This page last changed on Feb 21, 2007 by bjjordan.

Campus Network Infrastructure Directions

Strawman Architecture for 2010

Prepared by C&C's Network Infrastructure Directions Team

SUMMARY

If you don't know where you're going, any road will take you there.

OUTLINE

Part I (The Problem Space)

1.0 Background

1.1 Motivation

This "Goal-State Strawman" effort is about trying to define a (or a small set of) candidate UW campus network architecture(s), looking several years out. This is a high-level, broad-brush look at "most likely scenarios" in order to help guide our more detailed planning and design efforts. It involves speculating on actual user and institutional requirements, as well as Best Practice technology and design choices.

This document is one of the primary deliverables of the Network Infrastructure Directions (NID)team, which is charged with coming up with recommendations on network and network security service

requirements and architecture for consideration by our customers and C&C technical and policy leadership, e.g. the Service Directions Team (SDT) and Strategic Architecture Team (SAT).

In our strategic planning process, we have established "reliability and responsiveness" of our systems as a key goal. Additionally, we have a changing world of security threats and compliance requirements that dramatically impact our designs. Accordingly, we need to re-think our assumptions as we plan ahead.

The key drivers for the next-generation architecture are:

- Security requirements, e.g. traffic isolation
- Fault containment, e.g. infrastructure isolation
- Fault tolerance and resilience to errors
- Rapid fault diagnosis, via consistent and easily understood topology

We seek a "self-defending" network design, but when self-defense fails, one which minimizes the consequences of accidental or intentional faults, in terms of masking them via redundant/resilient design, reducing their duration via rapid diagnosis, and reducing their scope via fault containment and infrastructure isolation strategies.

1.2 Document Purpose

Document goals include:

- addressing the need to have combined engineering and strategy/requirements discussions documented in a way that can feed into customer requirements conversations, C&C's SAT and SDT processes, as well as serving as guides for design and implementation teams.
- addressing the need for a communication vehicle of requirements and architecture proposals, decisions, and rationale, so that when questions arise in the future there is a record of the thinking and alternatives considered.
- being a living document that records alternatives considered, their tradeoffs, and recommendations. (Some recommendations should have a "Use by" or "Revisit after" date associated with them.)

In particular, the Network Architecture "Goal-State Strawman" investigation is intended to clarify the technical direction we wish to pursue over the next several years. The idea is to create a living document that:

- Provides a general design target for UW networks, circa 2010.
- Defines key parameters and uncertainties.
- Assists in making tactical design decisions.
- Assists in communicating to and soliciting feedback from, our campus constituency.

1.3 Document Audience

- Network engineers within C&C
- Network Operations, Provisioning, and Customer Service units within C&C
- C&C Technical and policy leadership in C&C (e.g. SAT and SDT)
- Campus customers and IT oversight groups
- Other interested parties within C&C
- Other interested parties on campus
- Interested parties at other institutions

1.4 Design Process Roadmap

- Develop Strawman draft, focusing first on Assumptions and Questions
- Review approach internally
- Collaborate with Campus on refining Assumptions and Questions
- Revise Strawman based on feedback from Campus and others in C&C
- Expand and add detail

- Periodically review and update

1.5 Relationship to Other Documents

This is only one of a family of documents needed to describe network infrastructure directions. A complete document architecture would include white papers on:

- Best practices for datacenter networking design
- *

1.6 Network Mission and Constituencies

The UW campus network should support:

- knowledge workers (including faculty/staff/students): anytime/anywhere access to resources.
- researchers: "Extreme networking for extreme science"
- institutional infrastructure, e.g. building monitoring
- clinical systems
- advanced communication services, e.g. wireless VoIP

1.7 Cornerstone issue

The cornerstone issue upon which so many others rest is convergence, in several dimensions:

- geographic convergence (one network vs. multiple network zones)
- functional convergence (border routers, core routers, route servers vs. VRFs)
- traffic class convergence constrained by QoS and performance needs (e.g. VOIP)
- traffic class convergence constrained by security needs (e.g. quarantine, scada)

The Internet started out with a single traffic class, thus all services running on it were converged. Now it includes multiple strategies for de-converging traffic into separate classes, including multiple QoS queues, multiplexing mechanisms (MPLS, VLANs), and crypto isolation (SSL, SSH, IPSEC). Convergence is driven by saving money by reusing infrastructure that would otherwise need to be replicated. De-convergence is driven by the realization that convergence (sharing infrastructure) doesn't always lead to acceptable performance, reliability, or risk. (A case can be made that the personal lambda phenomenon is a direct result of this.)

1.8 Solution Space

Some of the key differentiators between candidate solutions include:

*the number of separate fault-zones (i.e. highly isolated networks) that are needed to achieve fault containment goals.
*the degree of service aggregation in terms of "box size", i.e. a few big routers and switches, or many smaller ones?
*the degree to which organizational/administrative rather than geographic topologies are embraced or promoted.

- what protocols are used in which parts of the network topology/hierarchy, e.g. Enet/VLANs(L2), MPLS (L2.5), IP (L3).
- topology: use of rings vs. star/hierarchy.

Security, survivability, and supportability are the key design drivers, and those goals translate directly to traffic and fault isolation requirements.

Policy issues that will influence definition of network zones and traffic isolation requirements include:

- security
- privacy
- performance
- manageability
- access rules and acceptable use
- desired connectivity
- risk management

2.0 Definitions

2.1 Zones

A network zone is an administrative domain, defined as a separate area/region/element of the network either because different policies
(e.g. security, acceptable use) may apply, or because of fault isolation needs. A zone may coincide with a geographic region, but that
would be fortuitous circumstance, not the defining characteristic. However, a zone does map to some subset of overall network topology.
Examples include: residence halls, medical centers, datacenter, UWB, UWT, perhaps wireless. Technically, a zone might be
distinguishable from surrounding network zones/regions/elements by as little as a router ACL, or as much as fully isolated physical
infrastructure and a distinct Autonomous System number.

- Network Zone: A group of L2 or L3 hosts or subnets sharing a common feature, or set of features.
- Administrative Zone: A groups of L2 or L3 hosts where the common defining feature is the administrative entity.
- Policy Zone: A group of L2 or L3 hosts where the common defining feature is a policy, or set of policies.
- Security Zone: A group of L2 or L3 hosts where the common defining features are the security requirements.
- Fault Zone: Hosts, zone or zones affected by a failure in the network.

