The phases of IPv4 phaseout at the Nordic Tier-1 for WLCG

NeIC NT1 Manager
Mattias Wadenstein
<maswan@ndgf.org>

2026-03-20
ISGC 2026
Taipei, Taiwan

NORDIC E-INFRASTRUCTURE COLLABORATION

Overview

- Need
- Constraints
- Plan
- Status

Background

- WLCG is transitioning from IPv4 to IPv6
- Current focus: Enabling sites to go IPv6-only
  - Will not force IPv6-only
- Previous efforts with deadlines:
  - Storage dual-stacked by 2019
    - Current status 98% done
  - Computing dual-stacked by 2024
    - Current status 74% done

Nordic Setup

- Distributed Tier-1 site
  - Storage (30PB disk + 30PB tape) and compute in 6 Nordic sites
- Associated Tier-2 sites in Slovenia and Switzerland
  - We integrate the storage for user convenience

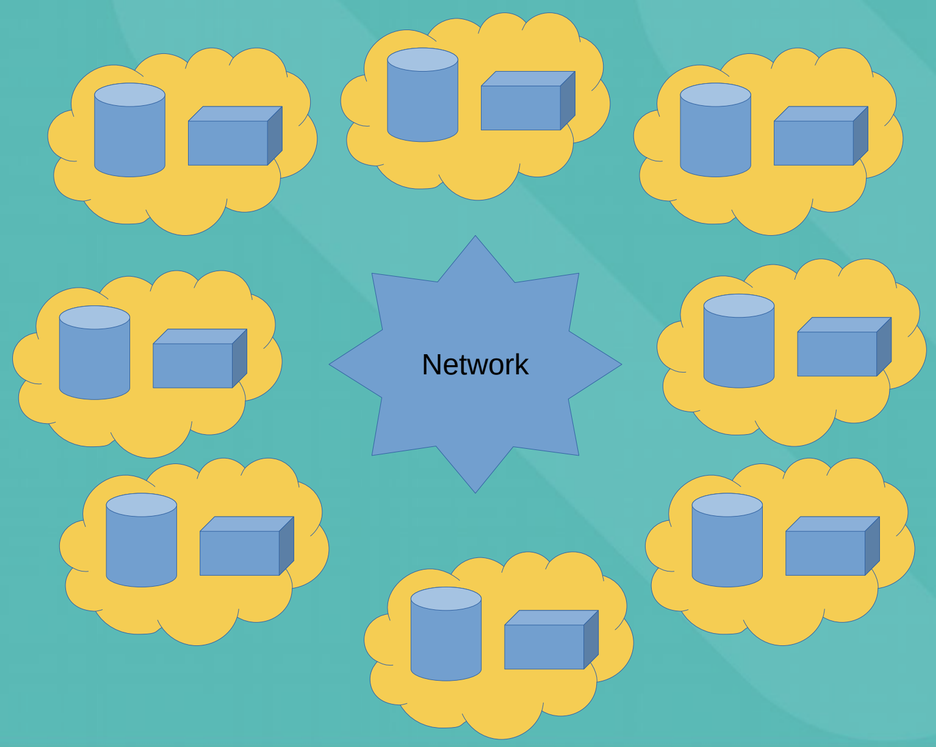

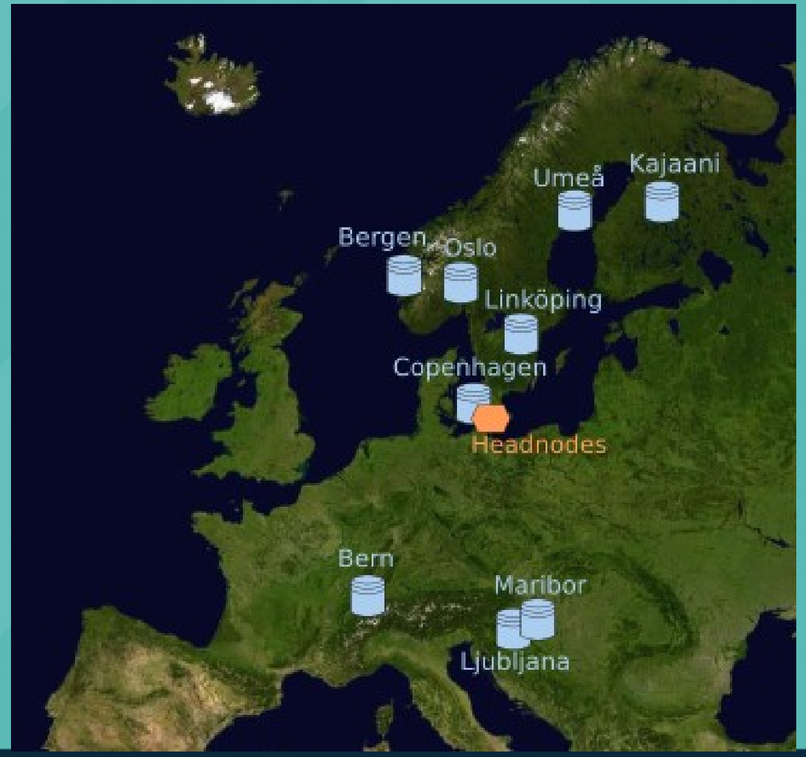

Our motivation

- We are facing IPv4 exhaustion
  - We have a /22 IPv4 LHCOPN split over 8-10 networks
  - Space and performance needed for run4 will likely require many more dCache pool nodes than today
    - And more ARC CE data staging nodes
  - Thus: Some of the networks are likely to get tight or run out
  - Renumbering might solve things temporarily, but is painful
- Maintaining two sets of IP ACLs is worse than one
  - Especially since the v4 one can’t be sensibly subsubnetted

Constraints

- Not all computing speaks IPv6 yet
  - Several years since the WLCG deadline for dual stack
  - Needs to be fixed in the next couple of years
- Some legacy code may need IPv4 storage access
- We can solve this by still having a legacy dCache door that proxies connections instead of redirecting
  - The normal door behavior is to issue redirects to pools for performance

