Implementation of a modern logging pipeline for Hadoop

CERN openlab Summer Student Programme Lightning talks

Jordie Biemold

Supervisor: Luis Pigueiras

18 August 2026

What is “logging”?

Short dated messages about important events to help developers and system operators understand what is going on in the system

Logs show:

- **When** something happened
- **How serious** (error or warning)
- **What** happened
- **Where** or **why** it happened

```
INFO  <2026-08-10T16:48:15.220+0200> <INFO> <authorize.ServiceAuthorizationManager>: <Authorization successful for [redacted] (auth:TOKEN) for protocol=interface org.
INFO  <2026-08-10T16:48:15.219+0200> <INFO> <hbase.Server>: <Auth successful for [redacted] (auth:TOKEN)>
INFO  <2026-08-10T16:48:15.210+0200> <INFO> <hbase.Server>: <Connection from [redacted] version=2.3.4, sasl=true, ugi=[redacted] (auth:TOKEN), service=ClientS
INFO  <2026-08-10T16:48:15.210+0200> <INFO> <authorize.ServiceAuthorizationManager>: <Authorization successful for [redacted] (auth:TOKEN) for protocol=interface org.
INFO  <2026-08-10T16:48:15.209+0200> <INFO> <hbase.Server>: <Auth successful for [redacted] (auth:TOKEN)>
WARN  <2026-08-10T16:48:15.003+0200> <WARN> <org.apache.hadoop.hdfs.DFSUtil>: <Unexpected value for data transfer bytes=8427 duration=0>
```

The problem

Identified:

- 10 applications/services
- 20 important log files
- 10 different log file formats

