FNAL Site Perspective on LHCOPN & LHCONE Future Directions

Phil DeMar (FNAL)
demar@fnal.gov
February 10, 2014

FNAL WAN Basics

- Aggregate WAN b/w:
  - 8x10GE migrating to 100GE + 3x10GE
  - All but 2x10GE allocated for “science data” movement
- LHCOPN = ~17Gb/s
  - 2x8.6Gb/s with addtl 3Gb/s for backup
  - Subjective evaluation: current b/w is adequate
- LHCONE = 10Gb/s
  - Subjective evaluation: current b/w is adequate
- E2E data circuits:
  - With 6 of 7 US T2s
  - Guaranteed 1Gb/s
    w/ scavenge to 10Gb/s
    - Routinely peak at 8-9Gb/s

FNAL WAN Security Model

- FNAL does not have a site firewall
- Site security based on wide spectrum of controls
  - Strong auth., vulnerability scanning, ACLs, IPS, web proxy, etc
- By default, science data must pass thru perimeter controls
- Bypass exception:

  "Known traffic from well-managed systems at trusted sites"

  - Exception based on risk analysis and acceptable residual risk
  - Bypass implementation = policy routing ACLs on the border
- LHCOPN & LHCONE traffic generally via policy route ACLs
- No reliance on security controls of external networks
  - Added layer protection is nice, but not essential

FNAL Tier-1 WAN Data Path(s) Today

- CMS Tier-1 integrated into campus network
- Routed IP traffic to T1 goes thru border routers
  - Bypass available for identified traffic
  - Security controls on the rest
- Separate border router for science data paths:
  - LHCOPN & LHCONE
  - E2E circuits

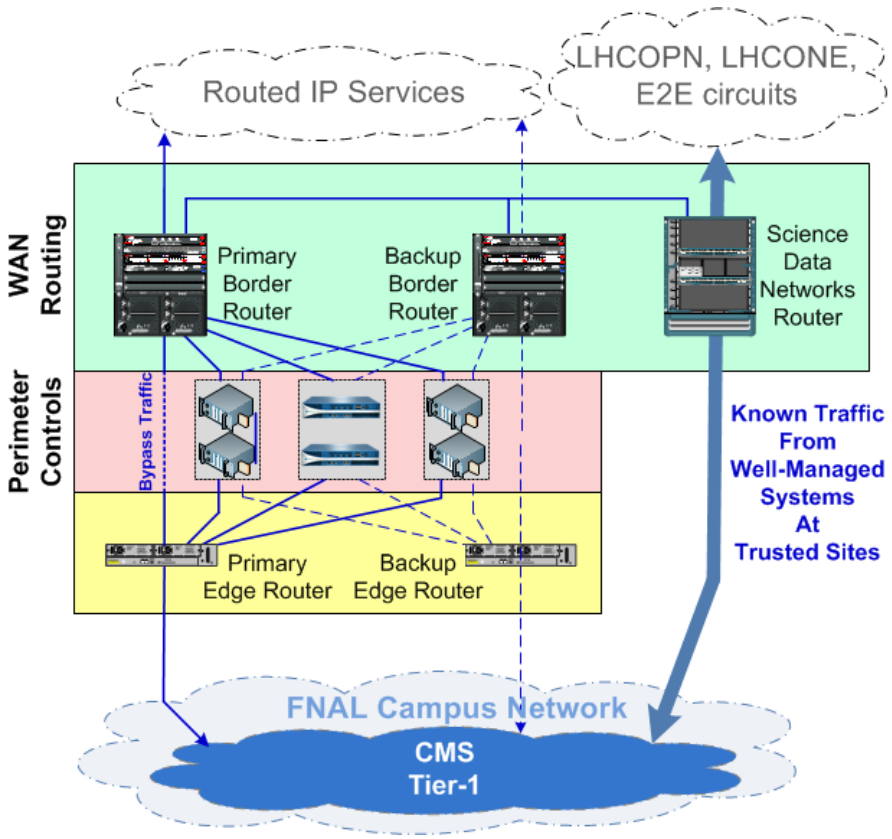

Bypass for Science Data Networks

- Internal policy-based routing (PBR) creates bypass path
- Incoming traffic not in PBR tables forwarded to border router:
  - Still gets into the Tier1
  - But passes through security controls
  - Also creates WAN path asymmetry
    - May cause firewall problem on remote end
    - Haven't run into actual cases of this yet
- Flow data monitored for non-bypass traffic from science data paths

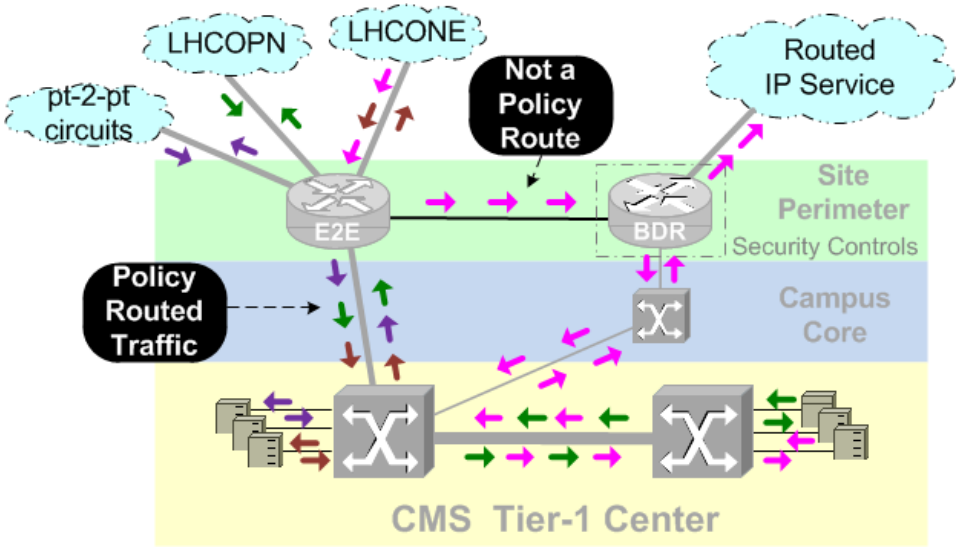

Coming Soon: New Perimeter Architecture

- 100GE costs necessitate consolidating bypass router functions into border & backup border routers
  - Actively working on 2x100GE configuration
- Expect to continue efforts to separate out science data:
  - But using virtual separation technologies, not separate physical infrastructure

