BM@N DAQ Data Center

BM@N Experiment at the NICA Facility
14th Collaboration Meeting
JINR, Dubna, May 13 – 15, 2025

ILIA SLEPNEV, JINR

DAQ Data Center

Outline

DAQ Data Center

- Mission & Design Principles
  - Purpose of the DAQ Data Center
  - Guiding Architecture Principles
- Data Flow & Network Fabric
  - End-to-End Data Path
- BM@N DAQ Network Topology
  - DAQ Network Performance & Reliability
  - High Availability & Storage Resilience
  - Readout Link Redundancy Constraints
- Compute, Storage & Virtualization
  - Distributed Storage Cluster (CephFS Layer)
  - Virtual & Bare-Metal Compute Tiers

LHEP Computing Resources

- LHEP data centers in operation

Monitoring

- Incident Resolution
- Node-RED Automation
- Grafana: DDC dashboard

Extra

- Infrastructure Management
- Monitoring Architecture
- Log Message Processing and Analysis

DAQ Data Center

Mission & Design Principles

Purpose of the DAQ Data Center

- Central hub for experiment data reception and archiving
- Decouples micro-second readout from second-scale processing
- Guarantees continuous acquisition during long physics runs
- Operates autonomously, no external IT dependencies
- Hosts online monitoring for quality-of-data checks
- Designed to evolve without disruptive rewiring or downtime

Guiding Architecture Principles

- All critical paths redundant, no single point of failure
- Software-defined everything: storage, network, virtualization
- Hardware chosen for low-latency performance
- Horizontal scalability—add nodes, no redesign required
- Observability first: fine-grained metrics, logs, alerts
- Security via dedicated VLANs, JINR SSO and firewall zoning
- Documentation-driven operations; infra declared in Git
- Emphasizes open-source, vendor-neutral technologies

DAQ Data Center

Data Flow & Network Fabric

End-to-End Data Path

- Detector front-ends push UDP streams via custom MStream protocol
- Hardware IP core in FPGA provides dat transfer and control
- First-Level Processor buffers, validates, formats packets
- Event builders assemble multi-detector fragments asynchronously
- Asynchronous layers prevent back-pressure on readout electronics
- FLP quality checks tag corrupted or partial events
- Event files streamed to high-availability Ceph storage
- Complete files are synchronized to Offline farm immediately

TOF Readout Board

- VXS 6U 160 cm form factor
- Front-Panel: detector I/O
- Backplane: sync, readout 1Gb/s

21-slot, VXS Chassis

- Power, Cooling, Remote control
- Backplane: Dual-Star, 10 Gb/s per slot

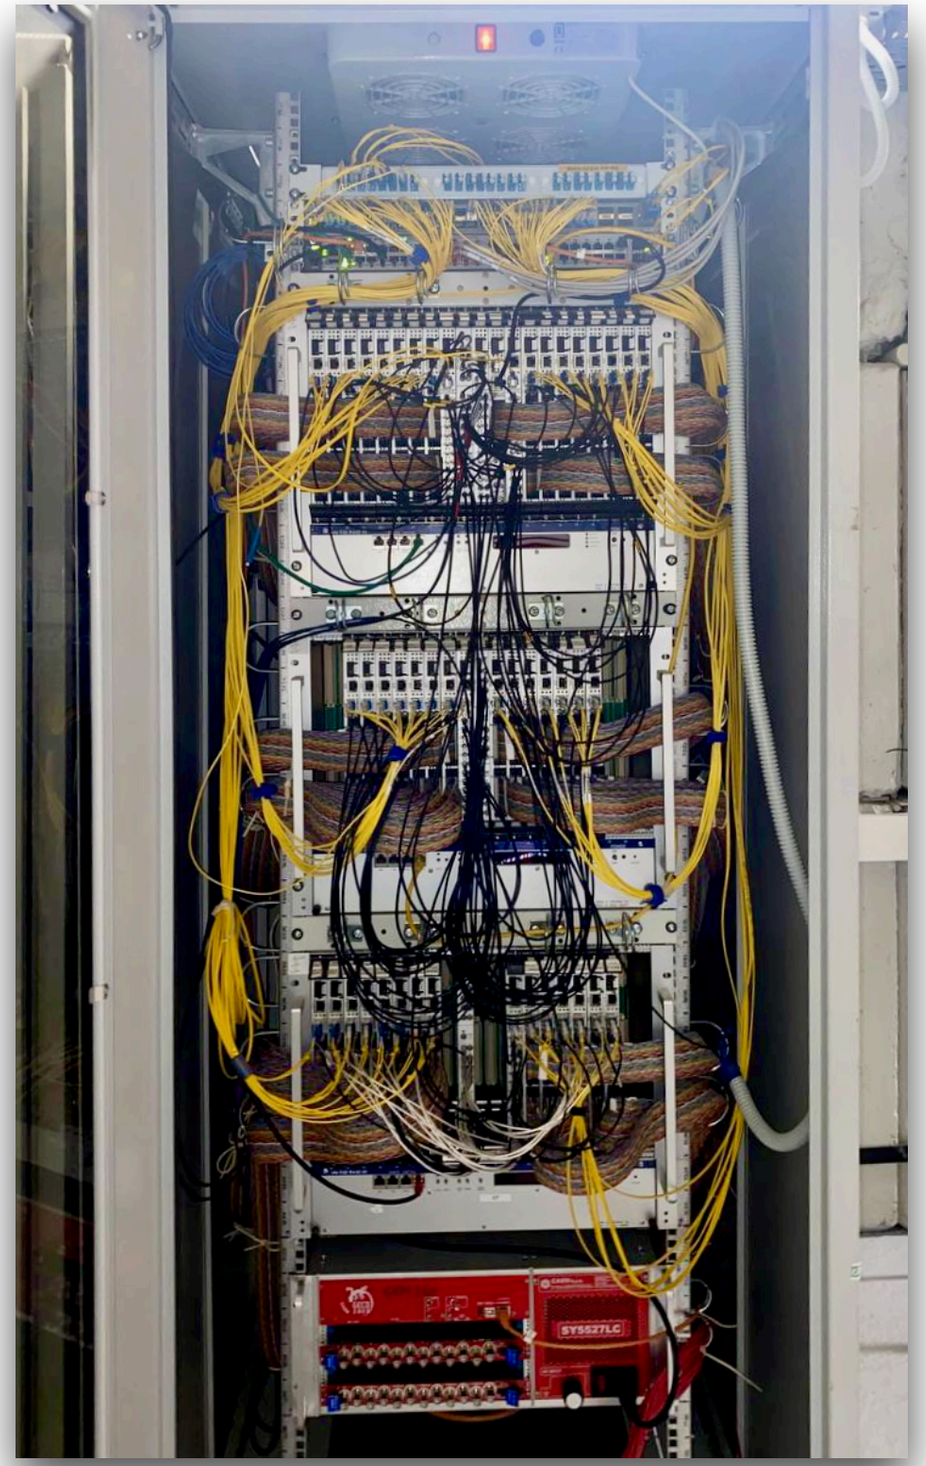

Readout Electronics Rack: GEM

- 3 × VME64x chassis with ADC64 boards
- Fiber-optical front panel readout links
- Ribbon cable detector I/O
- Top-of-Rack Cisco Nexus 93180 switch