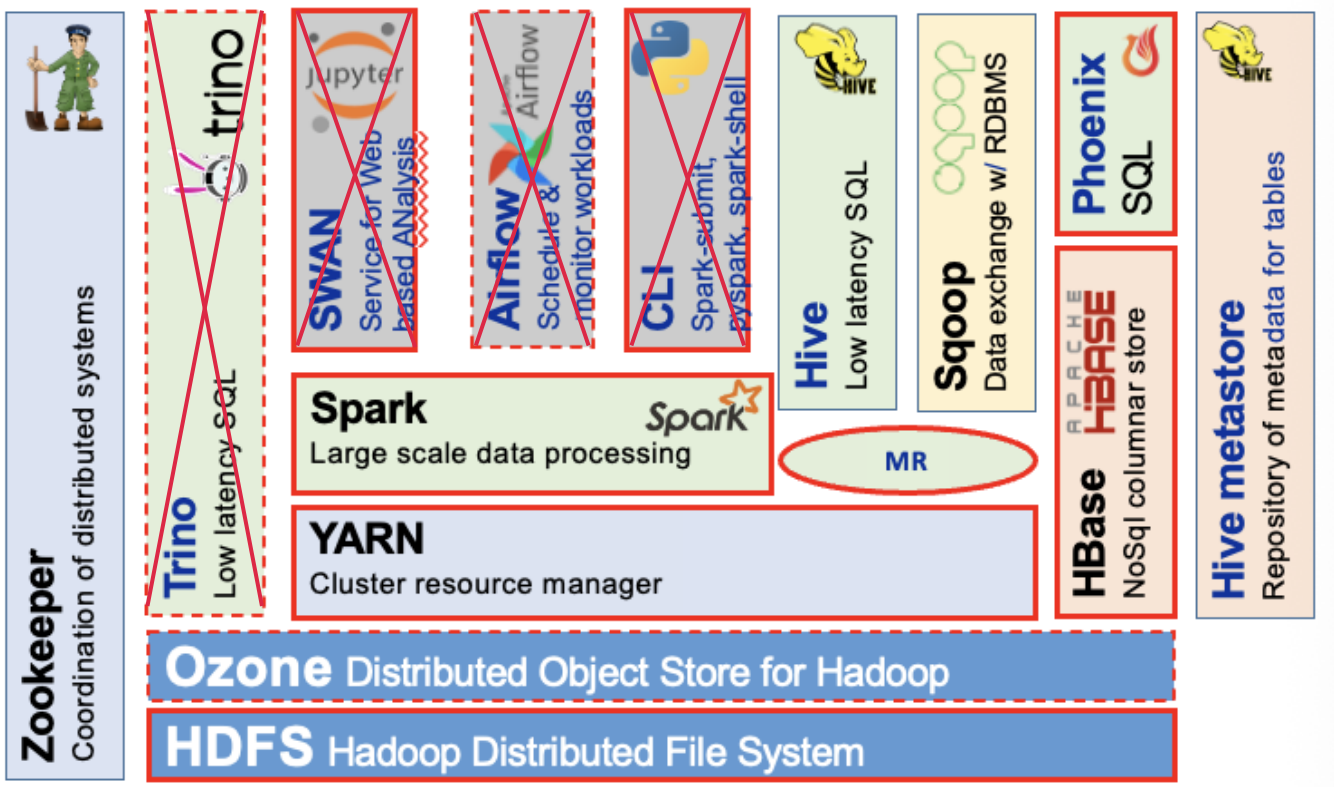

Imagine what happens when something goes wrong

Go through 5-20 files spread across servers (around 210) by hand!

