School of Engineering Network Design

Bradley R. Smith
brad@soe.ucsc.edu, x9-2370

January 16, 2004

1 Overview

This document presentes the design for the expansion of the School of Engineering's (SoE) networks to the new Engineering 2 (E2) building, and the SoE space in the new Physical Sciences Building (PSB). The general goals of this design are to provide capacity for three years of expansion in network bandwidth and port counts, to provide robustness in the context of a "cable-closet" failure mode, and to provide separation of campus NTS and SoE network administration zones (requested by NTS). The concrete impact of the design lifetime is in the following areas: router and switch throughput capacity, port count, and rack space; patch panel port count; and router-to-router and router-to-switch cable plant. The project growth rates used to derive capacity goals are a 25% per year increase in port density, and a 25% per year increas in per-port traffic load (for a total of 56.25% per year traffic load increase per router and switch). The baseline numbers for these lifetime projections are derived from current port count densities (stated in terms of ports per square foot), modified by current actual plans where available.

The cable-closet failure mode is based on expected failures only in the electronics (routers and switches), and connectivity in the SoE cable-closets. This design does not provide robustness in the context of faults external to these cable closets (e.g. back-hoe conduit breaks, cable-vault flooding, etc.). This design provides robustness in the event of a single-inter-building fiber failure (e.g. a bumped connector), and limits the affect of a router failure to the workstations downstream of that router. Specifically, each machine room is provided redundant connectivity to two routers, so failure of any one router will not compromise connectivity of any SoE machine room.

The separation of NTS and SoE network administration zones is accomplished by the deployment of two routers per building, and separate inter-building cable-plant. The NTS-managed inter-building cable-plant is composed of the cable-plant between E2 and the Communications building, BE and Communications, and the BE to PSB cable-plant included in the base bid for the PSB construction. Additionally, in the event of the sharing of space by SoE and one or more non-SoE units at the sub-cable-closet level (i.e. where the same cable closet serves both SoE and non-SoE users), separate switches will be installed in the cable closet for separation of SoE and NTS supported users.

The design is based on a VLAN model where all SoE subnets are distributed across all three buildings. The benefits of this design, compared with a fully routed design, include ease of administration (computer moves do not require operating systen or DNS reconfiguration), direct support of private (RFC1918) subnets, direct support of dedicated physical connectivity (e.g. if CBSE choose to own their own bandwidth), and more efficient use of the limited SoE IP address space. The cost of this VLAN-based design is in the requirement for a more sophisticated intrusion detection system (IDS), and the loss of DHCP's ability to control the default router assignment to the router local to a given workstation (although solutions to this last problem are being investigated).

The remainder of this document presents the new topology (both inter and intra-building), a summary of the equipment requirements, both inter and intra-building cable-plant requirements, a detailed description of the selected equipment, and a detailed design of the cable closets.

2 Topology

Figures 1 through 3 present the inter and intra-building topologies. All connections between SoE gear will carry VLAN trunks composed of all VLANs supported in the School. Links from SoE routers to NTS-managed routers will carry a point-to-point subnet for each pair (i.e. they will not carry VLAN trunks, which are not currently supported by NTS). In addition, all connections to SoE gear will be composed of one or more cables (fiber or copper) aggregated by the equipment at each end into a single logical interface. Use of link aggregation allows the scaling of bandwidth to support future needs of the School without the need to jump to 10Gbps technology which is currently prohibitively expensive. An important limitation of current link aggregation technologies is they allocate flows to cables within the aggregated link based on source/destination IP pairs; this results in a limit between any pair of hosts equal to the bandwidth of the cables composing an aggregated link. This limitation should be kept in mind as communications requirements grow.

The VLANs defined in the SoE networks fall into two classes called *public* and *private*. Public VLANs carry subnets that are advertised via routing protocols to the campus, and therefore reachable directly from the Internet. Private VLANs are non-routable (i.e. RFC1918) subnets that are not directly reachable from the Internet (i.e. Internet access to nodes on these networks requires traversing an appropriately configured firewall or network address translation (NAT) box). Table 1 lists the current SoE routed subnets that will be assigned VLANs.