DAQ Data Center

BM@N DAQ Network Topology

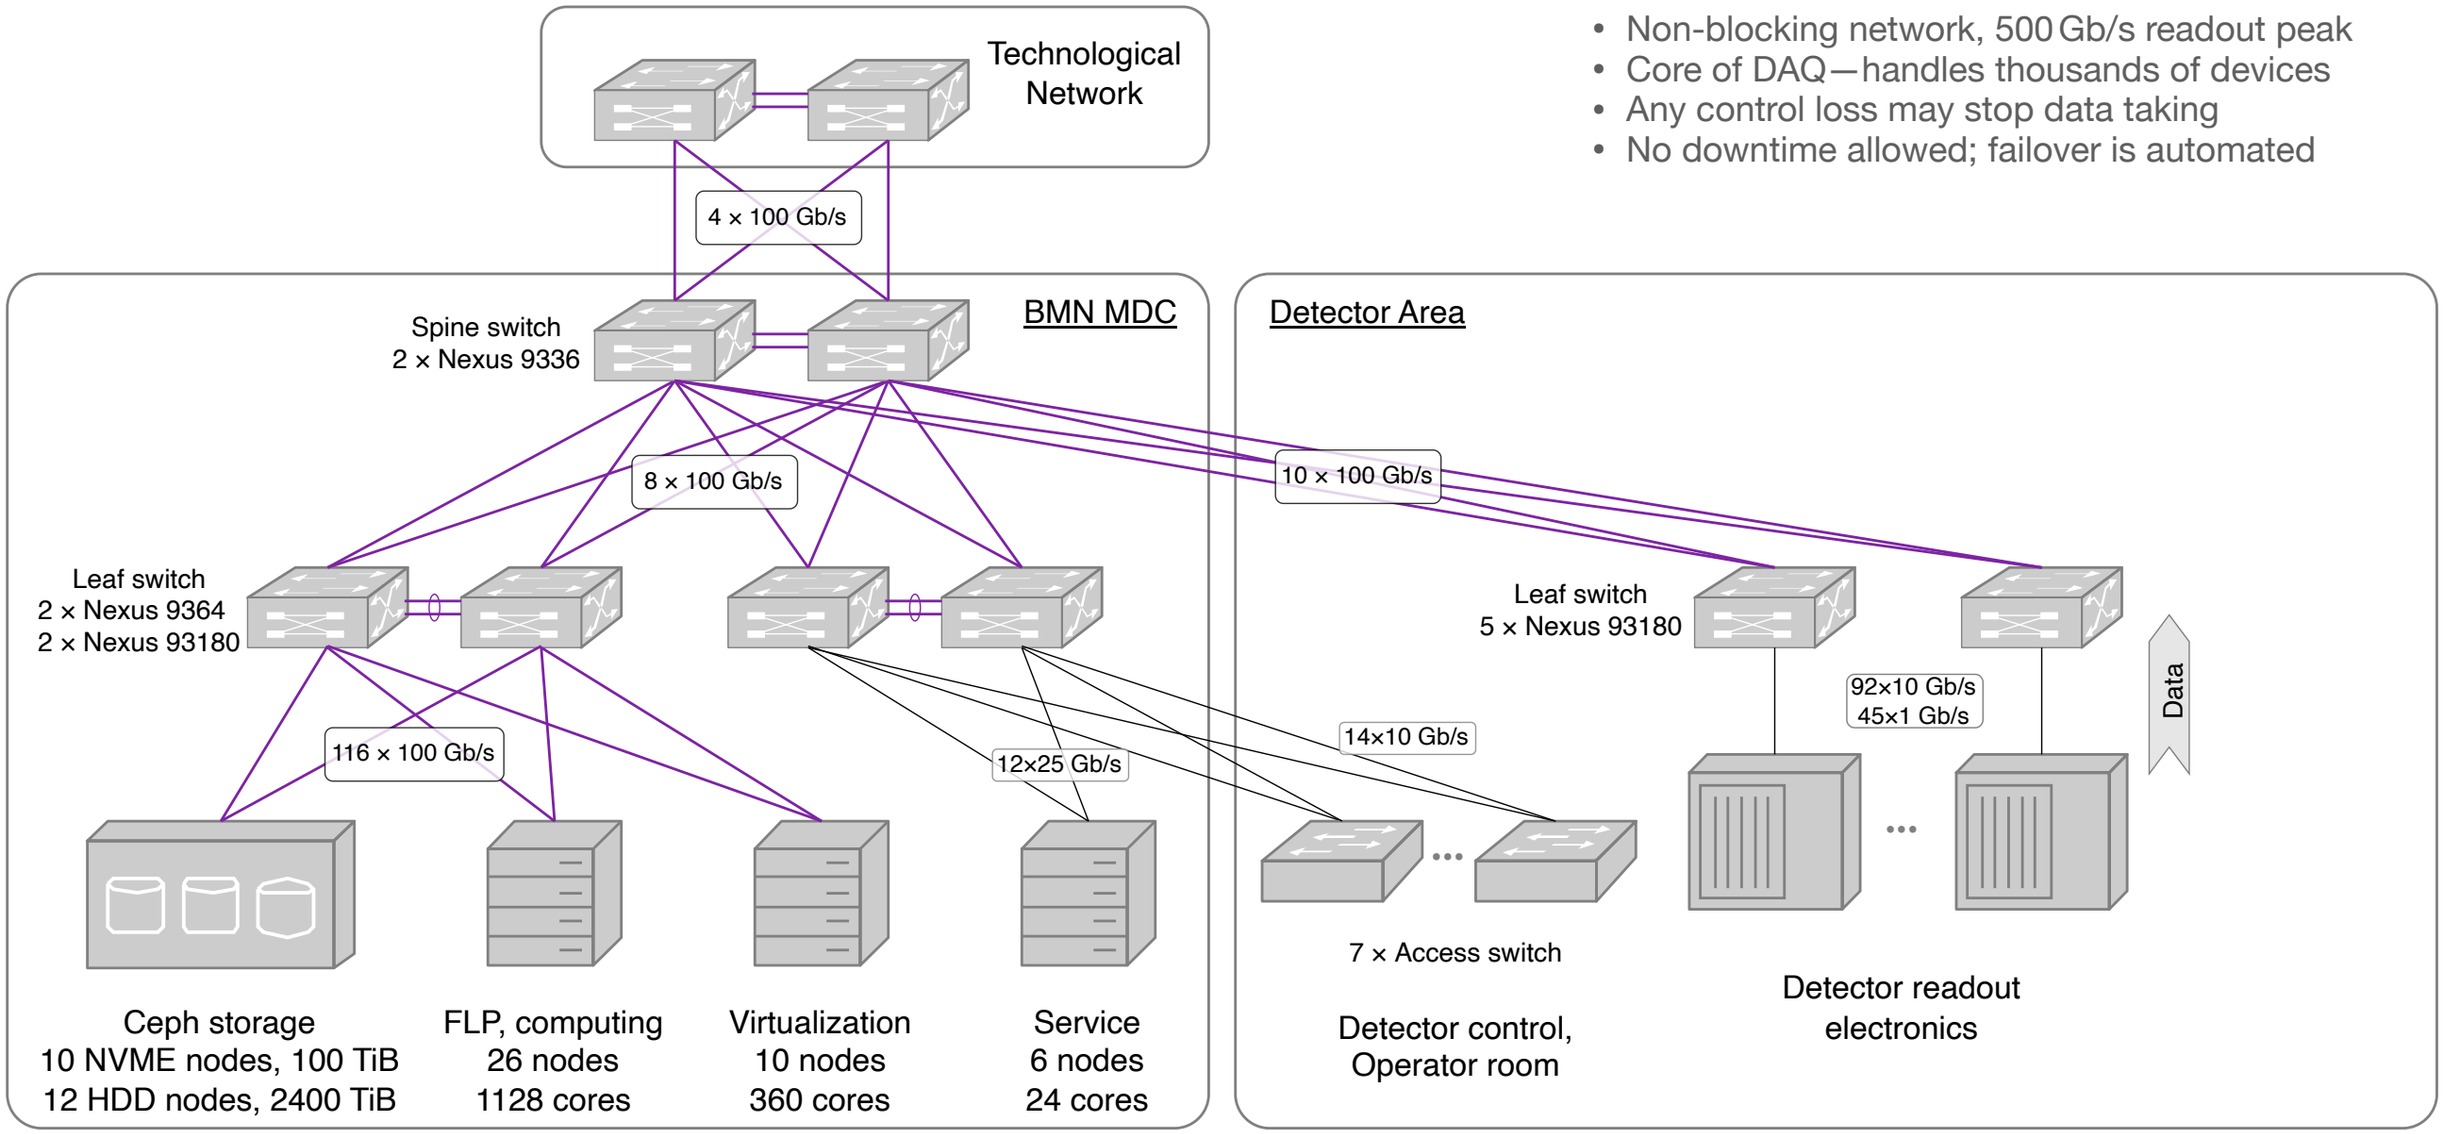

DAQ Network Performance & Reliability

- Non-blocking network, 500 Gb/s readout peak
- Core of DAQ—handles thousands of devices
- Any control loss may stop data taking
- No downtime allowed; failover is automated

High Availability & Storage Resilience

- Ceph ensures 24/7 access with auto-recovery on failure
- HDD pools use erasure coding, NVMe SSD pools triple-replication
- Storage remains online with no data loss or downtime
- External users can access data at all times

