
BNL Network/Storage Update

Bruce G. Gibbard

4 March 2006

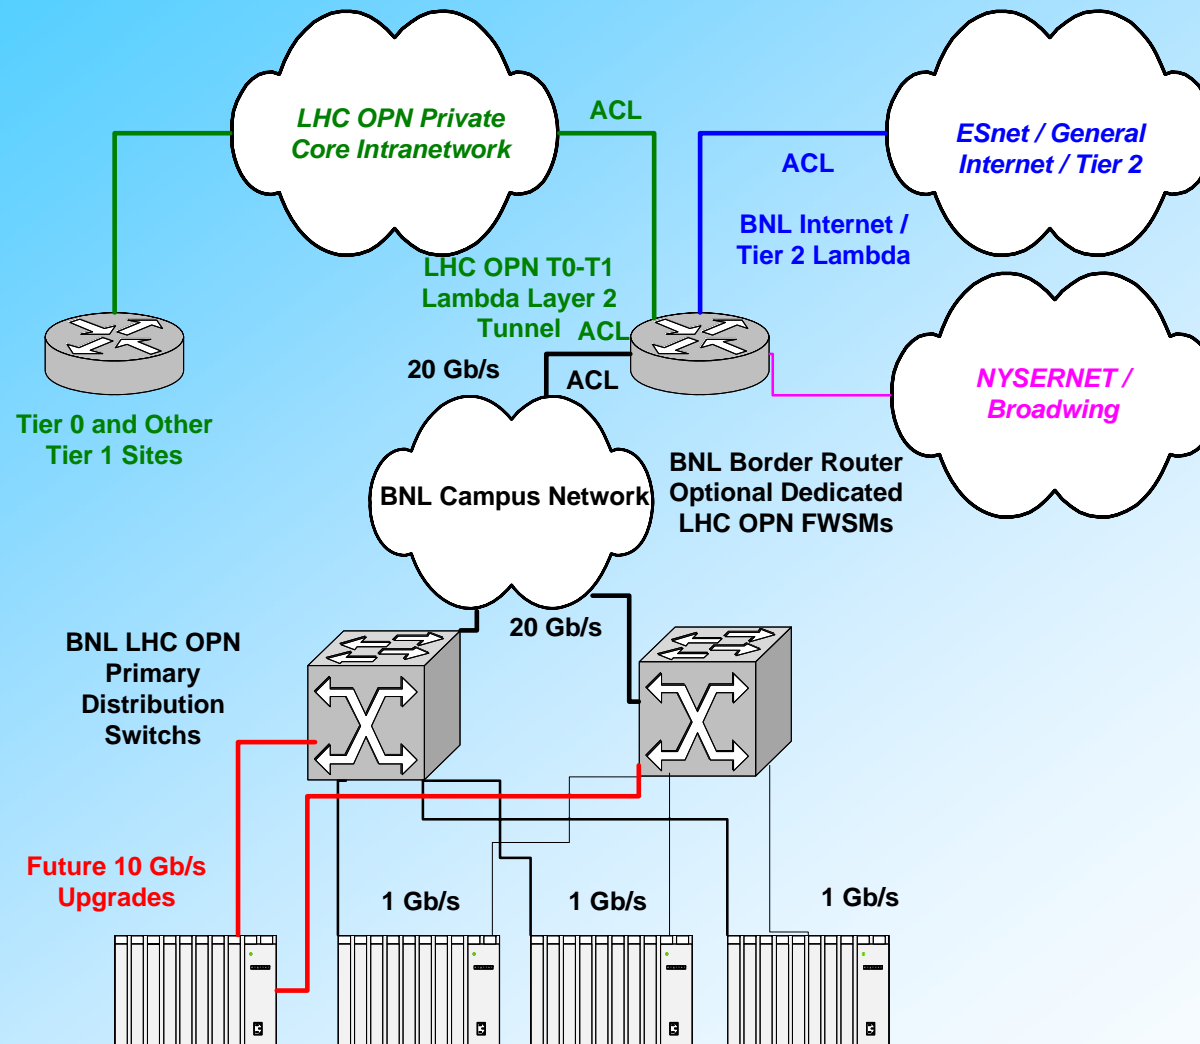
Brookhaven Science Associates
U.S. Department of Energy



