

# Hadoop

POW 2020

IT-DB-SAS

# Service in Numbers 2019

- > 99% of service availability
- HDFS size 11PB
  - 200MB/s of HDFS growth
  - 700MB/s writing rate, 500MB/s reading rate
  - 4kHz requests on namespace
- 1 GB/s ingest of structured data to HBase2 with NXCal
- ~ 1K jobs submitted everyday
- ~ 500 job containers running, ~20K finished every day

# Analytix (HDFS) - 2020

## ▼ HDFS Nodes Overview

Active Datanodes ▼


39

Dead Datanodes

0

## ▼ HDFS Storage

HDFS Usage



HDFS Capacity

13.4 PB

HDFS Used

6.049 PB

HDFS Growth Rate

424 MB/s

Number of Files Limit ▼



Number of Blocks

25.0 Mil

Block Write Rate

135 wps

# Analytix (HBASE) - 2020

## ↳ HBASE Nodes Overview

Active Region Servers


39

Dead Region Servers

0

## ↳ HBASE Database Size

Database Size Limit



Database Size


26.75 TB

Database Growth

14.4 MB/s

## ↳ HBASE Regions/Tables

Number of Regions Limit



Number of Regions

10.7 K

Number of Files

35.9 K

# NXCALS (HDFS) - 2020

ides Overview

Active Datanodes


30

Active Datanodes

0

Storage

HDFS Usage



HDFS Capacity

8.315 PB

HDFS Used

2.821 PB




HDFS Growth Rate

108 MB/s

Namespace

Number of Files Limit



Number of Files


72.8 Mil



Files Write Rate

346 wps

Number of Files Limit



Number of Blocks

63.1 Mil



Block Write Rate

436 wps

# NXCALS (HBASE) - 2020

## ▾ HBASE Nodes Overview

Active Region Servers

16

Dead Region Servers

0

## ▾ HBASE Database Size

Database Size Limit



9.8%

Database Size


9.754 TB

Database Growth

172 MB/s

## ▾ HBASE Regions/Tables

Number of Regions Limit



34%

Number of Regions

20.7 K

Number of Files

35.8 K

# Review of Activities 2019-2020

- Phase-out Cloudera [DONE]
  - Announce to users - a dedicated presentation/session needed?
  - Validation by users (QA, Hadalytic)
  - Perform the upgrade
- Migration to Hadoop3 [PARTIALLY DONE]
  - Analytix, Lxhadoop, QA - Hadoop3
  - NXCALS (Prod and Dev) [ONGOING]
- Extend the web portal with [PARTIALLY DONE]
  - compute resource requesting
  - integration with metrics
  - panel for administrators
- Security: deploy SSL and authorisation on all web UIs [PARTIALLY DONE]
  - Not Deployed but Urgent
- Move monitoring scripts to Nomad [NOT DONE]
  - Canary test
  - SLS
  - Shift notifications (including SMS)
- Consolidation of clusters [PARTIALLY DONE]
  - Merging LXHADOOP and ANALYTIX [DONE]
  - HBASE federation [MISSING]

# Review of Activities 2019-2020

- Test Scalable database [DONE]
  - OpenTSDB
  - Further test will be required
  - Identification of use cases
- HBASE Snaphots [DONE]
  - Missing Deployment on NXCALS
- Productionizing Apache Phoenix [USED BY EVENT INDEXING]
  - Software and configuration deployment
  - Monitoring
- Addapt and Report to OC11 [DONE]



# Short term plans 2020-2021

- **NXCALS**

- Install new Hardware
  - slots in critical area (Hbase) ?
- Scale Hbase
  - Does Hbase require to be in the critical area - if so find alternatives to scale.
- Isolate NXCALS
  - Specific protected environment for DEV and PROD
- Hbase Disaster Recovery Strategy
  - Snapshots
  - Replication (Master-to-Master, Master-to-Slave)
  - Copy tables
- Reduce Hbase Complexity
  - Number of Regions/Tables
  - Deploy automatic procedures
- Review and improve Monitoring
  - Current Dashboards
  - HBACK for inconsistencies
  - Log for Known issues

# Short term plans 2020-2021

- **NXCALS**

- Update to Hadoop3
- Update to Hbase 2.2.4 - 3.0...
- Implementation of Limits
  - HDFS Storage
  - HDFS Namespace Objects/Files
  - Hbase Regions, writes/read ops, size
- Review and implement (if required) a NXCALS dedicated QA environment
  - BE-CO will have ideally to target there a few ETLs processes
- Definition and Implementation of the Piquet Service Specific for NXCALS
- Structured operations guidelines
  - Many operations are not documented, we need to have a clear guide to apply recipes on common problems
  - Critical to reduce impact of potential problems and also to remove dependencies
- Intervention/Modification formal model definition

# Short term plans 2020-2021

- **General**

- Security: deploy SSL and authorisation on all web UIs
  - Not Deployed but Urgent
- Consolidation of clusters LxHadoop Analytic
  - HBASE federation
- Move QA environment from bare metal to Openstack
  - Reduce loads from IT-Monit and CERT
- Extend the web portal
  - compute resource requesting
  - integration with metrics
  - panel for administrators
- Move monitoring scripts to Nomad
  - Canary test
  - SLS
  - Shift notifications (including SMS)
- Productionizing Apache Phoenix
  - Software and configuration deployment
  - Monitoring

# Long term 2021 - ...

- Data Access Abstraction Layers
  - Presto,...
- Integrate Hadoop, Spark and Kafka services
- Federated namespaces vs object store
  - Evaluate Apache Ozon - object store on top of hadoop infrastructure
- Towards Scalable Database Service
  - Evaluate alternative HBase vs Ignite vs Redis vs Cassandra - as fast distributed database
  - Scalable Time Series Databases (alternative to Influx)
- Hive vs Iceberg - as data metastore
- High level services
  - Druid-like system? Automated Kafka -> HBase -> HDFS
- Re-Evaluate NXCALS Hbase model
- Integrate Cloud resources