

Application of InfluxDB2 in the monitoring system of the MICC MLIT JINR



Kashunin I., Mitsyn V., Strizh T.
June 2023




MICC – multifunctional information computing complex



General view of JINR MICC

Cooling system



Work nodes and disk servers

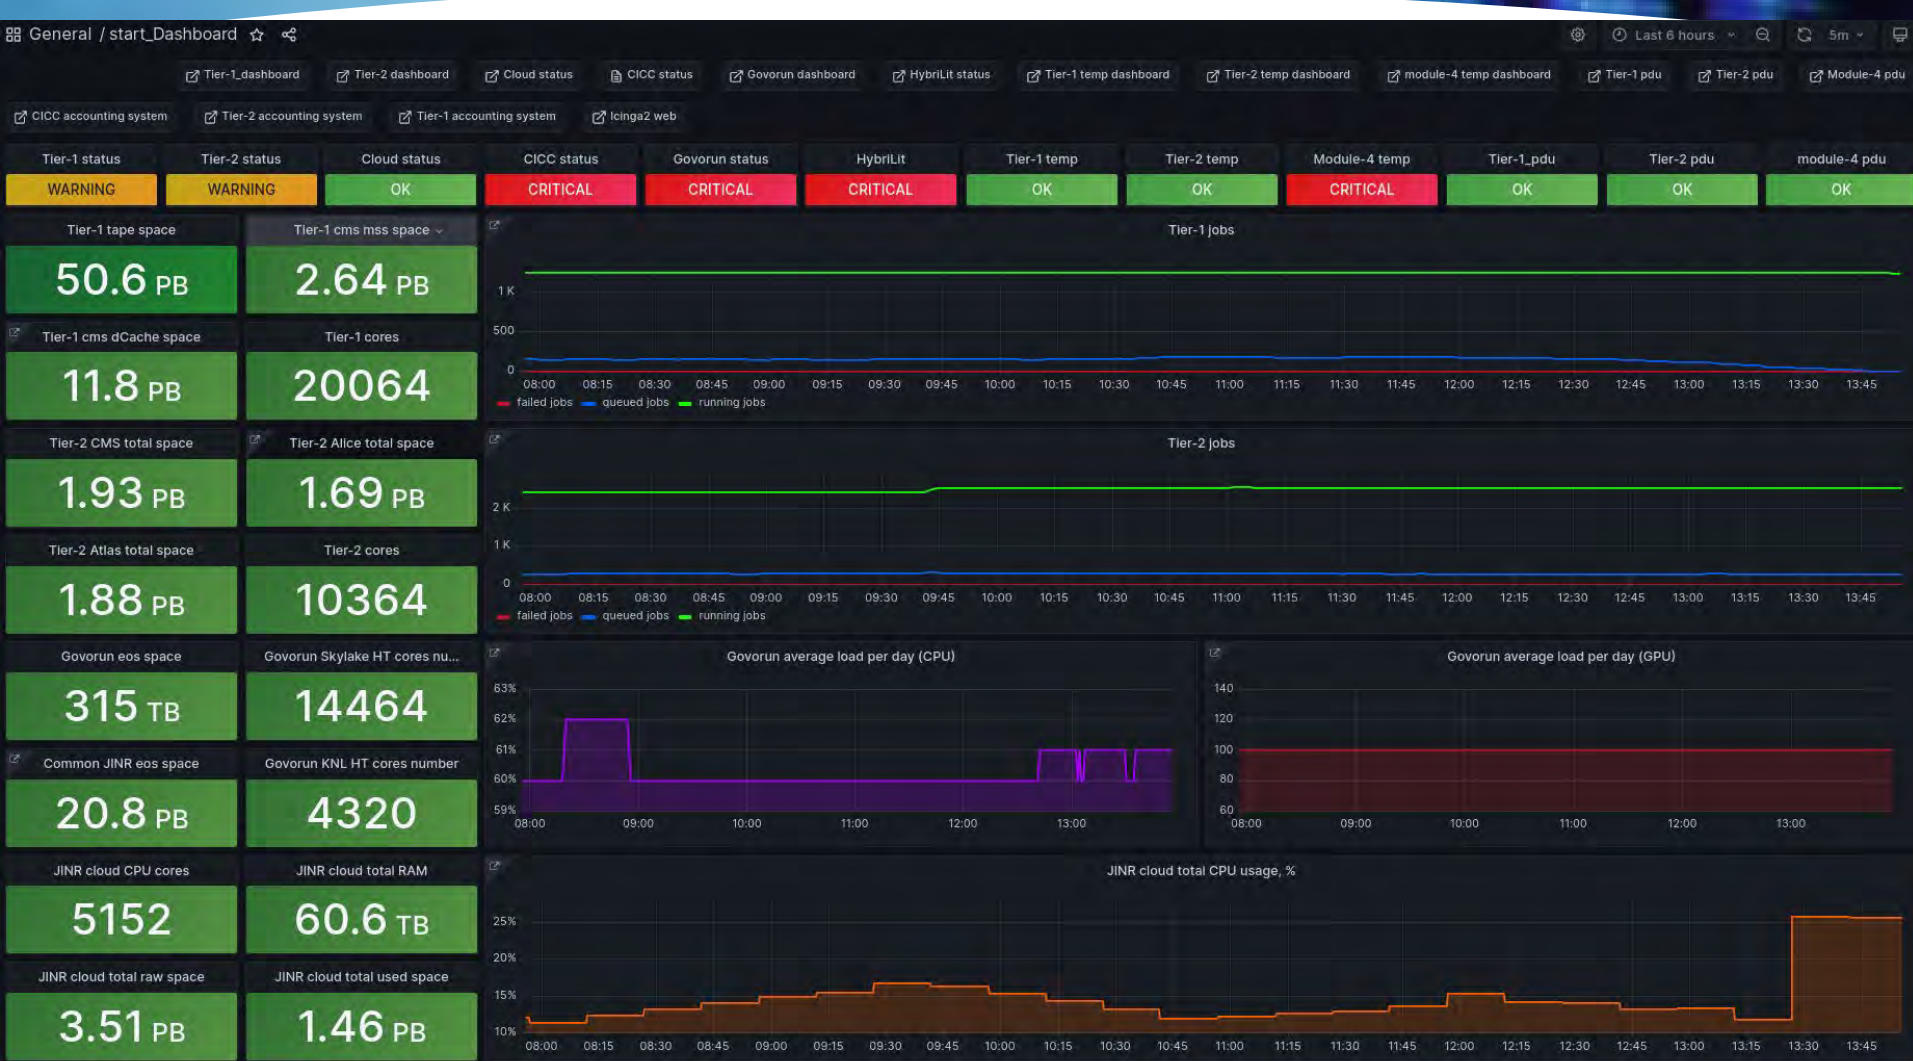


Tape library



UPS

LITmon – MICC monitoring system

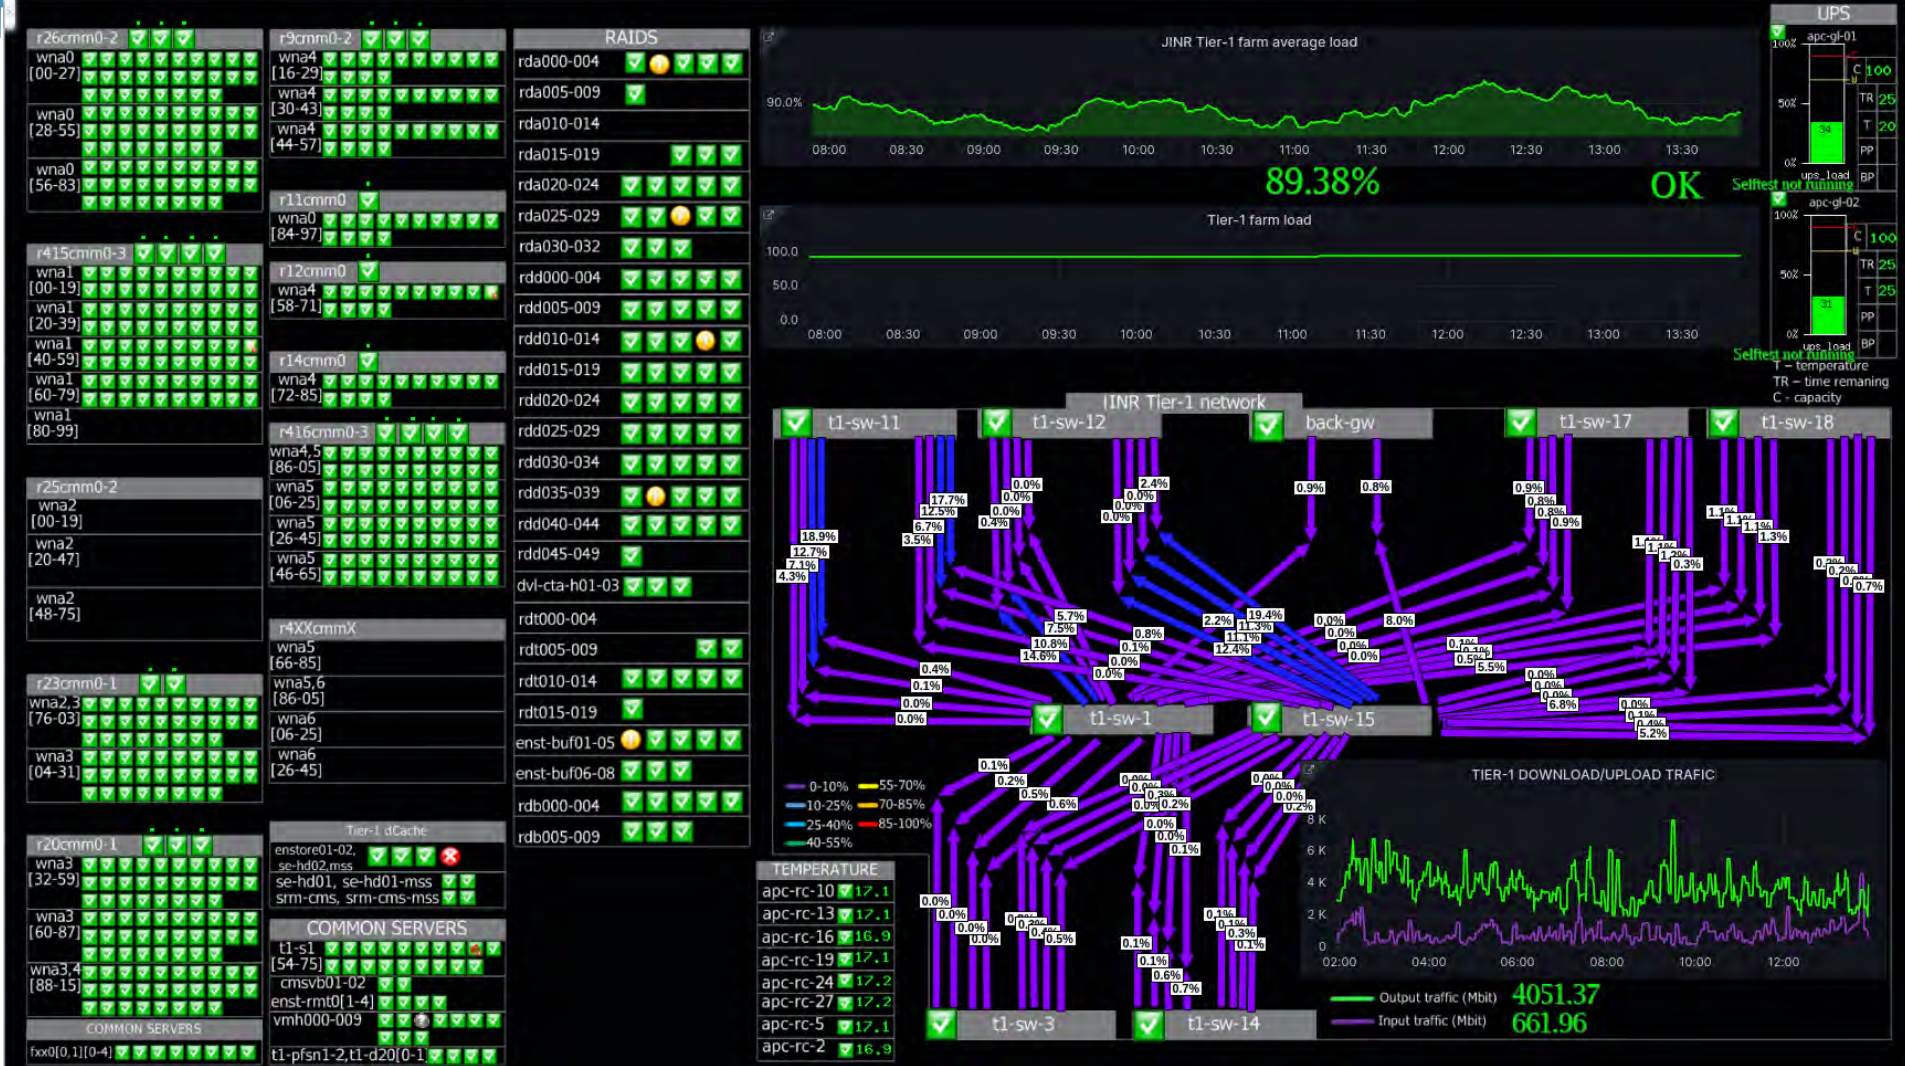


❖ 3 monitoring servers