Readout Link Redundancy Constraints

- No redundancy in readout links due to custom hardware
- Failover is complex, costly, and resource-limited
- Manpower and design constraints limit improvements
- Focus is on robustness and failure monitoring

DAQ Data Center

Compute, Storage & Virtualization

Distributed Storage Cluster (CephFS Layer)

- NVMe replicated pools for low-latency RBD workloads
- HDD pools with erasure coding for cost-effective capacity
- CRUSH algorithm enforces fault-domain-aware placement
- Self-healing scrubs verify checksums without operator action
- POSIX CephFS exports RAW data to FLP and event builders
- Grafana dashboards track PG-health, latency, rebuild rate
- Capacity expansion performed online by adding OSD nodes
- Data-at-rest encrypted; keys managed by vault-backed KMS
- Network-isolated background replication to off-site cluster
- Retention policy aligns with collaboration data mandate

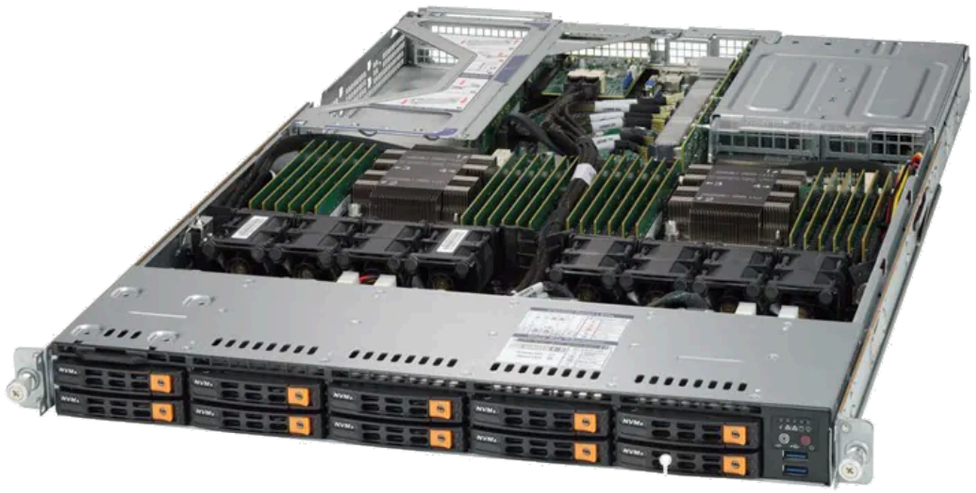

Ceph OSD node: 10 × 3.8 TB NVMe SSD

Ceph OSD node: 24 × 18 TB HDD

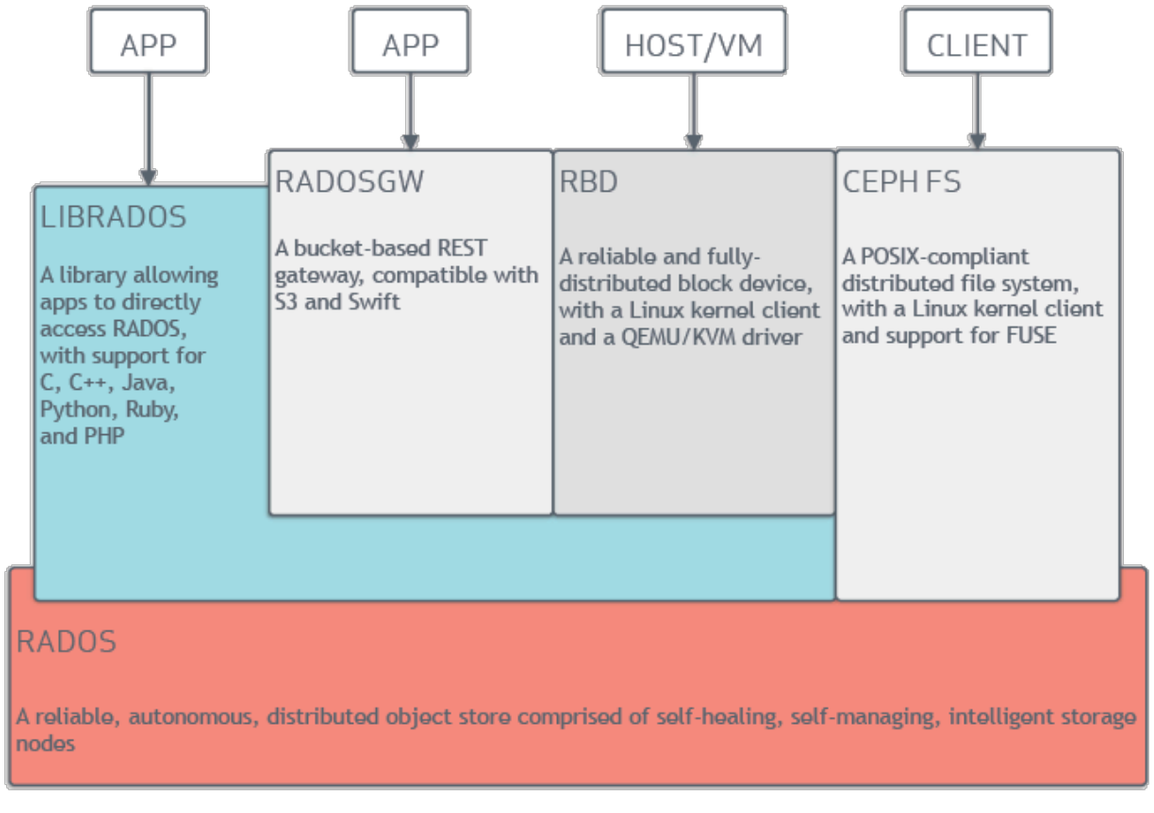

DAQ Data Center

Compute, Storage & Virtualization

Virtual & Bare-Metal Compute Tiers

- KVM / LXC cluster hosts control, monitoring, DB, web portals
- Bare-metal nodes reserved for FLP and heavy event builders
- Live-migration keeps services online during maintenance
- Snapshot roll-back guards against faulty software updates
- Resources re-balanced when DAQ idle to boost offline jobs
- Continuous incremental backups: hourly, daily, weekly, yearly
- SSO-protected self-service portal for VM lifecycle requests
- Performance tuned: CPU-governor, NUMA pinning, NIC interrupts