dCache design, redirect

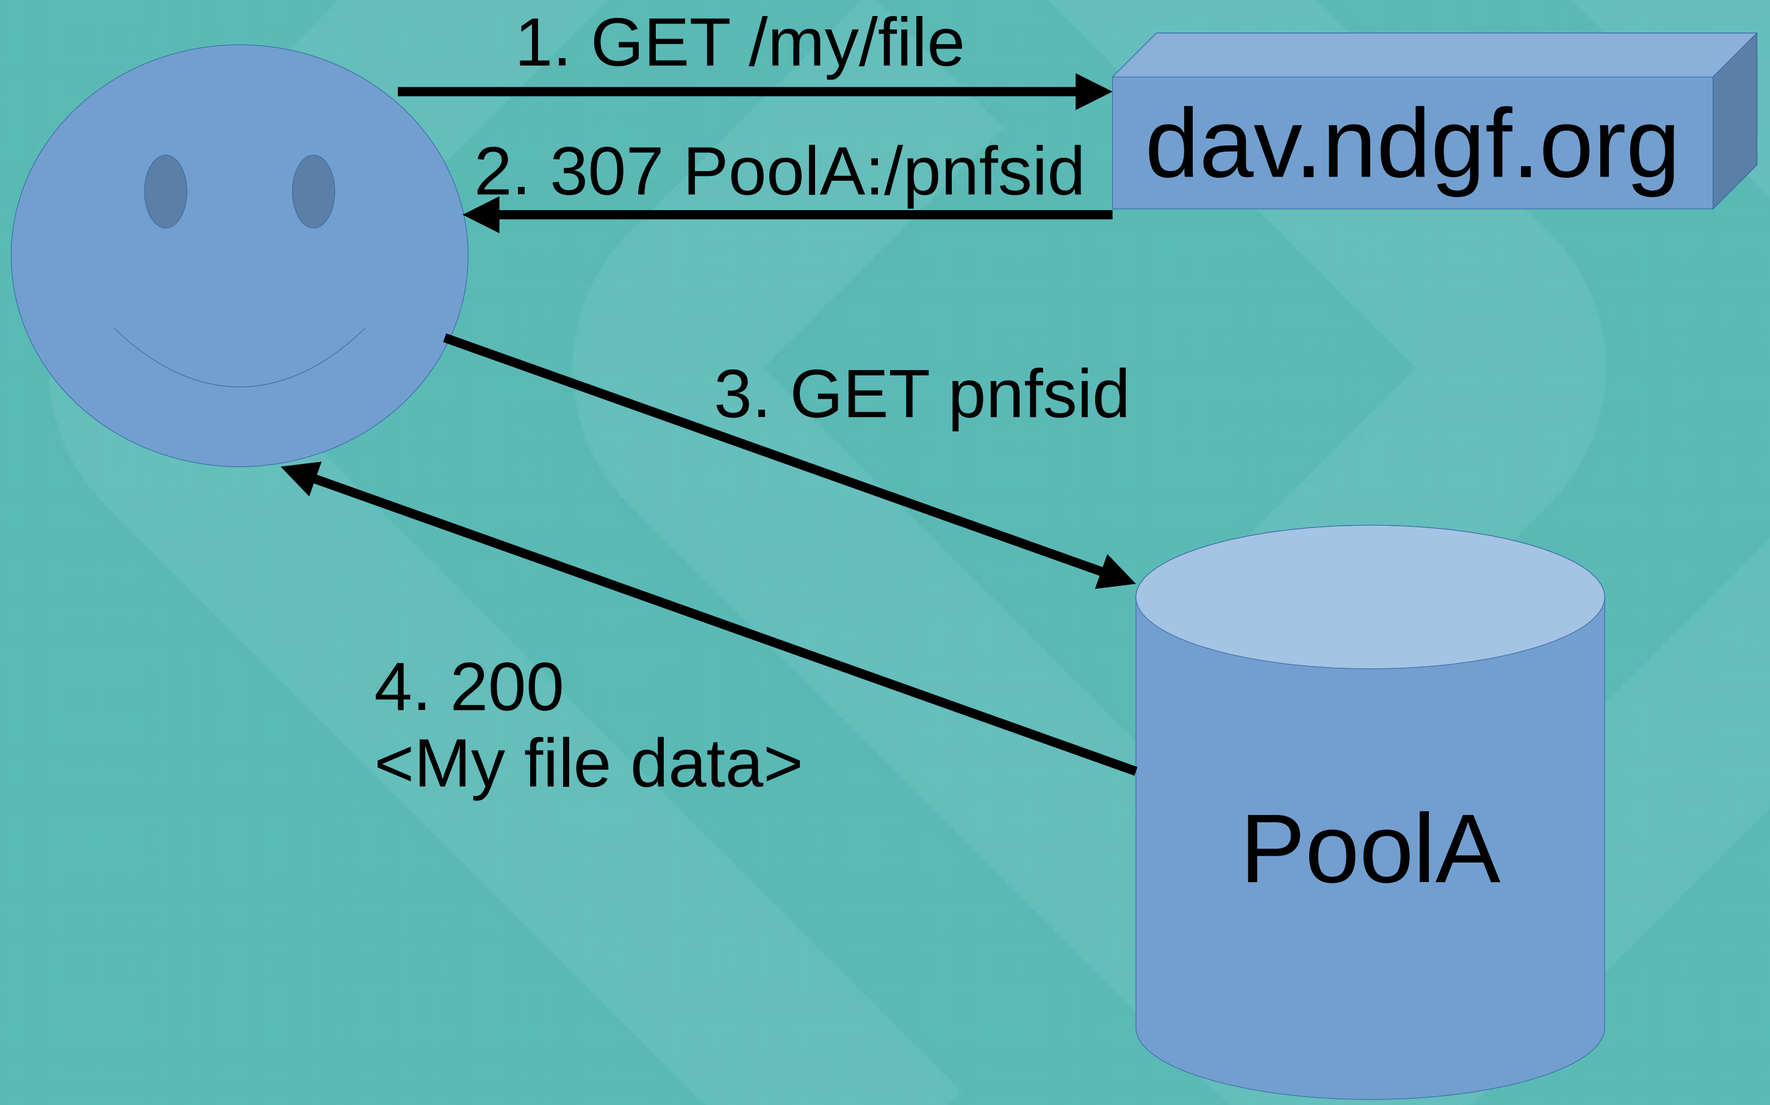

dCache design, proxy

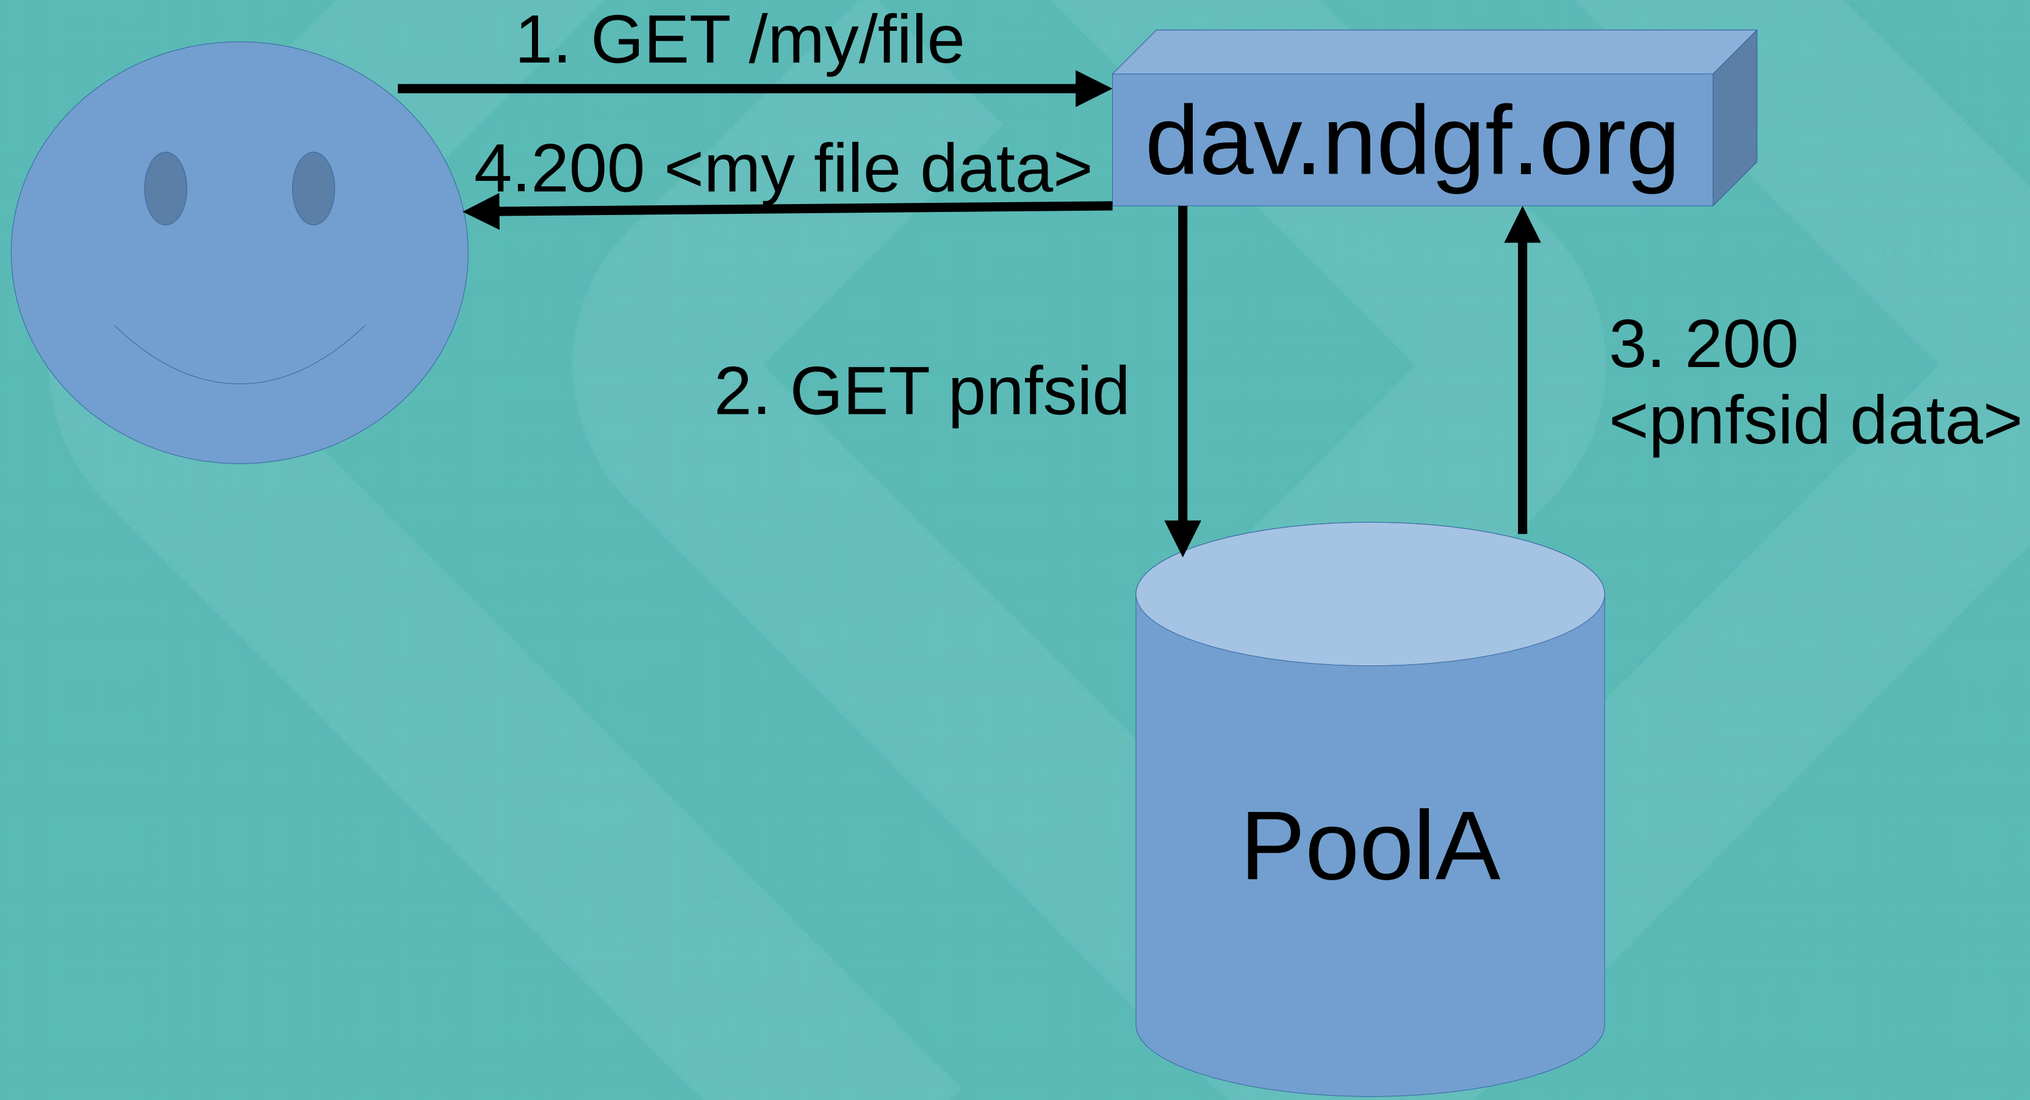

Downside of proxying

- Higher latency
- Lower throughput
  - The proxying doors will likely only have 10-40Gbit/s aggregate
- Higher network usage
  - Longer network path, especially when Pool and Compute is in the same location: Finland → Copenhagen → Finland
- More fragile connections
  - For a slow transfer or a direct read, a door glitch or restart will interrupt the transfer
  - A redirected transfer only relies on client, pool, and network

