



Network design

Exercise

Barbara Krasovec, EGI CSIRT, JSI

Split, June 2022

Objective

- The purpose is to think about network design for an HPC centre
- We will take HPC Vega as an example

Security integration in network design

- Integrating security early on prevents multiple security gaps later on
- Plan for growth
- Maintain network documentation and update it, when the topology changes
- Network redundancy - how much does it cost if services are down?
- Plan remote access to your internal network

HPC Network design

- HPC network usually fat-tree
- Two network fabrics: management network (administration) and servers network (compute)
- Top of rack switches (leaf switches) connected to core switches (top level switch - TLS)
- Compute network is usually private

Cluster interconnect

- Multiple options: ethernet, infiniband, slingshot, omniopath etc.
- Low latency and high throughput required (OpenMPI requires less than 1 μ s of latency)
- two dominant topologies: fat-tree and dragonfly (also torus, 2D/3D mesh, hypercube)

Network design gaining complexity

- where to put login nodes?
- adding storage to the cluster - scratch storage vs permanent storage?
- access to the permanent storage via login node or direct?
Authn/authz?
- some users need VMs for their UIs (cloud), access to storage?

Another consideration: IPv4 vs IPv6

- problems of IPv6 network design?
- DHCP IPv4 vs IPv6?
- dual stack, public IPv6 only, private IPv4/IPv6?

Our case

- 10 Login nodes
- 2 management nodes
- IPMI devices, monitoring tools, cooling management, power management services
- 30 S3 storage nodes + disks (eg. ceph)
- 30 storage nodes + disks for permanent storage (eg. ceph - dCache)
- 60 storage nodes for high speed scratch (eg. Lustre)
- 30 virtualization nodes for cloud services
 - VMs for backend services (DNS, squid, websites, monitoring tools, configuration management etc)
 - VMs for end users (private cloud)

Your network design

In groups of 5-6 discuss the possible network design:

- How many VLANs/subnets and argument why?
- Will you use NAT, if yes for which services and why?
- Which services will run on private, which on public network?
- Where would you put the firewall?
- Consider IPv6/IPv4 network, dual stack or exclusive?
- Estimate the number of public IPv4 addresses
- How will the storage be accessed? Directly, via login node, another jump host?
- How will you segment remote network?
- How to plan your cloud infrastructure? What are the considerations?