Dual-node compute server

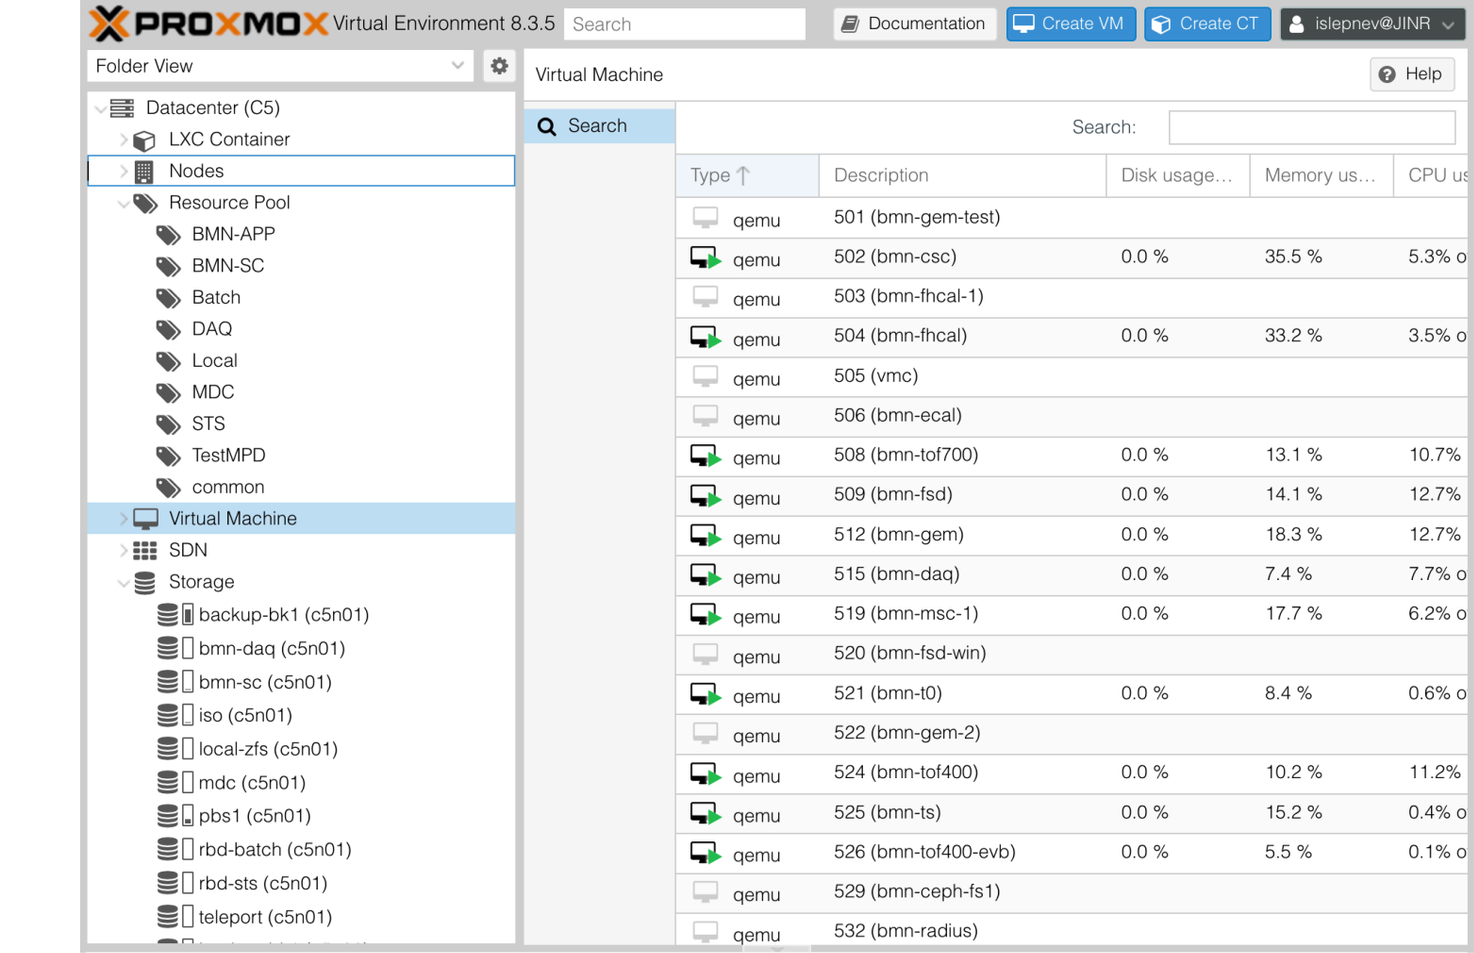

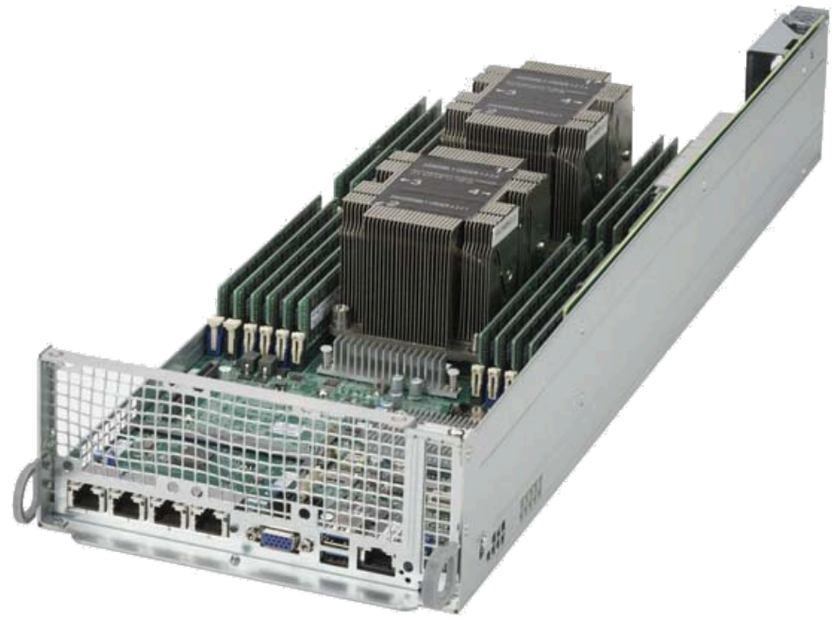

Compute node

RAM: 384 GB

CPU: Dual Xeon Gold 6154, 6342

LHEP Computing Resources

Data Centers in Operation

DDC — DAQ Data Center

- Data taking, online processing and monitoring
- Primary RAW data storage
- CephFS for immediate data access
- Secondary role: batch jobs outside DAQ periods

NCX — Offline Cluster

- Batch and interactive jobs at large scale
- Experimental and simulation data storage: EOS
- Shared by multiple experiments, 200+ users

	DDC (BMN)	DDC (MPD)	NCX
Location	Building 215, BMN area	MPD Hall	Building 216, room 115
Operating since / Last upgrade	2019 / 2023	2021 / 2024	2019 / 2023
CPU architecture	3.0 GHz Skylake 2.8 GHz Ice Lake	3.1 GHz Cascade Lake 2.8 GHz Ice Lake	2.6 GHz Broadwell-EP 2.0 GHz Skylake 2.5 GHz Cascade Lake
CPU cores — Total	1488	1760	4200
CPU cores — Batch	700 [1]	1000 [1]	3000 [2]
RAM, GB per CPU core	6–7.5	7.5	9.6–16
Node uplink, Gb/s	2×100	2×100	100
Local node storage (/tmp)	32 GB SSD	32 GB SSD	1 TB HDD
Shared storage (workspace, experimental and simulation data)	100 TB NVMe (CephFS) 2.5 PB HDD (CephFS)		124 TB NVMe (NFS+ZFS) 11 PB HDD (EOS)

(1) Additional CPU cores available when DAQ is inactive
(2) May be temporarily unavailable due to maintenance

CephFS

- Distributed POSIX-compliant file system on Ceph
- Supports high-throughput, parallel data access
- Used for RAW data from DAQ in DDC clusters
- Mounted natively on compute and DAQ nodes

EOS

- CERN-developed distributed file system
- Optimized for large-scale data access via XRootD
- Used in NCX cluster for simulation and analysis
- Suited for shared access by many experiments

