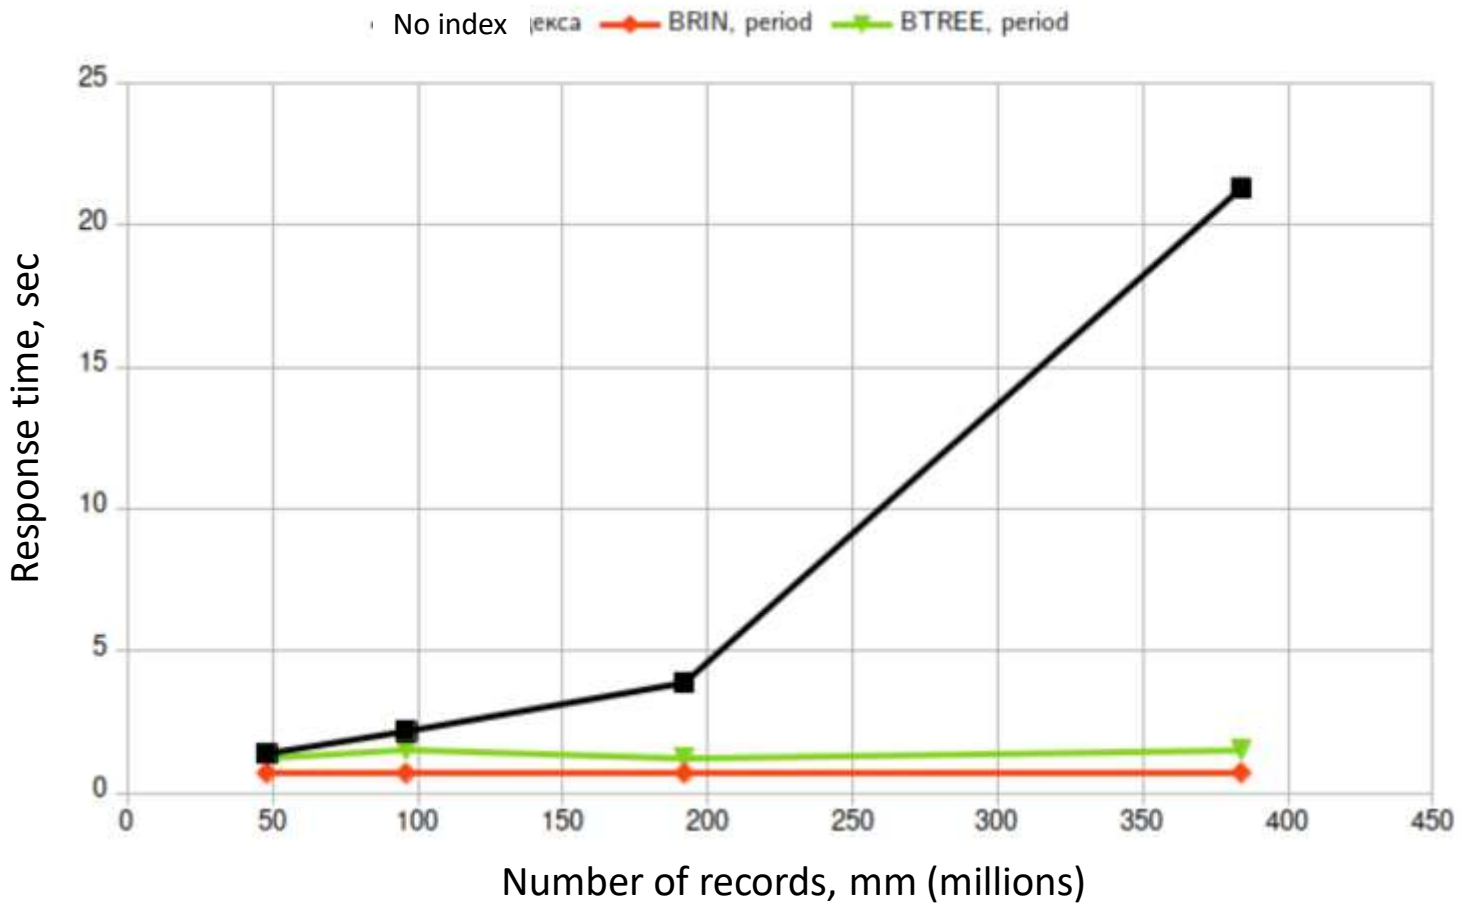

EMS for NICA. High Availability

- Implemented Master-Slave (RFC 2136) technology for all parts of EMS: DB, API, keepalive
  - Replication parameter = 2

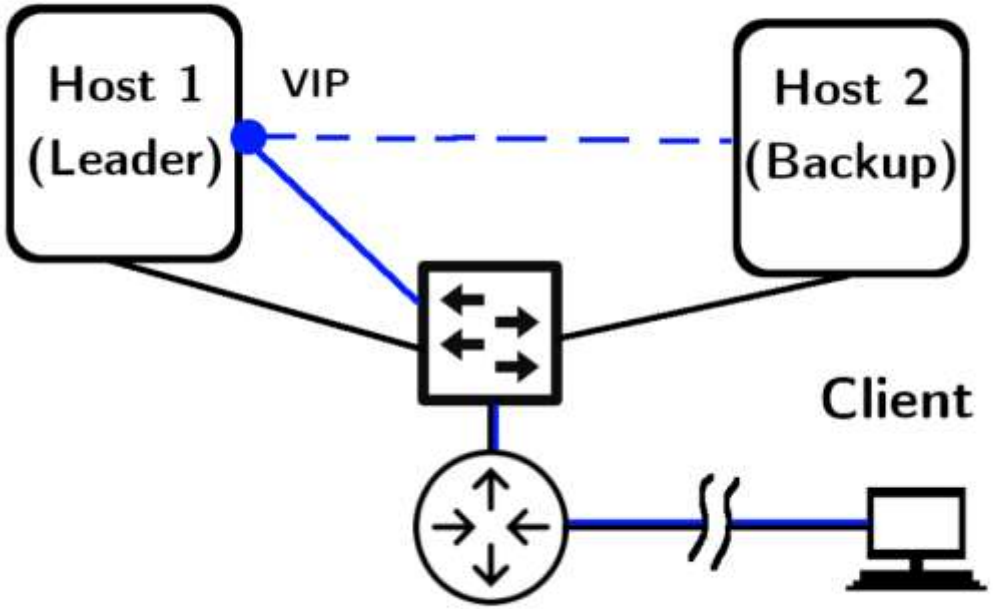

- Advantages:
  - Robust fault-tolerance (!)
    - Low downtime
    - Improved consistency
    - Data replication (safety)

EMS for NICA. Statistics Collection and Display

- Statistics is saved to the database and displayed in the Web interface
- Script run is scheduled periodically

EMS for NICA. Automated Deployment

- After running Ansible playbooks supports:
  - One-node deploy
  - Double-node deploy (on picture)

