ALICE

Erik Edelmann

erik.edelmann@csc.fi

NeIC / CSC

Pledges: CPU

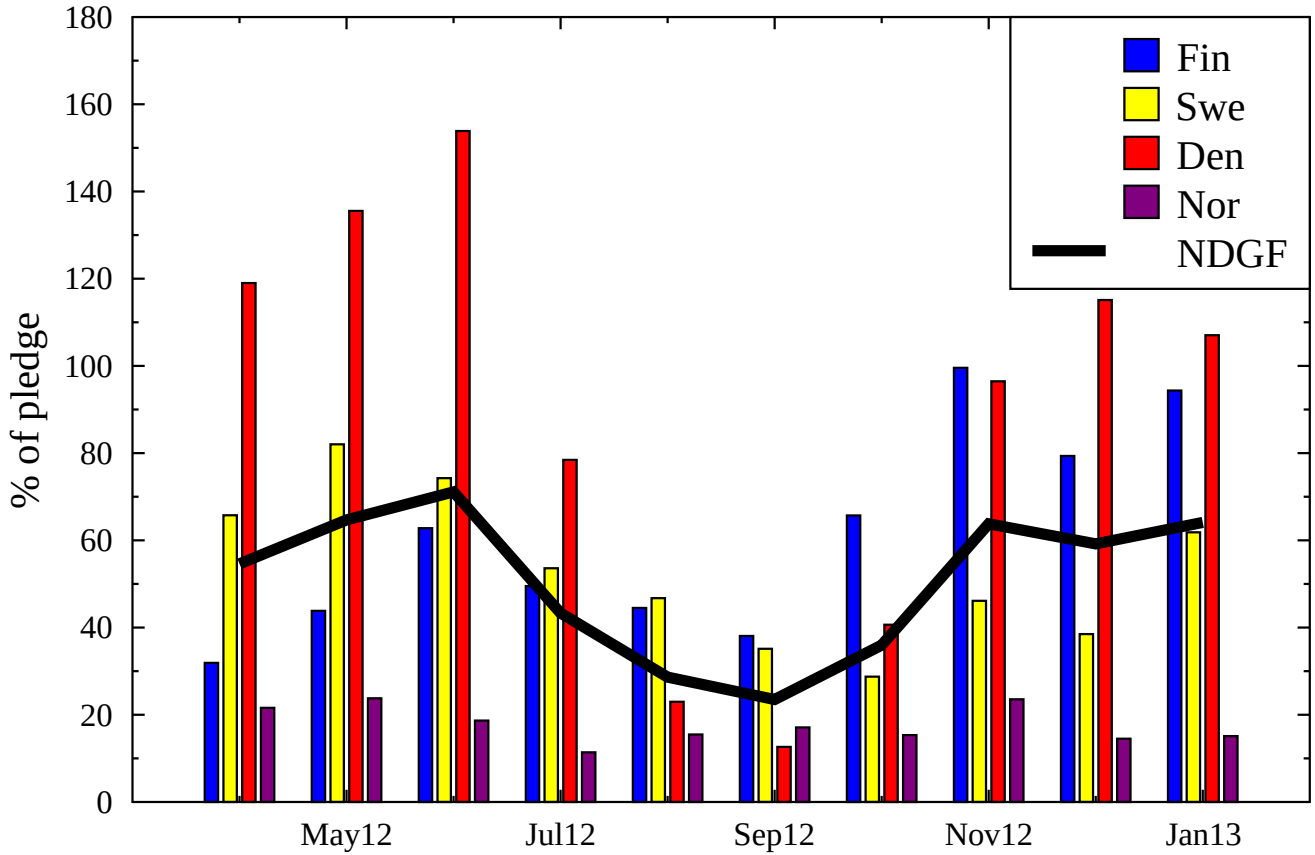
180
160
140
120
100
80
60
40
20
0
% of pledge
May12
Jul12
Sep12
Nov12
Jan13
Fin
Swe
Den
Nor
NDGF