Computing Resources

Data Centers in Operation

BM@N DAQ Data Center

NCX Offline Cluster

Monitoring

Incident Resolution

Detect
- Node-RED
- Zabbix

Alert
- Support engineers
- Management

Coordinate
- Phone call
- Messengers

Resolve
- Switch-On
- Verify
- Postmortem

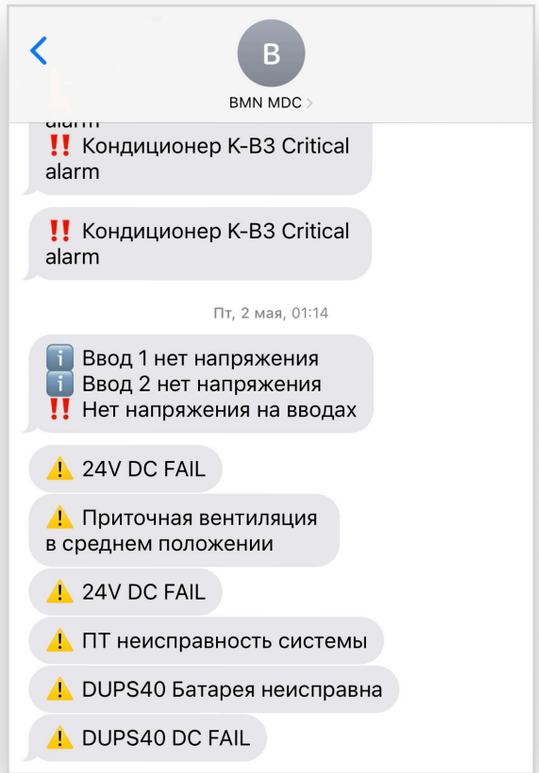

SMS Alert

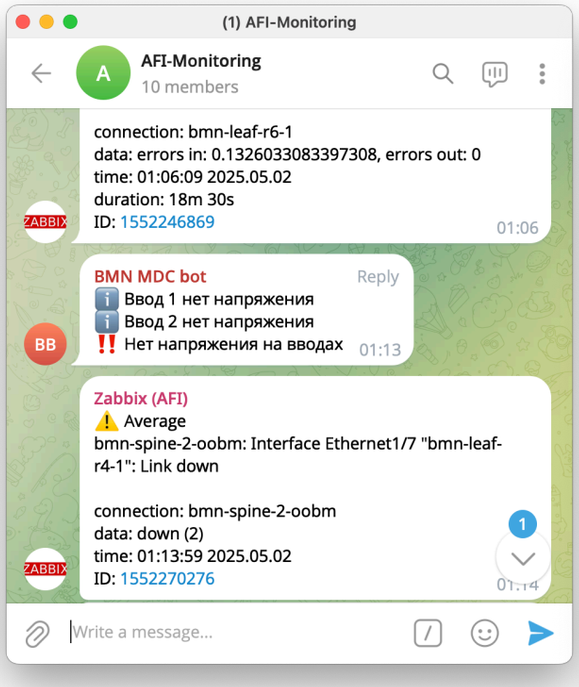

Telegram Alert

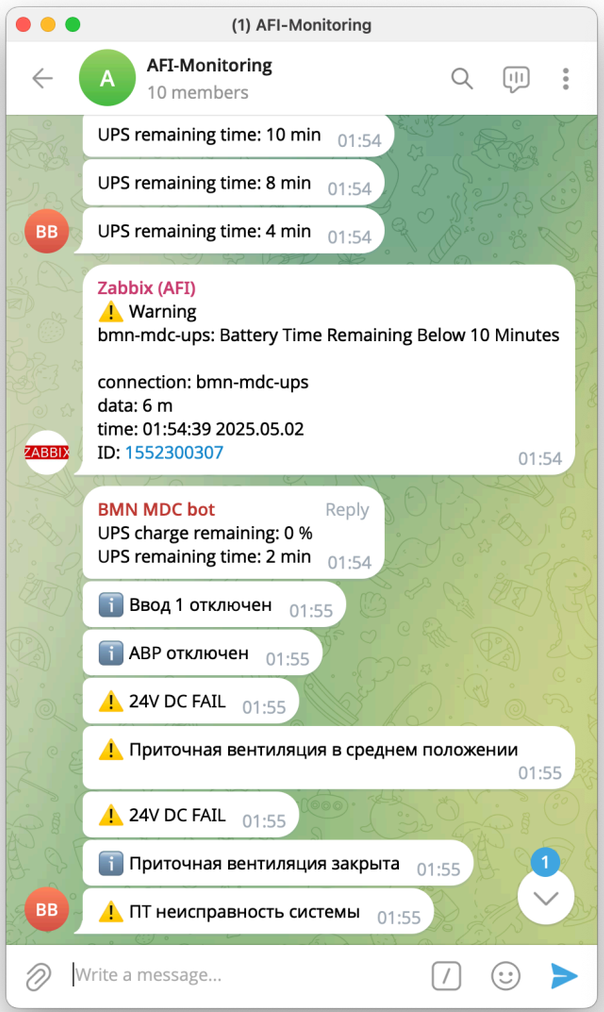

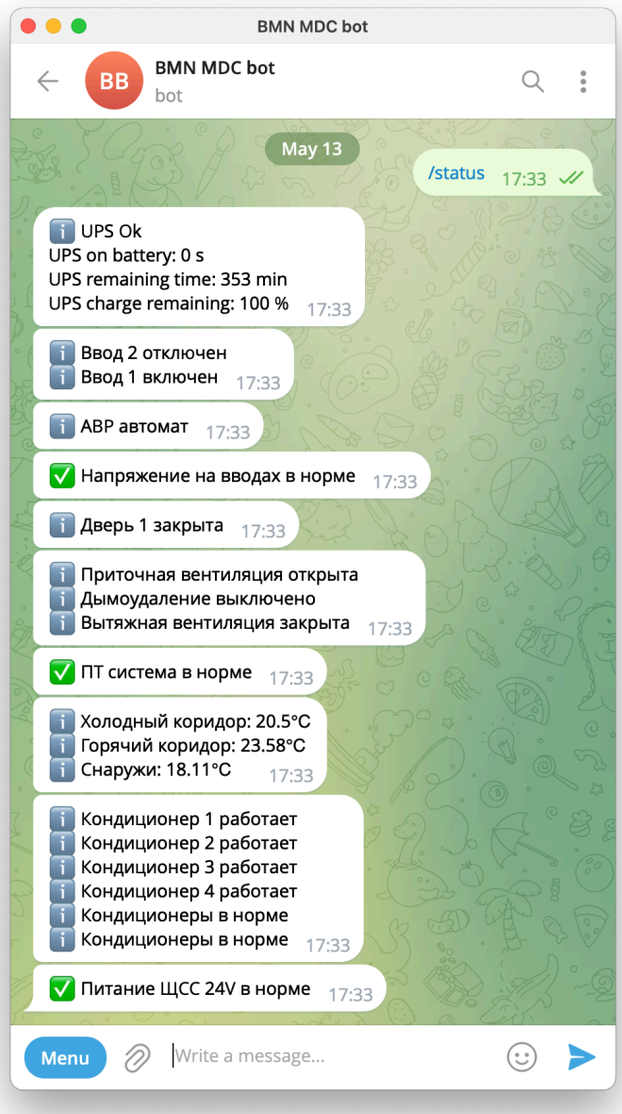

- Real-time alerts via SMS and Telegram bots
- Automated escalation from monitoring systems
- Critical events resolved via coordinated response
- Status recovery confirmed by bot interaction

Monitoring

Node-RED Automation

- Node-RED automates low-level alert processing
- Monitors power, HVAC, ventilation, security
- Custom logic triggers SMS and Telegram alerts instantly
- Each flow handles real-time decision-making
- Part of end-to-end automated monitoring system

