

Durham Forward Look

There and Not Quite Back Again...

Presented by Oliver Smith

30th June 2016

Our Environment

- Local group system
 - 80 Desktops ~320 Cores
 - 13 Physical / 12 Virtual 'Servers'
 - ~100TB Storage
 - 3 Dedicated Batch machines ~88 Cores
 - Virtual Hosts / Web Servers
- Grid Cluster
 - 108 Compute Nodes ~2616 Cores
 - ~220TB Storage
 - Virtual Hosts / UIs / Etc

Where We Were




Where We Are

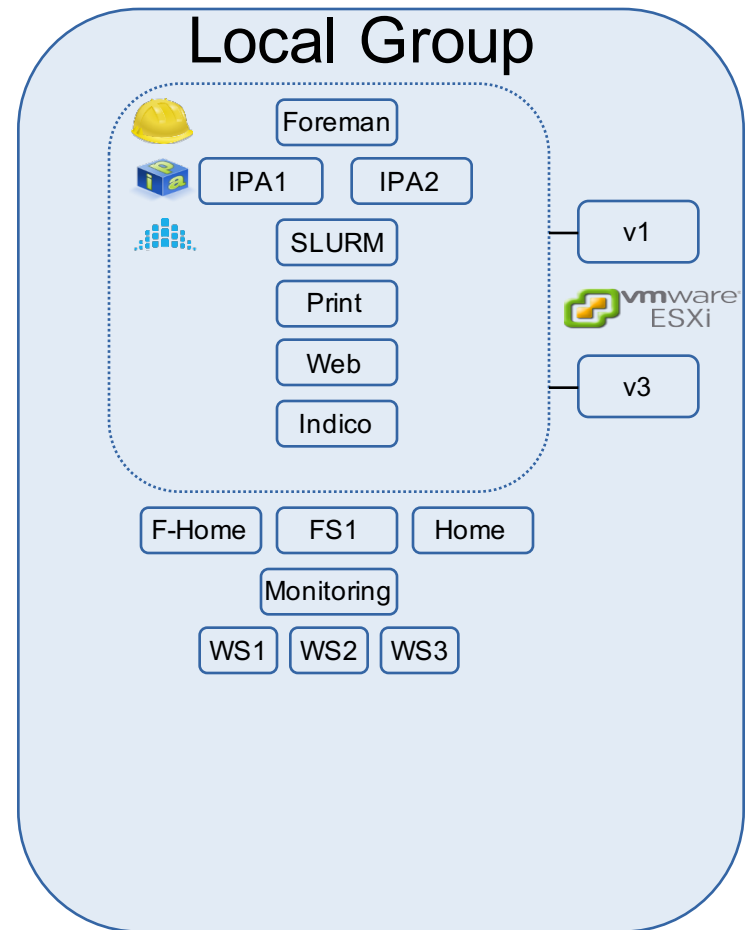
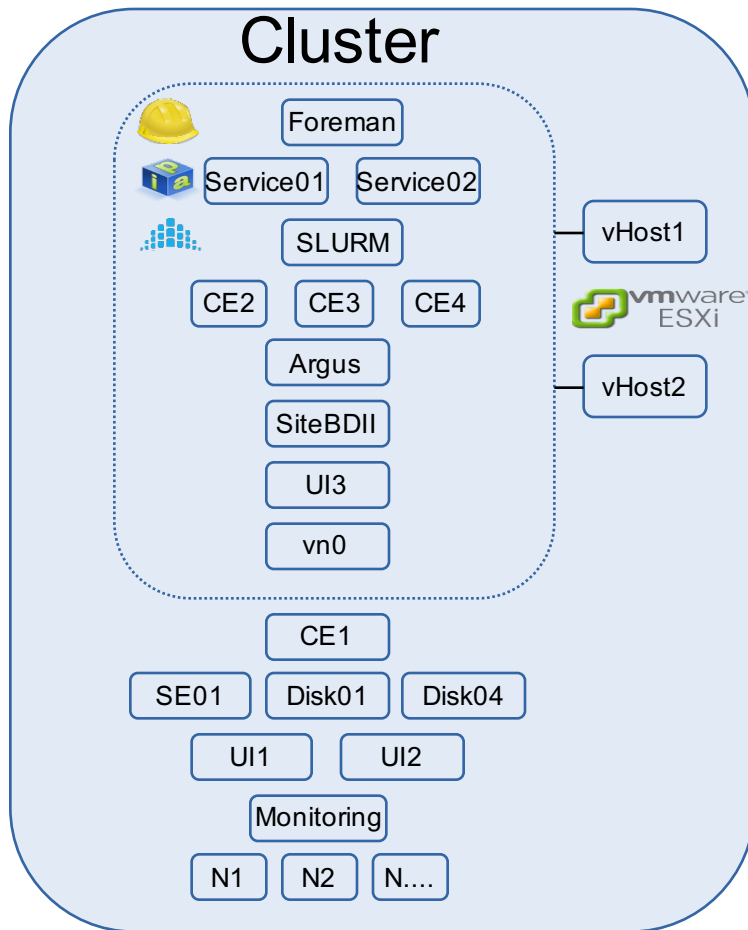


