The Analysis Facility at GSI

Sören Fleischer Raffaele Grosso

GSI

September 27, 2022

Outline

- The ALICE T2 at GSI
  - the cluster
  - software solutions
  - current issues
- Prospects for the ALICE Analysis Facility at GSI
  - AF requirements
  - current state

The Cluster - Hardware

Resources reserved on a shared cluster:

- $\sim$13k logical CPU cores (hyperthreading):

  ```
  PARTITION   CPUS   S:C:T    MEMORY    NODES
  grid        96     2:24:2   191762    134
  ```

  under a SLURM workload manager

The Cluster - Hardware

Resources reserved on a shared cluster:

- ∼13k logical CPU cores (hyperthreading):

PARTITION	CPUS	S:C:T	MEMORY	NODES
grid	96	2:24:2	191762	134

  under a SLURM workload manager
- 5.2 PB (in 2022 ⇒ 6.1 PB in 2023) disk storage under a Lustre distributed file system

The Cluster - Hardware

Resources reserved on a shared cluster:

- ∼13k logical CPU cores (hyperthreading):

PARTITION	CPUS	S:C:T	MEMORY	NODES
grid	96	2:24:2	191762	134

  under a SLURM workload manager
- 5.2 PB (in 2022 ⇒ 6.1 PB in 2023) disk storage under a Lustre distributed file system
- Network connection
  - internally 100 Gb/s EDR Infiniband
  - 10 Gb/s LHCONE, 2 Gb/s DFN

The Cluster - Hardware

Resources reserved on a shared cluster:

- ∼13k logical CPU cores (hyperthreading):

  ```
  PARTITION   CPUS   S:C:T    MEMORY    NODES
  grid        96     2:24:2   191762    134
  ```

  under a SLURM workload manager
- 5.2 PB (in 2022 ⇒ 6.1 PB in 2023) disk storage under a Lustre distributed file system
- Network connection
  - internally 100 Gb/s EDR Infiniband
  - 10 Gb/s LHCONE, 2 Gb/s DFN
- Memory limits imposed by SLURM via `cgroups`
  - limit is set on `PSS`, thus correctly taking shared memory into account (4.4 GB per physical core)
  - we don't set virtual memory limits

Software

Jobs run within Singularity containers.

- Host: minimal Red Hat Enterprise Linux compatible installation

Software

Jobs run within Singularity containers.

- Host: minimal Red Hat Enterprise Linux compatible installation
- Container:
  - used to be our own image
  - since the move to JAliEn-VOBoxes this is handled by JAliEn, using the Singularity runtime engine and image from `/cvmfs/alice.cern.ch/`.

Software

Jobs run within Singularity containers.

- Host: minimal Red Hat Enterprise Linux compatible installation
- Container:
  - used to be our own image
  - since the move to JAliEn-VOBoxes this is handled by JAliEn, using the Singularity runtime engine and image from `/cvmfs/alice.cern.ch/`.
- In the last months all jobs have been going through the GSI-8Core queue
  - JobAgents are run as SLURM jobs with 8 cores, which they fill with single- and multi-core AliEn jobs