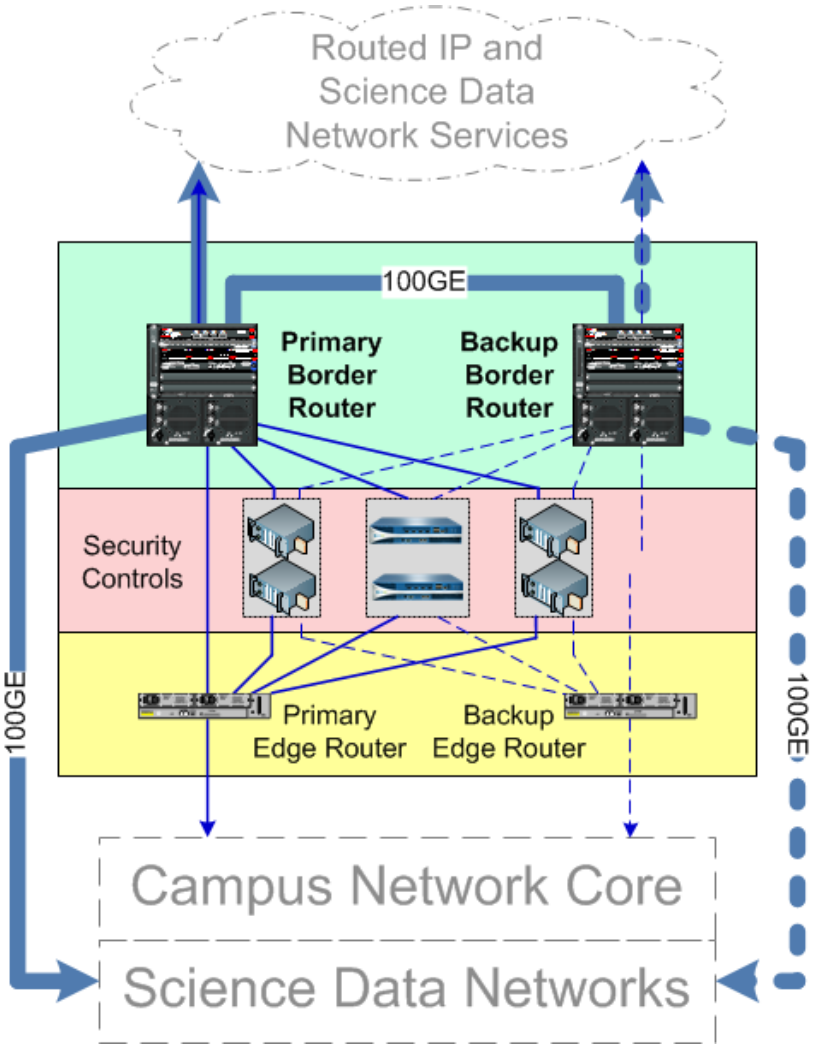

Fermilab

Tier-1 Perspective(s) - I

- T0 → T1 data movement (LHCOPN):
  - Raw data movement should continue to have “preferred handling”
  - But doesn’t need dedicated b/w (LHCOPN isn’t a distributed DAQ)
  - Goal is 48hrs to tape
  - Even for upcoming run, 10Gb/s would be more than sufficient…
- T1 ←→ T1 data movement (LHCOPN):
  - No “preferred handling” needs, just “adequate” b/w
  - Currently works well soaking up available LHCOPN b/w
  - Large flows very intermittent & not latency sensitive
- T1 ←→ T2 data movement (circuits, LHCONE, other)
  - No “preferred handling” needs, just “adequate” b/w
  - Circuits (static) to US T2s work very well
  - LHCONE & general R&E network paths to T2s vary considerably

Tier-1 Perspective(s) - II

- Potential changes to LHCOPN:
  - Keep “preferred handling”; don’t care about implementation
  - Would like to have any changes implemented by 1/1/15
- Building network-awareness into applications:
  - Willing to consider, if necessary…
  - But concerns about:
    - Troubleshooting would become extremely difficult
    - Ongoing maintenance another concern
  - For now, having capacious b/w available is working fine
- On/over the horizon WAN concerns
  - Impact of potential consolidation of tape archiving
  - How commercial cloud services would be supported
  - Firewall performance (100GE) issues at other sites

Summary

- Plan to keep current model of separating science data movement from general network traffic
  - Virtualized separation will be necessary (at least internally…)
- Would prefer to see LHC data carried on “LHC” networks
  - But not essential; LHC traffic on R&E routed paths will also be supported (ie., get bypass handling service)
- LHCOPN function should be preserved:
  - Implementation should evolve with technology
  - T2s on the LHCOPN?
    - This is not what LHCOPN was intended for
    - This is what LHCONE & general R&E networks are for
    - Don’t want to act as an ISP for T2s using LHCOPN
  - T1 traffic on LHCOPN has worked fine, but could be moved

Questions