Current Infrastructure

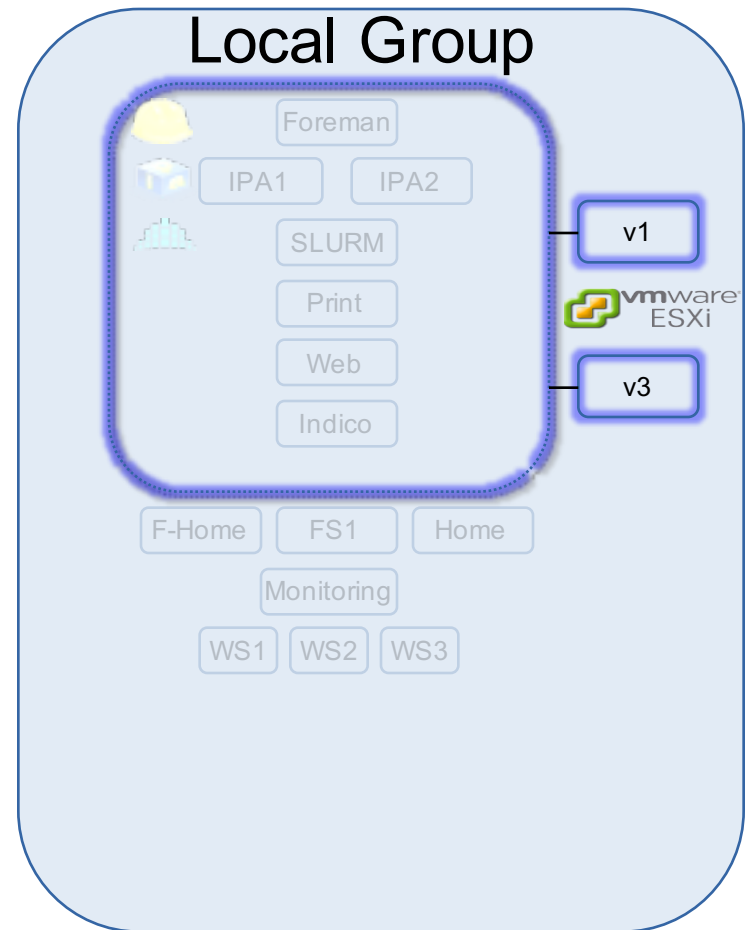
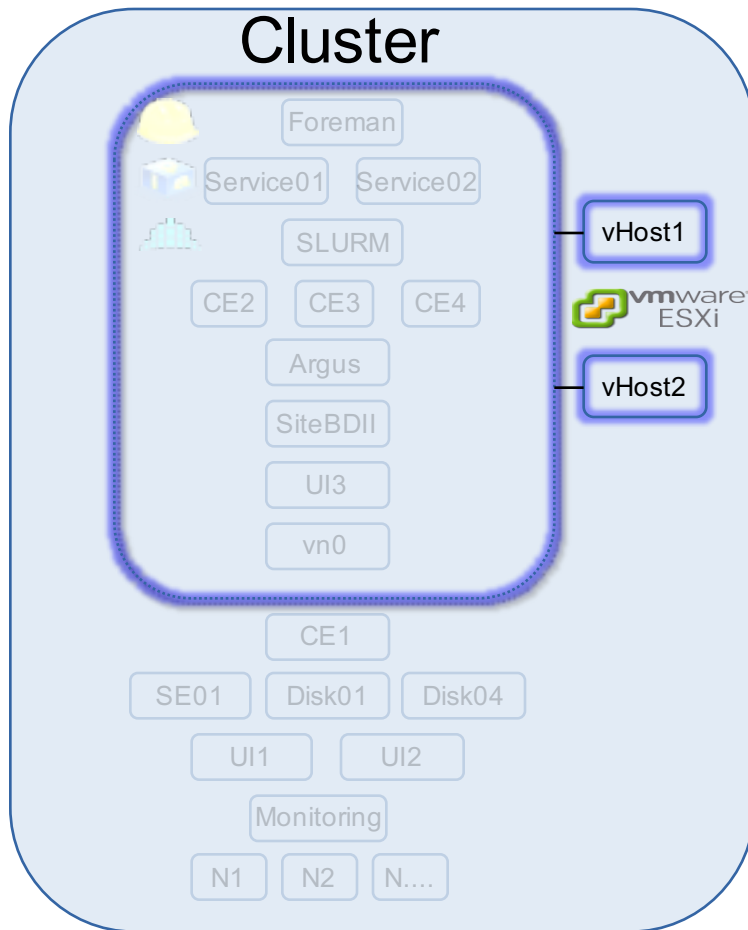
- Whole cluster setup scrapped and rebuilt in 2014
 - New physical space
 - New software stack with puppet and foreman
 - ~6 month very low usage period during rebuild
- Local group system rebuilt in 2015/2016
 - Modelled on the grid rebuild, simplify and automate everything
 - Still in the old space
- Both built to operate independently based on previous experience