2.2 Traffic Classes or Types

Traffic classes group different flavors of network traffic together
into categories that may be useful for implementing differing
policies, or for fault/interference considerations. These classes of
traffic will often coexist on the same shared network infrastructure,
but may be separated to achieve security or fault containment
objectives. Examples include:

- Voice (relatively low-bandwidth, latency sensitive)
- Video (potentially high-bandwidth, latency sensitive, often UDP)
- Real-time control (low-tolerance for outage/degradation)
- Real-time monitoring (ranges from life-critical to fairly tolerant of outages)
- Apps requiring extremely large (TB PB) file transfers
- Business-critical transactional flows (esp interactive apps)
- All other unclassified/best effort traffic (HTTP, HTTPS, FTP)

2.3 Network Service Classes

Network service classes differentiate possible product offerings in
the network service area. In particular, they distinguish different
(optional or departmentally mandated) connectivity semantics.
Examples:

- Standard globally-routed "open" Internet connection
- P172 campus-private network connection
- Subnet-firewalled connection
- Potential registered-host connection (MAC reg or 802.1x)
- Quarantine connectivity for suspected "threat" hosts
- Connectivity with/without Intrusion Prevention System (IPS) in path

- Connectivity for devices lacking self protection or where we want a high level of protection (Power Strips, Elevator, Fire Alarms, SCADA, etc.)

2.4 Levels of Isolation

For Layer 1 and 2 devices we will classifly the level of Isolation to:

- Type 1: Physically separate network infrastructure (individual fiber, independant copper runs, independant switches)
- Type 2: Common physical network infrastructure with full domain isolation (wave/time DM)
- Type 3: Common physical network infrastructure with logical domain isolation (VLAN's, EoMPLS, statistical DM)

For Layer 3 devices we will classify the level of Isolation to:

- Type A: Physically separate network infrastructure with constrained and strict interconnections (EGP).
- Type B: Physically separate network infrastructure with common routing (1 IGP) and management, with overt "hard" interconnections (physical interfaces with policy inflection devices between isolation areas)
- Type C: Common physical network infrastructure with logical isolation (MPLS, VRF's) with overt "hard" interconnections (physical interfaces with policy inflection devices between isolation areas)
- Type D: Common physical network infrastructure with traffic isolation (Policy-based routing, MTR, merged-VRF's) with "soft" interconnections (network software traffic combination and selection)
- Type E: No infrastructure provided isolation, total convergence, application based segregation (the Internet)

3.0 Architectural and Operating Principles

3.1 General architectural design principles:

- Good fences make good neighbors.
- The key to high-performance, high-MTBF, and world peace, is reducing contention for shared resources.

Design systems to

- embrace principle of least surprise
- have no silent failures
- contain faults
- enable measurability
- enable rapid diagnosability
- be self-defending; resistant to misconfiguration
- allow operators to discover problems before users do

<Ref other docs>

3.2 Operational Principles

This can be interpreted in two ways--maintaining and updating network or monitoring and troubleshooting the network. There was a bit of overlap in the two.

- Network devices must notify and be able to be polled for status changes.
- NOC must be able to join any service class offered
- Ability to sniff data either directly or remotely between any two points on the network.
- Limit number of technologies on campus. Thorough documentation, testing, training and integration into tools must be complete before going in to production.
- Processes in place to ensure network documentation is kept current
- "Living" documentation. Should be designed in a way that tools can be developed to

produce documentation on the network to indicate current topology/status.

- Self documenting. Devices should be labeled and placed in network in a way that make their use/function obvious.
- test plans to include interoperabiliyy with current network management systems

4.0 Assumptions

As noted above, key drivers for architectural change include evolving expectations around policy-based networking (especially "selective connectivity") and having adequate reliability/availability (which in turn drive requirements for improved redundancy and diagnosability.) We accept that "One size does not fit all", and even though one size might fit 80% of customer needs, we need to be responsive to the minority needs as well.

4.1 Security and Traceability:

- There is growing interest in "organizational" rather than the traditional "geographic" network topologies and boundaries, driven primarily by security considerations (specifically, fear of shared fate zones, and what neighboring hosts, if infected, might be able to do to one's own systems.)
- A number of departments desire (or are required by contract to have) more network isolation (selective connectivity) than in previous times.
- In general, customers desire a choice of connectivity options, e.g.
  - traditional "open" Internet connection
  - globally filtered Internet connection
  - connectivity to a semi-closed affinity group (dept'l net)
  - custom subnet firewall connection
  - connectivity to a "more trusted" class (perhaps 802.1x protected)
- At the same time, traditional perimeter defense strategies are becoming ever less effective as more traffic is tunneled over web ports (80/443) and encrypted end-to-end.
- End-to-end encryption at the OS-level (e.g. IPSEC) will become common.
- Attack vectors are changing, with new emphasis on attacks exploiting user behavior (e.g. Phishing).
- We do not really know what the demand will be for selective-connectivity options (as listed above).
- Some departmental administrators would like more visibility into the network traffic of their units, for IDS/IPS or diagnostic or traceability reasons.
- If we need to do network admission control (NAC), requiring daily logon, this would be limited to a specific zone (e.g. med ctr wireless) or to one of the optional service classes. (Unless CALEA etc force it.)
- Legal requirements for traceability and attribution may increase (e.g. CALEA).

4.2 Availability and Fault Containment:

- Dependency on the network for almost every aspect of one's work life is growing; hence, outages have increasing impact on the institution.
- Not all units have the same availability requirements, and it is difficult to know a priori where High Availability expectations exist, but where customers have articulated the requirement we will work with them to to investigate how their needs can be accommodated.
- The medical center will continue to require (and be willing to pay for) HA network design.
- It is a high priority for C&C to provide adequate availability for each use case, and we seek to establish a high standard for overall network reliability and availability.
- Fault-containment goals require different techniques at different service layers. Containing an electrical ground fault is different than containing a multicast storm. Even highly isolated networks

will generally have some form of interconnection, which can become a gateway through which certain kinds of faults propagate.

- Some services should be deployed on highly isolated physical infrastructure, e.g. patient monitoring, out-of-band network control/management, certain R&D needs such as "personal lambda" connections.
- DWDM and TDM are as good as separate fiber/copper for traffic isolation, but designs based on diverse copper/fiber/device infrastructure may still be needed in some cases for fault isolation.
- Zones or service classes requiring modest isolation should not span zones of high-isolation, although areas of similar semantics might exist in multiple isolated networks. For example, subnet zones (broadcast domains) would not be trunked across Layer-3 fault-isolation boundaries, but that does not preclude offering similar connectivity classes in different geographic areas.

