Network Engineering

[Building] Network Plan

Charlotte, NC

9/30/1998

Introduction

The [Building] is a new 18 story building upfit. The building will initially provide office space for approximately 1200 Associates, however, the building has the capacity to support up to 2700 Associates. [Building] will push the Campus LAN/WAN environment into a previously undeveloped area of [City].

According to preliminary information received in July/August 1998, basic Enterprise WAN access will be the predominant networking requirement. Some additional stand-alone networking will also be required, however, those requirements will be addressed by the individual client groups.

Network Engineering plans to include this facility on the standard Campus Backbones as well as incorporate the [Building] in all Campus infrastructure planning.

Base Building

Physical Plant

The [Building] has three basic floor footprints across eighteen floors. The Lobby through the 9th floor is a large 'L' shaped footprint of approximately 34,000 square feet each. The Lobby will be retail and no offices are planned at this time on that floor. The 10th through 15th floors (no 13th floor) are somewhat smaller floors in an octagonal shape directly above the vertex of the first nine floors. The remaining floors are smaller but continue this design.

Cableplant

The Main Distribution Frame (MDF, server room) will be located on the 2nd floor. A single IDF will be located on each floor, vertically stacked near the central core of the building. The IDF and MDF locations are not contiguous and therefore require riser conduit from the 2nd floor IDF to the MDF. Two stacked riser paths, one on either side of the building core, will be used for diverse riser cable routes.

The standard [Corporate] Cable plan will be implemented at the [Building] providing one jack plate with four cable drops from the local IDF to each workstation area. The physical cable selected for all drops in the building was an enhanced Level 5 cable. [Cable vendor] was contracted by Real Estate Services to bid and implement the cableplant.

In addition to the workstation Cabling, 400 pair of copper riser cable will be run from each floor to the PBX Frame. Either 50 or 100 pair of copper riser cable will be run from each IDF to the Data frame in the MDF. Two sets of fiber riser will be run from each IDF to the MDF riser patch facility. The risers will use diverse paths from the IDF to the MDF. Each of these will be 12 strands mm fiber and 6 strands sm fiber for a total of 24 strands mm and 12 strands sm fiber per floor.

Campus Infrastructure Access

Due to the distance from this building to the rest of the Campus, special emphasis and consideration has been placed on the Campus plan for the [Building]. In conjunction with the overall Campus Fiber plan, the [Building] will employ 48 strands of multimode and 48 strands of singlemode fiber in one innerduct. A second innerduct will be provided and left empty for future expansion. This physical infrastructure provides the most flexibility and ultimately the least expensive connectivity to virtually all current and future networking requirements. Fiber path diversity was not deemed technically required nor financially desirable for [Building].

Utilizing this fiber, the following Enterprise WAN backbones will be extended to the [Building]:

Campus Fast Ethernet backbone

Campus SNA Token Ring backbones

Development Network Campus backbone

With the extra fiber available, additional backbone and special network requirements can also be connected as requirements are learned.

The Campus fiber access to [Building] has been planned to permit additional buildings sufficient access to migrate to the new Campus Fast Ethernet Backbone.

Future Campus Building projects will be leveraged to provide diverse fiber paths.

Network Design Summary

Building Overview

The Building network design is based on Cisco equipment and VLANs which hub into two building backbone switches located in the MDF. Two routers (running in redundancy) provide interconnectivity between VLANs and to the Campus. Campus connectivity is discussed in the next section. Two Building backbone switches (Catalyst 5500 each equipped with 100bFX (fiber) and 10/100bTX (copper) interfaces) provide connectivity from the IDFs to the routers and other MDF equipment. An additional switch was specified to accommodate Development Network requirements via separate facilities. Two Cisco 7513 routers were specified to accommodate routing between VLANs and to the Campus. A separate Cisco 7206 router was specified for Development Network.