Solution: Fluent Bit to OpenSearch

Centralising the storage of logs

Fluent Bit
Collects the logs from application services running on servers

OpenSearch
Stores the logs in a central place

Grafana
Dashboards for insights

Architecture diagram

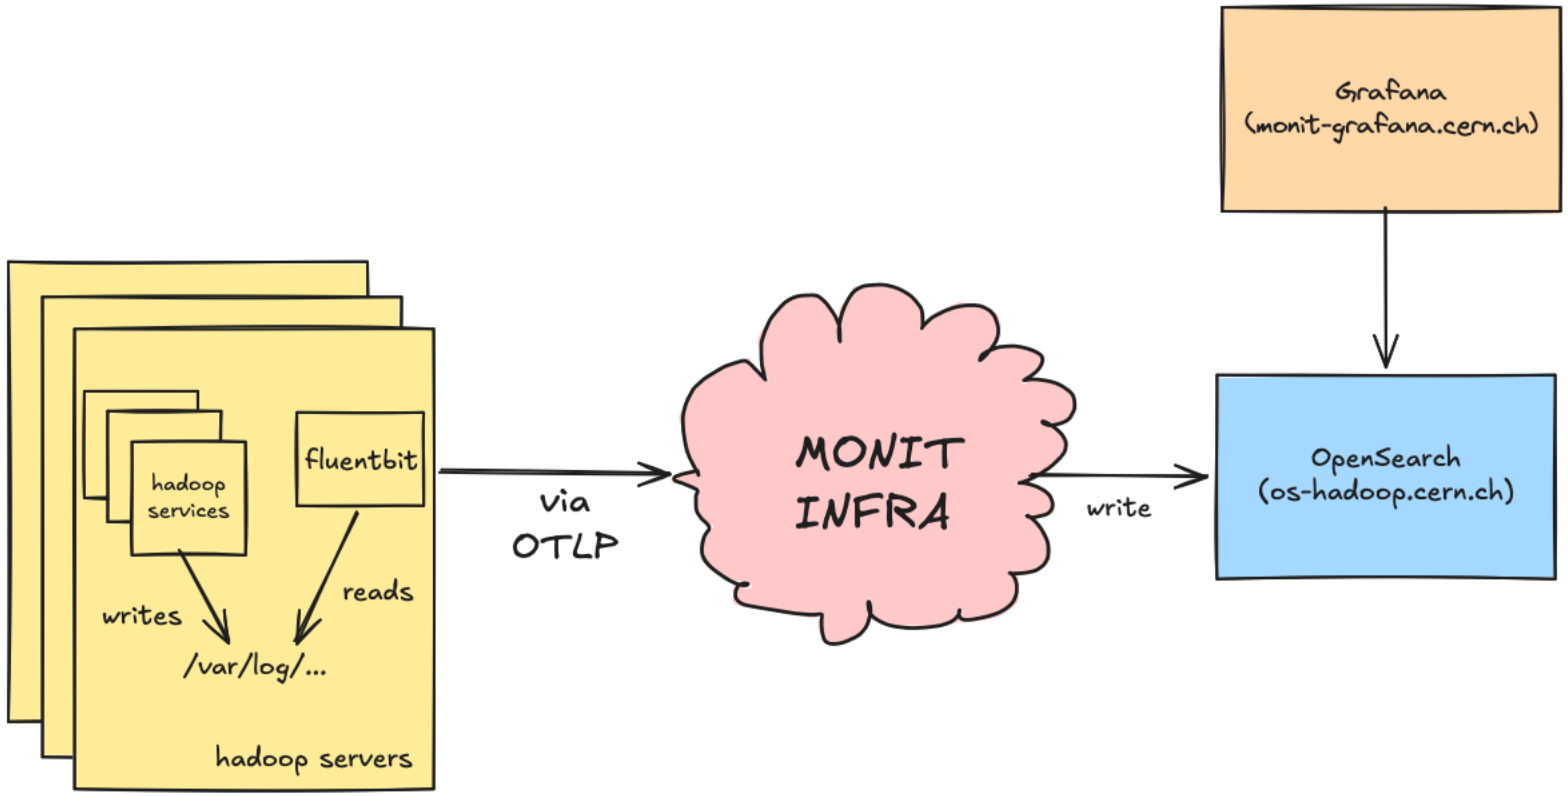

Example flow

```
[root@ithdp-hadoop-qa-10 ~]# ls -la /var/log/hadoop-hdfs
total 107132
drwxrwxr-x.  2 hdfs hdfs     4096 Aug 12 02:20 .
drwxr-xr-x. 22 root root     4096 Aug 11 05:25 ..
-rw-------.  1 hdfs hdfs  6325572 Aug 12 11:52 hadoop-hdfs-hdfs-datanode-ithdp-hadoop-qa-10.cern.ch.log
-rw-------.  1 hdfs hdfs 10063359 Aug  6 23:59 hadoop-hdfs-hdfs-datanode-ithdp-hadoop-qa-10.cern.ch.log.2026-08-06
-rw-------.  1 hdfs hdfs 15128600 Aug  7 23:59 hadoop-hdfs-hdfs-datanode-ithdp-hadoop-qa-10.cern.ch.log.2026-08-07
-rw-------.  1 hdfs hdfs  9713941 Aug  8 23:59 hadoop-hdfs-hdfs-datanode-ithdp-hadoop-qa-10.cern.ch.log.2026-08-08
-rw-------.  1 hdfs hdfs 15208922 Aug  9 23:59 hadoop-hdfs-hdfs-datanode-ithdp-hadoop-qa-10.cern.ch.log.2026-08-09
-rw-------.  1 hdfs hdfs  9626661 Aug 10 23:59 hadoop-hdfs-hdfs-datanode-ithdp-hadoop-qa-10.cern.ch.log.2026-08-10
-rw-------.  1 hdfs hdfs 15027399 Aug 11 23:59 hadoop-hdfs-hdfs-datanode-ithdp-hadoop-qa-10.cern.ch.log.2026-08-11
-rw-------.  1 root root      125 Jul 15 12:18 hadoop-hdfs-hdfs-datanode-ithdp-hadoop-qa-10.cern.ch.out
```