Tables 2 through 5 summarize the port and cable counts at various points in the infrastructure. In these tables the term *bundle* is used to identify a number of cables aggregated into a single virtual link using a link aggregation protocol as described above. Table 2 gives the count of intra-building cables attaching a building's core switch to machine room and cable-closet switches. Table 3 gives the count of inter-building cables and the cables to the SoE router. These counts include the riser counts from the previous table. Table 4 gives the count of cables appearing at each patch panel. These include the cables incident on the co-located core switch (Table 3), and any "cables" transiting the router closet.

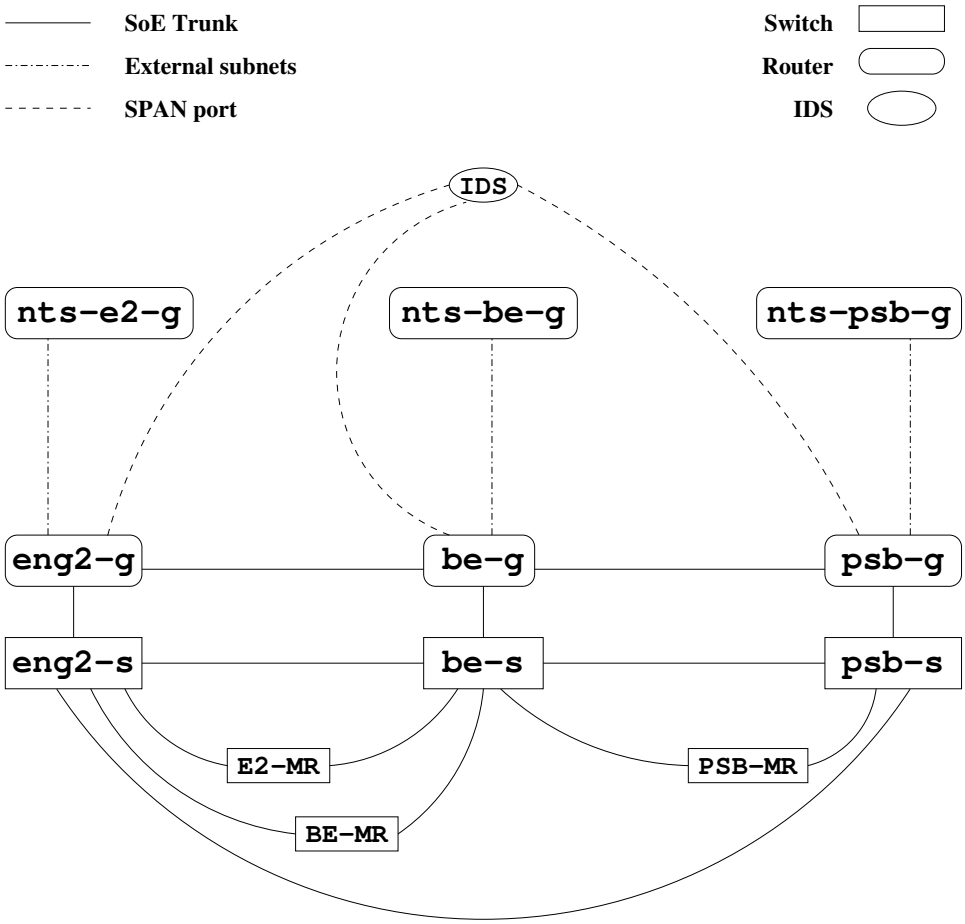

Figure 1: Inter-Building Layer 2 Topology