Each floor of the [Building] has physical space capacity to support up to approximately 360 work areas (in the most dense configuration). Initial plans are for around 100-140 users per floor. Allowing for printers, Conf. Rooms, and other equipment, a base of 240 ports per floor was chosen to meet all known requirements and most future requirements that may come about within the next twelve months. To meet this port density and allow for future requirements, two Catalyst 5500 switches have been specified for each IDF. All switches will utilize redundant Supervisory III modules. IDF switches will have fiber ISL modules installed on the Supervisory modules. In addition, a single 12 port 100bFX (fiber) module will be installed in slot 3. User ports will be on 24 port 10bTX modules (WS-X5013). 10/100bTX modules can be added by clients with higher bandwidth requirements at their own expense. A single VTP domain will be established for the building with the VTP root servers being the MDF switches. All VLAN trunking will be controlled from these switches.

From each IDF, Fiber trunks will be established from each switch to the two MDF switches via diverse risers. One riser will be connected to port 1/1 on the Supervisor module. The diverse riser will be connected to port 3/1 on the fiber card in slot 3. Separate trunks will be used for the Development Network as requirements dictate. Development Network trunks will be connected to ports 1/2 and 3/2 on diverse risers.

A single VLAN will be set up for each floor. Except under special circumstances, VLANs should not need to nor will they be designed to span floors. The IP Subnet mask for each floor will be 255.255.255.0. Development Network accessible VLANs may span floors as size and requirements dictate. Development Network subnet sizes will be determined as needed. These specifications will allow for approximately 240 connections per floor without having to expand the IP (DHCP) range or add additional ports.

The Visio document, BUILDING.VSD, details and helps visualize this design.

Campus Access

The Campus employs a number of backbones to offer separation and redundancy between IP Networks and technologies. Two SNA rings (ring 0 and ring 1) are used to provide Switched Major Node (PU) access from any building on the Campus. An FDDI Backbone exists between [list of eight buildings]. This ring is used exclusively on the legacy IP network. A Token Ring backbone remains in use between the original Campus Core buildings as a backup to the FDDI backbone and for Network Management devices. A new Campus dual Fast Ethernet Backbone was created in conjunction with the [Building]. This backbone is exclusively used for the 10.32 IP Network. Unlike the shared FDDI backbone, the new Fast Ethernet Backbone is switched in a dual hub/spoke design with layer 2 and layer 3 redundancy. Hub switches are located in [diverse buildings]. All routers on the Campus that require 10.32 network connectivity connect to these two switches via dark fiber between buildings.

The [Building] will be connected to the dual Fast Ethernet backbone, Development Network backbone and the SNA backbones just as the other Campus buildings. Two fiber connections from each router to the Campus Backbone switches will provide redundant access to the Campus and the rest of the Corporate WAN resources.

Future Requirements

The Network infrastructure and fiber access planned for the [Building] will provide the basis for future network requirements. Backbone requirements up to100Mb can easily run on multimode fiber and are significantly cheaper to do so (compared to singlemode fiber). Future backbone requirements, however, will require much higher bandwidth. Network Engineering is evaluating Gigabit Ethernet and SONET as migration directions for Campus backbone transport. These technologies are designed to run best over singlemode fiber. Singlemode fiber offers extended bandwidth and distance over multimode due to the nature of the light frequency and method of light propagation. Even with singlemode fiber, however, Gigabit Ethernet standards limit the point-to-point distance between nodes to 5 Kilometers. The [Building] will not employ these new technologies, however, the building and fiber plan have been designed with these technologies in mind.

Network Design Detail

Campus Access

The Campus network currently utilizes two distinct backbone architectures. These architectures are distinctly separated based on IP Networks. In both cases, the fundamental premise is to maintain routers ONLY on the backbone (strictly a transport for Campus traffic). Servers and other services should be connected through building backbones or other special segments (subnets). The legacy IP network routers use a shared FDDI backbone. The new 10.0.0.0 network routers are on a dual switched Fast Ethernet backbone. The [Building] will be connected to the ethernet backbone. The following diagram is a high level view of this connectivity.