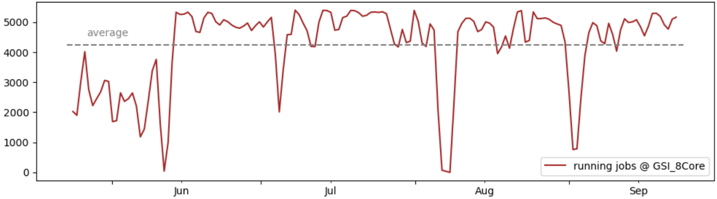

Average: $\sim$4.3k jobs from Hyperloop trains, max 5600 running jobs

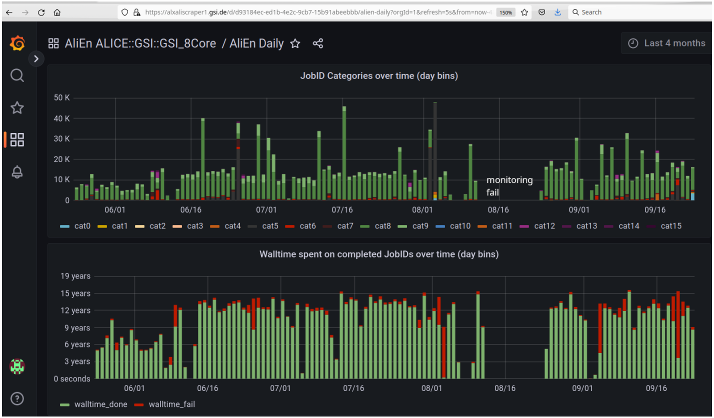

https://alxaliscraper1.gsi.de/d/d93184ec-ed1b-4e2c-9cb7-15b91abeebbb/alien-daily?orgId=1&refresh=5s&from=now-4
150%
Search
AliEn ALICE::GSI::GSI_8Core / AliEn Daily
Last 4 months
JobID Categories over time (day bins)
50 K
40 K
30 K
20 K
10 K
0
06/01
06/16
07/01
07/16
08/01
08/16
09/01
09/16
monitoring fail
cat0
cat1
cat2
cat3
cat4
cat5
cat6
cat7
cat8
cat9
cat10
cat11
cat12
cat13
cat14
cat15
Walltime spent on completed JobIDs over time (day bins)
19 years
15 years
12 years
9 years
6 years
3 years
0 seconds
06/01
06/16
07/01
07/16
08/01
08/16
09/01
09/16
walltime_done
walltime_fail

XRootD Plug-Ins

XRootD compiled with following self-developed plug-ins:

- symlink: locally creates an LFN-symlink pointing to the PFN ⇒ allows local access to files via LFN
- quota: calls `lfs quota` ⇒ allows MonaLisa to get the correct values for available storage from a shared distributed file system
- local redirect: see description in next slides ⇒ optimizes the I/O throughput of the analysis jobs