Grafana visualises data

Log files read by Fluent Bit (20+ files, 10 formats)

```
[root@ithdp-hadoop-qa-10 ~]# cat /etc/fluent-bit/hadoop-agent/input_hdfs-dn.conf
# Managed by puppet
[INPUT]
    Name              tail
    Tag               hdfs-dn
    Path              /var/log/hadoop-hdfs/hadoop-hdfs-datanode-*.log,/var/log/hadoop-hdfs/hadoop-hdfs-hdfs-datanode-*.log
    Path_Key          log.file.path
    Refresh_Interval  5
    Skip_Empty_Lines  On
    Skip_Long_Lines   Off
    DB                /etc/fluent-bit/hadoop-agent/db/hdfs-dn.db
    Multiline.Parser  java1
[root@ithdp-hadoop-qa-10 ~]# cat /etc/fluent-bit/hadoop-agent/filter_10-parse-hdfs-dn.conf
# Managed by puppet
[FILTER]
    Name  parser
    Match hdfs-dn
    Parser java1
    Key_name log
    Preserve_Key On
    Reserve_Data On
[root@ithdp-hadoop-qa-10 ~]# cat /etc/fluent-bit/hadoop-agent/filter_20-enrich-hdfs-dn.conf
# Managed by puppet
[FILTER]
    Name  modify
    Match hdfs-dn
    Add service.name hdfs-dn
    Add service.type hdfs
    Add monit_type hdfs-dn
```

… and the Fluent Bit OpenTelemetry output plugin

Fluent Bit sends data to OpenSearch

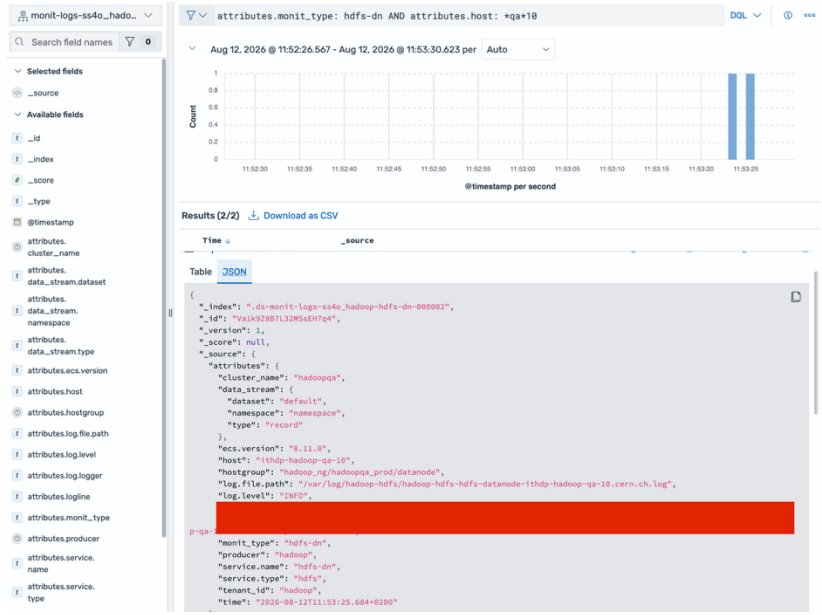

Grafana insights

Not all warnings and errors are important

Some will have to be filtered out (future task)