4.3 Service Trends (bandwidth and mobility):

- Bandwidth demands will continue to grow, due to increased utilization, more intensive apps (e.g. video conferencing), and technology evolution (e.g. networking stack optimization or improvements).
- We have some extreme bandwidth requirements on campus for advanced research efforts, e.g. 10Gbps links from UW labs to facilities in other parts of the world, and applications with very low tolerance for latency and jitter such as real-time HD conferencing.
- We assume that everyone will have a wireless computer and a mobile (cell, wifi, etc) phone; how soon those might obviate the need for desk phones and desk computers for many of our constituents is not yet clear.

4.4 Implementation:

- Diagnosability, measureability, and supportability in general are really, really important.
- By collaborating with others in our industry (including other institutions and vendors), we will deploy topologies which are inline with Best Common Practices at time of deployment. Today, this includes a layer-2 edge (building), layer-3 aggregation (routing center), layer-3 core, and layer-3 border.
- There will be a number of distinct network "zones", defined either by differences in policy (e.g. residence halls, wireless), or fault-containment (e.g. OOB device management), or non-interference (e.g. R&D or personal lambda), or security (e.g. clinical nets) needs.
- Most connectivity options must be implemented as software-defined features; Multiple types ("colors") of wall-jack doesn't scale, and needs to be limited to very special needs (possibly VoIP, or clinical A and B nets).
- In general, we assume that software-defined network configuration is preferable to having to physically reconfigure hardware and connections, e.g. having to move a user outlet to a different switch to offer a different network connectivity semantic. The exception would be when security or fault isolation objectives require maximum isolation between the different services.
- We anticipate more convergence of voice, video, and data services than previously, convergence plans will be constrained by best practices for fault isolation.
- We will continue to use Best-effort and over-provisioning but with a eye on QoS for applications for specific networks.

4.5 Institutional:

- In order to ensure the integrity of and proper operation of the network, it is important to have a single organizational entity be responsible for managing all active devices in the network from the customer walljack demarc to the ISP border.
- The UW is not likely to move toward an aggressive telco-charging model for the data network (per port service charges), though some form of charging for premium services is possible.

5.0 Requirements

We begin with some issues that must be considered when designing a network.
First operational threats that could interrupt service, and then policy constraints that could invalidate a particular design if/when they change. Note that these requirements apply to a network **service**, and therefore have implications not only for the technical design of the system, but also on operational procedures, training, documentation, etc.

5.1 Network Disruption (Operational) Threats that should be Considered

- Cross-connected edge ports
- End systems configured with an conflicting network address
- CPE router/host becoming DHCP server for subnet
- Router/switch total failure
- Router/switch partial failures
- Device (router/switch/firewall) mis-configuration
- Management/implementation proceedures that result in network disruption
- Broadcast/Multicast storm
- DOS attacks originating from outside
- DOS attacks originating from inside (on UW computers)
- Wireless DOS attack
- Bad software/firmware upgrade
- Link/device saturation of either data or control plane of routers and switches
- Disruption of environmentals (power, air conditioning)
- Disruption of links, e.g. fiber vs. back-hoe encounters
- Fire, flood, earthquake
- Physical security
- Device (OS and configuration) security

5.2 Policy/Traceability Considerations (Possible threats to the Design)

- Rogue wireless access points
- Lawful-intercept requirements, e.g. CALEA
- VOIP e911 requirements (for both office and mobile devices)
- Copyright infringement countermeasures requirements
- TOR (Onion router)techniques for masking identity vs. attribution requirements
- Billing separately for differentiated services
- Other billing requirements
- Fair network-utilization constraints (e.g. high-use penalty box)

5.3 Primary Goals

- A network that meets customer needsh1.
- Flexibility to enable deploying new services quickly and efficiently
- Availability sufficient for customer needs
- Limited consequences of any one failure
- Operationally supportable, including
    - rapid diagnosis
    - rapid fault mitigation
    - self-configuring wherever possible
    - localized maintenance possible
- Scalable
- Secure
- Measureable (See Success Metrics below)
- Cost-effective, including
    - support for incremental upgrade
    - avoidance of vendor lock-in

5.4 Success metrics

- Operational (pkt loss, MTBF, MTTD, MTTR, fault zone size)

- Deployment (percent of ports capable of 10/100/1000, etc)
- Features (percent capable of sFW, voip, multicast, v6, HD/IP)
- Strategic (percent subnets with own FWs, DNS, DHCP, etc)
- Wireless-specific (coverage, address exhaustion, roamability)
- Security (# of security-related outages)
- Flexibility (How often do we do one-offs or need to say "sorry")
- Usage (trend lines, percentage of population using service)
- Time spent fixing vs. time spent implementing

>>> NEED TO ALIGN 5.3 and 5.4 --every goal should have a metric

5.5 Anti-goals (These are things you really don't want to have happen)

- Number of user hosts affected by a single subnet outage > 250 ?? ''DISCUSS''
- Number of subnets affected by a single outage > ??
- Number of servers affected by a single outage > 5 ?? ''DISCUSS''
- Percentage of subnets without redundancy > ?? ''DISCUSS''
- Number of hosts in a shared broadcast domain > 250 ?? ''DISCUSS''

Proposed Matrix

Client Class	Connectivity	Outage	Users Affected (> 60 secs)	Comment
Basic Host	Single Homed	Agg Switch	1000	Single Aggregation Switch
Basic Host	Single Homed	Agg Router	4000	Single Routed Interface
Premium Host	Dual Homed	Agg Switch	1000	Single Aggregation Switch
Premium Host	Dual Homed	Agg Router	0	Dual Routed Interface
HA Host	Dual Homed	Agg Switch	0	Dual Aggregation Switch
HA Host	Dual Homed	Agg Router	0	Dual Routed Interface
Basic Server	Single Homed	Agg Switch	25	Single Aggregation Switch
Basic Server	Single Homed	Agg Router	100	Single Routed Interface
Premium Server	Dual Homed	Agg Switch	25	Single Aggregation Switch
Premium Server	Dual Homed	Agg Router	0	Dual Routed Interface
HA Server	Dual Homed	Agg Switch	0	Dual Aggregation Switch

HA Server	Dual Homed	Agg Router	0	Dual Routed Interface
Wireless	Single Homed	Agg Switch	2000	
Wireless	Single Homed	Agg Router	0	Dual Routed Interfaces
Personal Lambda	Single Homed	Agg Switch	100%	L2 infrastructure and edge router are single points of failure
Personal Lambda	Single Homed	Agg Router	100%	L2 infrastructure and edge router are single points of failure

- Mean-time-to-awareness of problem > 2 min
- Mean-time-to-diagnosis > 15 min
- Mean-time-to-repair > 60 min for Size/Impact severity 1
- Mean-time-to-repair > 180 min for Size/Impact severity 2
- Mean-time-to-repair > one business day for Size/Impact severity 3
- Inability to connect to consoles in an emergency
- Users or operators surprised by policy enforcement into thinking the network has failed (e.g. silent firewall-policy-induced failures)

5.6 Diagnostic Infrastructure Goals

Without an initial reference point and ongoing traffic pattern analysis, it can be difficult to track and diagnosis network anomalies. Therefore, a diagnostic infrastructure strategy should include the task of developing a baseline state of the network from which variations can be measured. This should be coupled with a process for ensuring ongoing comparative analysis.