❖ About 1800 nodes

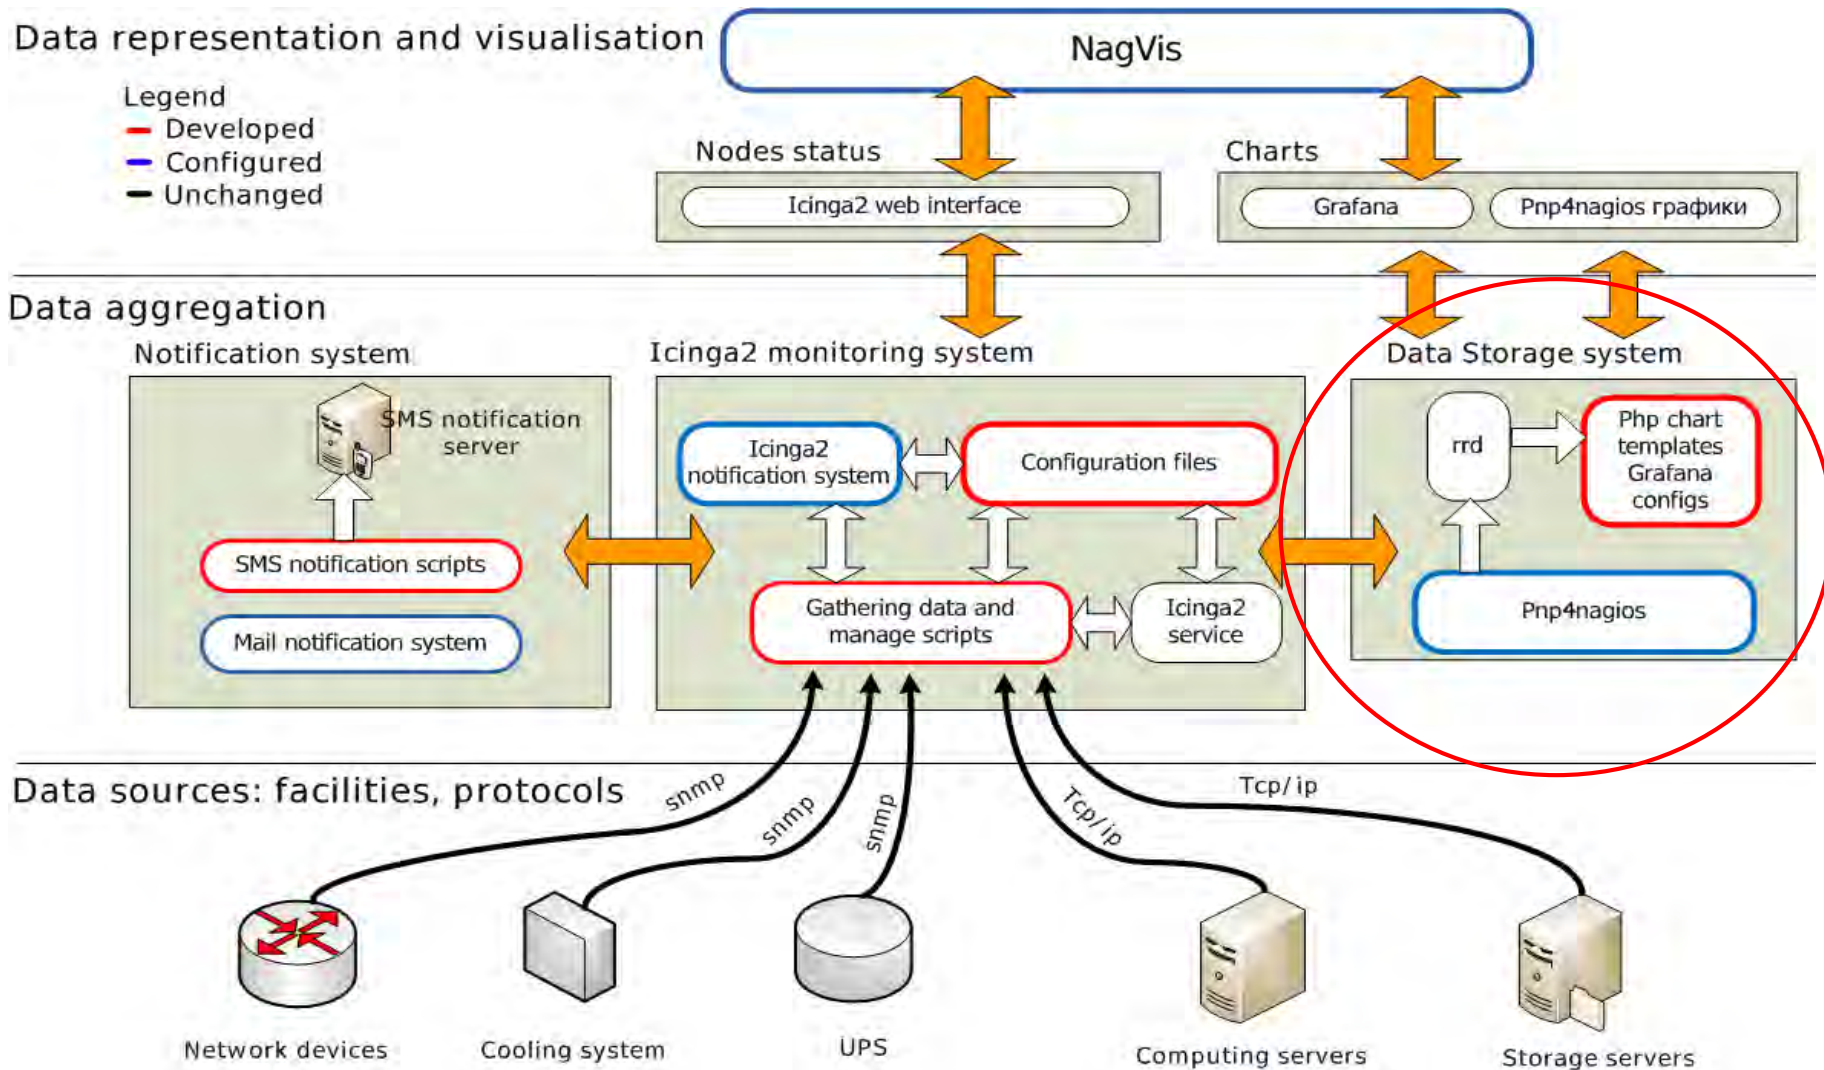
❖ About 16000 service checks

General view of the JINR CMS Tier-1 dashboard



Structural scheme of the monitoring system

LITmon




Monitoring data system analysis: The original system based on RRD



365 Gb data for 4 years + 495 GB spool files

Monitoring data system analysis : cons



More than 20 Gb RAM

About 5 Gb RAM

- ❖ The need for excess disk space and RAM to store charts.
- ❖ Low performance compared to competitors.
- ❖ Installing extra software.
- ❖ There is no option to enable charts automatically.
- ❖ Inability to replicate and save data.

