



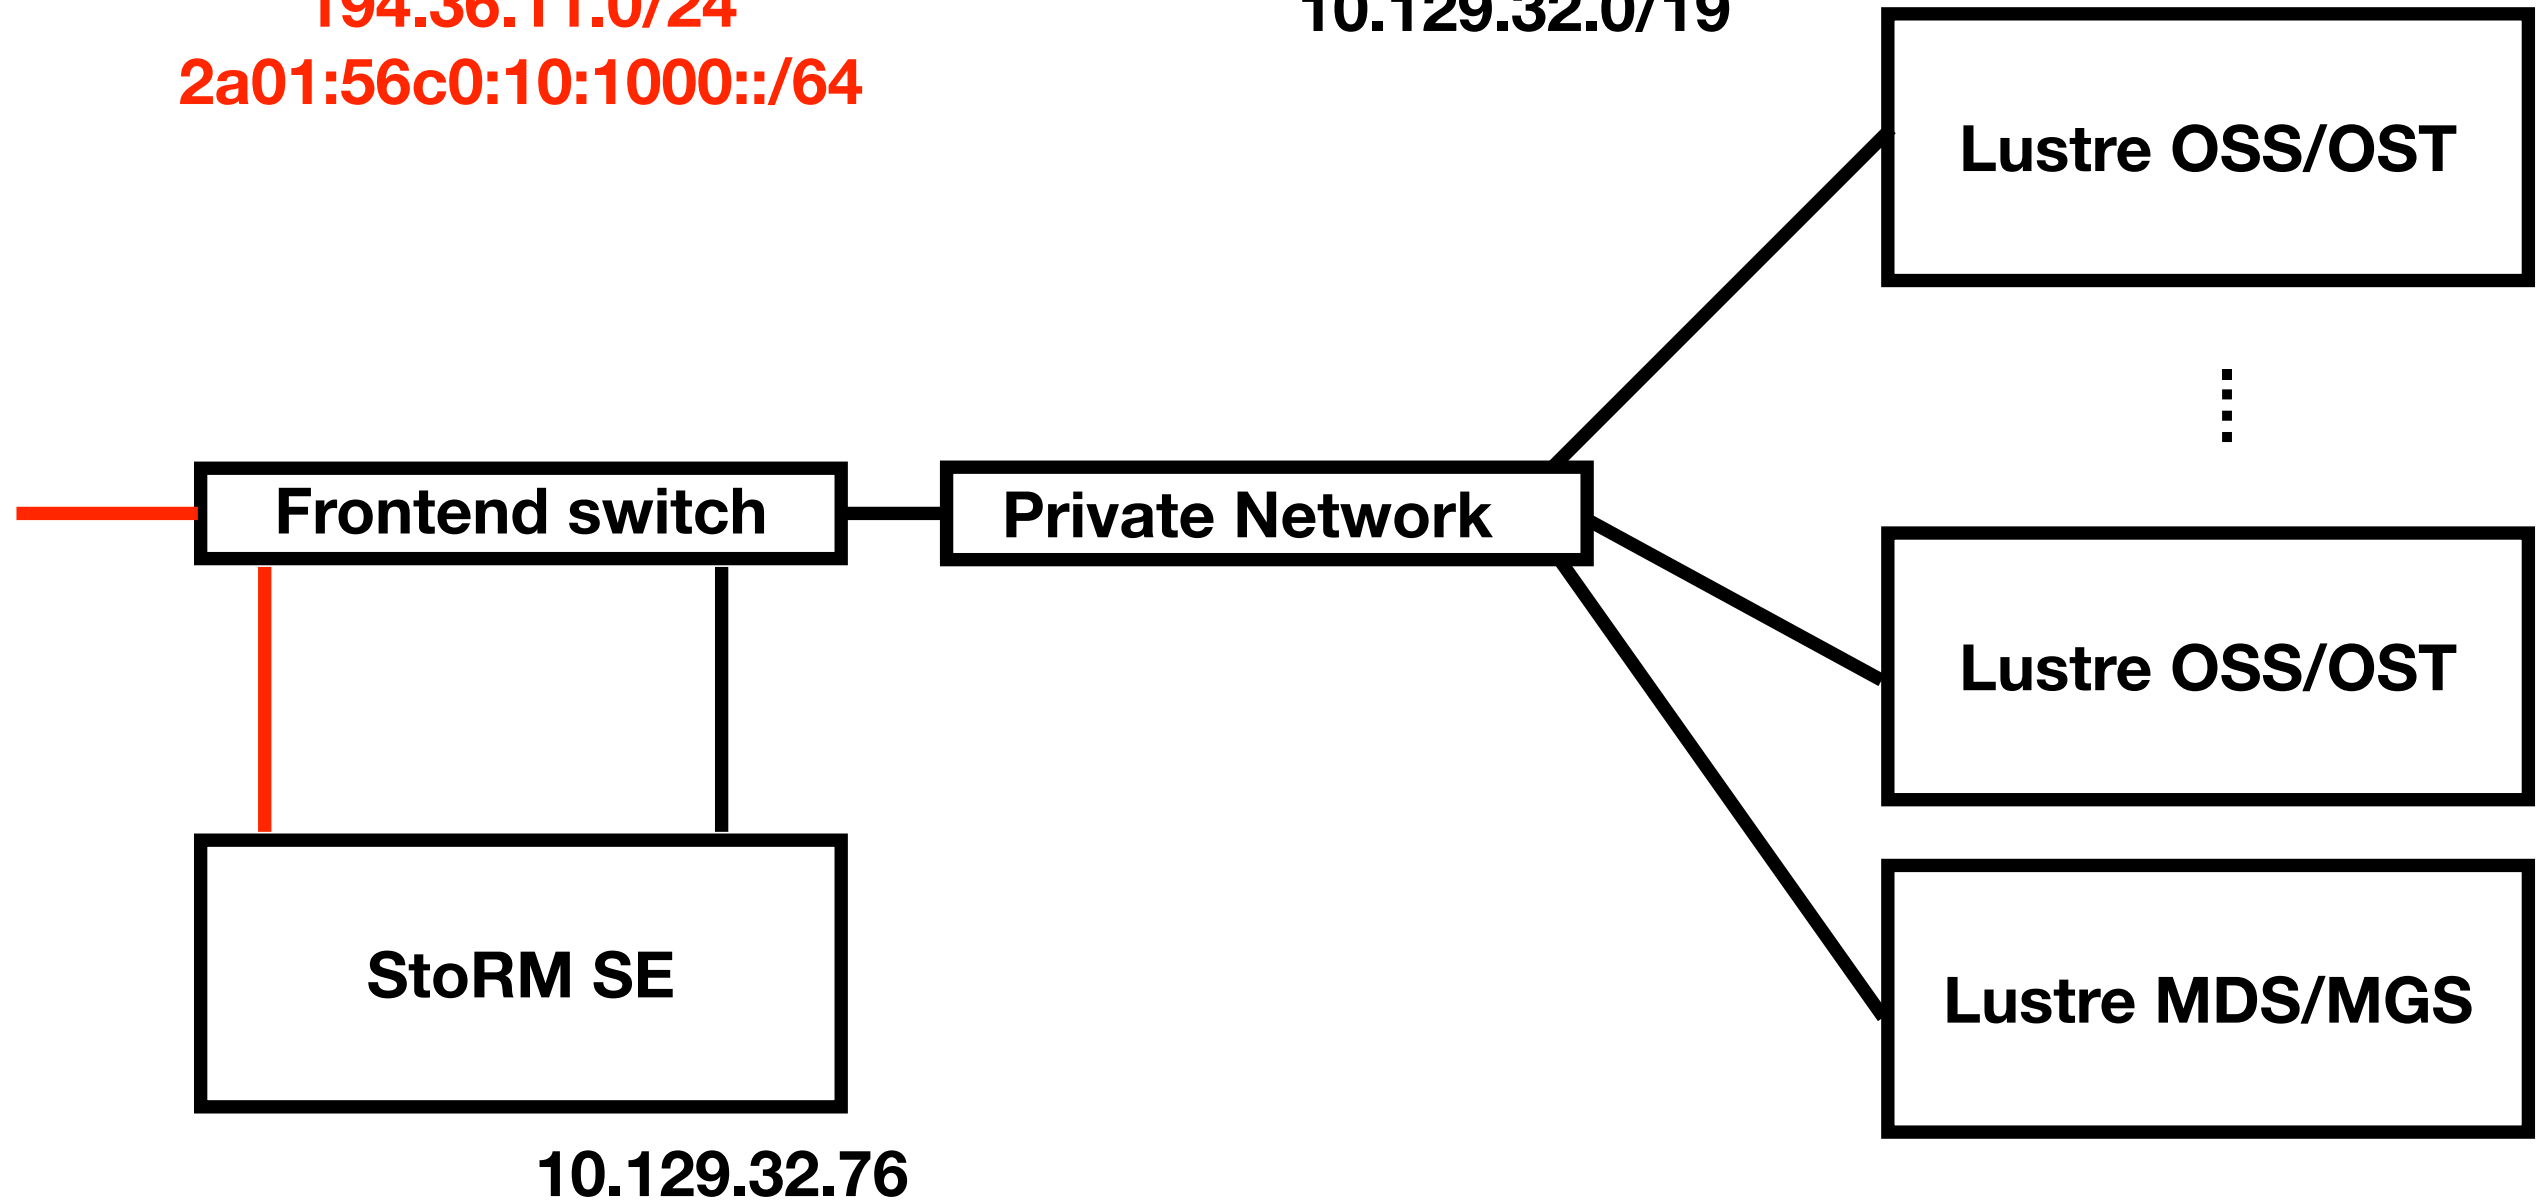
QMUL IPv6 status

- Given recent discussion expanding IPv6 on wlcg to more than storage (ULA/RFC1918/routing ...), I thought it worth while to show how a we (QMUL) have set up its IPv6.
- Having had a network professional (Terry Froy) on the team helps get the right answer.
- Need to make sure ordinary sysadmin can maintain and evolve the service in the long term.