Pledges: CPU

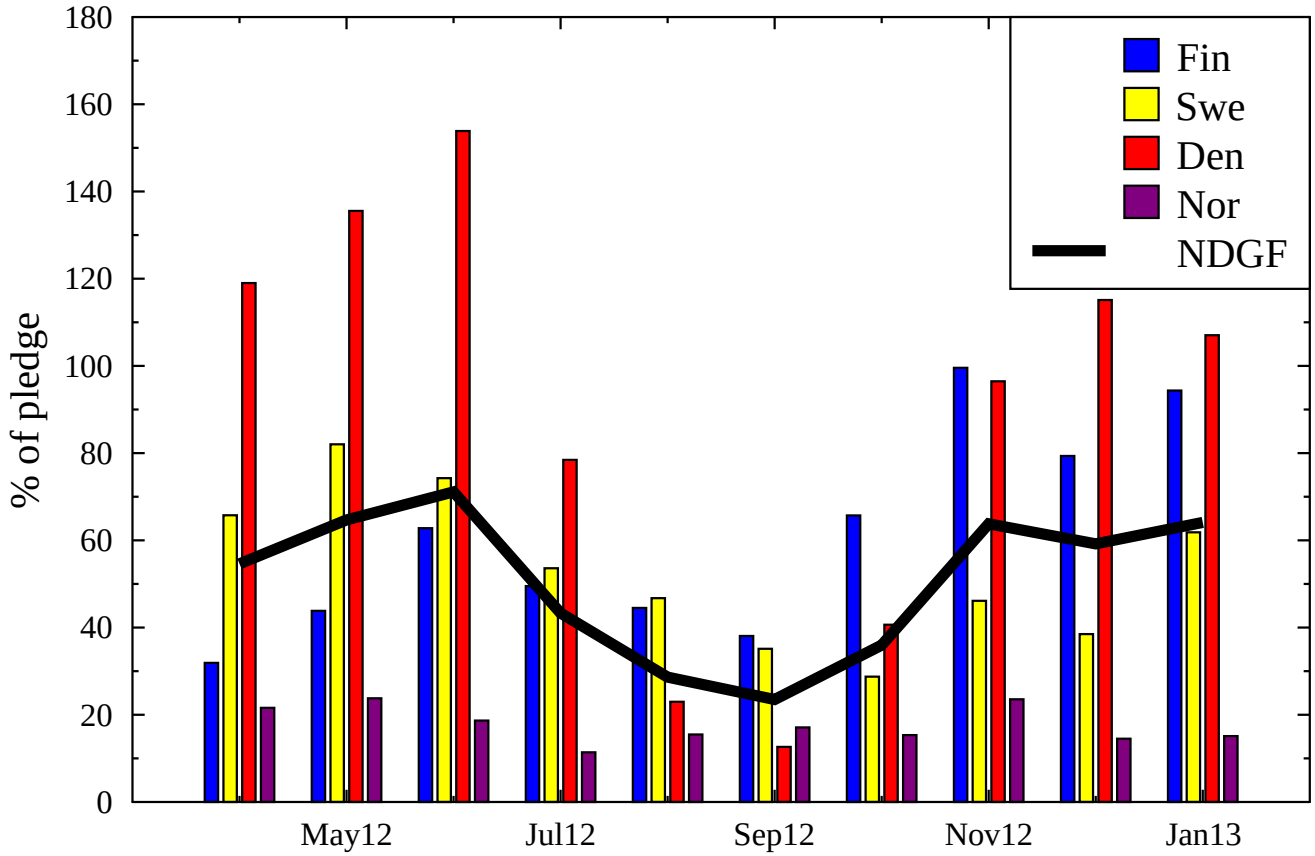

In 2013 the pledge increases from 12535 to 13383 kHEPSPEC06

Pledges: Storage

- Disk
  - 2012 Pledge: 1.325 PB (NDGF-T1) + 400 TB (SNIC-T2)
  - Installed: 711 TB + 400 TB
  - In use: 617 TB + 125 TB
- Tape
  - 2012 Pledge: 1.761 PB
  - Installed: 2.4 PB (according to bdii1.ndgf.org)
  - In use: 216 TB

Pledges: Storage

- Disk
  - 2012 Pledge: 1.325 PB (NDGF-T1) + 400 TB (SNIC-T2)
  - Installed: 711 TB + 400 TB
  - In use: 617 TB + 125 TB
- Tape
  - 2012 Pledge: 1.761 PB
  - Installed: 2.4 PB (according to bdii1.ndgf.org)
  - In use: 216 TB

In 2013 disk pledge increases to 1480 TB, Tape to 2276 TB

Job efficiency & storage architecture

ALICE has been infamous for bad job efficiency ...