High-level block diagram

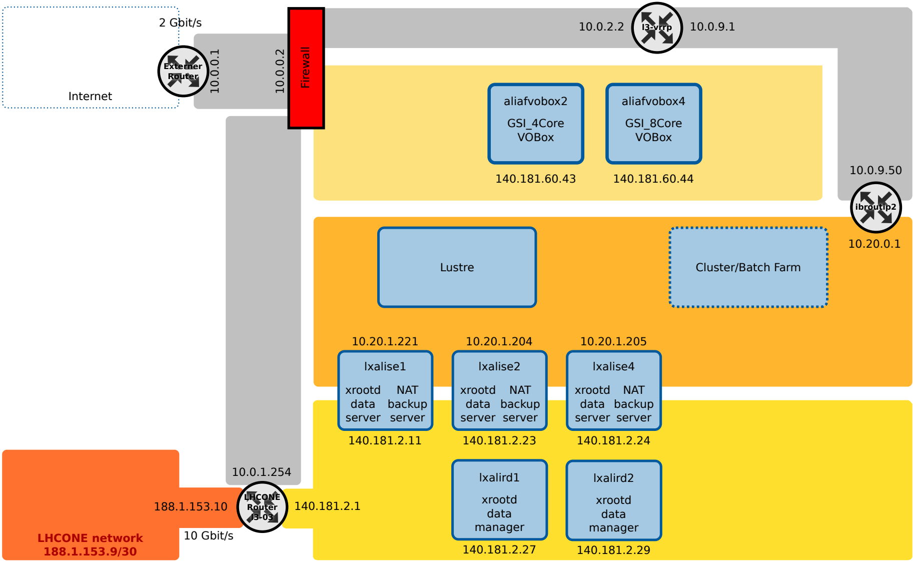

xrootd Local Redirect Plugin

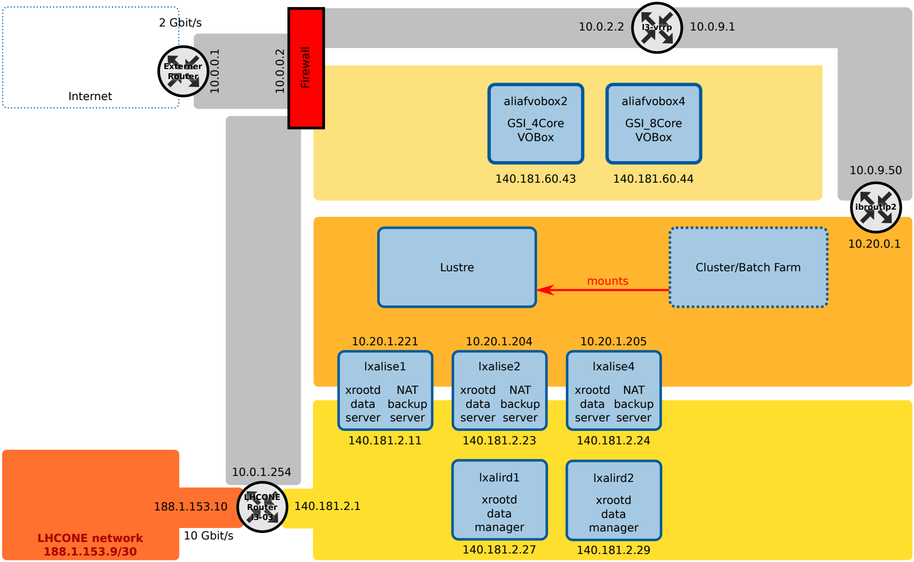

xrootd Local Redirect Plugin

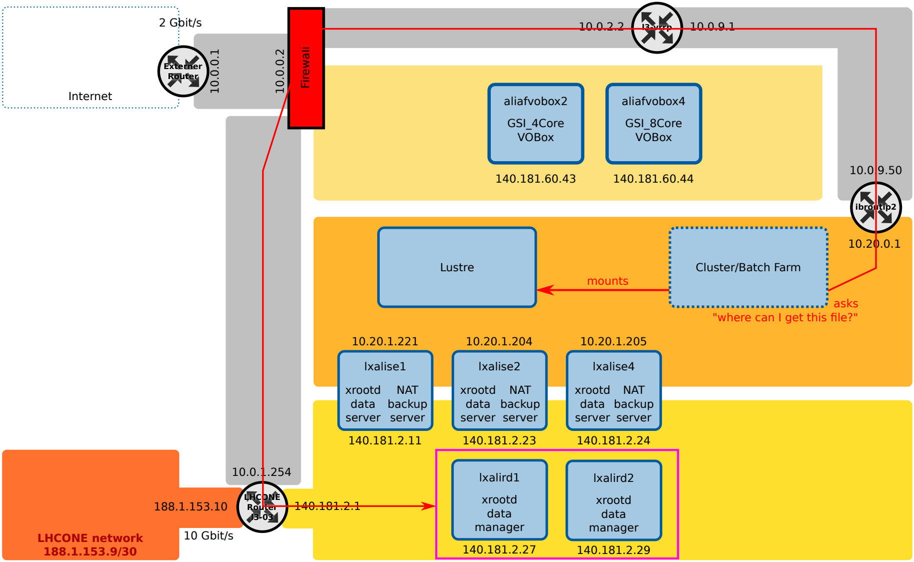

xrootd Local Redirect Plugin

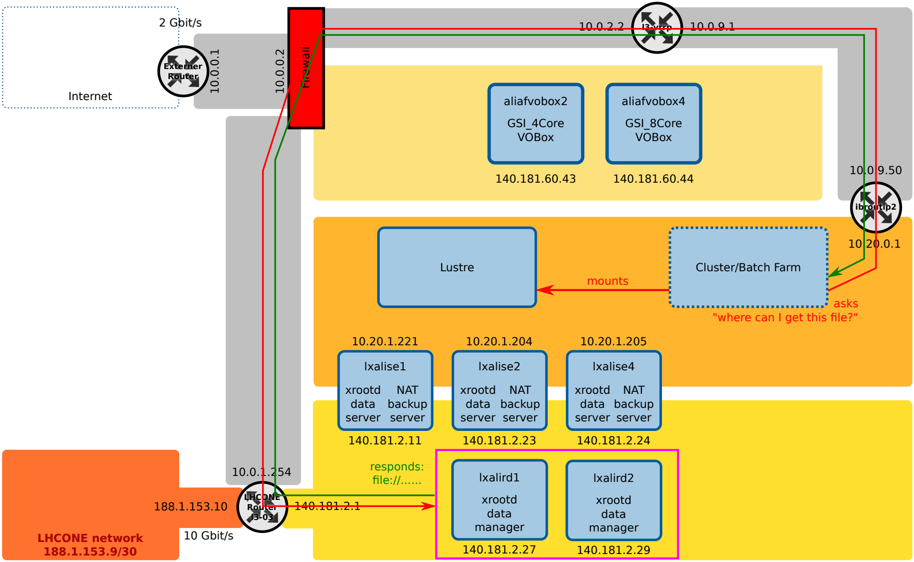

xrootd Local Redirect Plugin

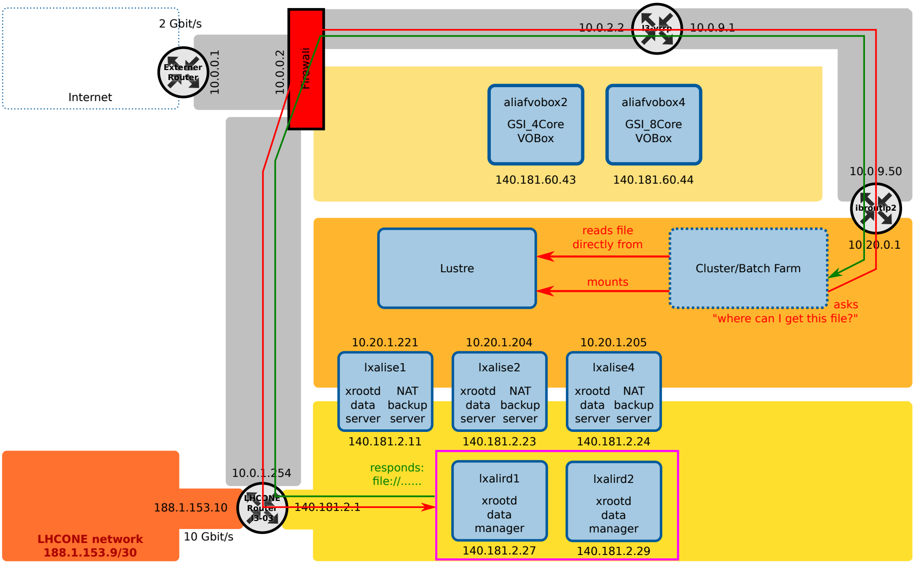

2 Gbit/s

Externer Router

Internet

10.0.0.1

10.0.0.2

Firewall

10.0.2.2

l3-vrrp

10.0.9.1

aliafvobox2

GSI_4Core VOBox

140.181.60.43

aliafvobox4

GSI_8Core VOBox

140.181.60.44

10.0.9.50

ibroutip2

10.20.0.1

Lustre

reads file directly from

mounts

Cluster/Batch Farm

asks "where can I get this file?"

10.20.1.221

lxalise1

xrootd data server

NAT backup server

140.181.2.11

10.20.1.204

lxalise2

xrootd data server

NAT backup server

140.181.2.23

10.20.1.205

lxalise4

xrootd data server

NAT backup server

140.181.2.24

10.0.1.254

responds: file://......