IPv6 on Storage

External Public network
VLAN 843
194.36.11.0/24
2a01:56c0:10:1000::/64

Internal private network
VLAN 4000
10.129.32.0/19



194.36.11.44
2a01:56c1:10:1000::c224:b2c

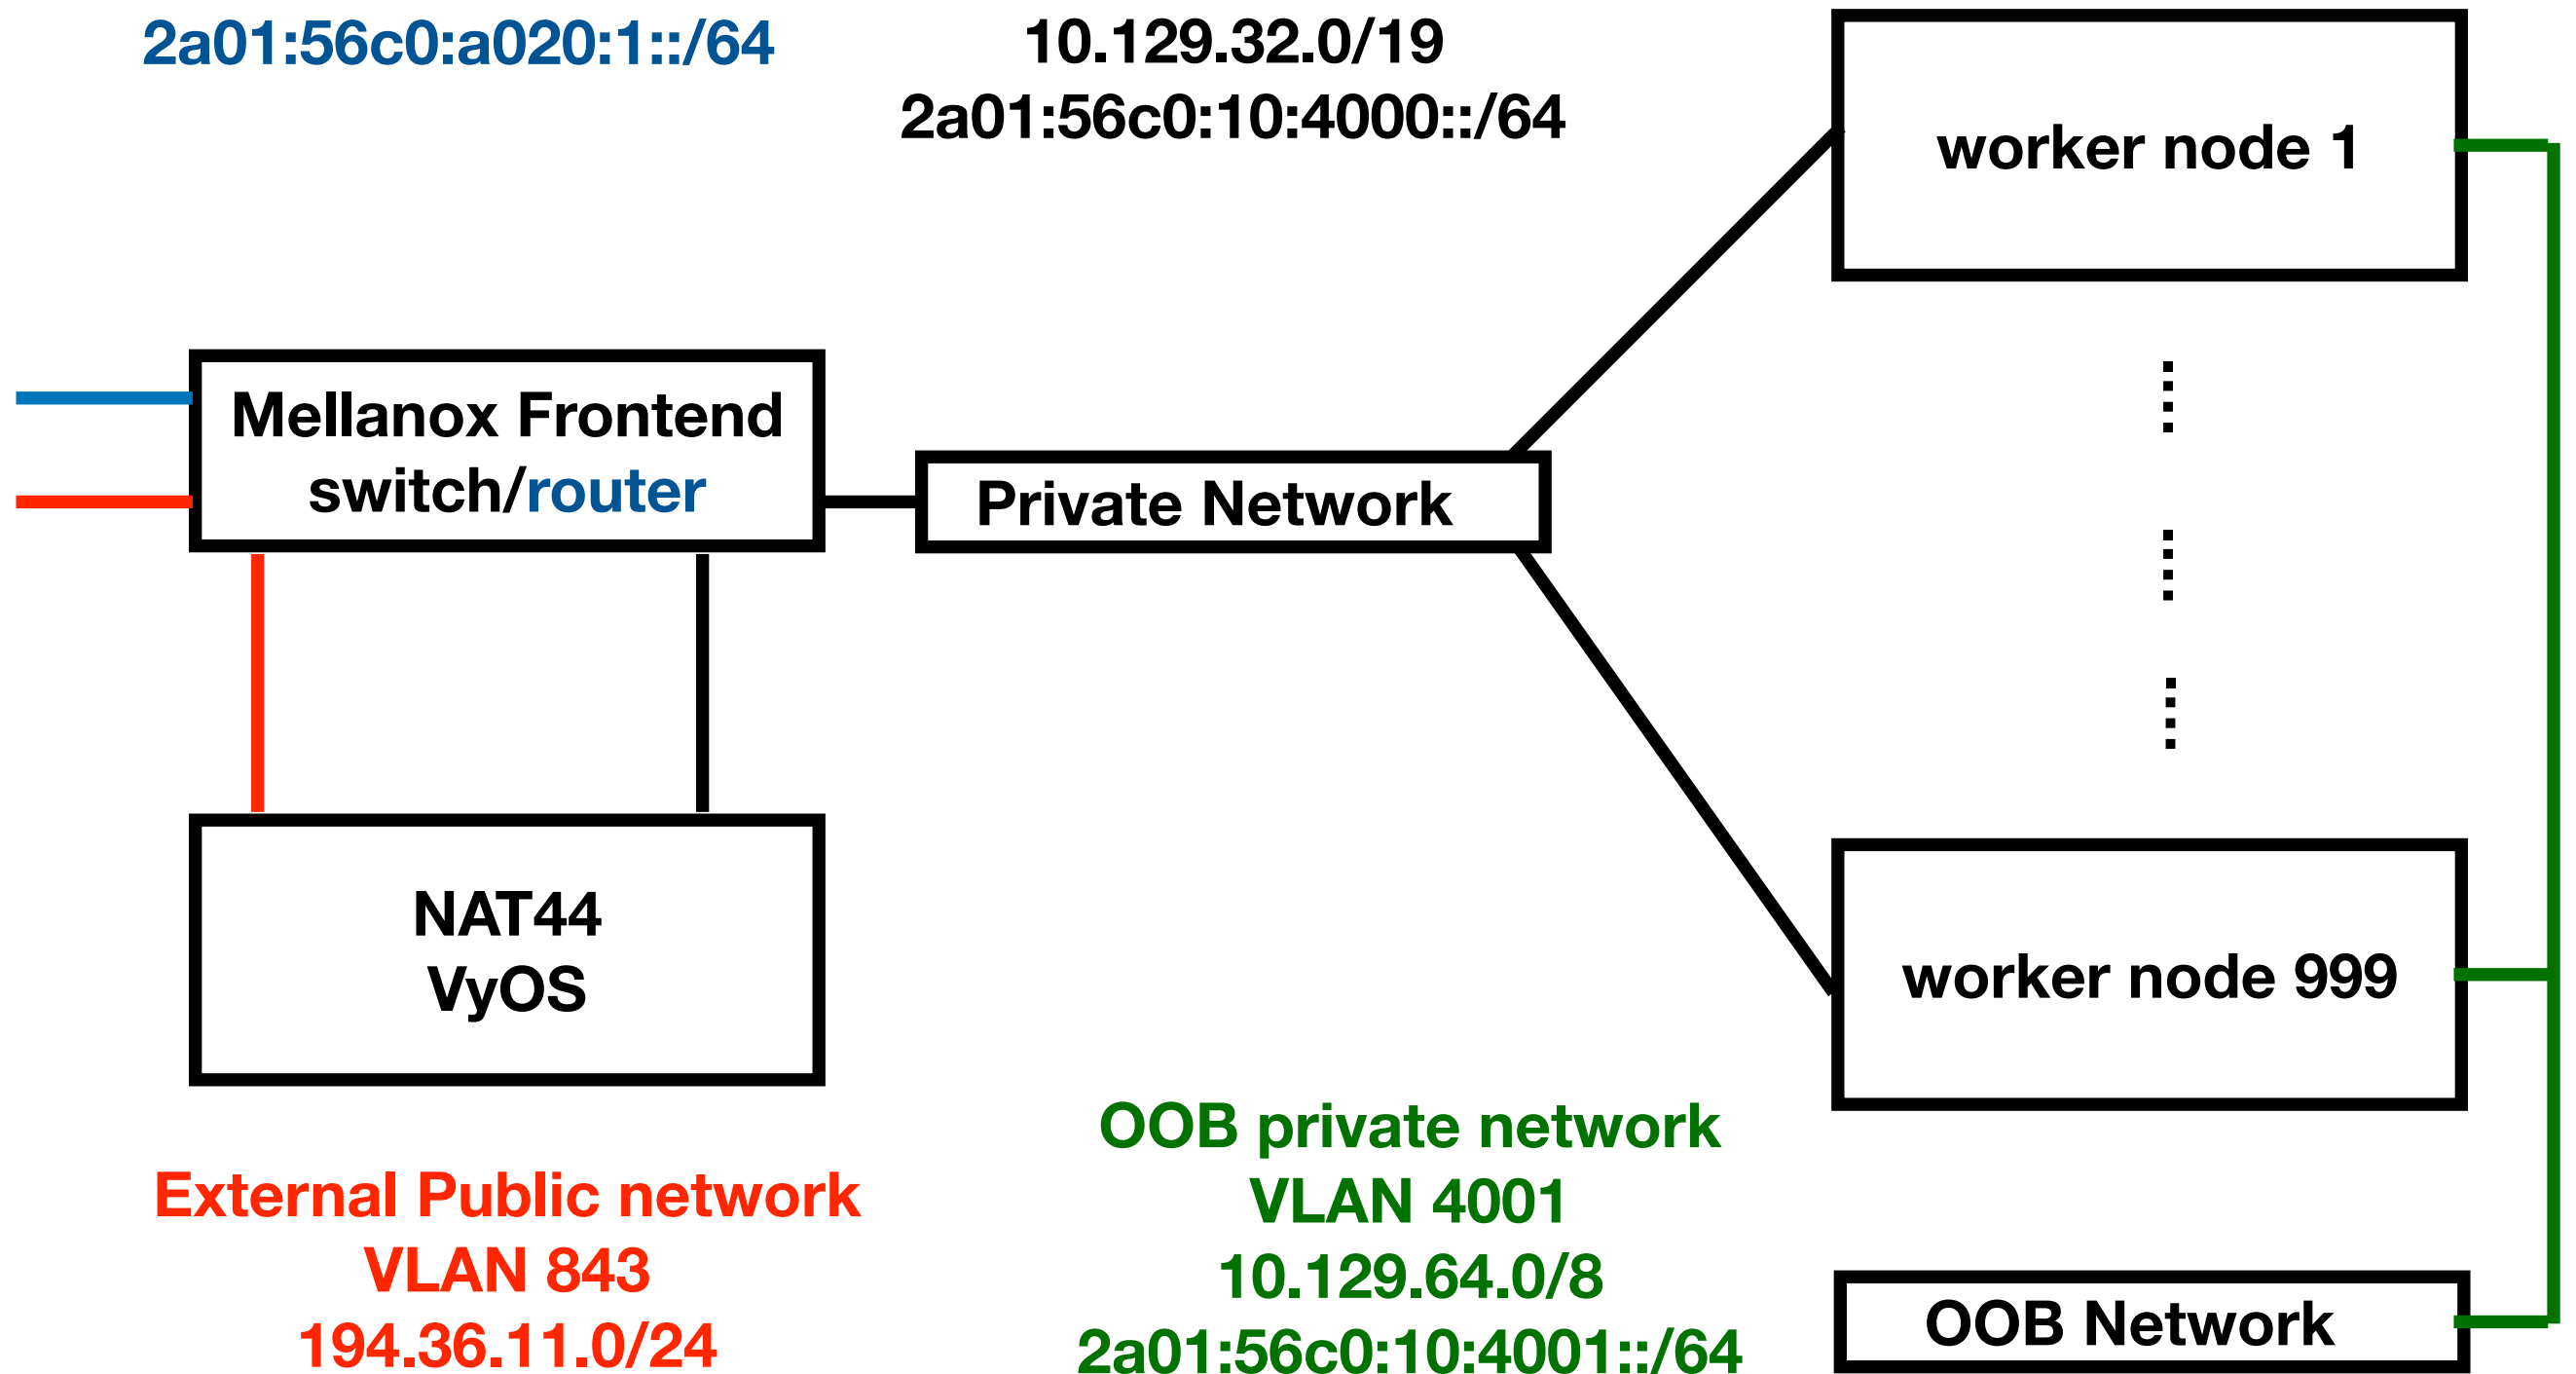
Notes

- Just added IPv6 /64 to VLAN 843.
- Originally internally used default VLAN (1) and 10.1.0.0/16. After discussion with central networks team moved to 10.129.0.0/16 and VLAN 4000 to 4010 (now reserved for our use only across campus).
- VLAN 4000 and 843 are physically separate (separate network ports). Historic but kept for practical and performance reasons.
- If a worker node writes to local SE we have LUA records to redirect to SE via private interface.
- A call to getaddrinfo (glibc) might return multiple answers. StoRM SE requires gai.conf file to correctly sort IP address (RFC 3484)
- IPV6 addresses use node's IPv4 address as the end octant. This is in DNS. Assigned statically.