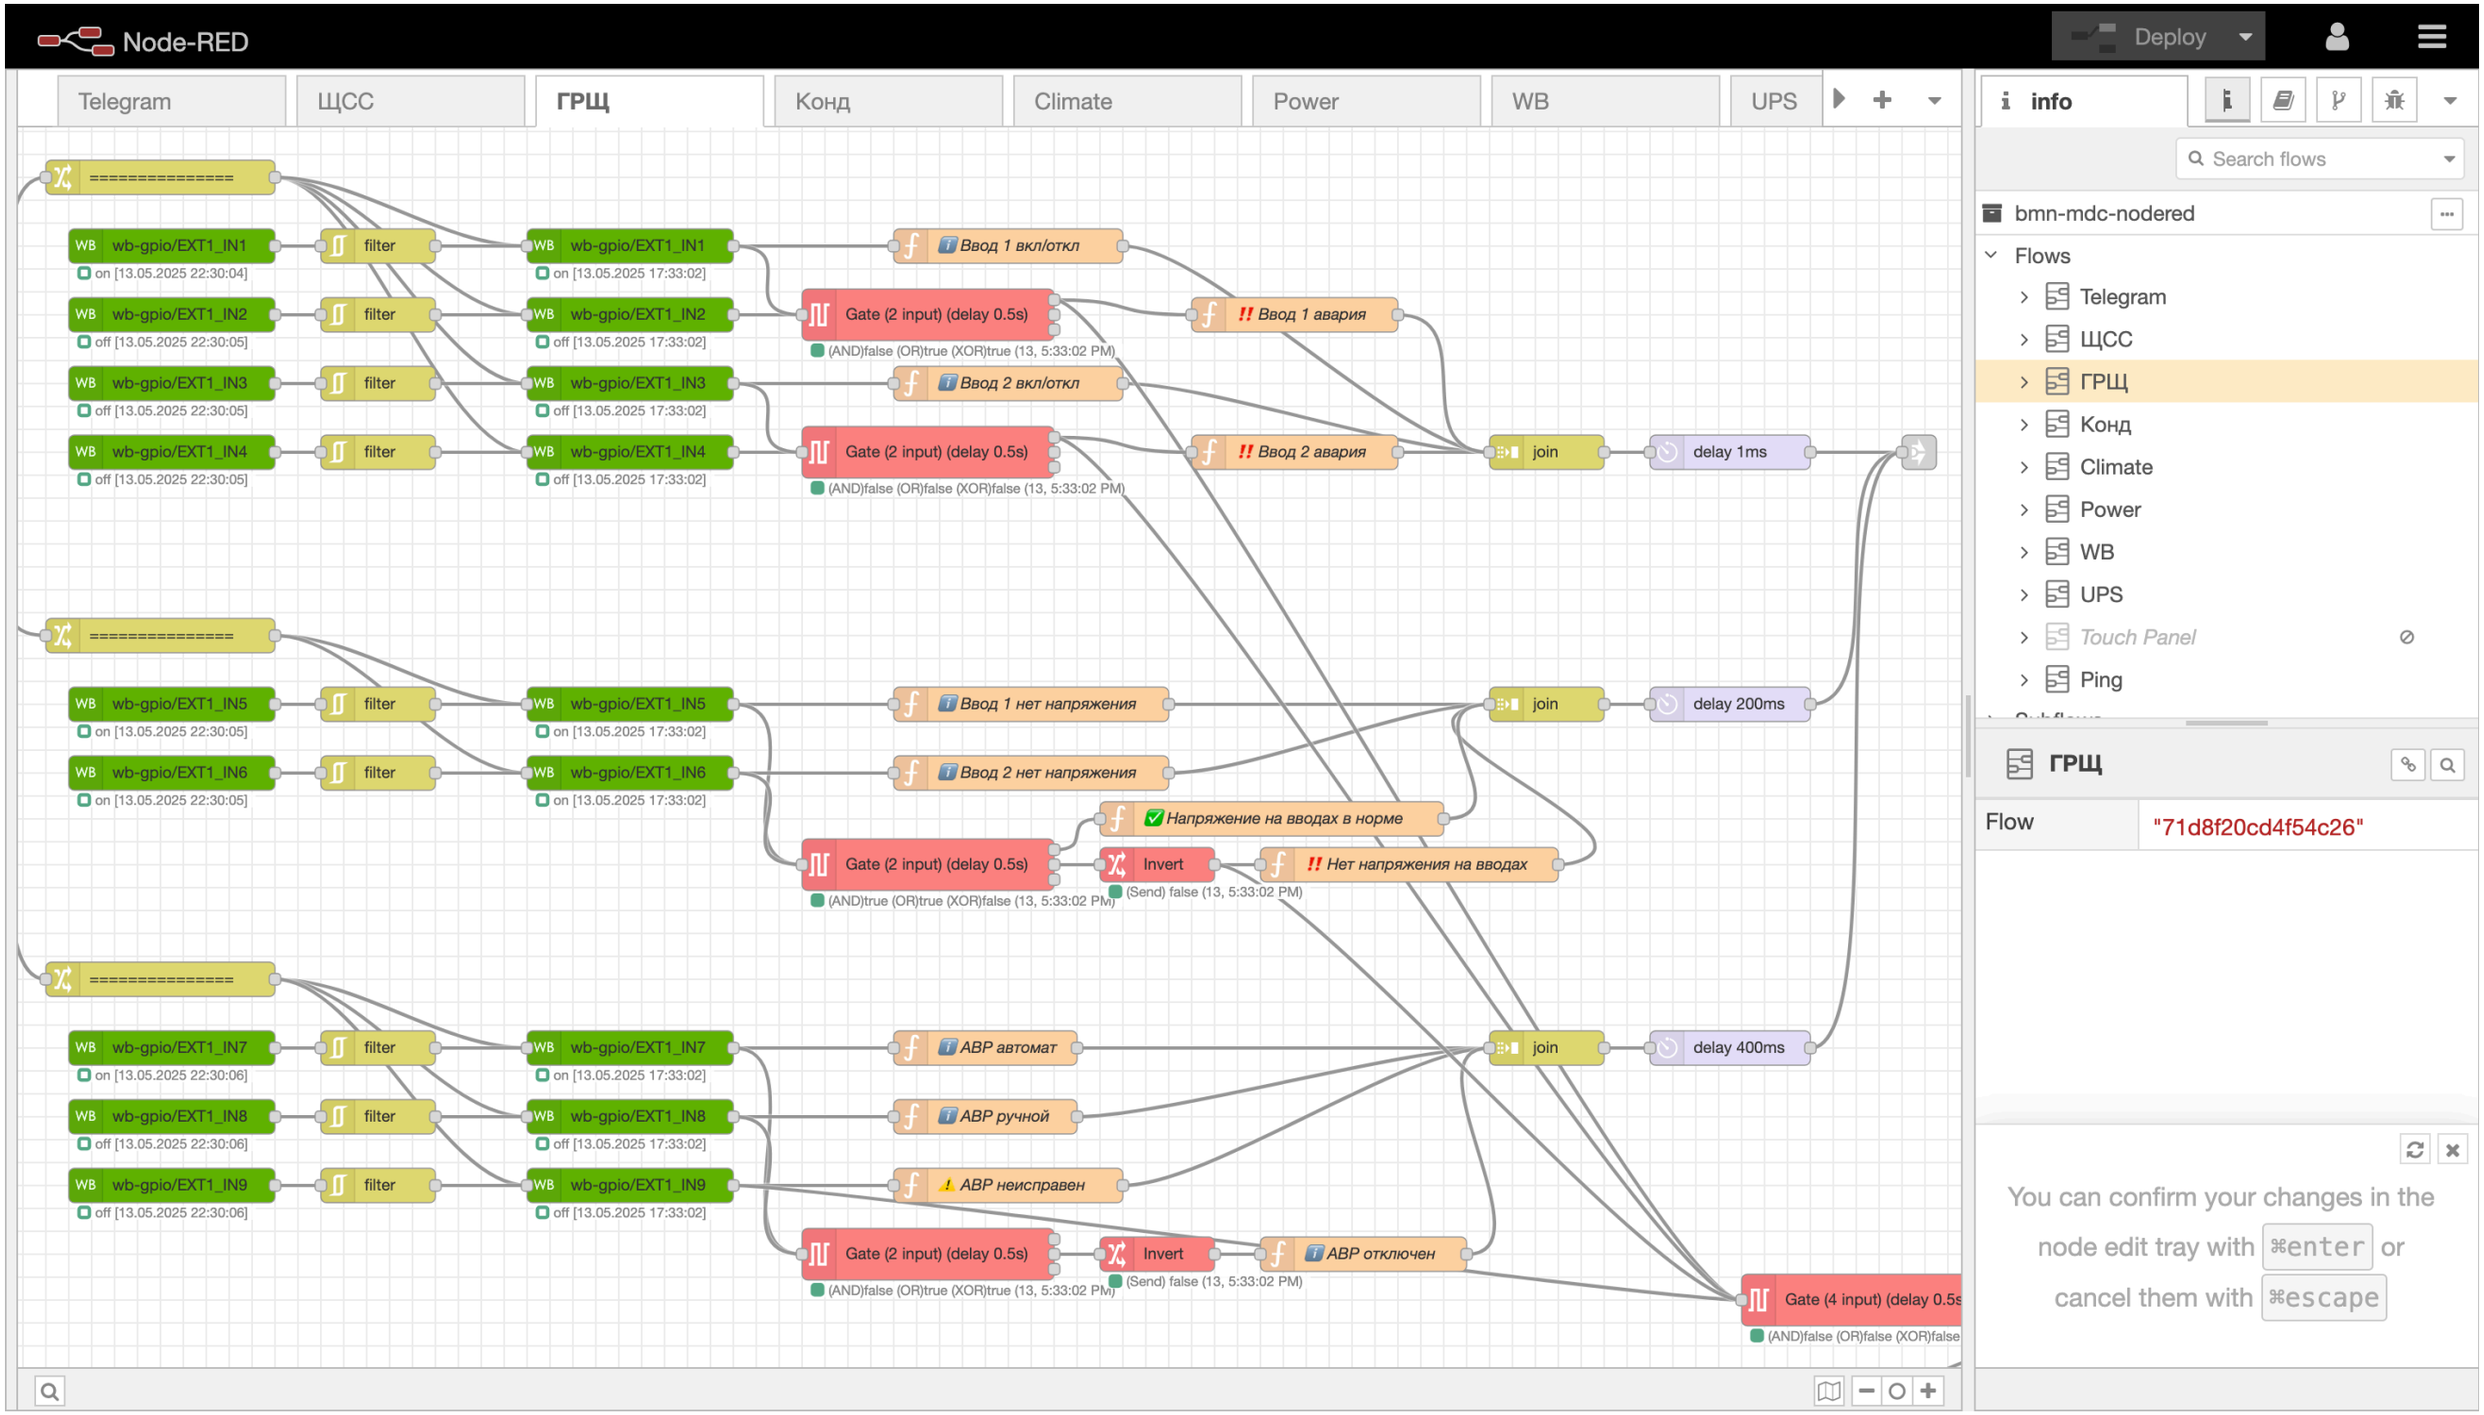

Monitoring

Grafana: DDC dashboard

Acknowledgements

DAQ Data Center

Sergey Bazylev

Andrey Egorov

Alexander Fediunin

Ivan Filippov

Sergey Kuklin

Andrey Shchipunov

NCX Offline Cluster

Ivan Slepov

Your contributions made this work possible.