LHCONE Router l3-03

188.1.153.10

140.181.2.1

10 Gbit/s

LHCONE network 188.1.153.9/30

lxalird1

xrootd data manager

140.181.2.27

lxalird2

xrootd data manager

140.181.2.29

xrootd Local Redirect Plugin

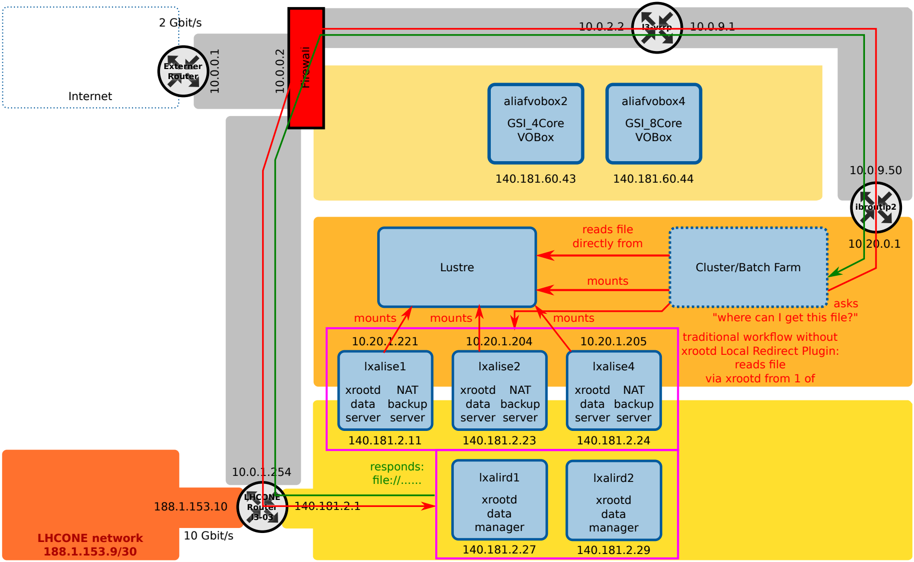

alinat

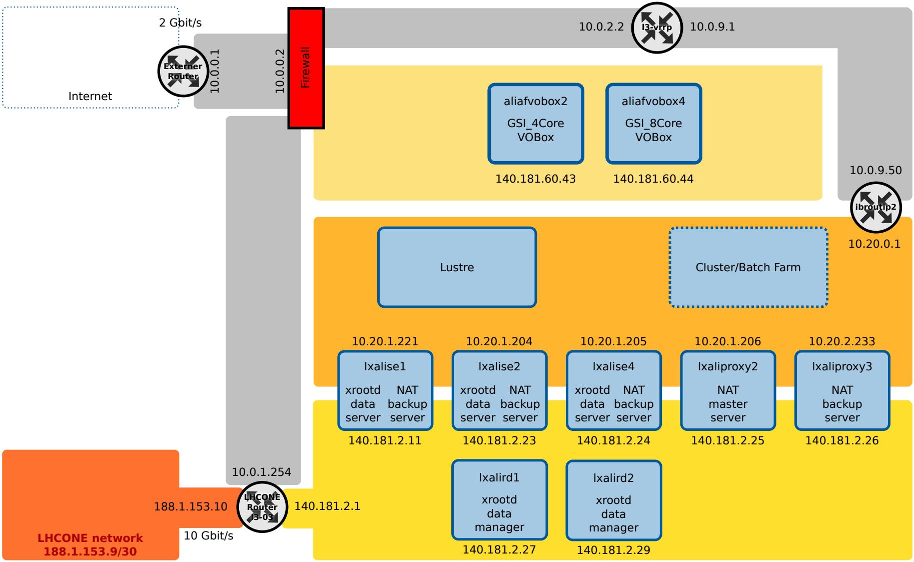
2 Gbit/s
Externer Router
Internet
10.0.0.1
10.0.0.2
Firewall
10.0.2.2
l3-vrrp
10.0.9.1
aliafvobox2
GSI_4Core VOBox
140.181.60.43
aliafvobox4
GSI_8Core VOBox
140.181.60.44
10.0.9.50
ibroutip2
10.20.0.1
Lustre
Cluster/Batch Farm
10.20.1.221
lxalise1
xrootd data server
NAT backup server
140.181.2.11
10.20.1.204
lxalise2
xrootd data server
NAT backup server
140.181.2.23
10.20.1.205
lxalise4
xrootd data server
NAT backup server
140.181.2.24
10.20.1.206
lxaliproxy2
NAT master server
140.181.2.25
10.20.2.233
lxaliproxy3
NAT backup server
140.181.2.26
10.0.1.254
LHCONE Router l3-03
188.1.153.10
140.181.2.1
10 Gbit/s
LHCONE network
188.1.153.9/30
lxalird1
xrootd data manager
140.181.2.27
lxalird2
xrootd data manager
140.181.2.29

alinat

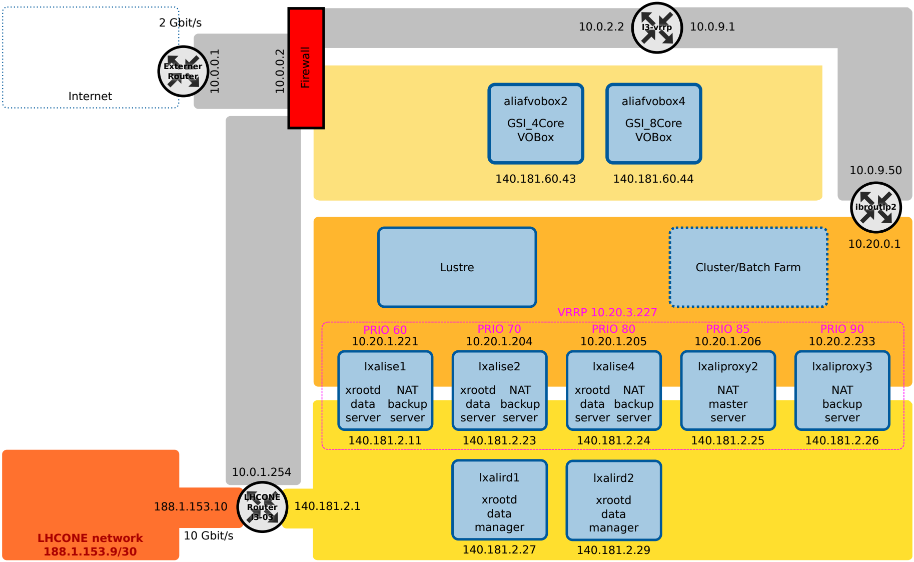

2 Gbit/s
Externer Router
Internet
10.0.0.1
10.0.0.2
Firewall
10.0.2.2
l3-vrrp
10.0.9.1
aliafvobox2
GSI_4Core VOBox
140.181.60.43
aliafvobox4
GSI_8Core VOBox
140.181.60.44
10.0.9.50
ibroutip2
10.20.0.1
Lustre
Cluster/Batch Farm
VRRP 10.20.3.227
PRIO 60
10.20.1.221
lxalise1
xrootd data server
NAT backup server
140.181.2.11
PRIO 70
10.20.1.204
lxalise2
xrootd data server
NAT backup server
140.181.2.23
PRIO 80
10.20.1.205
lxalise4
xrootd data server
NAT backup server
140.181.2.24
PRIO 85
10.20.1.206
lxaliproxy2
NAT master server
140.181.2.25
PRIO 90
10.20.2.233
lxaliproxy3
NAT backup server
140.181.2.26
10.0.1.254
LHCONE Router l3-03
188.1.153.10
140.181.2.1
10 Gbit/s
LHCONE network
188.1.153.9/30
lxalird1
xrootd data manager
140.181.2.27
lxalird2
xrootd data manager
140.181.2.29