Current Infrastructure


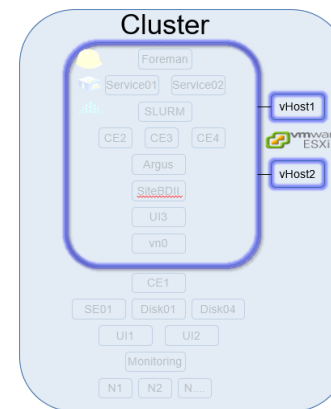


Issue 1 – Virtual Hosts

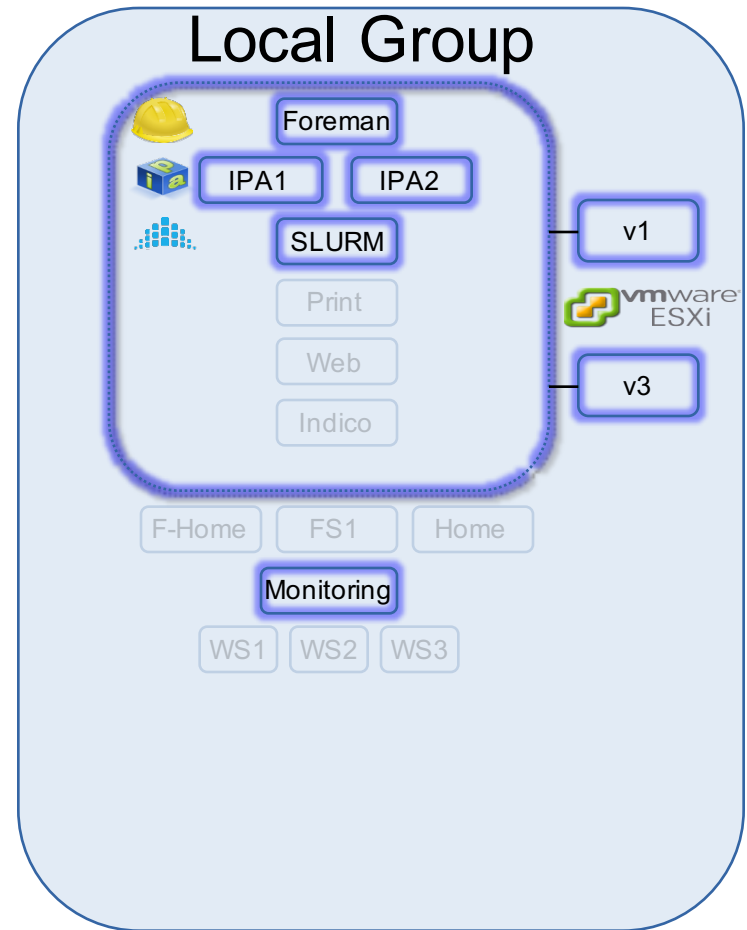
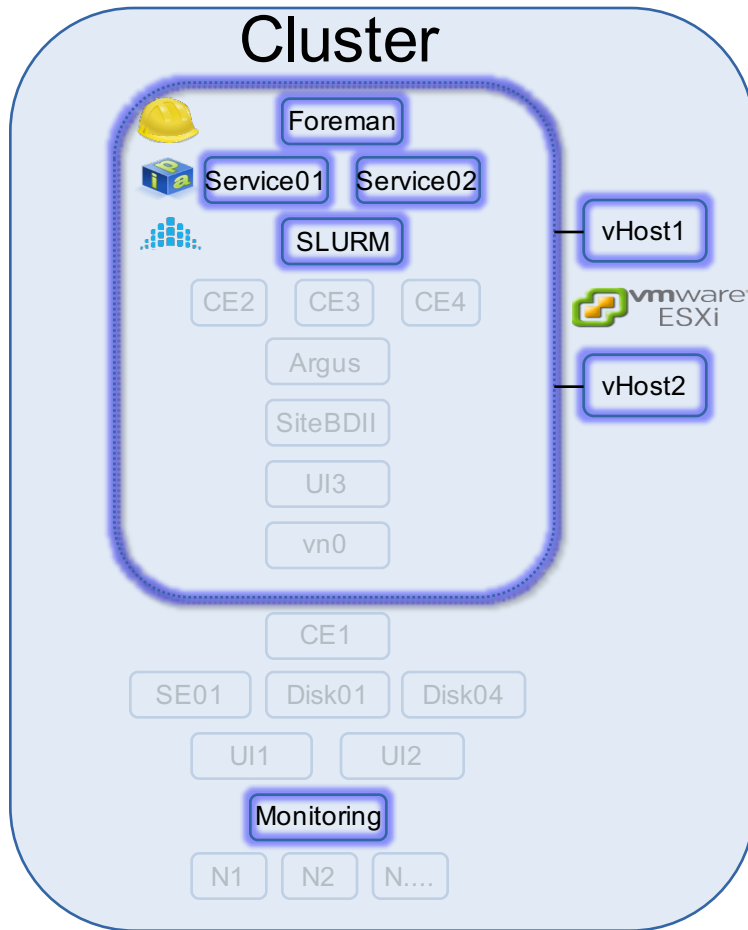


Issue 1 – Virtual Hosts

- Two ESXi hosts per system
- Very solid and stable platform with a great feature set...
- But we have the basic (free) licence missing all the exciting functionality
 - No shared storage – Poor IO to the local disks
 - No redundancy for host loss – can you guess why we have 18 CEs?
 - No stable backups
 - Very few changes to running VMs
 - Upgrading hosts live is not possible
 - No access for non sysadmins