Job efficiency & storage architecture

... but it's not that bad anymore ...

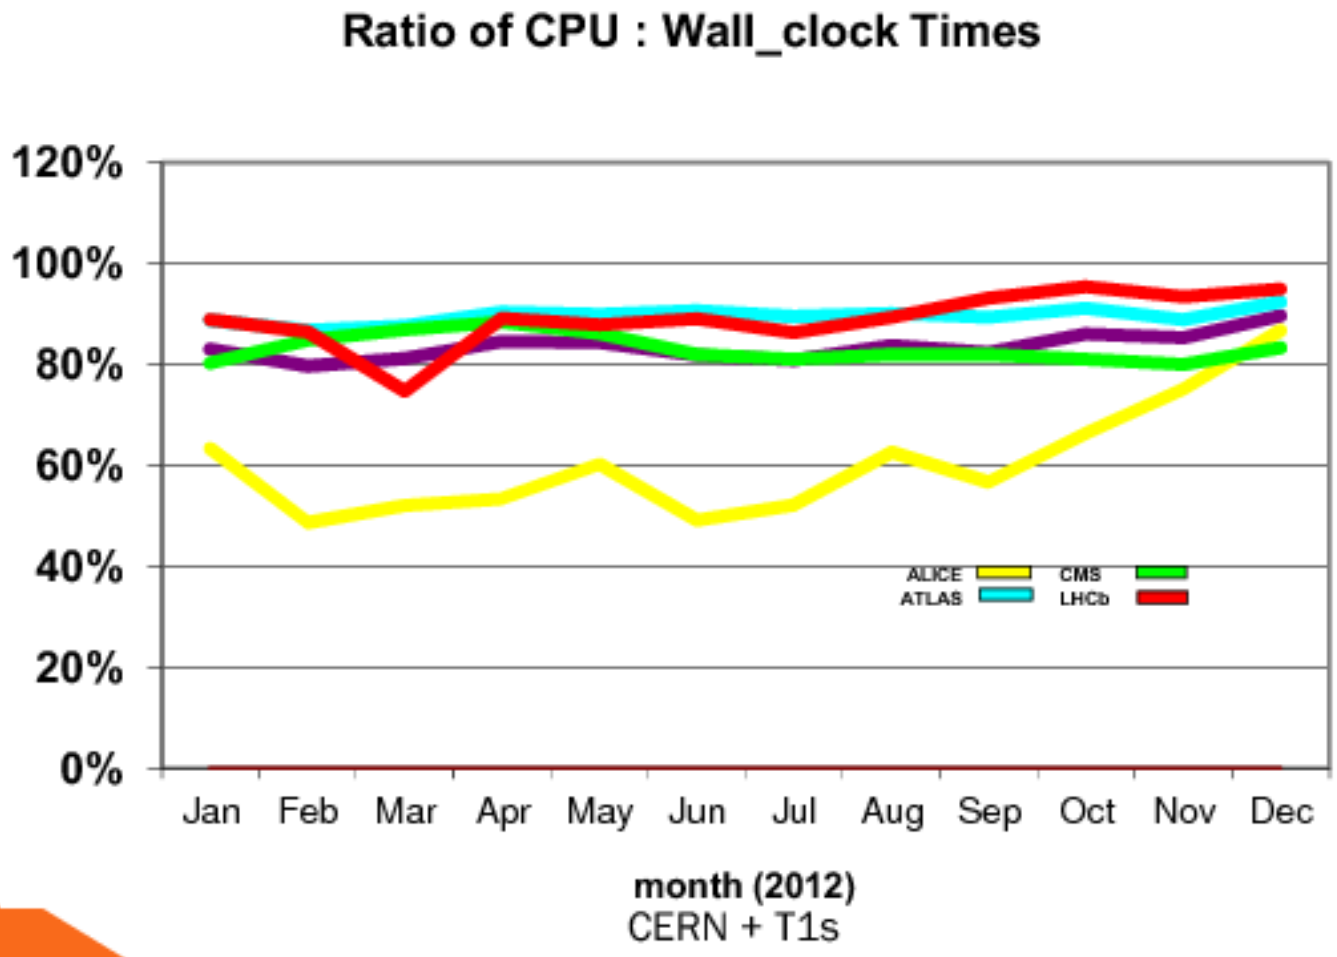

Job efficiency & storage architecture

... except for NDGF

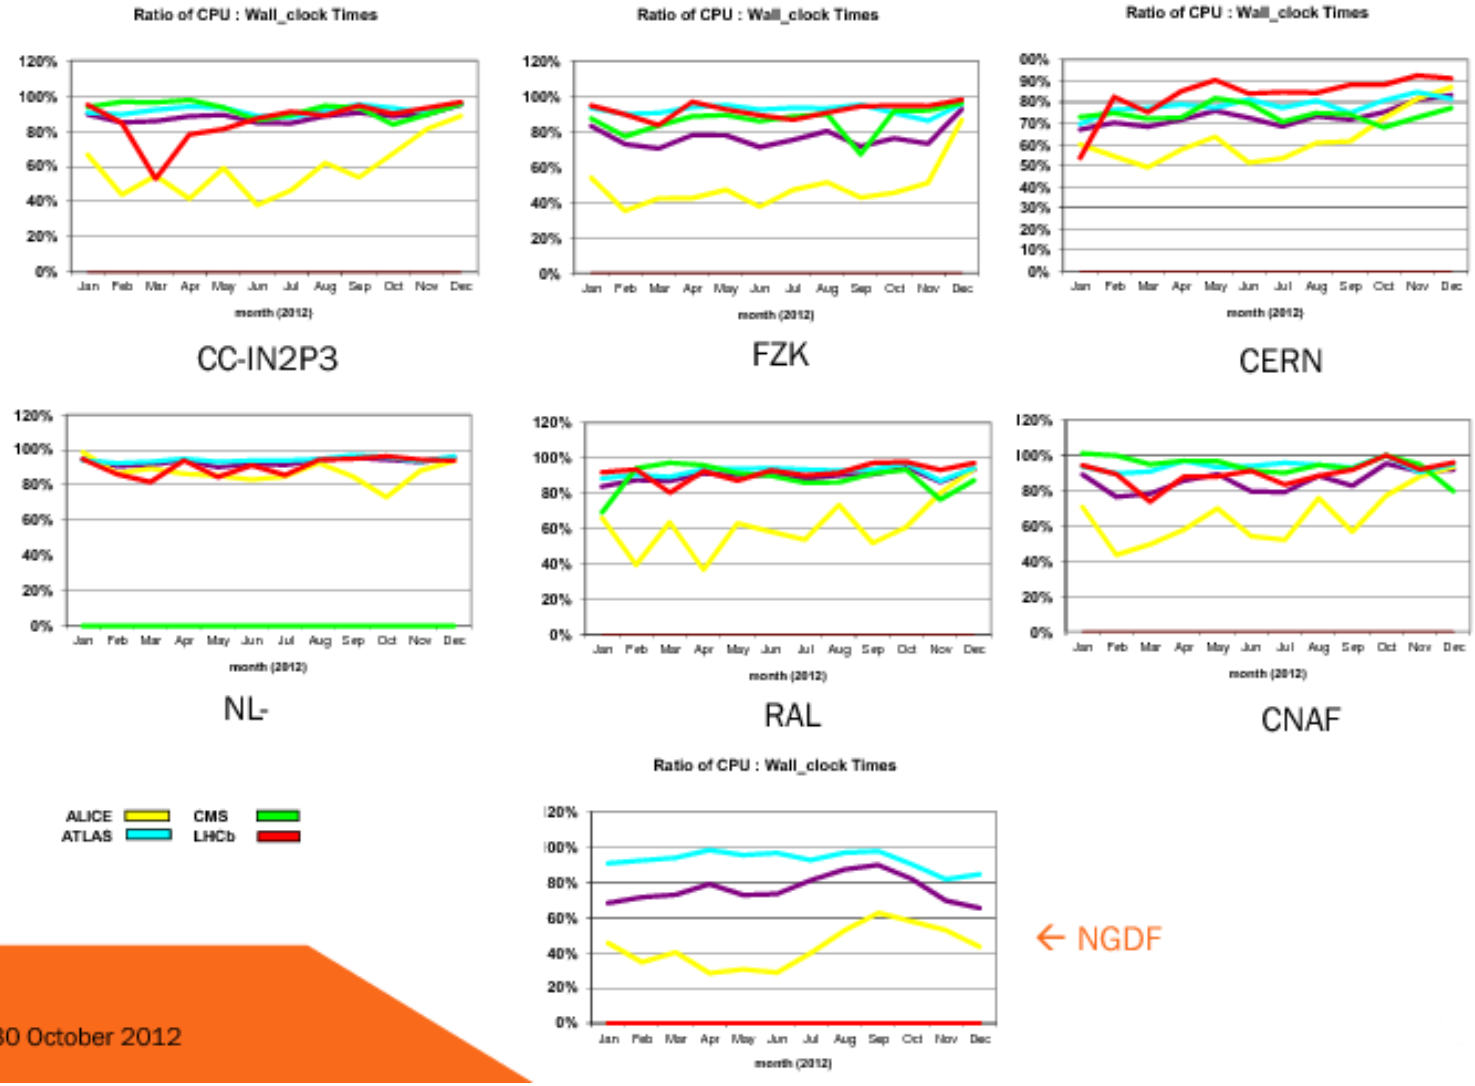

Job efficiency & storage architecture

- Distributed storage not necessarily a problem
- but could it be made more efficient?
  - Caching?
  - ’hints’ of where the data is located internally?

SW installation: past, present and future

- Past: VO-box installed onto disk shared with the WNs

SW installation: past, present and future

- Past: VO-box installed onto disk shared with the WNs
- Present: Jobs install onto local disk using a BitTorrent based system.
  - Shared disk not required anymore.

SW installation: past, present and future

- Past: VO-box installed onto disk shared with the WNs
- Present: Jobs install onto local disk using a BitTorrent based system.
  - Shared disk not required anymore.
- Future: CVMFS

SW installation: past, present and future

- Past: VO-box installed onto disk shared with the WNs
- Present: Jobs install onto local disk using a BitTorrent based system.
  - Shared disk not required anymore.
- Future: CVMFS
  - When? ”They’ll tell us beforehand”