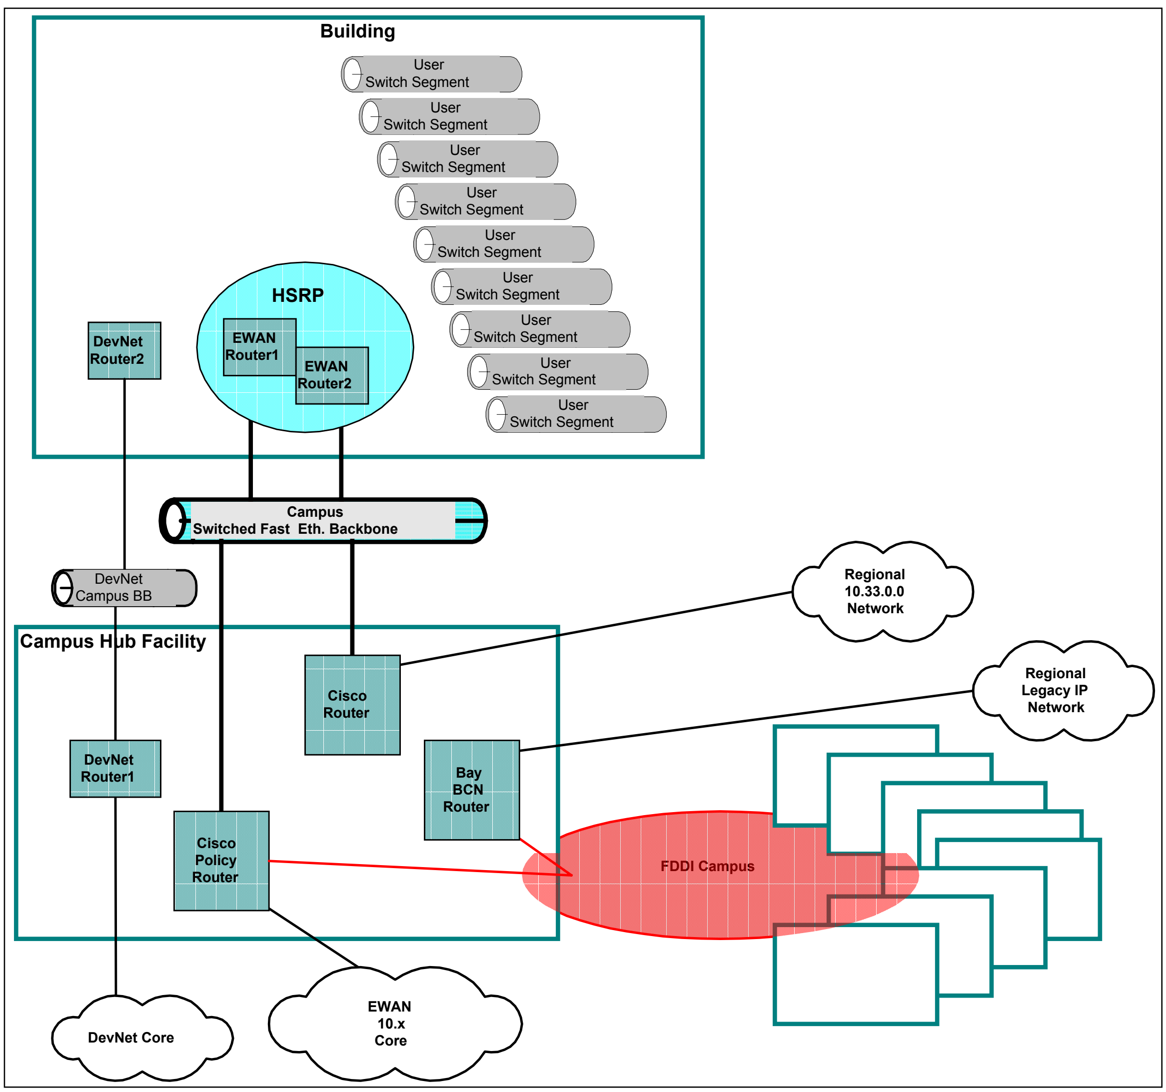

IP Address Plan & VLAN Assignments

The IP address plan for this region is in the midst of change. Originally, the change was from the legacy IP network to the 10.32.0.0 network. Future changes in Layer 3 architecture will dictate another change to the [new] registered IP network. Due to support constraints and timing, the [Building] will follow the 10.32.0.0 layer 3 network plan and then undergo re-addressing in 2000.