Thank you!

Extra slides

Infrastructure Management

Infrastructure-as-Code

Tool	Method	Approach, our usage	Tasks
Puppet	Pull	functional (declarative)	configure services, settings
Ansible	Push	procedural (imperative)	updates, one-time tasks
—	—	manual admnistration	other complex tasks

- Machine-readable, version-controlled configuration files (YAML, Ruby)
- Puppet modules:
  - provision, configure, manage OS and application components
  - supported by community or our custom solution
- Hierarchical design: roles, profiles, classes are assigned to groups of computers. Dev and Prod environments.
- Documentation of IT Infrastructure configuration

daq.yaml 2.34 KB

```yaml
---
classes:
  - apel
  - apel::testing
  - autofs
  - profile::service::cephfs_automount
  - sysctl::base
  - ssh::client
  - ssh::server

apel::testing::enabled: '1'
daq_vncserver::home_manage: true
daq_fedora::homedir::desktop_bg: '#1b3324'
desktop::desktop: 'LXDE'

autofs::mounts:
  net:
    mount: '/net'
    mapfile: '-hosts'
  ceph:
    mount: '/-'
    mapfile: '/etc/auto.ceph'
    options: '--timeout=120'
```

init.pp 344 Bytes

```
#
class cvmfs(
  String $package_release,
  String $package_release_url,
  String $package_ensure,
  Boolean $package_manage,
  Array[String] $package_name,
) {
  contain cvmfs::repo
  contain cvmfs::install
  contain cvmfs::config

  Class['::cvmfs::repo']
  -> Class['::cvmfs::install']
  -> Class['::cvmfs::config']
  ~> Service['autofs']
}
```

```
[root@bmn-evb ~]# puppet agent -vt
Info: Using environment 'production'
Info: Retrieving pluginfacts
Info: Retrieving plugin
Info: Retrieving locales
Info: Loading facts
Info: Caching catalog for bmn-evb.he.jinr.ru
Info: Applying configuration version '1684344730'
Notice: /Stage[main]/Autofs::Service/Service[autofs]/ensure: ensure changed 'stopped' to 'running' (corrective)
Info: /Stage[main]/Autofs::Service/Service[autofs]: Unscheduling refresh on Service[autofs]
Notice: Applied catalog in 9.74 seconds
[root@bmn-evb ~]#
```