Gathered information can be used as a basis for pro-active planning, such as, trend analysis, performance tuning, and capacity planning. Additionally, it can be used in a re-active manner for tracing anomalous, and service impacting behavior across the network.

Baseline in this context could be in the form of quantitative data, based on network flow data and log files, or established network use policies.

5.6.1 Re-active Diagnosis (Event Isolation)

This is driven by the need to rapidly identify, localize and quantify service or network issues.

Authorized C&C staff should be able to inspect all switched and routed network traffic at any network locations of their choosing, as we currently do at the campus border.

The intended uses for this are:

- Helping to diagnose client connectivity problems
- Investigating undesirable network traffic (with ethereal, WireShark, tcpdump, snort, etc.)

It would also be desirable to occasionally be able to temporarily connect one of our test systems to a remote network as if the system was local. This allows us to test and fix problems without having the customer on the phone in the loop. It also allows us to interact with hosts on the customer's network, which may not be willing or able to talk to devices outside their broadcast domain (such as rogue DHCP servers or unauthorized wireless gateways).

5.6.2 Pro-active diagnosis (trending, capacity planning)

This is driven by the need to discover potential network congestion points, identify deviations from baseline, budget justification for infrastructure, and development of a traffic-engineering plan.

The data store should be designed such that it provides flexibility for varying information retrieval and data presentation systems.

The intended uses for this are:

- Statistical traffic sampling, monitoring, and trend analysis
- Network capacity planning and performance optimization.

In addition to capacity planning, a clear representation of current and historical traffic characteristics would be particularly useful in the creation and fine-tuning of QoS policy, security event correlation and trending, and as reference to gauge the impact of introducing new technology into the network.

5.7 Perimeter Defense Paradigm

We want to offer choices... but not so many that everyone goes crazy trying to support them. Possibilities include:

- Optional firewalls at subnet and/or inter-cloud boundaries,
- Choice of several general connectivity options at edge (e.g. P172),
- Firewalls at border of separate zones (and/or logical networks.)
- IPS --between zones, and potentially wherever a firewall might go.

5.8 High-Availability Building Network Model

Although we can't afford to deploy HA designs everywhere, we want to have a vision for one or more levels of HA design that we can deploy as circumstances require and/or funding becomes available.For example:

- Two aggregator switches, with separate fiber to separate routers in different router centers.
- Singly-connected (non-replicated) edge switches (except for extreme HA needs).
- Two Enet jacks per outlet, each one going to a separate edge switch and aggregator switch.
- Role of STP to be determined (whether to embrace active/intentional loops for HA).

6.0 Questions

6.1 General Arch/Strategy Questions/Issues

- Convergence (subject to risk, reliability, regulatory and support requirements)
  -geographic (how many zones? how big?, network sizing)
  -multi-function (use of VRF rather than separate boxes)
  -multi-service (data, voice, video)
  -traffic isolation classes building control, console-access, patient monitoring...)
- Perennial recurring questions
  -How big should networks be? (More/fewer Zones/sectorization?)
  -How big should subnets be?
  -Where L2 vs. L3?
- Framework
  -Multiple networks for different types of traffic?
  -Multiple networks for scalability and fault isolation?
  -Network utility via a software-defined-net, seperate physical connection (**Need clarification on meaning of "network utility"
- Network self-defense and fault isolation strategies
  -STP, storm control, ACL - L2 (port, VLAN), L3 (subnet, aggregation, campus edge), MDIX-off, etc
  -Partitioning into zones
  -Partitioning/protecting control plane (e.g. SIP signalling, device CPU protection)

-Partitioning management plane (e.g. OOB console access, seperate OOB network)
-High Availibility
-IPS/IDS
- Advanced services
-Streaming, multicast and unicast design (scopes/fault containment/failure modes)
-QoS or overprovision? (depends on convergence answer)
-Wireless
-Policy Routing
-IPV6
-Security
- Security strategies:
-Traffic isolation (e.g. Quarantine)
-Traffic visibility, capture and baseline
-Firewalls, Traceability, Host-health-assurance
-Admission control (e.g. 802.1x)
- Supportability
-One size fits all (e.g. do switch ports need indiv. configs?)
-Diagnostic framework
-CALEA
-Packet capture (near impossible on high-speed segments)
-OOB management

6.2 Convergence Questions/Issues

Issues that grow out of the convergence question include the following:

- How many semi-autonomous or partially isolated network zones should we have in order to meet fault containment goals? (TEAM#2: Valid)
- How many different traffic classes, with what level of isolation, do we need to implement? (TEAM#2: Valid)
- What technologies should be used to implement the different traffic and infrastructure isolation levels? (TEAM#2: Valid)
- How much do we want to converge functions that previously we've used separate routers for? i.e. Do we sacrifice the "good fences make good neighbors" design principle in order to reduce cost (and potentially complexity)? (TEAM#2 Invalid, combines the previous questions and seems redundant. Vote to remove)

In particular:

- How isolated should the following network be (dynamic list)?
    - Medical Center(s)
    - UWB Campus
    - UWT Campus
    - UWS Campus
        - Dorms
        - DataCenter
        - Wireless
        - VOIP
        - Building/elevator control
        - R&D
        - Emergency backup
        - SAN
        - Backdoor

- How many **kinds** of Enet jacks should a typical office have? (TEAM#2: Can services be combined/ tracked on existing wall Jacks?)
    - One: any categorization or isolation is software defined
    - One by two: two equivalent s/w defined ports, but on different nets for redundancy/resiliency
    - Two: "Normal Data" and "VoIP"
    - Three:
        - Campus: "Normal Data", "VoIP" and "Video"
        - Hospital: Clinical net A, Clinical net B, Open Net
- What types of software-defined network service should be offered? (TEAM#2: Valid Examples)
    - "Open" Internet and/or P172

