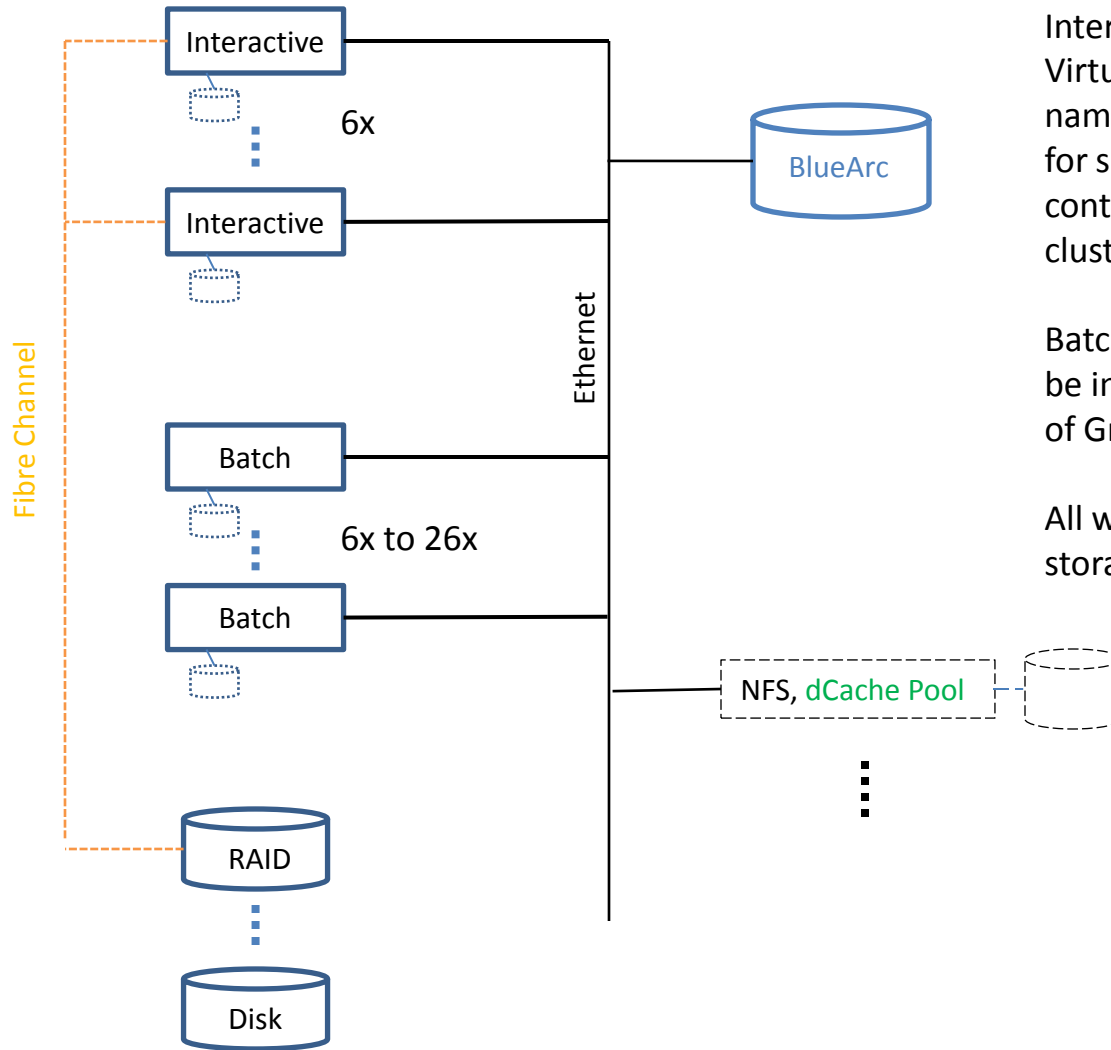


GPCF Status

20-Jan-2010

- Philosophy
 - Affordable
 - Leverage power of large purchases
 - Supportable
 - Standardized, high quality, supported systems
 - Minimal number of varieties
 - Professional and sustainable system administration
 - Expandable
 - Add capacity while maintaining other criteria
 - Flexible
 - Allocate capacity as needed
 - Upgrade / restructure hardware as needed
 - Performant
 - Attention to processing, storage, and I/O needs