alinat

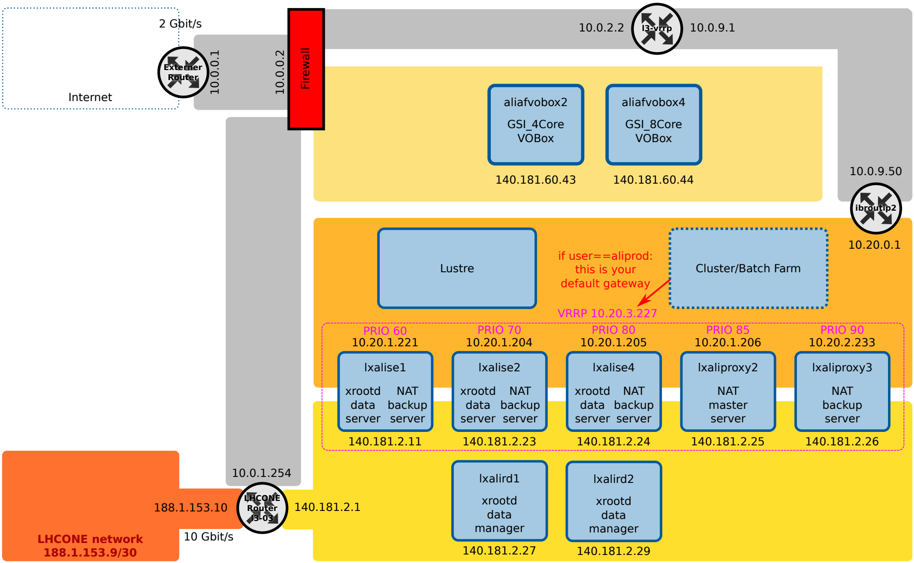

2 Gbit/s
Externer Router
Internet
10.0.0.1
10.0.0.2
Firewall
10.0.2.2
l3-vrrp
10.0.9.1
aliafvobox2
GSI_4Core VOBox
140.181.60.43
aliafvobox4
GSI_8Core VOBox
140.181.60.44
10.0.9.50
ibroutip2
10.20.0.1
Lustre
if user==aliprod: this is your default gateway
Cluster/Batch Farm
VRRP 10.20.3.227
PRIO 60
10.20.1.221
lxalise1
xrootd data server
NAT backup server
140.181.2.11
PRIO 70
10.20.1.204
lxalise2
xrootd data server
NAT backup server
140.181.2.23
PRIO 80
10.20.1.205
lxalise4
xrootd data server
NAT backup server
140.181.2.24
PRIO 85
10.20.1.206
lxaliproxy2
NAT master server
140.181.2.25
PRIO 90
10.20.2.233
lxaliproxy3
NAT backup server
140.181.2.26
10.0.1.254
188.1.153.10
LHCONE Router l3-03
140.181.2.1
10 Gbit/s
LHCONE network
188.1.153.9/30
lxalird1
xrootd data manager
140.181.2.27
lxalird2
xrootd data manager
140.181.2.29