    - Authenticated (802.1x) + Firewalled access
    - Firewalled access
    - Quarantine
    - Choice of organizational affinity group, plus firewall
- If security requirements preclude sharing subnets, how far do we go toward organization-based networking topology? (TEAM#2: Valid)
    - subnet splits (to provided dedicated subnets) only
    - each org unit in a different location gets similar semantics equivalent to being behind the same firewall
    - each org unit in a different location gets same broadcast domain
- What technologies should be used to implement different degrees of service and affinity-group isolation? (TEAM#2: Valid, some examples but not all solutions.)
    - None within the network infrastructure (e.g. App defined)
    - L3: IPSEC, TOS/DiffServ
    - L2.5: MPLS
    - L2: VLANs
    - L1: DWDM/CWDM, separate fiber/copper
- If we offer subnet firewalls, does this imply org topologies, or at least non-shared subnets? (TEAM#2: remove or merger with security requirements above.)
    - yes
    - no
- What form of auth and traceability solutions need to be deployed? (TEAM#2: Valid, some examples but not all solutions)
    - None
    - log mining, e.g. pubcookie, arp cache data
    - captive portal
    - 802.1x
- If something more than log-based attribution is needed for wired nets, for whom? (TEAM#2: merger with question above)
    - clinical nets?
    - dorm nets?
    - campus nets?
- How large should a subnet be? (TEAM#2: Valid, answered in 5.5)
- How should we support needs for "extreme networking" (e.g. personal lambdas)? (TEAM#2: Valid)
- Active/active or active/passive redundancy (cf. DNS rotaries)? (TEAM#2: Valid)

6.3 Misc. Questions

- What/How many types of connectivity/isolation service?
  (e.g. besides patient monitoring, how many colors of wall jack?)
- NAC? vs. traceability and host-health-assurance?
- Multiple autonomous network zones?
- Self-defending network req'ts
- How much convergence? (cost v. failure v. support)
- Performance? QoS?
- Personal lambdas?
- Network security? Where? Delta-T?
- Gold/silver/bronze nets (gold=C&C-run DHCP, DNS, etc)?
- Usage-based pricing (dorms?) v. p2p clamping
- Organizational toplology?
- DRBR: how much geographic diversity for paths and equipment?
- Bandwidth management (dorms and elsewhere)
- v6?
- SSM?
- CALEA? (and other regulatory reqts)
- Advanced App/R&D net?
- Wireless assumptions (WWAN vs. .11 vs. UWB vs. ??)

6.4 Near-Term

- Spanning Tree (where/when, and active vs. passive) (TEAM#2: Remove, this has been answered)
    - Pervasive use of passive STP on Access Routers
    - Selective use of active STP for subnets with HA firewall config

    - Pervasive use of L2 passive spanning tree
    - Pervasive use of L2 active spanning tree (Intentional L2 loops for redundancy)
- Degree of convergence for VoIP (TEAM#2: Valid)
- Consensus on ideal HA building design (TEAM#2: Valid)
- Consensus on optimum fault-zone size; degree of convergence (TEAM#2: Invalid, this is a long term question)
- Best way to provision personal lambdas (TEAM#2: In-progress of being answered by DSinn)
- If we go from one network with one class of service to multiple network instances and service classes with mediated interconnection between nets/service classes, what are the unintended/ undesired consequences of that level of isolation? (TEAM#2: Already being considered in 6.2, should be removed)
- Best strategies for quarantine, security isolation and fault zone reduction. (TEAM#2: Valid)
- Rogue AP control; traceabilty of clients (TEAM#2: Valid)
- We don't yet **really** understand how to diagnose wireless nets. (TEAM#2: Should be combined with question above)

6.5 Longer-Term

- How soon might **most** user workstations be wireless laptops or handhelds?
- How soon might desktop VoIP phones be superfluous (due to cell phones or wifi-voip phones)?
- Is there any hope of getting back to a zero-conf, near-zero-exception L2 infrastructure within buildings?
- Answers to all the policy requirements questions implied in section 5.2 above

---

Part II (The Solution Space)

7.0 Alternative Solutions

This summary provides a quick overview of several different campus/enterprise network approaches:

7.1 Status Quo

- Multiple separate (Type A) core networks:
    - Medical Center clinical net
    - Medical Center patient monitoring net(s)
    - VoIP
    - Main data net for all campuses
    - UWS "backdoor" and/or "maintenance" nets
    - UWS R&D
    - (soon) UWS NLR connection(s)
- No MPLS, no VRF (yet), no edge VLANs
- Six classes of service, except for subnet firewalls
  TP protected, Medical Center, VoIP, Un-protected, Dorms, Wireless
- No TOS/DiffServ
- Separate Enet jacks for VoIP
- L3 only at border and access/agg/Level-1 routers
- L2 at edge, building agg, campus backbone
- IPS detection at the perimeter with limited isolation of hosts known to be infected
- NOTE: near-term changes
    - core moves from L2 to L3
    - add VRF's (but only in the core routers to support IPS needs)

7.2 Most Probable/Promising Designs

7.2.1 Common Elements

- All of the existing core networks continue to be Type A.

- Potential core network changes:
  - Add new Type A OOB-Management net
  - Med Ctr may want additional Type A patient monitoring nets
  - Decide whether existing "backdoor" & "maintenance" nets should be changed
  - Decide whether an isolated SAN is needed
  - Decide whether dorms, datacenter, wireless, or UWB/UWT, should become physically isolated
- Network zones listed above that are not isolated as separate Type A physical nets would be isolated by separate VRFs (Type C logical nets)
- Multiple service classes would be overlayed on any given Type A core network by using VRF's
- Access and border layers would be L3
- Move from (newly deployed) L3 core to an MPLS (L2.5) core, supplemented by L3 function to link different VRF/connectivity classes.
- Lambda network (e.g. NLR) connectivity multiplexed to campus and then distributed locally.
- Buildings would remain L2, but with dynamic VLANs to permit software-defined isolation/class of service, and with higher-availability design using redundant aggregator switches as a goal.
- Dynamic L2 VLANs to the edge (or edge sw ACLs?).
- Robust isolation/quarantine of "unhealthy" hosts, e.g. via VLANs or ACLs in edge switches.
- Consider migrating from separate VoIP jacks to two (separate and more independent) general-purpose data jacks per outlet.
- Most likely some form of MAC registration, and/or possibly 802.1x/ae as one of the service class options.
- Security protection (IPS, IDS, and/or firewall) between logically separate networks is possible.
- TOS/DiffServ may be used for broad classes of service in the core (e.g. VoIP in a converged access layer).