Monitoring data system analysis : end of support

github.com/linge/pnp4nagios/blob/master/README

master pnp4nagios / README

Go to file ...

linge retired

Latest commit 5e09f53 on Jan 16, 2022 [History](#)

1 contributor

7 lines (4 sloc) 184 Bytes

Raw Blame   

```
1 This project is no longer active.  
2  
3 Thank you for the many years we have worked together on this project.  
4  
5 Feel free to take over the project.  
6  
7 This repository is no longer maintained!
```

ADIOS

Optimization of data aggregation and processing layer

Analysis of the monitoring data storage system

Evaluation of actual software solutions

Organize your data migration plan

Creation of data transfer software

Integration into the monitoring system


Evaluation of relevant solutions

Software name	Icinga2 support	Query language	Grafana support	Key features
Prometheus	+ -	PromQL	+	The monitoring system + time series database
Graphite	+	-	+	Time series database
InfluxDB 1.8	+	InfluQL (SQL)	+	Time series database
InfluxDB 2.6	+	Flux	+	Time series database
Elasticsearch	+	DSL	+	Text information database
Opentsdb	+	SQL	+	Time series database

Evaluation of relevant solutions: results

Product name	Cons
Prometheus	Lack of a developed query language. Icinga2 does not support by default.
Graphite	Low performance in read/write operations and query execution time (14 times worse than Influxdb2)
InfluxDB 1.8	An older version of Influxdb that will be deprecated over time.
InfluxDB 2.6	+
Elasticsearch	Low query execution speed for operations with the output of points, focus on working with text.
Opentsdb	Low performance in read/write operations and query execution time (9 times worse than Influxdb2)

Data migration



```
import rrdtool
```

```
from influxdb_client import InfluxDBClient, Point  
from influxdb_client.client.write_api import SYNCHRONOUS
```

```
[root@litmon-01 pnp4nagios]# ./rrdflux2.py -h  
Usage: rrdflux.py [-u|-m] -f <RRD FILE> [-H <INFLUXDB HOST>] [-p <INFLUXDB PORT>] -d DATABASE [-U user] [-P password] [-h]  
Updates or dumps passed RRD File to selected InfluxDB database  
-h, --help                Display help and exit  
-u, --update              Only update database with last value  
-m, --dump                Dump full RRD to database  
-f, --file                RRD file to dump  
-H, --host                Optional. Name or IP of InfluxDB server. Default localhost.  
-p, --port                Optional. InfluxDB server port. Default 8086.  
-d, --database            Database name where to store data.  
-U, --user                Optional. Database user.  
-P, --password            Optional. Database password.  
-s, --start                Start days ago.  
-e, --end                End date ego, e.g. now  
EXAMPLE: ./rrdflux2.py -H wna000.jinr-tl.ru -s -20d -e now
```

Operation algorithm of the monitored data storage system




Old

- ❖ Disk space - 860 Gb used
- ❖ RAM - 30 Gb required
- ❖ The need to create a new unix socket when adding a node
- ❖ Lack of data replication capability



influxdb




Grafana

New

- ❖ Disk space – 16 Gb used
- ❖ RAM – 5 Gb used
- ❖ Automatic creation adding new nodes to the database
- ❖ DB replication

Replication data system



```
object Influxdb2Writer "influxdb2" {  
  host = "159.93.227.19"  
  port = 8086  
  organization = "jinr"  
  bucket = "icinga2"
```

```
object Influxdb2Writer "influxdb2_backup" {  
  host = "159.93.227.250"  
  port = 8086  
  organization = "jinr"  
  bucket = "icinga2"
```

```
api checker icingadb ido-mysql influxdb2 influxdb2_backup
```

Optimization of the algorithm of the data storage system

Optimization of the amount of stored data:
intermediate stages of data processing are excluded

Data Security Guaranteed

Optimized data distribution algorithm, implemented
comfortable access to user monitoring system data

**A new system for storing traceable data by
the LITmon monitoring system has been put
into operation**



Thank you for your attention!



Abstract

One of the most important components of the monitoring system LITMon MICC LIT JINR is the data storage system. Initially, it was based on the RRD database and a special pnp4nagios plugin, support for which ended in 2022. Required features no longer work. The RRD database is morally obsolete and has ceased to meet performance requirements and has begun to consume more computing resources of the monitoring system server in comparison with analogues. Migrating data to a database based on InfluxDB software will solve these problems.