Monitoring Architecture

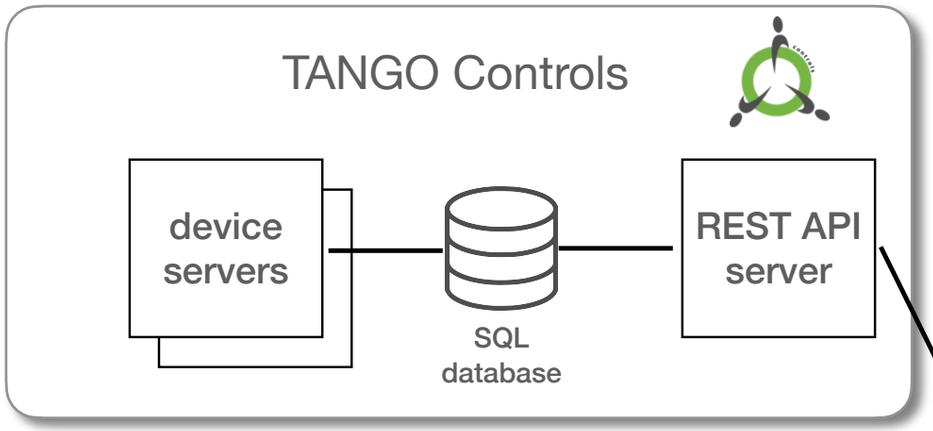

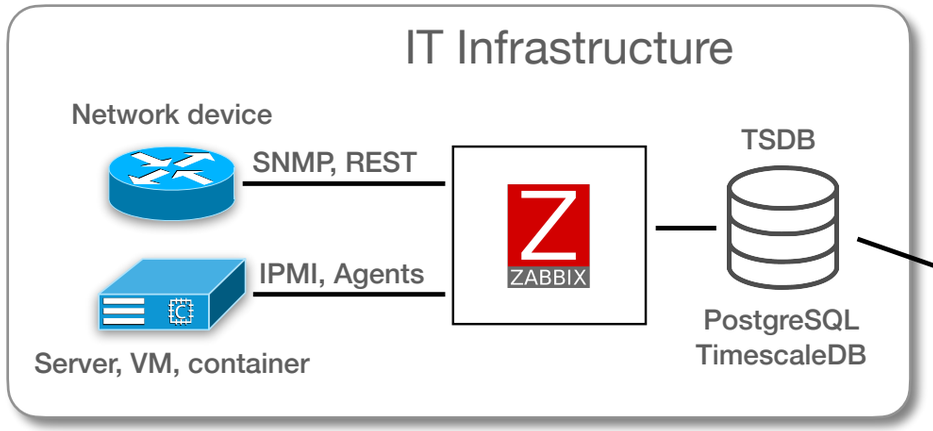

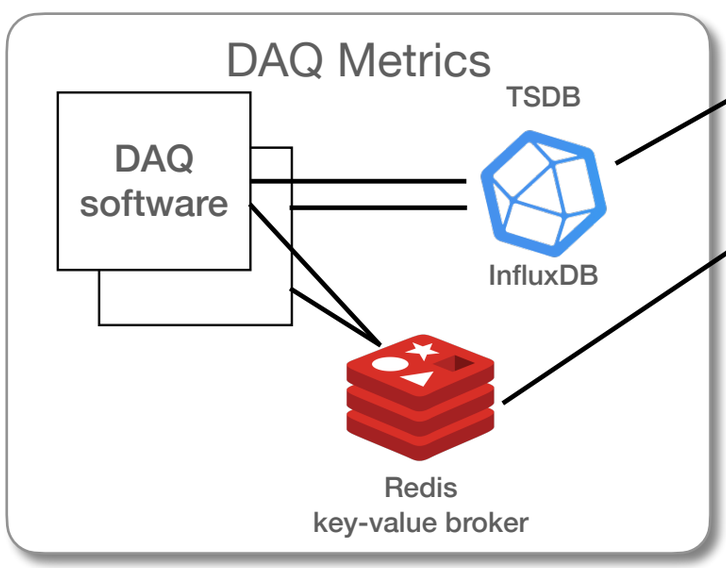

Time-Series Data panels

Query data
Visualise
Analyse
Alert

Unified dashboards
Data source plugins

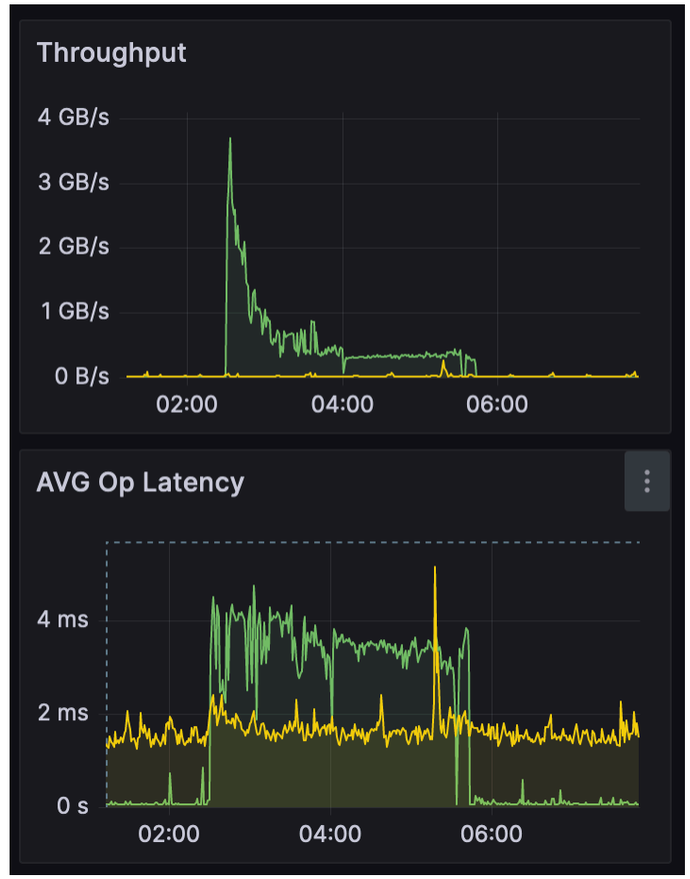

Plotly panel for arbitrary data
Graphs and Histograms

Log Message Processing and Analysis

Elasticsearch, Logstash, Kibana

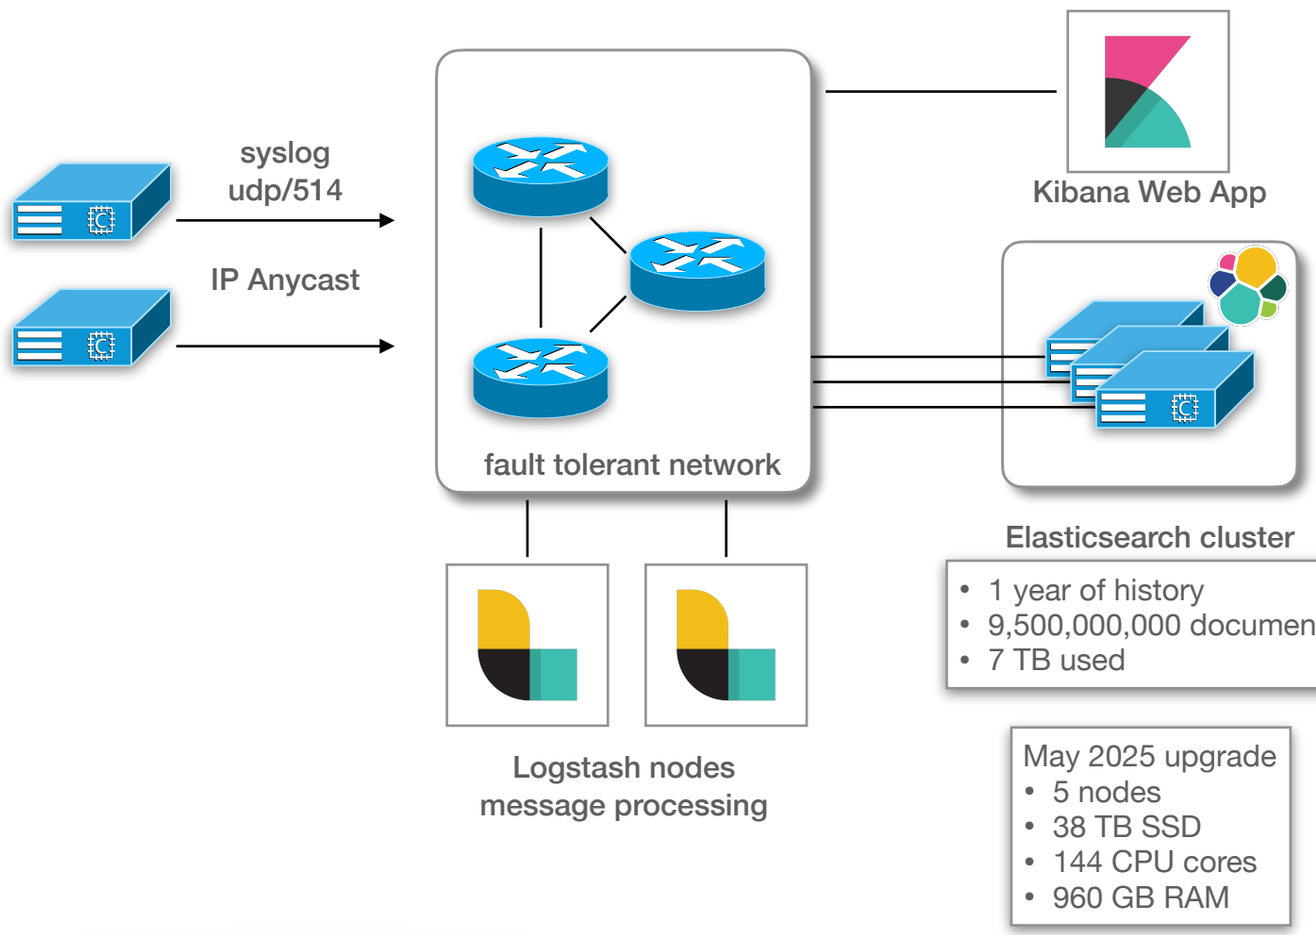

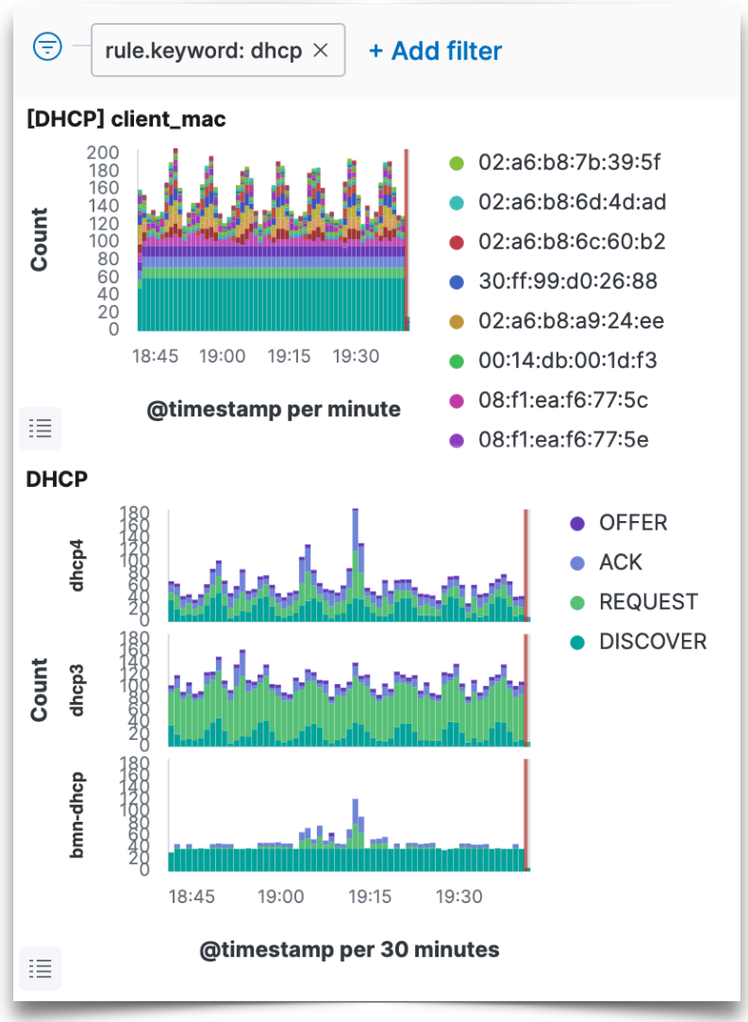

>	May 16, 2023 @ 19:41:04.000	bmn-dhcp	error	02:a6:b8:0d:84:c2	DISCOVER	network 10.18.88/24: no free leases
>	May 16, 2023 @ 19:41:03.000	dhcp3	informational	00:01:7f:54:00:01	REQUEST	BOOTP from dynamic client and no dynamic leases
>	May 16, 2023 @ 19:41:03.000	bmn-dhcp	error	02:a6:b8:6f:29:2c	DISCOVER	network 10.18.88/24: no free leases
>	May 16, 2023 @ 19:41:02.000	bmn-dhcp	error	02:a6:b8:0c:3b:f6	DISCOVER	network 10.18.88/24: no free leases
>	:02.000	dhcp4	error	00:13:95:21:0c:2a	DISCOVER	network 10.18.16/24: no free leases

Exit full screen