7.2.2 Highly Converged; few service classes; geographic topology

- Minimal use of VRFs (e.g. just for tipping point)
- No use of TOS/DiffServ

7.2.3 Multiple service classes; moderate convergence; geographic topology

- Small numbers (approx 10) Type C nets
- Pervasive VRRP
- Redundant FWSMs with active spanning tree

7.2.4 Organizational topology; Low service convergence

- Potentially hundreds of "Type C" networks, one for each department or related group of departments.
- Firewalls between logical networks; no active spanning tree needed.

7.2.5 Contrast: Many Separate Networks (Best fault containment, worst cost)

- 10-30 maximally separate (Type A) network zones
- Additional isolation via VRFs for approx 10 classes overall, each contained within an individual physically-separate network

7.2.6 Contrast: Just one (or two) Separate Networks (Best cost, worst fault containment)

- One maximally separate network zone (or two, if a customer demands and funds).
- Isolation primarily provided via VRFs for at least 30 classes overall.

7.3 Roads Not Taken... Extreme/Avante Garde Alternatives

- DWDM or MPLS campus core; MPLS building aggregators and L3 edge routers.
- DWDM or MPLS campus core; L3 building aggregators and L2 edge switches.
- DWDM or MPLS pushed to the edge for R&D/personal lambda needs.
- Multiple campus-wide Enet VLANs, for multiple service classes, linked by border routers.
- Three or four different Type A networks appearing on each standard Enet wall outlet, as compared to Enet jacks whose semantics are fully software defined.
- Departmental L2 sub-networks spanning multiple buildings.

- L2 core with trunked VLAN's for differing services.

8.0 Recommended Solution

9.0 Conclusions and Recap

Part III (Appendices)

Appendix A: Analysis of Key Design Issues

Architectural Issues Template

Issue:

Description:

Alternatives:

Tradeoffs:

Recommendation:

Issue: Minimizing Mean Time To Diagnosis

How do we ensure that network problems can be diagnosed quickly?

<Ref Network Goal/Requirement... 15-30 min target for MTTD>

MTTD depends on at least three elements:

- Training
- Instrumentation
- System complexity

In order give ourselves the best shot at rapid diagnosis, we apply Einstein's words to the network architecture:
"Things should be made as simple as possible, but no simpler." Once that objective is achieved, and knowing
that contemporary networks are not inherently simple, we will need to define instrumentation requirements to
give the staff, even well-trained staff, a fighting chance at meeting the MTTD objective.

Part of the evaluation of each architectural strawman is to identify a set of typical fault symptoms, and
examine what must be true in order to pin point the cause of the problem within the target MTTD.

Issue: Optimizing Fault containment zone size

Description: This concept can be applied to different classes of fault, at different layers of the system. For
example, subnet boundaries represent natural fault containment zones for broadcast storms. Thus the size of a
subnet defines the scope of an outage due to a broadcast storm. At

the highest and most general level, the size
of "the network" or a "network zone" defines the scope of other kinds
of faults, such as multicast storms,
routing misconfigurations or failures, etc.

Fault-containment Strategies, by Network Layer

Tools: diversity, redundancy, isolation

- L0 -phys infrastructure:
- L1 -physical: fiber connections rather than copper; separate devices on separate power
- L2 -link: Ethernet and broadcast protocol faults
- L3 -network: e.g. multicast storm
- app layer:

Issue: Optimizing Subnet size

Description:

In data centers we have chosen to deploy appropriately sized subnets dependent on the number of grouped hosts of
a given customer or customer class, so as to constrain the broadcast domain fault zone. This approach also maps
well to our deployment model of a single L2 switch per subnet.

In contrast on campus we have chosen to deploy /24 subnets as the norm so as to provide a consistent numbering
scheme for customers and operators to expect.

There is increasing pressures from customers and internally to re-assess this approach. Some of the reasons
are:

- Redundancy - Moving to a /23 subnet would allow us to offer router redundancy to existing /24 subnets at not
  additional costs
- Scaling - We have a highly variable host-density on a per-subnet basis and do not garner any of the benefits
  of fewer infrastructure IP addresses that larger subnets would give us
- Consistency - We do not have a consistent mapping of IP addresses for infrastructure devices across differing
  subnet sizes (.100 for /24's, first address for others, etc.)

Alternatives:

- Status Quo

While we may garner benefits from changing our allocation model, the inertia of making such a change would be
too great, so leaving things as they are would be best.

- Engineered subnets

In this model we expand our approach for the data-center to be all subnets. Subnet size would be derived from
the expected host-density, across the range from /20 down to /31's.

- Single subnet per device

In this model we dedicate a subnet per L2 switch based on it's port size. This would reduce the number of
neighbors within a single broadcast domain, but may (depending on connectivity model) result in more

infrastructure (fiber and router ports) being consumed.

Tradeoffs:

Recommendation:

Engineered Subnets

Issue: network zone size

Description:

A network zone is defined as a topological region of the network that exhibits some level of fault and traffic
isolation beyond that afforded by a normal subnet boundary. This includes completely separate routing domains
with administrative routing protocols in place to constrain the interactions (BGP, dual OSPF processes); wholly
independent networks that do not interact with the others (backend network); to segments of the network that
share in the overall routing of others, but with differing client policies (residence halls)

Alternatives:

To date, we have focused on having as close to a single zone for all elements of the enterprise as we can.
A contrasting alternative would be to define multiple enterprise elements as distinct zones separated by
protocol, firewall, and/or physical infrastructure boundaries. For example:

- Medical center(s)
- Data center(s)
- Residence halls
- Wireless
- UW Seattle main campus
- UW Bothell
- UW Tacoma
- Backend
- VoIP
- Out-of-band Management

Tradeoffs:

Introducing boundaries increases some types of complexity and reduces others. For example, more zones means more diverse topology to manage, and to keep monitor. Conversely, such boundaries can --in addition to reducing the impact of a failure-- be simpler to
manage by constraining the probable source
of a problem due to the reduced scale of the network, and support a more rapid diagnosis via divide-and-conquer
methods. Moreover complexity can make the network more costly to operate and manage, as managing many identical
things is easier than managing many things having different characteristics and behavior.

