Experience with IPv6 deployment at FZU AS CR in Prague

Marek Eliáš, Lukáš Fiala, Tomáš Kouba
Jiří Horký, Jiří Chudoba, Jan Kundrát, Jan Švec
elias@fzu.cz

HEPiX Fall 2011 Workshop, 26. October 2011

Institute of Physics AS CR, v. v. i. (FZU)

Outline

1. Our current problems with IPv4
2. Setup of our central router, DNS and DHCPv6
3. DHCPv6 vs autoconfiguration
4. PXE in IPv6
5. Accessibility of Grid CRLs through IPv6
6. Future plans

Our motivation for IPv6 transition

Fyzikální ústav
Akademie věd ČR, v. v. i.

Lack of IPv4 addresses

- We have one /24 subnet, we currently use 3 times more addresses
- Workernodes are in a private address space

Routing problems

- DPM does not support multihoming (workernodes have to use the same address to access a DPM pool node as the outside world)
- Traffic between workernodes and DPM pool nodes must be routed through the central router
- New 10 Gbps infrastructure in FZU last year $\Rightarrow$ routing has become unsustainable

Our motivation for IPv6 transition

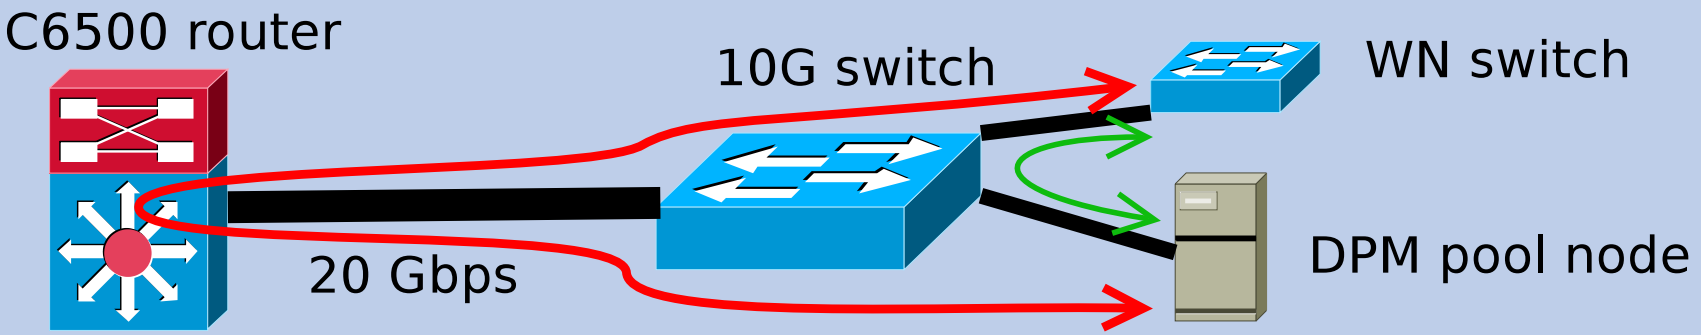

Currently deployed solution

- Suggested by Maarten Litmaath
- DPM pool nodes are connected directly to both, private and public networks

```
inet 147.231.25.45/24 brd 147.231.25.255 scope global eth0.25
inet 172.16.0.45/16 brd 172.16.255.255 scope global eth0.172
```

- Workernodes from private 172.16/16 network have a static route to access public IP of each DPM pool node directly without routing

```
ip route add 147.231.25.45 dev eth0
```

Setup of central router

Fyzikální ústav
Akademie věd ČR, v. v. i.

Firewall problems

- We are using a Cisco C6500 with FWSM in transparent mode
- FWSM does not support filtering of IPv6 traffic in transparent mode
- Possible solutions: switching the FWSM to routed mode or enable multiple context mode and filter IPv4 in transparent and IPv6 in routed mode.

Possible temporary workaround

Use the ethertype ACL rules to have the FWSM pass all IPv6 traffic unfiltered:

```
access-list outside_ether_access_in remark IPv6
access-list outside_ether_access_in ethertype permit 86dd
access-list inside_ether_access_in remark IPv6
access-list inside_ether_access_in ethertype permit 86dd
```

Setup of central router

Fyzikální ústav
Akademie věd ČR, v. v. i.

Setup of local network

- No IPv6 connectivity in production environment
- IPv6 is only in a separate VLAN
- Routing advertisements were tested
- Stateless DHCP on a cisco router for distributing addresses of DNS resolvers works well

Cisco IOS vulnerability

- Security advisory published on 28. September[1]
- An attacker can cause a router to reload by sending malformed IPv6 packets to the right interface of the router
- Nearly all versions of IOS are vulnerable
- Fixes are available from Cisco

[1] http://www.cisco.com/warp/public/707/cisco-sa-20110928-ipv6.shtml

DNS and DHCP

All management machines are installed with SL 6.1

DNS Setup

- No problems with running a master server and a resolver
- One of our masters has IPv6 address which is in GLUE record of a parent zone
- We have one IPv6 only zone with a master without IPv4 connectivity
- IPv6 only resolver is not a good idea for now (not everybody trying IPv6 have IPv6 enabled masters)

DHCPv6 and Routing Advertisments

- RA setup on a Cisco router
- Cisco routers support stateless DHCP: to be used in combination with RA to advertise DNS resolvers
- ISC DHCPv6 implementation from SL 6.1 works well
- Stateful DHCP: ISC dhcp and dibbler tested for client & server

DHCPv6 vs Autoconfiguration

Basic requirement

- We definitely want a fixed IPv6 address to be configured on each node
- IPv6 address which changes in time does not make sense for production services nor workernodes