Recent BNL LAN Upgrades

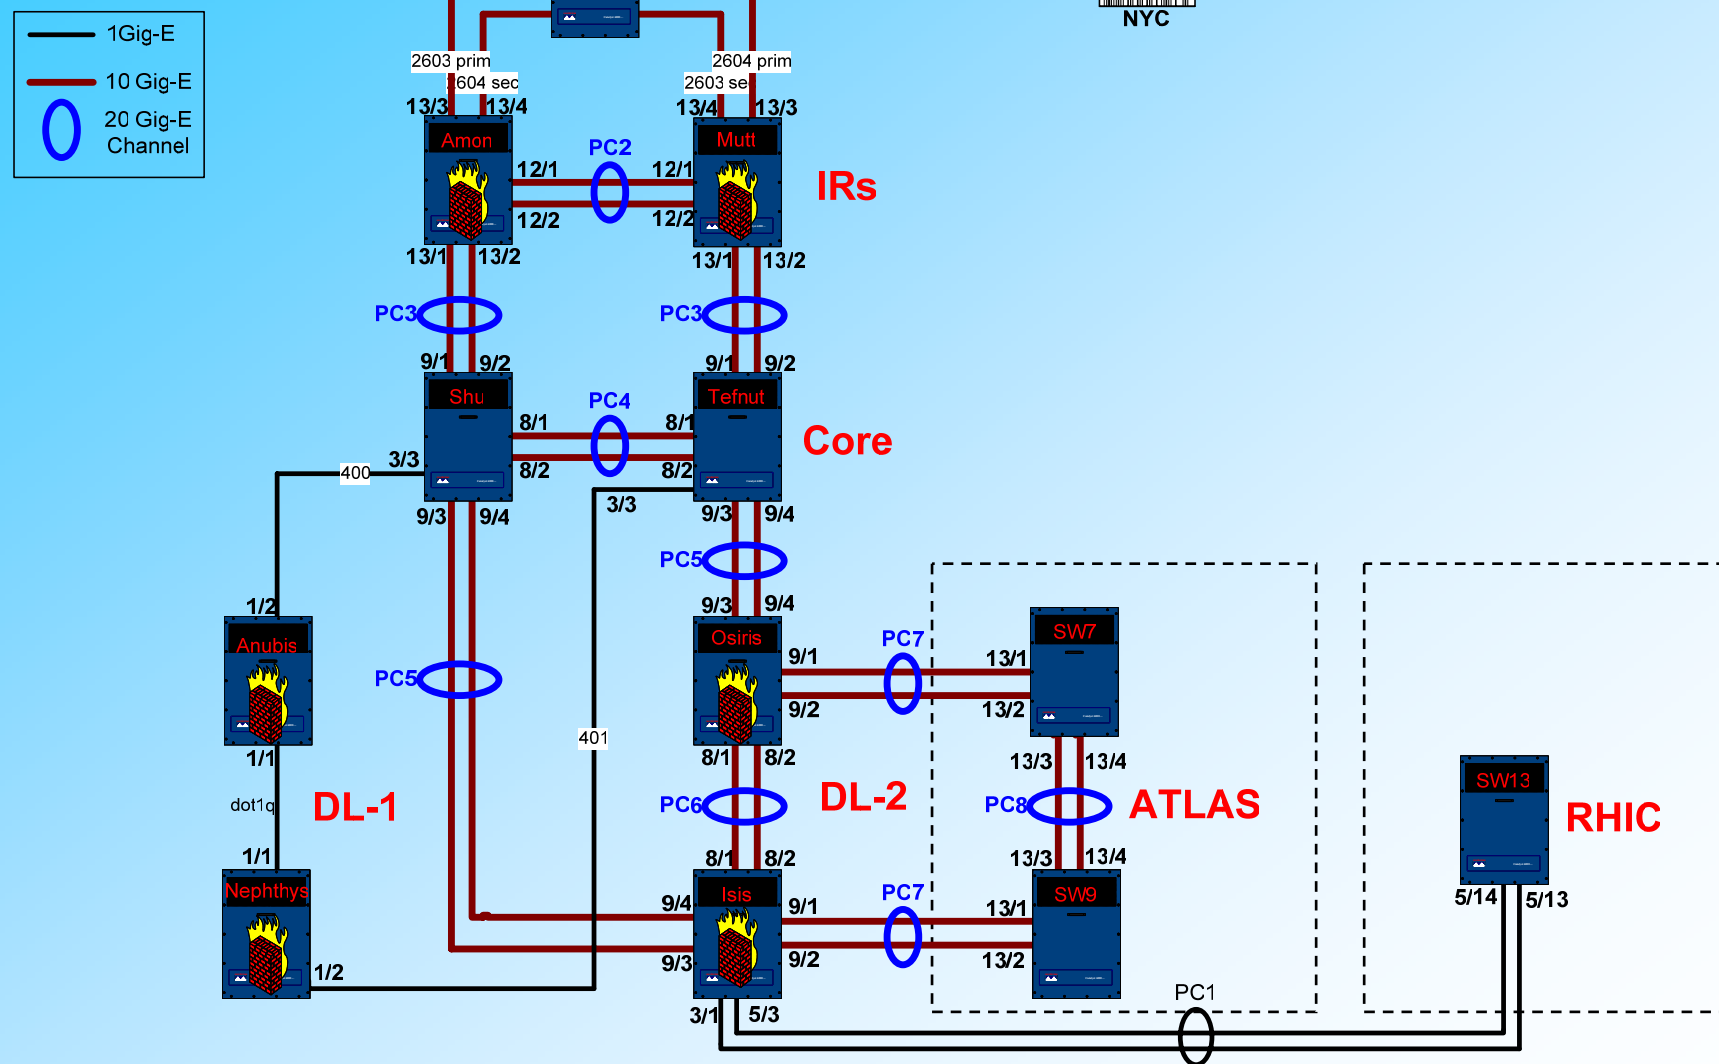
- Established dual 10 Gb/s Interfaces as 20 Gb/s EtherChannel 802.1q trunks
 - Hashing by source & destination (address and ip port)
- Established separate physical paths for primary and secondary VLAN traffic using spanning tree costing
 - Different VLAN's have different primary and secondary paths to load balance the architecture
- Established 10 Gb/s Generic Internet access via ESnet