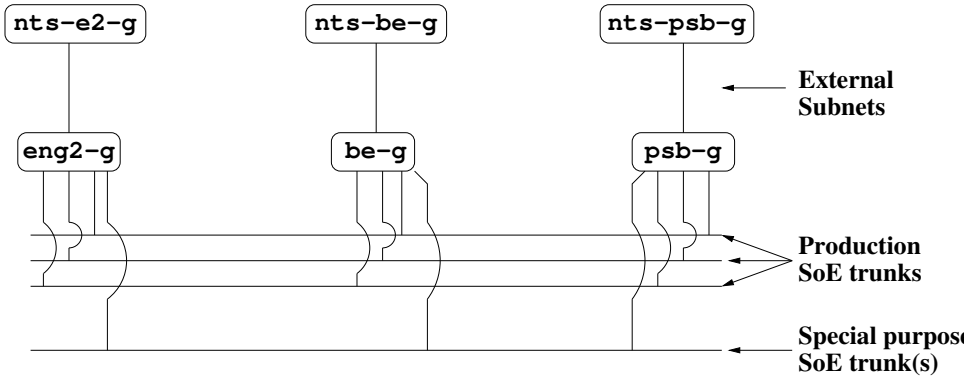

Figure 2: Inter-Building Layer 3 Topology

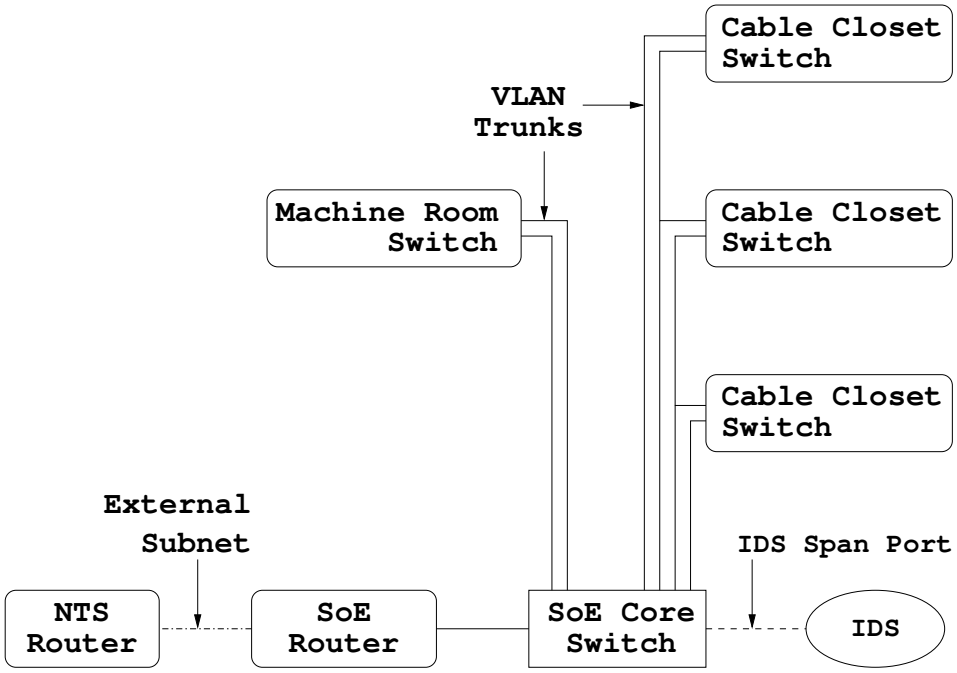

Figure 3: Intra-Building Layer 2 Topology

Subnet	Description	Address Range
128.114.48.0/21	"Secure net"	128.114.48.0 - 55.255
128.114.56.0/22	Unallocated	128.114.56.0 - 59.255
128.114.60.0/24	Wireless	128.114.60.0 - 60.255
128.114.61.0/24	Research networks	128.114.61.0 - 61.255
128.114.62.0/24	Instructional labs	128.114.62.0 - 62.255
128.114.63.0/24	"Insecure net"	128.114.63.0 - 63.255
128.114.5.0/24	Sinsheimer	128.114.5.0 - 5.255
10.1.0.0/16	Private CBSE KiloKluster	

Table 1: School of Engineering Subnets

Building	Bundle Count	Cable Count			Description
		Yr 1	Yr 2	Yr 3	
E2	11	TBD	TBD	TBD	10 cable closets and one machine room
BE	10	TBD	TBD	TBD	9 cable closets and one machine room
PSB	TBD	TBD	TBD	TBD	