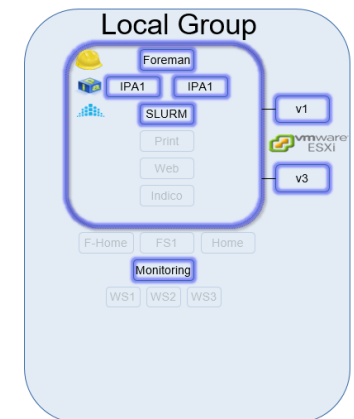
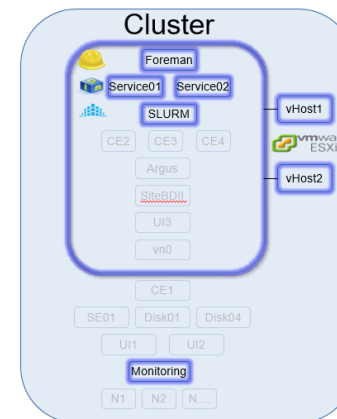


Issue 2 – Duplication

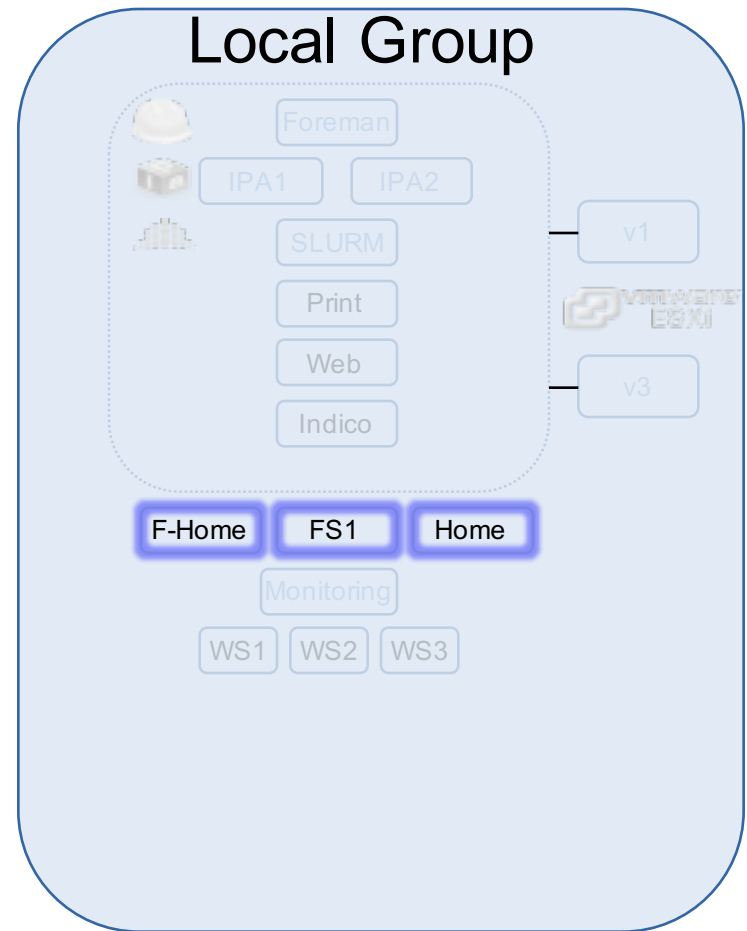
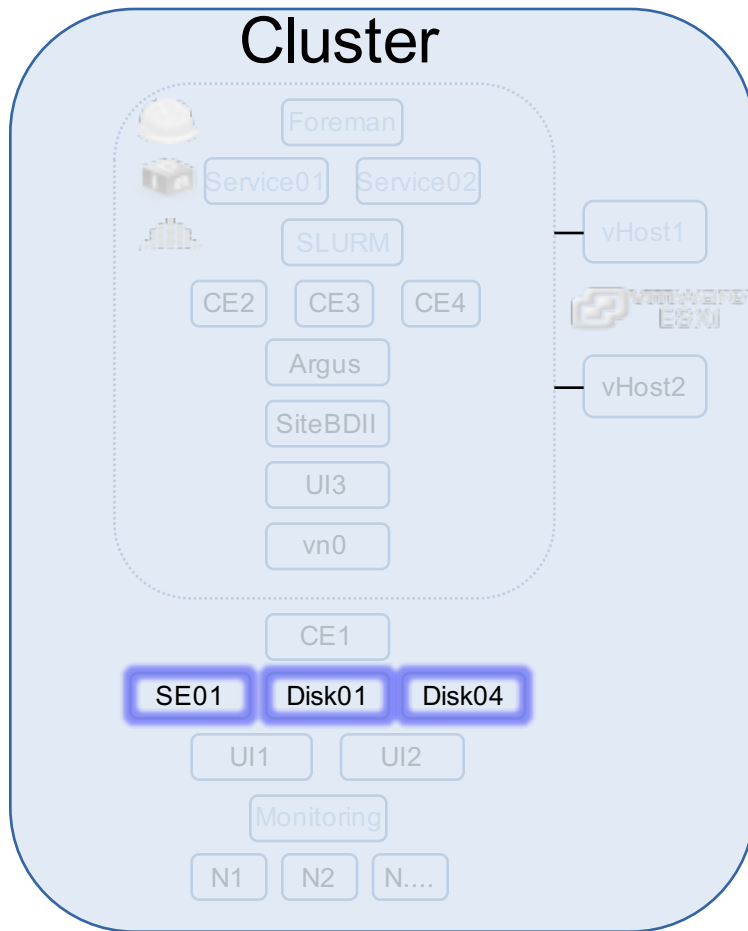