GPCF Configuration - Base

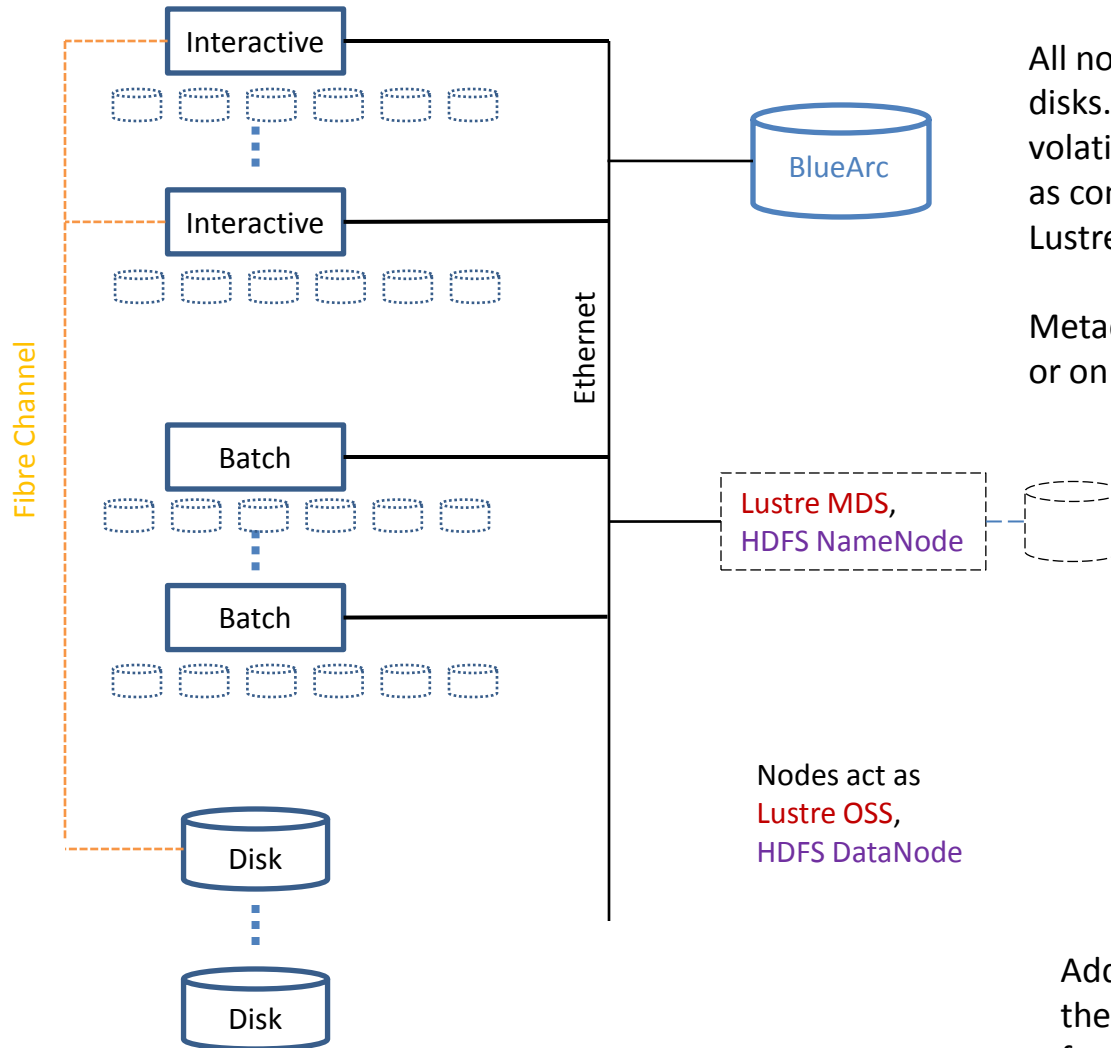


Interactive nodes are virtualized. Virtual machines will have static names/addresses. Local disk is for scratch use. SAN array will contain VM images plus provide clustered resilient storage.

Batch nodes not virtualized. Will be in configuration representative of Grid nodes.

All will have access to networked storage.

