

IPv6 Deployment Experience at the GridKa Tier-1 at KIT

A. Petzold

STEINBUCH CENTRE FOR COMPUTING - SCC



IPv6 @ GridKa

- several iterations of IPv6 deployment plan in 2017
- complex internal routing setup for IPv4
 - servers have two interfaces w/ internal and external IP, WNs have only internal IP
 - VRF w/ two virtual routers for internal and external traffic
 - allows NAT for hosts w/ only internal interface
 - requires route leaking between VRFs to allow communication between WNs and servers
- different requirements for IPv6
 - no NAT
 - no separation of internal/external traffic on servers
 - no virtual routers for internal/external traffic

Resulting IPv6 Config

- IPv6 should be deployed on internal NIC
 - some servers get separate tagged VLAN interface for IPv6

- typical file server

```
ens2: inet 10.65.124.128/24
```

```
    inet6 2a00:139c:5:417c::80/64
```

```
ens2.533@ens2: inet 192.108.68.34/26
```

internal IPv4

IPv6

external IPv4


Deploying IPv6 for CMS dCache

- testing and preparation on pre-production dCache instance
 - scripts for deploying IPv6 IPs on correct interface
 - checked dCache internal communication
 - what AAAA records need to be deployed and where
 - local and remote transfer tests on IPv4 and IPv6
- IPv6 deployed for all involved CMS server in minutes
 - IPv6 IPs on all servers
 - additional AAAA records in internal and external DNS
 - confirmed that dCache knows about all IPs
- What could go wrong?

Fallout

■ FTS transfers

- transfers by dual-stack FTS GridKa \Leftrightarrow IPv6 enabled sites worked
- transfers by dual-stack FTS GridKa \Leftrightarrow IPv4-only sites **broken**



➔ dCache gridftp door expects a true dual-stack setup, i.e. IPv6 and external IPv4 address on same interface

➔ immediate change of IPv6 network setup required!

- dual-stack setup for the doors within 24h
- required many unprepared routing + firewall changes internally & externally

■ at first everything seemed fine ...

■ ... reports about unreachable ports on gridftp doors

- were closed in the firewall with IPv4-only setup
- indicates transfers going through door??
- dCache pools are still using dual-homed setup

Results

- dCache (actually gridftp) requires true dual-stack setup, can't us dual-homed
- will test xrootd soon (just because)
- everything works now, but not the way we want it to
- ➔ Complete change of IPv6 deployment setup at GridKa required
 - still want to retire VRF
 - perfsonar servers first candidates
- Many thanks to CMS and dCache.org for bearing with us and helping with debugging!
- **Recent feedback from Paul Millar at dCache Workshop**
 - **problem of the gridftp protocol**
 - **workaround could in principle be implemented in dCache**