Powered by

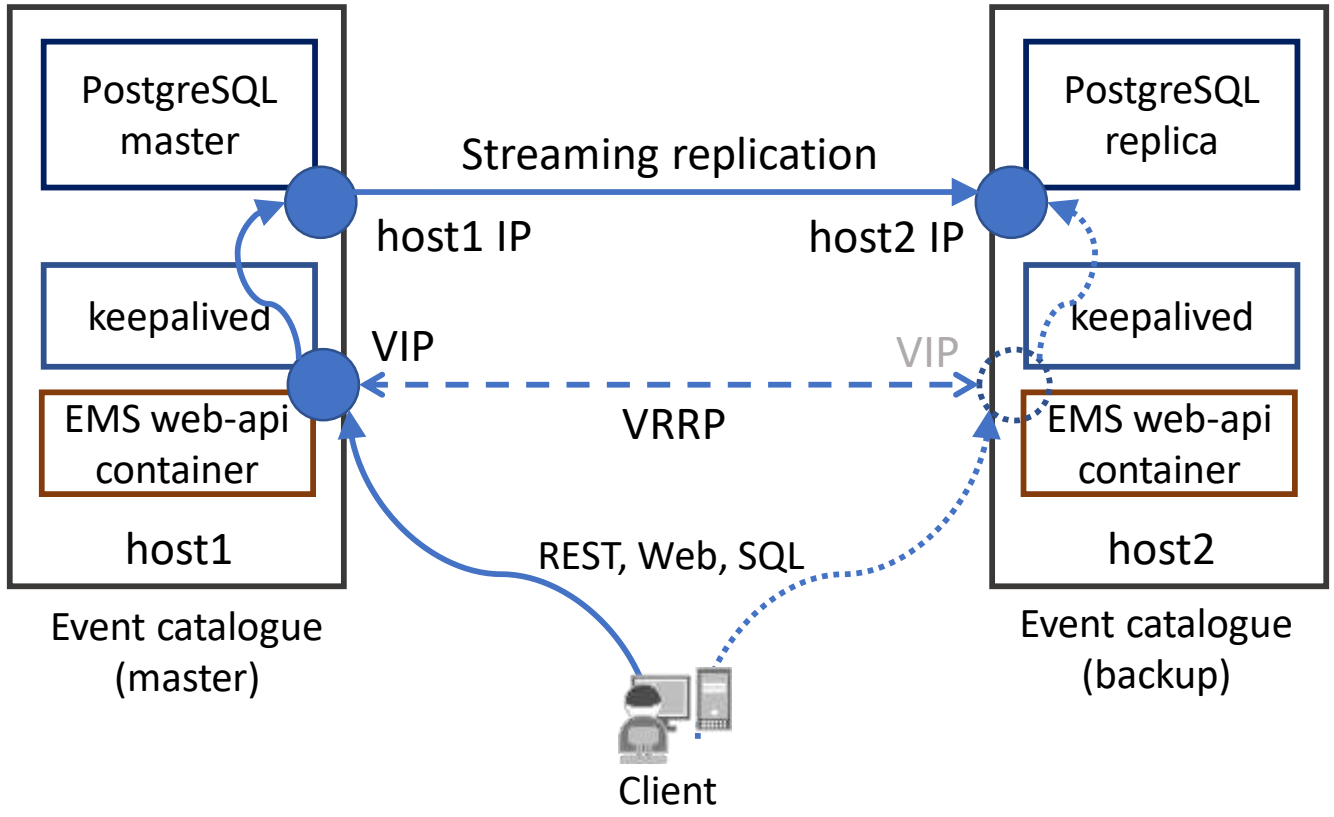

EMS for NICA. Integration with KeyCloak-based authorization system

- Authentication and authorization in EMS
  - FreeIPA / LDAP support has been dropped
  - KeyCloak token-based authentication and authorization is now supported
  - Database-based authentication is supported as before

```
keycloak_auth:
  server_url: "https://bmn-user.jinr.ru"
  realm: "BMN"
  client_id: "emd_api"
  client_secret: "*****"
  writer_group_name: "bmneventwriter"
  admin_group_name: "bmneventadmin"

# database_auth: True
```

EMS for NICA. Planning

EMS for NICA. Plans

- Integration with the VisionForge system

- Control the EMS by the Monitoring Service (Telegraf – InfluxDB - Grafana)

Thank you for your attention!

Abstract

Event indexing, or event metadata systems (EMS) are common for particle physics experiments. Their main goal is to keep a searchable catalogue of experimental events from which a subset of data can be extracted based on given filtering criteria. The BM@N experiment's EMS has been designed, developed and deployed previously and is now being improved to increase its performance, convenience for users, as well as fault tolerance. In this talk, the architecture of the current version of the BM@N EMS is reviewed and some recent improvements that have been applied are discussed.