Table 2: Riser Cable Counts

Building	Bundle Count	Cable Count			Description
		Yr 1	Yr 2	Yr 3	
E2	16	TBD	TBD	TBD	Router, riser, E2-BE, E2-PSB, E2-BEMR, E2-IDS
BE	10	TBD	TBD	TBD	Router, riser, BE-E2, BE-PSB, BE-E2MR, BE-PSBMR, BE-IDS
PSB	TBD	TBD	TBD	TBD	Router, riser, PSB-E2, PSB-BE, PSB-IDS

Table 3: Core Switch Cable Counts

Building	Bundle Count	Cable Count			Description
		Yr 1	Yr 2	Yr 3	
E2	16	TBD	TBD	TBD	Core switch, E2MR-BE
BE	20	TBD	TBD	TBD	Core switch, E2-PSB, PSB-E2, BEMR-E2, E2-IDS, PSB-IDS
PSB	TBD	TBD	TBD	TBD	Core switch, PSBMR-BE

Table 4: Patch-Panel Cable Counts

Building	Bundle Count	Cable Count			Description
		Yr 1	Yr 2	Yr 3	
E2-BE	5	TBD	TBD	TBD	E2-BE, E2MR-BE, E2-BEMR, E2-PSB, E2-IDS
BE-PSB	4	TBD	TBD	TBD	BE-PSB, BE-PSBMR, PSB-E2, PSB-IDS

Table 5: Inter-Building Cable Counts

2.1 Intrusion Detection System (IDS)

The intrusion detection system is connected to the network by SPAN ports of the "external subnet" ports connecting SoE and NTS routers. SPAN ports are in effect clones of the target port, duplicating all traffic sent out the target port onto the SPAN port. The IDS system uses this functionality to get an aggregated view of traffic entering and leaving the SoE network for use in assuring the security of SoE networks. An artifact of the VLAN design is the potential for "asymmetric" traffic flows into and out of the SoE networks. Concretely, it is inevitable that some external traffic flows (traffic flows with one endpoint connected to an SoE network and the other not) will traverse the SoE network boundary by one router outbound, and a different one inbound. The result being there will be flows where no single router will see all traffic. Since this capability is critical to the function of the IDS, this precludes the use of a separate IDS for each router, instead requiring the use of a single IDS which monitors traffic fed from all three routers.

2.2 Configuration of External Connectivity

The two challenges of this topology relate to how external connectivity is configured. In specific, how is inbound or outbound traffic routed such that the load is balanced over the three pairs (SoE and NTS) of routers, and traffic does not traverse a given node (router or switch) twice on its way into or out of the SoE domain? The solution chosen for this design involves different technology for inbound and outbound traffic. Inbound traffic uses traditional L3 routing solutions for balancing traffic by assigning primary and secondary routers for each SoE subnet. Outbound traffic uses a "first hop redundancy protocol" with functionality equivalent to Cisco's Gateway Load Balancing Protocol (GLBP). In specific, this functionality assigns a single next hop IP address to each subnet that is shared by all three SoE routers, assigns multiple MAC-layer (i.e. Ethernet) addresses for each individual router, and coordinates ARP responses to balance offered load over the three routers.

3 Equipment Requirements

Routers:

- support a "first hop redundancy protocol" with functionality equivalent to Cisco's Gateway Load Balancing Protocol (GLBP)
- support a link aggregation protocol also supported by the switches
- have the software features and bandwidth capability to support default deny firewall functionality (including stateful functionality)
- provide SPAN port functionality
- dual power supplies
- **port count and bandwidth TBD**

Switches:

- support a link aggregation protocol also supported by the routers
- dual power supplies
- **port count and bandwidth TBD**
- **Optional:** 802.1x dynamic VLANs

IDS:

- assymetric sessions (i.e. the inbound and outbound streams of a flow may be traverse different routers)
- **port count and bandwidth TBD**
- **terminal server for remote management?**

Network Servers:

- DHCP and DNS servers
- **terminal server for remote management?**

3.1 External Connectivity

4 Cable Plant

5 Equipment Selection