Infrastructure boundaries to achieve fault containment may increase cost, so this has to be analyzed and weighed
against the benefits. On the other hand, if we start with goals (or anti-goals) that say it is unacceptable to
have a total network failure, save for cases of major regional

disaster, then it follows that there must be one or more fault boundaries in the system. The less certain we are about our ability to anticipate and defend against specific failures, the more important it is to implement the best, most comprehensive infrastructure isolation mechanisms we can think of.

Said differently, in order to embrace the single network zone design, we need to come up with astonishingly good answers to how we will achieve the goals relating to maximum fault zone size. For example, how many subnets or wireless access points can go down due to a single element failure? Establishing this number is a network requirement, rather than architecture, activity which will necessarily include both political and economic components. The requirement could also be defined in terms of the product of people affected and time-to-repair, in which case one might tolerate a larger fault zone if you could guarantee that the diagnosis and repair could be accomplished very quickly (but even then, the political cost of an outage to an IT organization is a function of **both** number affected and duration.)

We have a clear requirement from the medical centers to maximize the survivability of their network in the face of any problems on campus; this leads clearly to a partitioning strategy to implement a separate zone for them, as we have.

Recommendation:

We need more than one network zone; the only question is how many, and what criteria should be applied in the future to decide when a zone becomes too big and/or how small of a zone we will support (I.E. how to decide that we don't want too many zones).

Specific strawman suggestion is to adopt the list above as (ten) distinct zones, then do the impact analysis on that scenario with regards to cost and supportability for different degrees of isolation between zones. Definition of the plausible mechanisms for different degrees of isolation between zone is to be determined.

This topic overlaps the question of traffic class isolation (below) within each zone, where we will attempt to define the classes and probable grouping therein. For example, a separate out-of-band management network could potentially span multiple (primary) network zones without compromising the fault containment objective.

Issue: Traffic Isolation Classes

Description:

Type A: Physically separate infrastructure; maximum isolation

In this model, we would have a separate physical infrastructure (routers, links, switches), though "hard" multiplexing (wavelength or TDM) might offer sufficient isolation from adjacent channel interference (though raising the need to focus on common element failures across differing networks). Example would be a dedicated patient monitoring network, or even more critically, a patient

treatment system that relied on network
connectivity.

Type B: Physically separate infrastructure with shared routing, management, and "soft" interconnections

In this model, we would have some independent physical infrastructure, but with some measure of physical
isolation, yet shared routing and management paths. For example our current "conservative convergence" strategy
for VOIP, which mixes a separate "blue" network and separate "blue" Enet jacks with many VOIP phones plugged
into the regular data network.

Type C: Infrastructure provided logical isolation (MPLS, VLAN, VRF's), with strict, "hard" interconnections

In this model, we would have a converged physical infrastructure, but with logical isolation via VRF's coupled
with MPLS and/or VLAN's to create the logical separation.

Type D: Infrastructure provided traffic isolation (Policy-based routing, MTR)

In this model, we would have a converged physical infrastructure, but with administrative protocols such as
policy-based routing and multi-topology routing providing the traffic separation.

Type E: No infrastructure provided isolation, total convergence, application based segregation (the Internet)

In this model, we would have a converged physical infrastructure with no isolation of the underlying traffic.

Today we have a mix of A and B class networks within our scope. The bulk of our network is a class B network
with multiple, differing interconnections for the various isolation classes. For the Medical Center we use a
dual OSPF routing process, for the VoIP network they are separate devices within a single OSPF area off of the
campus network, for the back-end network they have totally independent routing.

Nesting. Within a maximally isolated Type A network, there can be multiple instances of Type C networks, and
within any give Type C network, there can be multiple instances of Type D or E networks.

This diagram illustrates the concepts and inter-relationships of the different Issolation classes.
[[http:staff.washington.edu/dmorton/files/NID/traffic_iso_main.jpg]]

Alternatives:

- Status Quo

Due to the limited number of zones defined above and since we do not foresee the need for that list to grow, we
accept that building physically separate zones with the associated interconnections is the best approach.

- More isolated networks

As we foresee the need for many more zones that will only function best with full physical isolation, we decide

that we need to build more physically separate networks. As an example, the School of Medicine becomes its own network with individual routers and is provided with a dedicated border relationship to the other networks.
Similar individual networks may be built out for the other zones detailed above.

- Hybrid physical & logical

Since the list of zones will continue to grow and we feel that a consistently deployed physical infrastructure supporting multiple logical networks is operationally attainable, we decide to deploy a set of physically separate networks that consistently support the class of logically isolated zones required within. So the Medical Center and Main Campus would be deployed as entirely physically separate infrastructure, yet each would support via logical separation their respective child zones. For example Dorms and Wireless being logical isolation for Main Campus physical infrastructure.

Tradeoffs:

Recommendation:

Type A (Physically separate infrastructure; maximum isolation)

- Patient-monitoring network(s)
- Clinical data network(s)
- Out-of-band management network
- Primary campus data network family
- R&D network
- Datacenter SAN
- Emergency access/Backup ISP network

Type B (Hybrid of A and C)

- VoIP

Type C (shared physical infrastructure, distinct routing domains via VRF)

- UW Seattle main campus
- UW Bothell
- UW Tacoma

We should look at the above three Type C networks and decide if they are truly separate, since the near term plan is to have 7600 class routers at Bothell & Tacoma, and thus are really just separate aggregation routers that connect to the "Primary campus data network family" and then have the set of required below Type C networks instantiated at their site... - dsinn
---

Within each of the above Type A networks or Type C zones, a separate Type C network zone could be defined for any of the following service categories:

- Data center(s)
- Remote SANs?
- Residence halls
- Wireless
- Building access

- Elevator control
- Backend private network?

Type D (Policy-Routing)

- Quarantine?
- Trusted?

Type E (no infrastructure-based isolation)

- One of the options within each network zone

Network Topology: Geographic or Organization-based?

Security requirements, and to a lesser extent performance and application protocol concerns, have led to increased pressure to offer organiztion-based network topologies. There is a continuum of options which include:

- Full VLAN: organizational units get a distinct VLAN, with a common broadcast domain, even if that means transporting broadcast domains across a Layer 3 backbone. Experience at other institutions suggests this is a problem-prone scenario.
- VRF domains: organizational units get a distinct Virtual Router Facility on shared access routers, linked via MPLS or VLANs or separate fiber/copper. There is a questiohn of how many VRF domains can be reasonably supported.
- Limited VLANs: organizational units get a distinct VLAN, but not a common broadcast domain. Each access router offers semantically equivalent connectivity services to corresponding VLANs (i.e. departments), which means having equivalent firewall rules for traffic going outside the department, and special firewall rules to permit distinct (usually low-impedance) rules among hosts in the same department (but on different subnets).
- Subnet splits: often it is sufficient to give departments on a shared subnet their own separate subnet so they can avoid fear of hostile machines in other depts accomplishing layer-2 attacks (e.g. arp spoofing), and/or implement their own firewall rules.
- Purely geographic: no separation by org unit; subnets are defined by geography alone.