- Central networks did not turn off router advertising (SLAAC).

Worker Nodes with IPv6

Linknet to campus router
VLAN 1039
2a01:56c0:a020:1::/64

Internal private network
VLAN 4000
10.129.32.0/19
2a01:56c0:10:4000::/64



Mellanox Router Config

##

L3 configuration

##

```
vrf definition internet
vrf definition internet rd 10.129.88.37:2
ip routing vrf internet
interface vlan 843 vrf forwarding internet
interface vlan 1039 vrf forwarding internet
interface vlan 4000 vrf forwarding internet
interface vlan 843 mtu 9000
interface vlan 1039 mtu 9000
interface vlan 4000 mtu 8252
ipv6 routing vrf internet
```

##

IPv6 configuration

##

```
interface vlan 4000 ipv6 nd ra lifetime 3600
interface vlan 4000 ipv6 nd ra interval max-period 300
interface vlan 843 ipv6 address 2a01:56c1:10:1000::1/64 primary
interface vlan 1039 ipv6 address 2a01:56c0:a020:1::a/64 primary
interface vlan 4000 ipv6 address 2a01:56c1:10:4000::1/64 primary
interface vlan 4000 ipv6 nd ra dns-server 2a01:56c1:10:4000::53:1
interface vlan 4000 ipv6 nd ra dns-server 2a01:56c1:10:4000::53:2
interface vlan 4000 ipv6 nd prefix 2a01:56c1:10:4000::/64
ipv6 route vrf internet ::/0 2a01:56c0:a020:1::1
```

#

ACL configuration

##

```
ipv6 access-list private-cluster-network
ipv6 access-list private-cluster-network bind-point rif
ipv6 access-list private-cluster-network seq-number 1 deny tcp any 2a01:56c1:10:4000::/64 ack 0 syn 1
ipv6 access-list private-cluster-network seq-number 65535 permit ip any any
interface vlan 1039 ipv6 port access-group private-cluster-network
```

Notes

- University has been allocated /29 from RIPE.
- GridPP has been assigned a /44. We have then assigned /56 to the GridPP cluster, other GridPP projects will get separate /56 each.
- Maintaining physical barrier between world and cluster by using the router ability of the Mellanox switch. The router has no complex requirements. Small number of MAC addresses to deal with.
- We route 2a01:56c0:10:4000::/64 externally via VLAN 1039.
- We do not route 2a01:56c0:10:4001::/64, this is the Out Of Band (IPMI, etc...) IPv6 allocation. Also physically separate from main network.

Future Plans

- SLURM, our batch system, can use IPv6 instead of private IPv4. Set “CommunicationParameters=EnableIPV6”. Need IPv6 addresses in DNS.
- Lustre has active project to enable IPV6. Production LTS version with Lustre IPv6 likely a few years away.
- Will then be able to drop main internal private IPv4 network.
- OOB(IPMI etc): only a few nodes can't do IPv6, Can move over to IPv6 only soon but not a priority.
- PXE booting?