Issue 2 – Duplication

- Both systems are standalone
 - No risk of one breaking the other like we've previously had
 - Selling anything interconnected is hard
- But we now have two of everything
 - Foreman – Provisioning/Puppet
 - IPA x 2 – LDAP/DNS/DHCP
 - SLURM
 - Monitoring – Ganglia/Grafana/Icinga
- Duplication of effort
- Waste of hardware resources
- Confusion for users



Issue 3 – Storage





Issue 3 – Storage

- Not enough storage on the cluster
 - Cluster disk servers have to support multiple uses
 - DPM storage
 - Home space for local cluster users
 - Job session storage
- Too much storage in the wrong places
 - GlusterFS on compute nodes is our first attempt to solve this
 - It could potentially expose ~300TB of our underutilised storage but can affect node performance and access is slow due to the structure of DPM
 - Makes managing the nodes a nightmare
- No 'easy' access to grid storage

Issue 4 – Compute

- ARC and SLURM are very stable with few issues
 - Submission time can be a problem
- Never enough compute to meet ‘Pheno’ requirements
 - Usage patterns are very bursty
 - Aiming to double the cluster CPU capacity in the next three years
- Resources available on Local Group are idle the majority of the time
- Spare resources always available on our virtual hosts



Issue 5 – Network

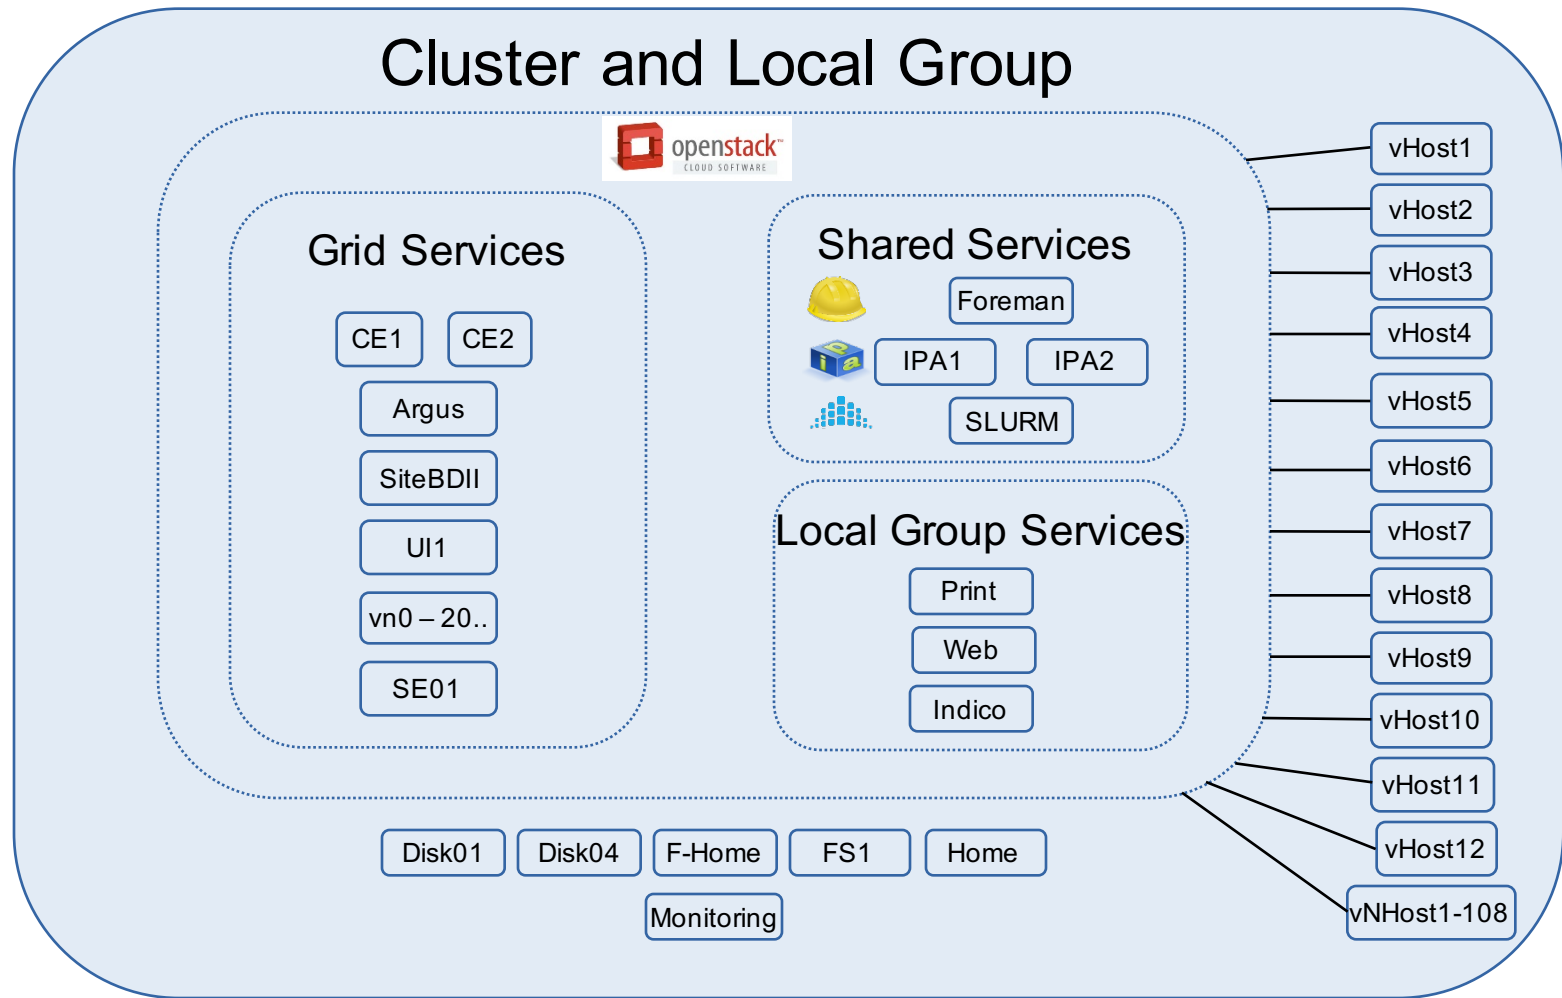
- Durham's internet connection – The Wet string™
- Cluster connection to Durham core network
 - Close to the edge but not close enough
 - Single point of failure
- More support from central IT service is necessary for any improvement
 - IPv6 network available since Feb 2015 but no rDNS for the foreseeable future
 - Slow connectivity between local group and cluster affects analysis and backup