Layer 1 Service Requests

We are beginning to get requests for wavelength or dark fiber connectivity. In some cases, this may be the right answer; in others, it may be possible to find a way to leverage primary infrastructure and still meet customer requirements. The continuum of requests includes:

- Networks supporting network research. Example: Tom Anderson's participation in the NSF GENI project. It is plausible that attempting to build an overlay network to support clean-sheet next-generation Internet design might have undesired side-effects on the research. If so, a "personal lambda" network would be appropriate.

- Networks supporting vast amounts of real-time streaming data. Example: John Delaney's Neptune project to instrument volcanically active parts of the sea floor is one example.
- Networks supporting vast amounts of bulk data transfer. Harvey Freeman at Caltech and others involved in the CERN Large Hadron Collider project would be examples, as would Tony Tyson's Large Synoptic Survey Telescope (LSST) -~~UW is a partner~~- that has a 3 billion pixel camera and generates 30 terabytes of data each night.
- Networks that require high bandwidth and low latency, and are mission critical. Jim Loter in UW's College of Engineering is interested in dark fiber (or ?) to provision a remote SAN installation on campus. In due course, one can imagine considerable demand for remote SAN installations, on and off campus.

In categorizing these requests, the Tom Anderson case may have the strongest rationale for avoiding layer 3 infrastructure; however, the other cases place a heavy burden on the advocates of a normal shared layer 3 infrastructure to demonstrate that performance, latency, predictability, reliability, and diagnosability goals can be met without providing unshared wavelengths.

L3 Isolation Technologies

- Over-provisioning and Rate-limiting

In this model we would combine an abundance of bandwidth in the core with high-level rate limiting of all other traffic. Customer traffic requiring assurances would need to ingress on a router where we have a large amount of bandwidth toward the core, along with commensurate core bandwidth to reach its ultimate destination. Further, we would need to limit the bandwidth of all other traffic on the edge router such that it's cumulative load can not choke the uplinks to the detriment of the preferred traffic.

- Prioritization via DiffServ

In this model we would leverage prioritization within the layer-3 components in our network to resolve contention issues on a per-port and per-hop basis. Traffic from a customer requiring assurances would need to have the DiffServ code-point (DSCP) verified and/or imposed upon at the closest C&C managed point to the customer as we can provide. These DSCP bits would then allow other devices within our network to appropriately reserve buffering space and queue the customers traffic. Because of this preferential treatment being reliant upon DSCP bits, which can be set by any users of the network for their own traffic, we would need to insure that for the non-assured traffic the DSCP markings are reset to a default value. Failure to do so could result in the other users of the network co-opting the preferred queuing that we have implemented. This would result in the traffic requiring assurances from not receiving the actual prioritization that we had intended since the non-prefered traffic was now intermixed with it.

- Bandwidth guarantees

In this model we would need to have some logical separation of the end

user traffic at the edge which would then
allow us to send the preferred and non-prefered traffic each over
their own path. We would thus be able to
dedicate links for each class, and insure that for the preferred
traffic we had provisioned and appropriate
level of capacity.

The Bandwidth Guarantees model will need to be researched further to
determine exactly how we would implement
the path separation and will be very tied to the logical separation
model we choose.

Description of "Many Networks" Scenario

This model takes the approach of building many smaller, fairly
identical network infrastructures each supporting
a unique class of traffic. Each of these networks is built to an
identical model, scaled relative to the
requirements, and functions as an atomic unit. There is then a
higher level network in the hierarchy that meshes
the many small networks together.

So, for example, we might take the planned Medical Center topology of
two border routers with point-to-point links
to four medical center aggregation routers and use this as the basic
atomic unit for the "many" networks needed
to support the rest of our clients. This building block would then be
replicated for each of the following
areas we see as needing a individual building block:

- VoIP
- Wireless
- Medical Center
- UW Bothell
- UW Tacoma
- NW UW Seattle Campus
- SW UW Seattle Campus
- E UW Seattle Campus
- Residence Halls
- Streaming
- Data Center

Each would be on a class of hardware as is appropriate for the
specific area, yet we would not exceed four
aggregation routers.

There would then need to be an interconnection model to allow each of
the atomic units to talk to one another.
At the interfaces for this interconnection we would be able to place
traffic inflection devices to impose the
policy appropriate for the traffic class (so firewall's, IPS's, IDS's,
etc.). There would then be an overall UW
border router pair that would then support the connectivity to outside
of the University.

Within each of these atomic units, there may be the need for
differentiated traffic classes (say in the Medical
Center where there is an interest in setting up clinic centric
networks), but these differentiation would not
extend outside of the atomic unit.

Description of "Unified Networks" Scenario

This model takes the approach of focusing on building a small hand-full of physical networks which we then virtualize our traffic classes over. Most likely these would remain the current three that we have, but that we would move out of the model of having functional separation on a given router in a network and move to virtual separation on all of the routers in the given network.

For example, within the Seattle Campus network today Residence Halls, Wireless, Remote Access, Streaming, Tipping-point protected and non-Tipping-point protected traffic classes are defined based on the router that the segment is terminated upon; in this new model all six of these traffic classes would be available on any of the routers within the Seattle Campus network, as is required by the customer needs where the router is. We would also be able to more-easily deploy other traffic classes as required via soft configuration of the new class, instead of dedicating hardware as is required in our current model.

Similarly the differing traffic classes needed within the Medical Center and the VoIP networks would similarly exist within those atomic networks.

Interconnecting the three physical networks can either continue over the Seattle Campus on the non-Tipping-Point protected class or any of a number of other options: At the UW Border Routers, over a new traffic class, or on separate hardware.

Description of "Extreme/Aggressive" Scenario

This model takes the approach that most service providers have taken with their network, that being a focus on a single network infrastructure that supports all of their traffic classes. We would build one overall physical network infrastructure that would carry all of the required traffic classes as virtual networks within. There would no longer be a separate and distinct Medical Center network, nor a isolated VoIP network. Both would be logically isolated on the same physical hardware as all of the campus networks. Further, we could leverage L2 VPN technologies to allow us to deliver the personal lambda's from the edge of our network to where the customer is.

END OF DOCUMENT