BNL LHC OPN Conceptual Block Diagram



Storage servers multi-homed to bridge LHC OPN to US ATLAS Tier 1 network

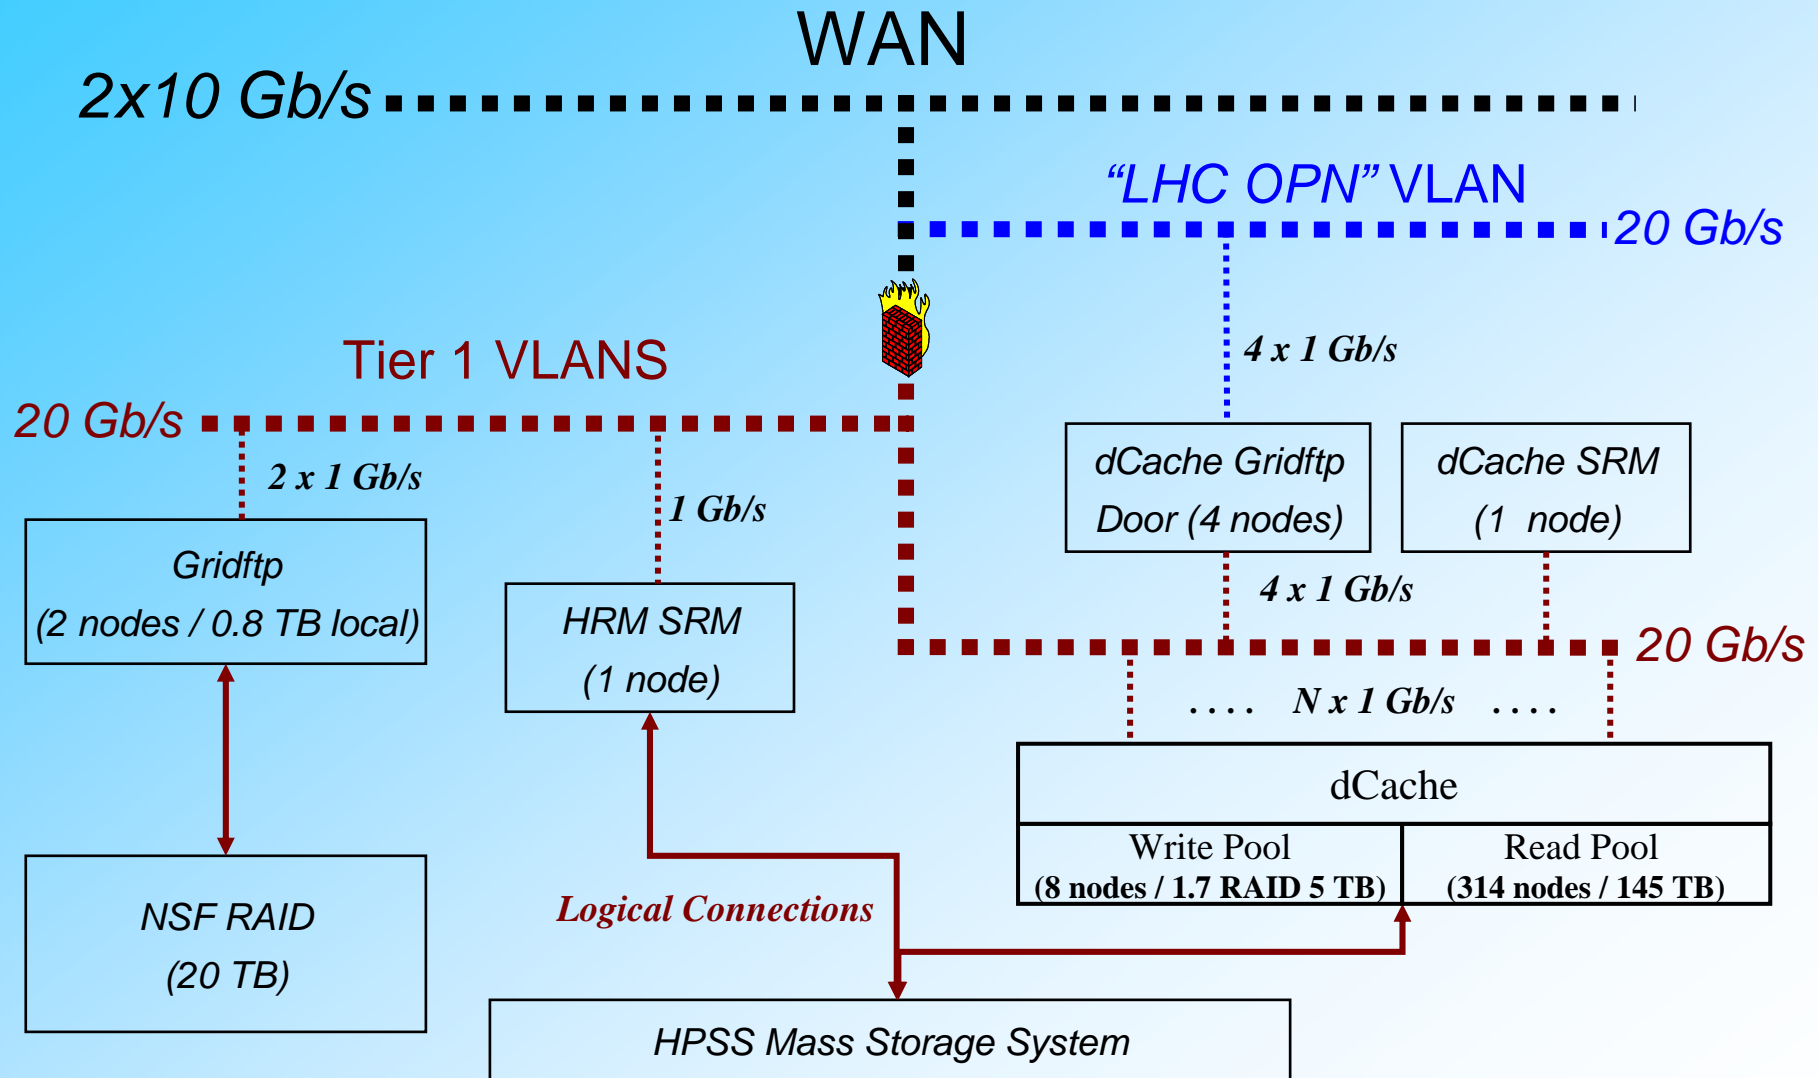
Current BNL 20 Gig-E Architecture



BNL Tier 1 dCache / Linux Farm

- US ATLAS Tier 1 at BNL uses its Linux farm
 - ... as its primary compute resource and
 - ... as its primary online storage resource via dCache
 - 100% overlap between standard worker nodes and dCache read pool servers
- Special function servers such as gridftp doors, write pool servers, etc., are single function and not used as worker nodes
- ATLAS Linux farm contains 300+ nodes with ~200 TB of local disk, most of which is part of the production dCache system
- The dCache system is backed by a Mass Storage System managed by the HPSS HSM
 - “Permanent” ATLAS storage is realized by migration through and staging back through the dCache system

BNL Tier 1 WAN Storage Interfaces



Upgraded BNL Tier 1 Mass Storage System

- HSM – HPSS Version 5.1.1
- Robotics – StorageTek/Sun SL8500
 - Physical capacity: ~ 6500 cartridges (2.6 PB for LTO Gen 3)
 - Current licensed capacity: 1448 cartridges
- Tape Drives – 10 x LTO Gen 3
 - 400 GB/cartridge uncompressed
 - 80 MB/sec streaming transfer rate
- Mover/Servers – (5 Tape + 4 Disk) x Dual Processor Linux Nodes
 - 4 x 1 Gb/sec EtherChannel connection per server
- Disk Cache – 9 TB Data Direct RAID 3
 - 4 port x dual channel controller connects to 4 mover/servers via FibreChannel
 - Theoretic bandwidth 1600 MB/sec
- Projected actual maximum ingest rate to tape \leq 400 MB/sec