GPCF Expansion – Add disks



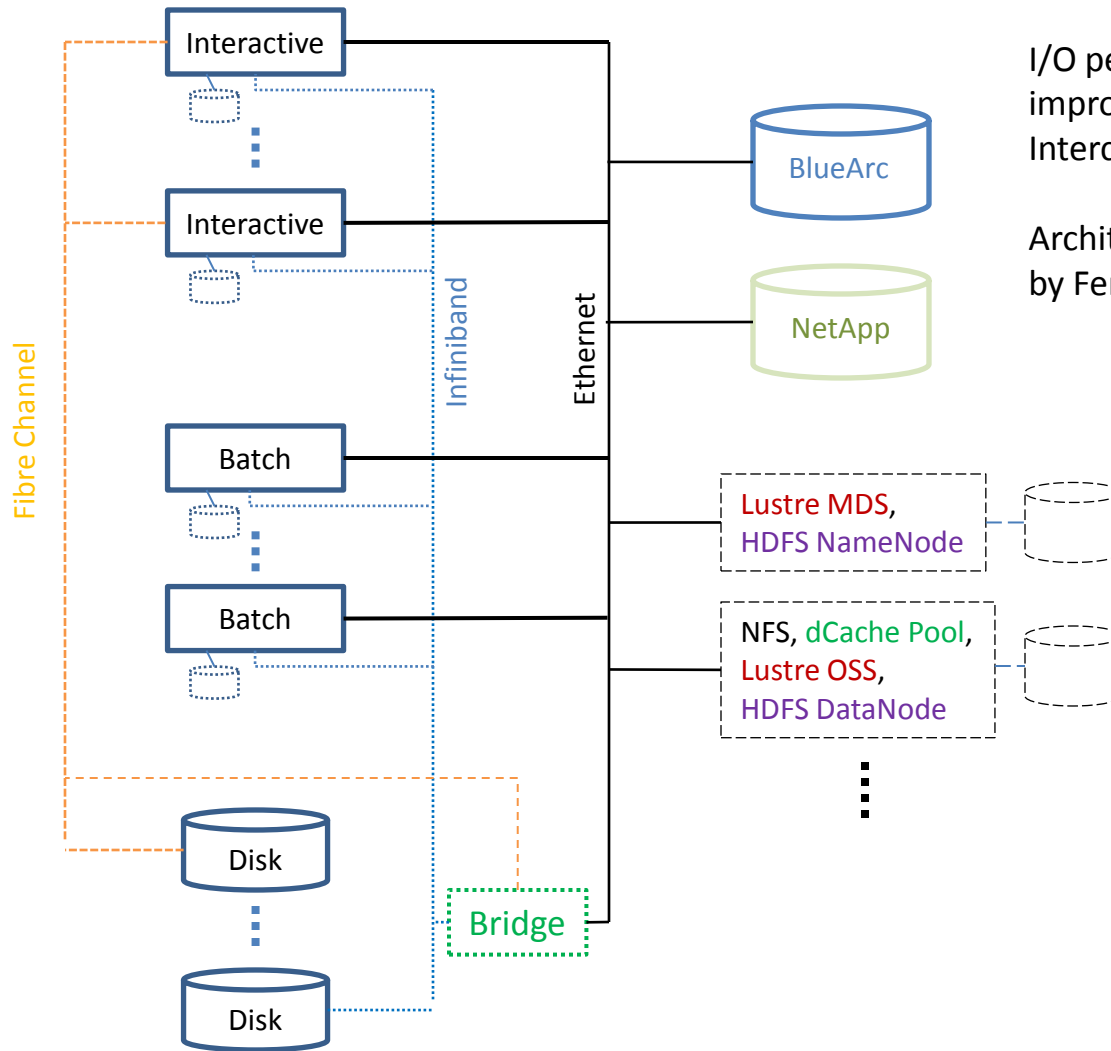
All nodes can have additional data disks. These could be configured as volatile local storage, or more likely as components of a large distributed Lustre or HDFS file system.

Metadata can be on one of the nodes or on dedicated metadata server(s).

Nodes act as
Lustre OSS,
HDFS DataNode

Additional servers can be added to the GPCF to provide specialized functions, possibly virtualized.

GPCF Expansion – Add interconnect, etc



I/O performance is potentially improved by upgrading the node Interconnect, e.g. with Infiniband.

Architecture decisions will be guided by FermiCloud studies.

GPCF Purchase Status

- Base purchase of 6 interactive + 6 batch nodes
 - Dual quad-core Intel E5540
 - 24 GB memory
 - RAID
 - Mirrored 250GB system disks
 - (1) 2TB data disk, slots for 5 more
 - Fibre channel HBA on interactive nodes
 - ~ \$5K per node
 - This HW requires SLF5 (though SLF4 can reside in virtual machines)
- Option for 20 more batch nodes and peripherals
- Fibre channel SAN
- 30+ TB Nexsan Raid array, dual controllers, w/ FC interfaces
- Racks, PDUs, console server, management network
- ~ \$58K remains for other storage
- ~ \$15K remains for other servers
- ~ \$10K remains for network needs
- Reqs just leaving CD
 - Estimate late March delivery

Other activity

- Decisions to be made in advance of HW arrival
 - Final choice of virtualization software
 - User & group management schemes
 - Batch system
 - Resource allocation process (with REX)
- Implementation plans
 - Meetings to settle on physical location
 - Network configuration issues:
 - Physical connections
 - Address space (for VMs)
 - Determine allocations
 - Determine VM needs
 - “individualized” VMs
 - Storage / file system details

More on the “to do” list

- Plans for migrating:
 - fnalu functionality
 - Access to existing file systems
- Other relevant activity
 - FermiCloud purchase is proceeding:
 - 18 nodes, with option for additional 18
 - Each very similar to GPCF node, but have Infiniband and 6 x 2TB data disks
 - Will influence future GPCF storage directions