alinat

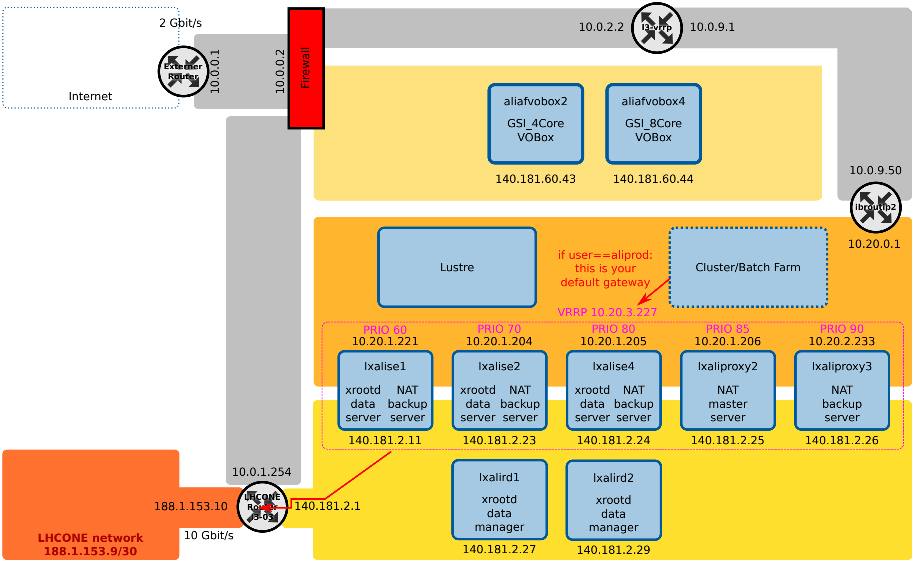

Internet

2 Gbit/s

Externer Router

10.0.0.1

10.0.0.2

Firewall

10.0.2.2

l3-vrrp

10.0.9.1

aliafvobox2
GSI_4Core
VOBox

140.181.60.43

aliafvobox4
GSI_8Core
VOBox

140.181.60.44

10.0.9.50

ibroutip2

10.20.0.1

Lustre

if user==aliprod:
this is your
default gateway

Cluster/Batch Farm

VRRP 10.20.3.227

PRIO 60 10.20.1.221	PRIO 70 10.20.1.204	PRIO 80 10.20.1.205	PRIO 85 10.20.1.206	PRIO 90 10.20.2.233
lxalise1 xrootd data server NAT backup server	lxalise2 xrootd data server NAT backup server	lxalise4 xrootd data server NAT backup server	lxaliproxy2 NAT master server	lxaliproxy3 NAT backup server
140.181.2.11	140.181.2.23	140.181.2.24	140.181.2.25	140.181.2.26

lxalird1
xrootd
data
manager

140.181.2.27

lxalird2
xrootd
data
manager

140.181.2.29

10.0.1.254

LHCONE Router l3-03

188.1.153.10

140.181.2.1

10 Gbit/s

LHCONE network
188.1.153.9/30

alinat

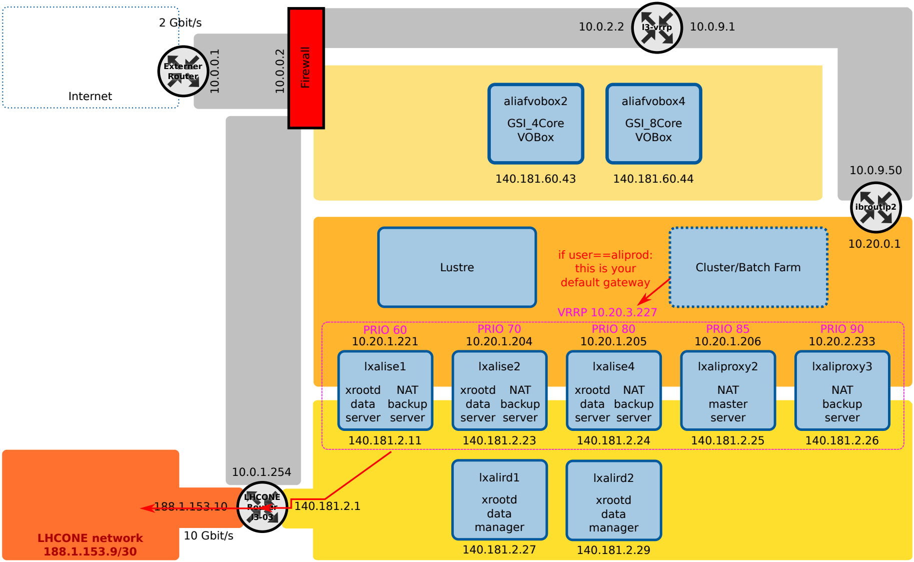
2 Gbit/s
Externer Router
Internet
10.0.0.1
10.0.0.2
Firewall
10.0.2.2
l3-vrrp
10.0.9.1
aliafvobox2
GSI_4Core
VOBox
140.181.60.43
aliafvobox4
GSI_8Core
VOBox
140.181.60.44
10.0.9.50
ibroutip2
10.20.0.1
Lustre
if user==aliprod:
this is your
default gateway
Cluster/Batch Farm
VRRP 10.20.3.227
PRIO 60
10.20.1.221
lxalise1
xrootd data server
NAT backup server
140.181.2.11
PRIO 70
10.20.1.204
lxalise2
xrootd data server
NAT backup server
140.181.2.23
PRIO 80
10.20.1.205
lxalise4
xrootd data server
NAT backup server
140.181.2.24
PRIO 85
10.20.1.206
lxaliproxy2
NAT master server
140.181.2.25
PRIO 90
10.20.2.233
lxaliproxy3
NAT backup server
140.181.2.26
lxalird1
xrootd data manager
140.181.2.27
lxalird2
xrootd data manager
140.181.2.29
10.0.1.254
LHCONE Router l3-03
188.1.153.10
140.181.2.1
10 Gbit/s
LHCONE network
188.1.153.9/30