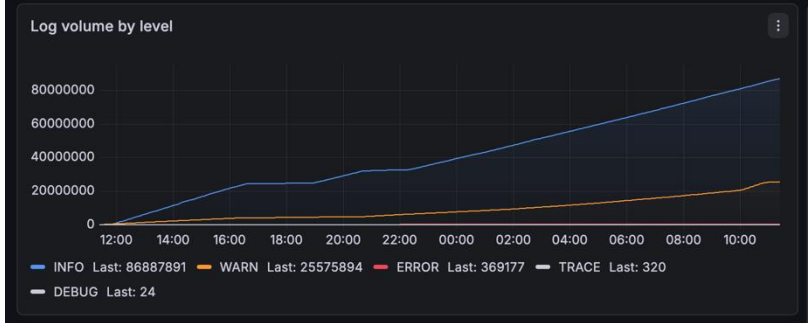

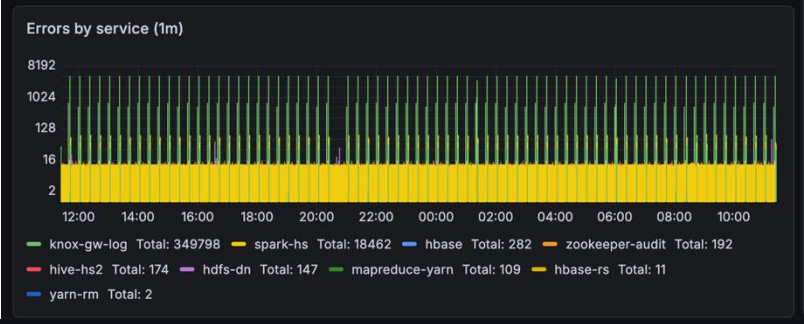

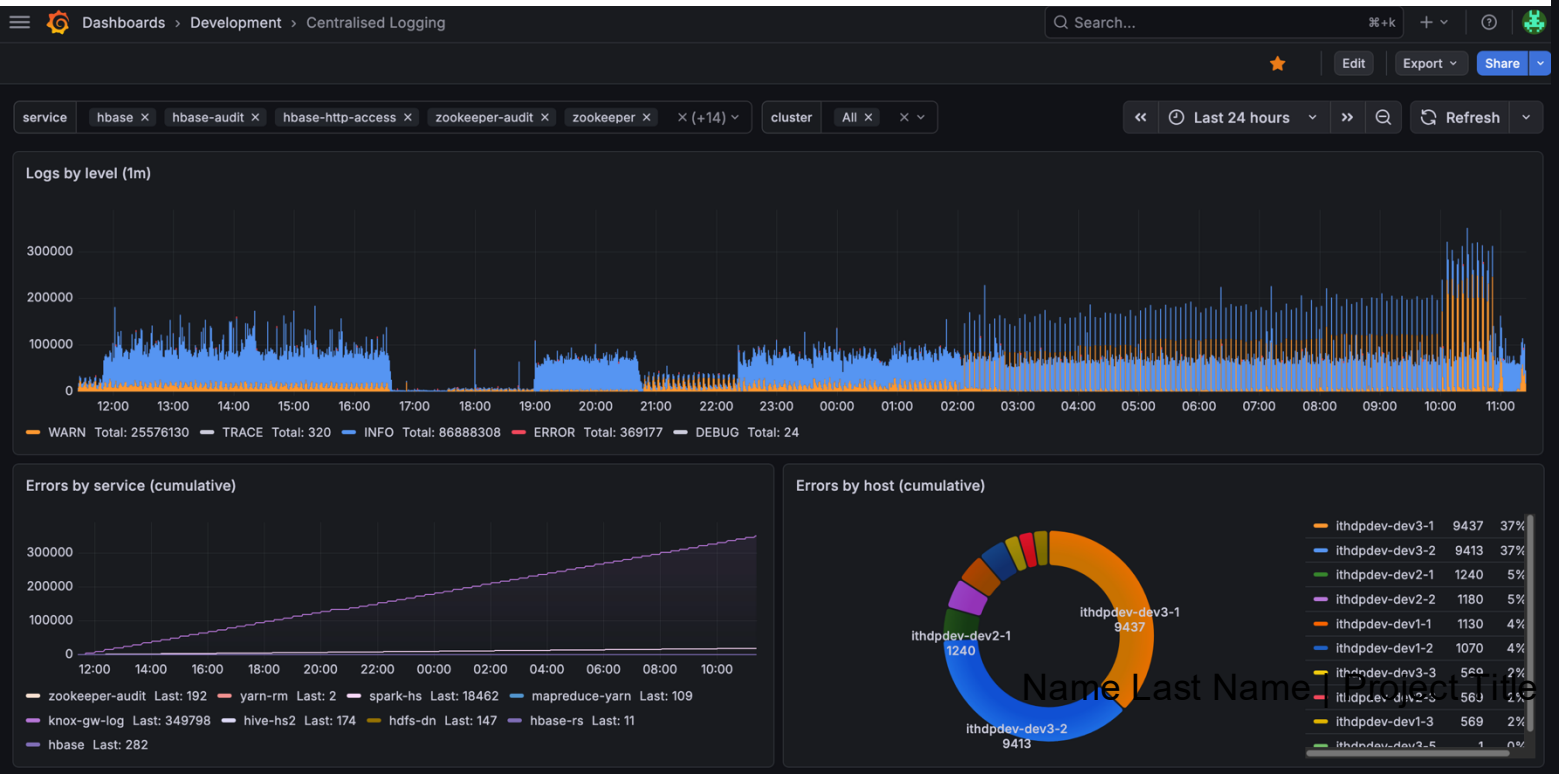

Top erroneous loggers

Service	Level	Count	Logger
hdfs-dn	WARN	25081448	org.apache.hadoop.hdfs.DFSUtil
hdfs-dn	WARN	456432	org.apache.hadoop.hdfs.server.datanode.DataNode
knox-gw-lo	ERROR	349371	knox.gateway
hdfs-dn	WARN	22551	org.apache.hadoop.hdfs.server.datanode.fsdataset.impl.FsDatasetIm
spark-hs	ERROR	17072	FsHistoryProvider
hbase	WARN	4547	snapshot.SnapshotFileCache
knox-gw-lo	WARN	2130	service.metadata
knox-gw-lo	WARN	1719	knox.gateway
hbase-rs	WARN	1713	regionserver.HeapMemoryManager
hbase	WARN	1478	zookeeper.Login
spark-hs	ERROR	1390	ServerRuntime$Responder
hdfs-dn	WARN	486	org.apache.hadoop.fs.CachingGetSpaceUsed
hbase	WARN	468	server.HttpChannelState
hbase	WARN	468	server.HttpChannel

1270+ logs/second

In development environment (30 servers)

After 3 days OpenSearch was full

Too many log messages 😰

[redacted] 🌴 09:27

don't worry about hadoop, I'll address it

5

[redacted] 09:32

> what do you mean?

[redacted] 🌴 10:06

>
>
Time	alertname	cluster	bundle	host	severity	description	rity	alert_group	value	mountpoint
3 hours ago	Disk full		[redacted]		critical	Detected mountpoint with less than 10% free	1	node_metrics	7.45	/var/lib/opensearch/hadoop1_master
40 minutes ago	Disk full				critical	Detected mountpoint with less than 10% free	1	node_metrics	9.65	/var/lib/opensearch/hadoop1_master
3 hours ago	High disk watermark	hadoop1			critical	High disk watermark is reached on the following data nodes.		opensearch	92.4	
an hour ago	High disk watermark	hadoop1			critical	High disk watermark is reached on the following data nodes.	1	opensearch	90.9	