Future Infrastructure

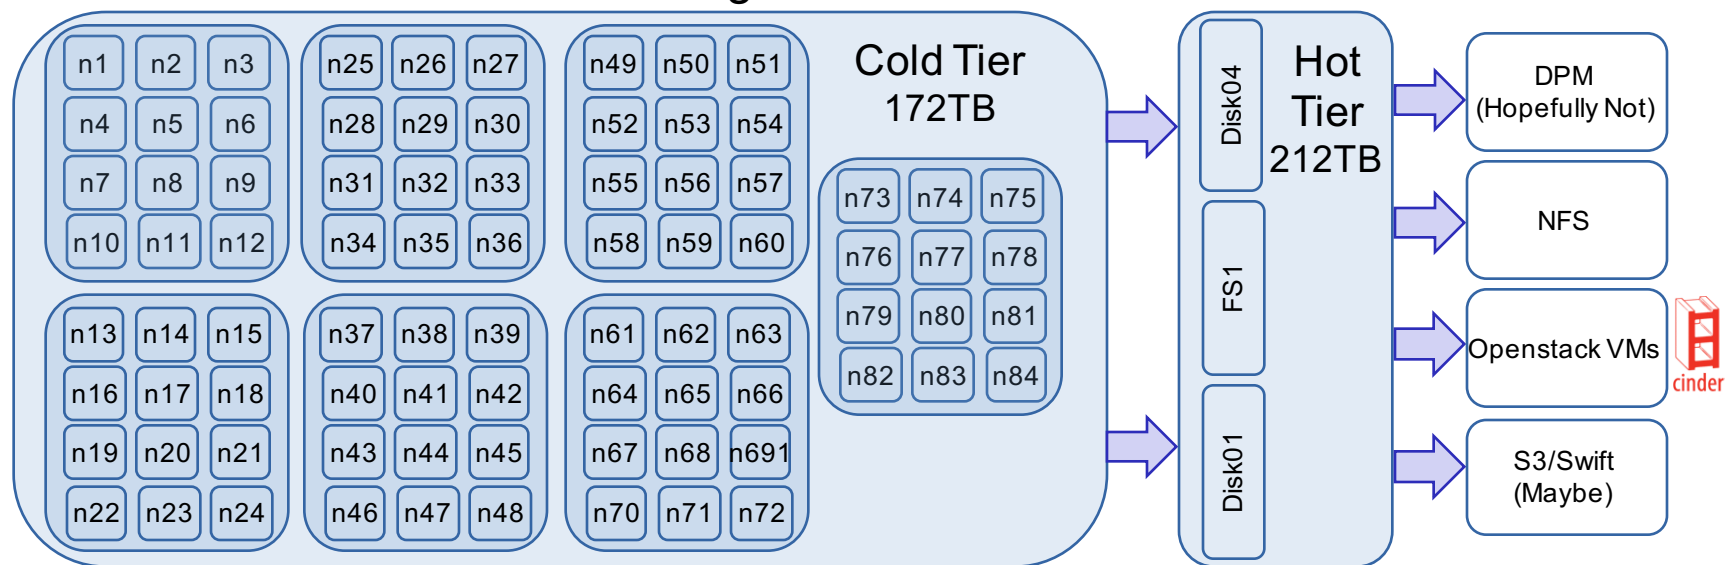
- There and back again – Merge the two systems where it makes sense whilst increasing reliability
 - Build a shared core that can handle the requirements of both systems
 - Merge the duplicated services
- Simplify and automate even more
- Share hardware, physical space and network access

Future Infrastructure – Core



Future Infrastructure - Storage

- Merge home space for both Grid users and Local Group users, consolidate to one physical fileserver per internal group
- Expand GlusterFS to provide a single, tiered storage system spread across all nodes and existing file servers



Future Infrastructure - Compute

- Reduce CEs to two – allow submission for all VOs on each
- One SLURM system
 - Queue for desktop batch with existing desktops in and all nodes (Low Priority)
 - Other queue for grid cluster with all nodes (High Priority)
- Nodes all in Openstack Nova
 - One VM per node with all resources allocated – Fixed allocation in SLURM
- Spare resources in Openstack to be provisioned as cloud nodes on demand scheduled by SLURM

Questions?