Means of configuring fixed IPv6 address

- Manual address configuration: difficult to manage especially for IPv6
- Stateful DHCP: client is identified using DUID instead of MAC
- Stateless autoconfiguration: yes, using EUI64, a node choses the same IPv6 address generated from MAC address of active interface

DHCPv6 vs Autoconfiguration

Stateful DHCP

- DUID is usualy stored in `/var/lib/...`
- But how do we know the DUID of a freshly installed system?
- And what is the DUID of an uninstalled system?
- In our IPv4 setup the DHCP server says which host configuration should be installed on which hardware

Autoconfiguration

- Autoconfiguration seems to be simpler for implementation in hardware, some special devices may lack support of stateful DHCP
- If the active interface is replaced or is changed by failover, IPv6 address is changed too

$\rightarrow$ Problems with DNS (cached entries), firewall, system logs

PXE in IPv6

Fyzikální ústav
Akademie věd ČR, v. v. i.

support in DHCPv6

- No `next-server` option in DHCPv6, server directly specifies url of image to be loaded through option `boot-file-url`
- Described in RFC 5970 from september 2010
- RFC 5970 is missing on dibbler's list of implemented RFCs; ISC dhcp 4.2.2 doesn't seem to support it

gPXE

- Supports IPv6, but has no means of automatic boot available
- When IPv6 is enabled manually, it can setup IP addres with RA and get DNS resolvers from stateless DHCP
- Subsequent manual download of boot image from http server works well including resolving IPv6 address from DNS
- Its configuration interface works a bit strangely

PXE in IPv6

SGI Altix 340

- PXE implementation does not care about IPv6
- It simply tries DHCP for IPv4 and then gives up

IBM iDataPlex dx360

- Node with UEFI
- PXE implementation does not care about IPv6
- Possibility to load custom EFI application performing boot
- We are not aware of such an EFI application which would support booting from IPv6 network

Availability of Grid CRLs

Fyzikální ústav
Akademie věd ČR, v. v. i.

IPv6 EUGridPMACrlChecker

Check date: 19.10.2011 text report (take a while)

CA Subject CA Crl url	IPv6 DNS	IPv6 GET
CN=CERN Trusted Certification Authority, DC=cern, DC=ch		
http://ca.cern.ch/ca/CRL/CERN Trusted Certification Authority.crl	NO	
CN=CESNET CA 3, O=CESNET CA, DC=cesnet-ca, DC=cz		
http://crl.cesnet-ca.cz/CESNET_CA_3.crl	YES	YES
CN=CESNET CA Root, O=CESNET CA, DC=cesnet-ca, DC=cz		
http://crl.cesnet-ca.cz/CESNET_CA_Root.crl	YES	YES
CN=CESNET CA, DC=cesnet-ca, DC=cz		
http://www.cesnet.cz/pki/crl/cn=CESNET CA,dc=cesnet-ca,dc=cz.crl	YES	YES
CN=GRID2-FR, O=CNRS, C=FR		
http://crls.services.cnrs.fr/GRID2-FR/getpem.crl	NO	
CN=CNRS2-Projets, O=CNRS, C=FR		
http://crls.services.cnrs.fr/CNRS2-Projets/getpem.crl	NO	
CN=CNRS2, O=CNRS, C=FR		
http://crls.services.cnrs.fr/CNRS2/getpem.crl	NO	
CN=CyGridCA, O=HPCL, O=CyGrid, C=CY		
http://cygrid.org.cy/CyGridCA/afe55e66.r0	NO	
CN=Grid-Ireland Certification Authority, O=Grid-Ireland, C=IE		
http://www.cs.tcd.ie/Grid-Ireland/gi-ca/1e43b9cc.r0	YES	ERR
CN=Grid Canada Certificate Authority, O=Grid, C=CA		
http://www.gridcanada.ca/ca/bffbd7d0.r0	NO	
CN=HellasGrid CA 2006, OU=Certification Authorities, O=HellasGrid, C=GR		
http://crl.grid.auth.gr/hellasgrid-ca-2006/82b36fca.pem	NO	
CN=HellasGrid Root CA 2006, OU=Certification Authorities, O=HellasGrid, C=GR		
http://crl.grid.auth.gr/hellasgrid-root-ca-2006/28a58577.pem	NO	

Availability of Grid CRLs

Fyzikální ústav
Akademie věd ČR, v. v. i.

Availability of certificate revocation lists is needed by most components of gLite. Our tool[2] checks availability of CRLs of certificate authorities from lcg-CA bundle.

- Checks presence of AAAA record of corresponding web server
- If AAAA record is present, checks wether CRL is downloadable
- Current score is not optimistic:
  - Only 7 of 94 authorities have CRL on a server with AAAA record
  - Only 6 CRLs are really downloadable through IPv6

What to do next?

- Ask CA's to support IPv6 (who is competent to do this?)
- Implement the same check in EUGRID nagios? (check itself is trivial)

[2] `http://www.particle.cz/farm/admin/IPv6EuGridPMACrlChecker/`

Future plans

Fyzikální ústav
Akademie věd ČR, v. v. i.

- Prepare some nodes for joining HEPiX IPv6 testbed
- Prepare a working system for node installation (maybe using private installation network)
- Try to run a CFEngine instance on IPv6
- Prepare our monitoring tools for IPv6 and SNMPv6, IPv6 version of netflow and so on
- Find a reasonably secure solution for accessing managemet interfaces of machines (which usualy do not support IPv6) from remote client through IPv6

Thank You

Marek Eliáš

elias@fzu.cz

http://www.farm.particle.cz

Work partially supported by CESNET, z. s. p. o.
project number 416R1/2011