[redacted] 10:07

> ahh nice thanks

System 10:07

You were **added to the channel** by [redacted]

[redacted] 10:07

5

Commented on [redacted]s message: **don't worry about hadoop, I'll address it**

> @Jordie Biemold is testing to the limit 😄
>
> 🤪 1
>
> 🥠

[redacted] 🌴 10:26

I'll change the retention to 3d instead of 30d until we have all the logs and we can properly estimate how big of a cluster we will need

👌 1 🐡 1

severity	description
critical	Detected mountpoint with less than 10% free
critical	Detected mountpoint with less than 10% free
critical	High disk watermark is reached on the following data nodes.
critical	High disk watermark is reached on the following data nodes.

Task for the system operators: configure the logging granularity

It takes time to figure out which logs are useful to keep around, and for how long

Lessons learned

Challenges

- Working with new technologies and the systems at CERN
- Learning the Puppet deployment workflow (different from Terraform which I used before)
- Figuring out how to put everything together in an environment I haven't setup myself

What I learned

- Working in a complex system where a lot of services are interconnected
- Fluent Bit (advanced) and OpenSearch
- Creating dashboards in Grafana
- Some basic particle physics 😄

Thank you!

(and many thanks to my supervisor and colleagues)

Jordie Biemold

18 August 2026

Copyright

- CERN Hadoop overview from presentation “Introduction to Hadoop, Spark and Big Data at CERN” (2025)
- Logos courtesy of Fluent Bit by CNCF, OpenSearch Software Foundation, Grafana Labs
- Picture: pt [dot] memedroid [dot] com/memes/detail/287212