One VO-box to rule them all?

- With the ”shared disk” requirement gone, could we have only one VO-box that submits through ARC?

One VO-box to rule them all?

- With the ”shared disk” requirement gone, could we have only one VO-box that submits through ARC?
- Probably, but do we want to?

One VO-box to rule them all?

- With the ”shared disk” requirement gone, could we have only one VO-box that submits through ARC?
- Probably, but do we want to?
  - Depends on what we do with storage

One VO-box to rule them all?

- With the ”shared disk” requirement gone, could we have only one VO-box that submits through ARC?
- Probably, but do we want to?
  - Depends on what we do with storage
  - Pros:
    - Fewer VO-boxes to keep track of.
    - Easier to add new clusters.
    - Simplified accounting (Using ARC everywhere would be enough for this)

One VO-box to rule them all?

- With the ”shared disk” requirement gone, could we have only one VO-box that submits through ARC?
- Probably, but do we want to?
  - Depends on what we do with storage
  - Pros:
    - Fewer VO-boxes to keep track of.
    - Easier to add new clusters.
    - Simplified accounting (Using ARC everywhere would be enough for this)
  - Cons:
    - We would replace something that works with something untested.
    - ”Political” issue: On the Alimonitor map, one nordic country would have a huge cluster, the others nothing.

Long shutdown plans

- For us, mostly business as usual:
  - Reprocessing of 2011 data, MC, user analysis ... $\Rightarrow$ jobs running, needs access to data.

Long shutdown plans

- For us, mostly business as usual:
  - Reprocessing of 2011 data, MC, user analysis ... $\Rightarrow$ jobs running, needs access to data.
- Sites may be asked to upgrade to SLC6 (”if OK with the sites”)

Long shutdown plans

- For us, mostly business as usual:
  - Reprocessing of 2011 data, MC, user analysis ... $\Rightarrow$ jobs running, needs access to data.
- Sites may be asked to upgrade to SLC6 (”if OK with the sites”)
- IPv6
  - CERN aims to have IPv6 fully deployed by the end of LS1 (autumn 2014)
    - Publicly accessible services should be reachable on dual-stack.
  - ALICE would like to start testing their software with IPv6 already now
    - Would be nice if VO-Boxes could be dual-stacked ASAP.

Summary

- We need more resources
- We’ll have to think how we do storage and organizes VO-boxes
- We’ll need IPv6.