alinat

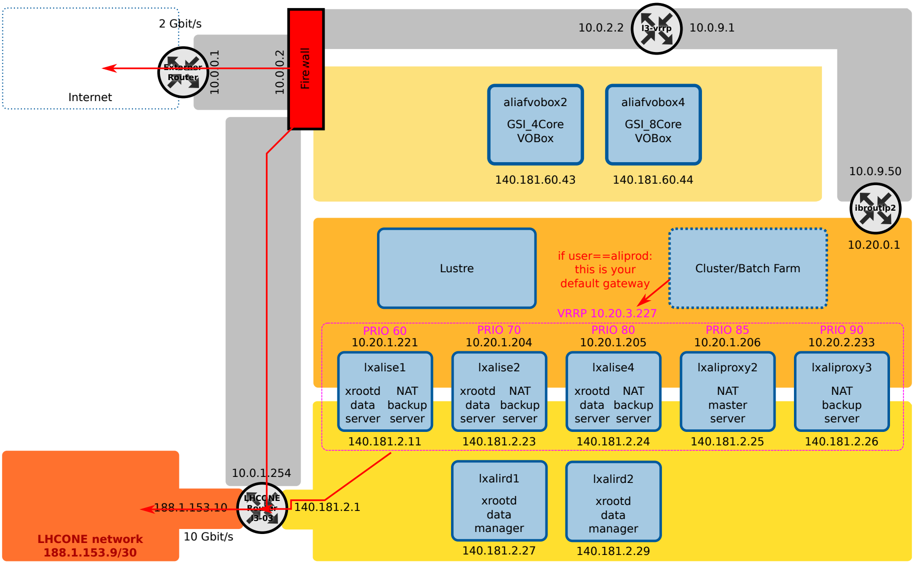
2 Gbit/s
External Router
Internet
10.0.0.1
10.0.0.2
Fire wall
10.0.2.2
l3-vrrp
10.0.9.1
aliafvobox2
GSI_4Core
VOBox
140.181.60.43
aliafvobox4
GSI_8Core
VOBox
140.181.60.44
10.0.9.50
ibroutip2
10.20.0.1
Lustre
if user==aliprod:
this is your
default gateway
Cluster/Batch Farm
VRRP 10.20.3.227
PRIO 60
10.20.1.221
lxalise1
xrootd data server
NAT backup server
140.181.2.11
PRIO 70
10.20.1.204
lxalise2
xrootd data server
NAT backup server
140.181.2.23
PRIO 80
10.20.1.205
lxalise4
xrootd data server
NAT backup server
140.181.2.24
PRIO 85
10.20.1.206
lxaliproxy2
NAT master server
140.181.2.25
PRIO 90
10.20.2.233
lxaliproxy3
NAT backup server
140.181.2.26
10.0.1.254
LHCONE Router l3-03
188.1.153.10
140.181.2.1
10 Gbit/s
LHCONE network
188.1.153.9/30
lxalird1
xrootd data manager
140.181.2.27
lxalird2
xrootd data manager
140.181.2.29

Current Issues

- CE “freezes”: is running but unable to spawn threads, in particular `squeue` commands ⇒ the VOBox silently stops spawning new JobAgents
    - temporary solution: monitor log file to promptly restart ce
    - long term solutions: fixes in JAliEn:
        - CE must not fail silently
        - understand why the critical condition is reached
- permissions: getting 0600 instead of 0644 with TPC transferred files
  Solution: being discussed with XRootD developers
- hard-coded list of Singularity bind mounted directories
  Solution: being discussed with JAliEn developers

Analysis Facilities in ALICE computing model

- AFs supposed to provide 50% of CPU share for analysis
  - receive AODs from O^2 farm and T1/T2s
  - produce histograms and trees
- 10% of sampled AODs for quick analysis and cut tuning

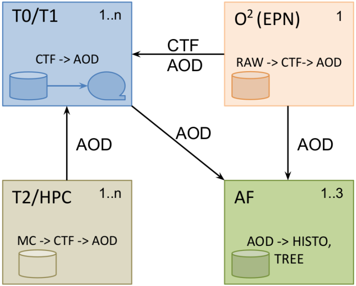

Requirements:

- serve 6-8k job slots with ∼15 MB/s/core[1]
- aggregate throughput of 100 GB/s
- be able to digest more than 5 PB of AODs in a 12-hour period

[1] S.Piano, ALICE Week June 6th

Satisfying the requirements

- Already close to the required number of cores.
- Benchmark showed linear increase of throughput vs number of concurrent jobs up to 30 GB/s for 2000 jobs. Challenges to rely on linear increase up to 100 GB/s still to be faced.
  - improve utilization of the shared Lustre file system
  - improve and enforce algorithms for distribution of files over OSTs

Thanks for your attention!

Questions?