The Campus backbone uses a dual switched Fast Ethernet architecture. This design separates and protects Campus connectivity from failures at layers 1, 2 and 3. Both [Building] routers will be connected to both Campus switches (VLANs 500 and 501). The remainder of the building will use IP addresses in the block from 10.32.144.0 through 10.32.159.0.

All network devices will be managed via vlan 1, subnet 10.32.144.0 / 24.

The Building backbone, used for general building resources, has been assigned vlan 20 on subnet 10.32.145.0 / 25.

Each floor will use a single subnet masked at 255.255.255.0. The production EWAN vlan number will be the floor number (eg. vlan 2 – 2nd fl., vlan 3 – 3rd fl., etc.). The IP subnet assignments for the building follows:

FL.	VLAN	SUBNET	MASK	Comments / Description
2	1	10.32.144.0	255.255.255.0	Management vlan
2	20	10.32.145.0	255.255.255.192	Building Backbone – Gateways, global access servers.
2	13	10.32.145.128	255.255.255.192	Integration Testing.
2	n/a	10.32.145.240 thru ..145.254	255.255.255.255	Router Loopback.
2	2	10.32.146.0	255.255.255.0	User connectivity.
3	3	10.32.147.0	255.255.255.0	User connectivity.
4	4	10.32.148.0	255.255.255.0	User connectivity.
5	5	10.32.149.0	255.255.255.0	User connectivity.
6	6	10.32.150.0	255.255.255.0	User connectivity.
7	7	10.32.151.0	255.255.255.0	User connectivity.
8	8	10.32.152.0	255.255.255.0	User connectivity.
9	9	10.32.153.0	255.255.255.0	User connectivity.
7	999	10.32.154.0	255.255.255.128	DevNet – Staff Support.
n/a	free	10.32.154.128	255.255.255.128	Unassigned.

All Subnets (except vlan 1) will support IPX. The IPX network address will be the hexadecimal equivalent of the IP subnet address.

Additional proprietary IP configuration details have been removed.

Router Configuration

Two production EWAN routers will be installed in the [Building]. For redundancy, all user subnets will be assigned to Hot Standby Router groups based on vlans on common router interface ports. All router interfaces will be 100bFX (mm fiber). As a compromise between price and performance, each Fast Ethernet router port will service at least two vlans, but not more than four. Port level VLAN density should be managed to approximately 500 users (10bTX connectivity) per 100Mb router interface. If a large number of users require 100Mb to the desktop, this number will have to be lowered to maintain performance expectations. The basic performance ratio rule of thumb for standard campus resource usage and office automation (on local servers) is approximately 50:1 (users:router interface). This means a 50x oversubscription rate on the router interface. This has been determined to be acceptable at this time and will be monitored for future performance impact.

EWAN Router port assignments have been designated as follows:

Slot/Port.vlan	ISL / VLAN	SUBNET / MASK	Standby	Comments / Description
0/	n/a			unassigned.
1/	n/a			unassigned.
2/	n/a			unassigned.
3/	n/a			unassigned.
4/0/0.2	2	10.32.146.0 / 24	2	
4/0/0.4	4	10.32.148.0 / 24	2	
4/1/0	500	10.32.2.0 / 25	n/a	Campus Backbone
5/0/0.3	3	10.32.147.0 / 24	3	
5/0/0.5	5	10.32.149.0 / 24	3	
5/1/0	501	10.32.2.128 / 25	n/a	Campus Backbone
6/				RSP/4
7/				RSP/4
8/0/0.6	6	10.32.150.0 / 24	6	
8/0/0.8	8	10.32.152.0 / 24	6	
8/1/0.1	1	10.32.144.0 / 24	1	Mgt Vlan
8/1/0.20	20	10.32.145.0 / 26	1	Building Backbone
9/0/0.7	7	10.32.151.0 / 24	7	
9/0/0.9	9	10.32.153.0 / 24	7	
9/1/0.999	999	10.32.154.128 / 25	9	DevNet Accessible vlan.
9/1/0.13	13	10.32.145.128 / 26	9	Integration Testing
10/				unassigned.
11/				unassigned.
12/				unassigned.