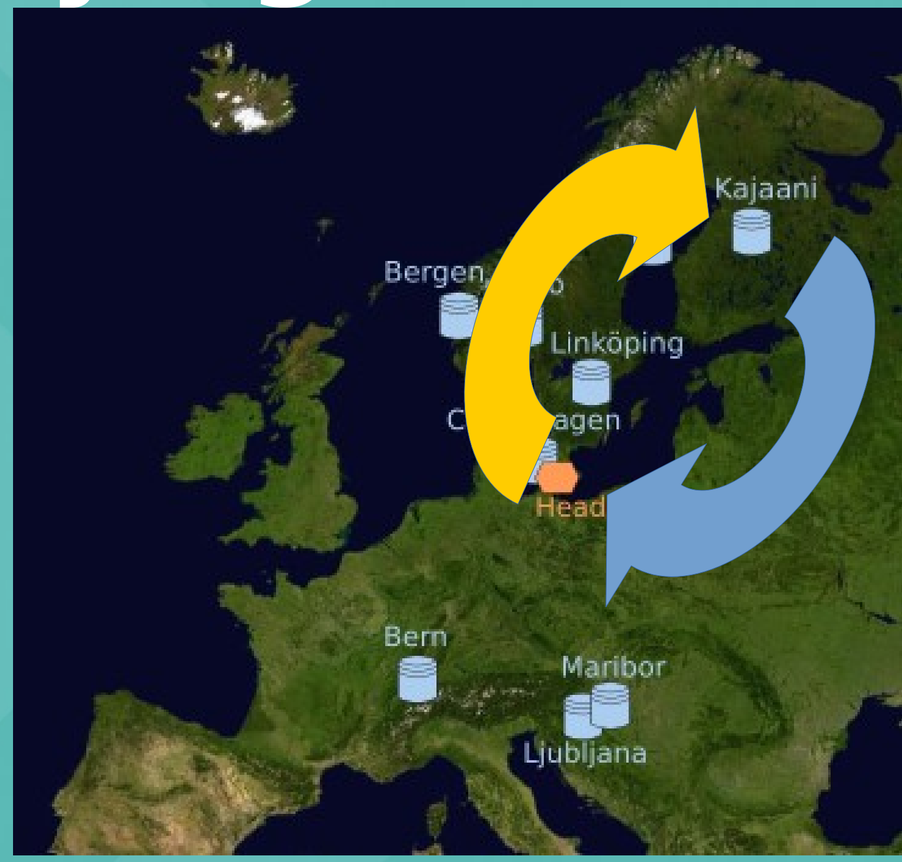

Constraints

- Experiments would rather not see such a big change during data taking
  - We should get this done during the long stop, 2026-2029
- Getting ready by the next data challenge in 2027 would be great
  - But not set in stone, in case of issues found

Experiment response

- ALICE
  - Some computing is IPv4 only, but a proxying door should be fine
  - Storage all dual-stacked, fine
- ATLAS
  - Some computing is IPv4 only, but a proxying door should be fine
  - Not all storages (2 T2, 5 T3), concerned about third-party copies
    - [Numbers as of January 2026]

Plan

- Set up proxying xrootd and webdav door for legacy IPv4 access
  - Including HA features so we can do upgrades smoothly
- Drop IPv4 from pools in pre-production dCache
  - Check if all of our monitoring and management still works
  - Check if ATLAS is fine with this
  - Time: This year
- Repeat in production after issues sorted
  - Time: In time for DC27

Known Worries

- Third party copies
  - HAproxy and dCache can detect if a client uses IPv4 or IPv6 and handle them by different paths
  - Third-party copies means that the “client” is CERN FTS, but the other destination might be a v4-only storage element
- External monitoring
  - Experiment, WLCG R&A, etc
- Most of our own deployment, management, and monitoring is untested in IPv6-only
  - Fixable, but might take time by an already busy team

Current status

We are here

Questions?

NORDIC E-INFRASTRUCTURE COLLABORATION