In general, all odd vlans will use router 1 as their primary HSRP router and all even vlans will use router 2 as their primary HSRP router.

The following IPX SAP filters are required on each interface:

```
ipx output-sap-filter 1000
…
access-list 1000 deny ffffffff 233
access-list 1000 deny ffffffff 30C
access-list 1000 permit ffffffff
```

In addition to the production EWAN routers, one DevNet router will be installed in the [Building]. The DevNet router will be used for tier 2 DevNet connectivity requirements in the [Building]. DevNet Accessible subnets use the same IP address scheme as standard production. The default gateway assignment, however, points to the DevNet router (as primary). DevNet Tier 2 routers connect to tier 1 (core) routers and learn routes to all 172.16.0.0 (DevNet) resources. They also learn routes to production networks, however, production networks are not advertised into DevNet. DevNet vlans start at 999 and work backwards with each new requirement. DevNet segments are usually small in requirements and therefore are typically subnetted with a 25 bit mask.

Additional proprietary Router configuration details have been removed.

Switch Configuration Design

The [Building] Ethernet switch design calls for two switches in each IDF (user floors) to provide approximately 240 user connections per floor. Two MDF switches have been specified to receive trunks from all the IDF switches and provide central connectivity for MDF equipment (routers, servers, etc.). The MDF switches provide physically diverse connectivity in the MDF for equipment that is capable of supporting such diversity. All local router 1 connections go to switch 1, all local router 2 connections go to switch 2. Between this and HSRP, connection diversity is supported for switch, cable, and router failures.

VLAN 1 connectivity from the MDF to each IDF is accommodated via a single 10bTX connection from each IDF switch to the MDF switches. This reduces the vlan 1 impact on trunks carrying production data traffic. For VLAN 1, management IP address, IDF switches will use the floor number for the tens position and the switch number for the units (eg. 10.32.144.31 for switch 1on the 3rd floor, 10.32.144.52 for switch 2 on the 5th floor). Refer to the ip address spreadsheet for full address assignments.

The VTP Domain for the building will be [xxxxxxx]. All switches in the building will participate in this domain. The two MDF switches will be the root servers for all production VLANs. A separate DevNet MDF switch will be the root server for DevNet vlans. The primary production VLAN for each floor will match the floor number (eg. vlan 2 – 2nd fl., vlan 3 – 3rd fl., etc.). DevNet VLANs start at 999 and work backwards.

Each IDF switch will be connected to both MDF switches thus providing a direct connection to resources on both switches. Since the building design calls for only one vlan per floor and vlan 1 will run on a separate trunk, IDF trunks will be set for the specific vlan on that floor (i.e. not set for trunking). This will reduce the overhead on the trunks and limit potential spanning tree problems on vlan 1. DevNet vlans will connect to the DevNet MDF switch on separate trunks.

By using two switches on each floor, the number of users per trunk is limited to 120. It is not recommended to exceed 200 users (at 10Mb each) on any single 100Mb trunk. This equates to a 20:1 oversubscription ratio. A 10:1 ratio is desired for cost/performance. In the event that 100bTX user connectivity requirements arise, trunking ratios must be determined based on projected user traffic load and destination resource location (local server Vs. Campus access).

All user switch ports (10bTX) will be set to Half Duplex and Port-Fast enabled. All trunks will be set to Full Duplex (no Port Fast). MDF devices (routers, servers, etc.) will be set for Full Duplex and Port Fast enable.

All trunks and MDF switch port descriptions will reflect the name and port of the connected device. User ports will not be named.

The following is a sample diagram of connectivity from an IDF to the MDF:

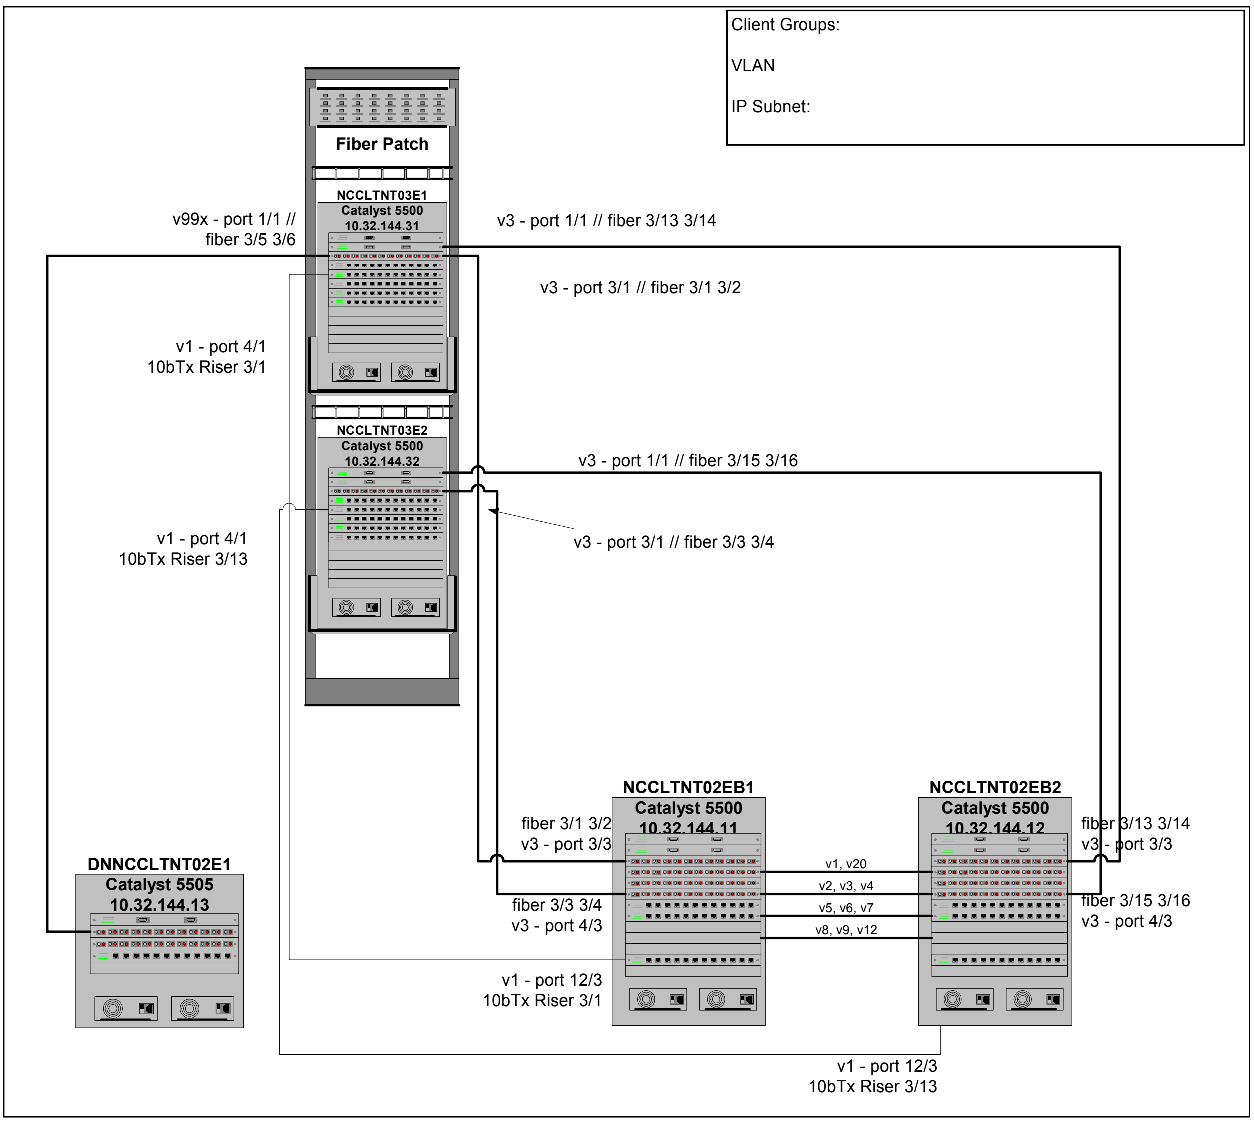

Additional proprietary configuration details and diagrams have been removed.