



# Presentation to the OIT Tech Meeting

Michael Scarpellino

Manager, Network Architecture and Applications

2 May 2007

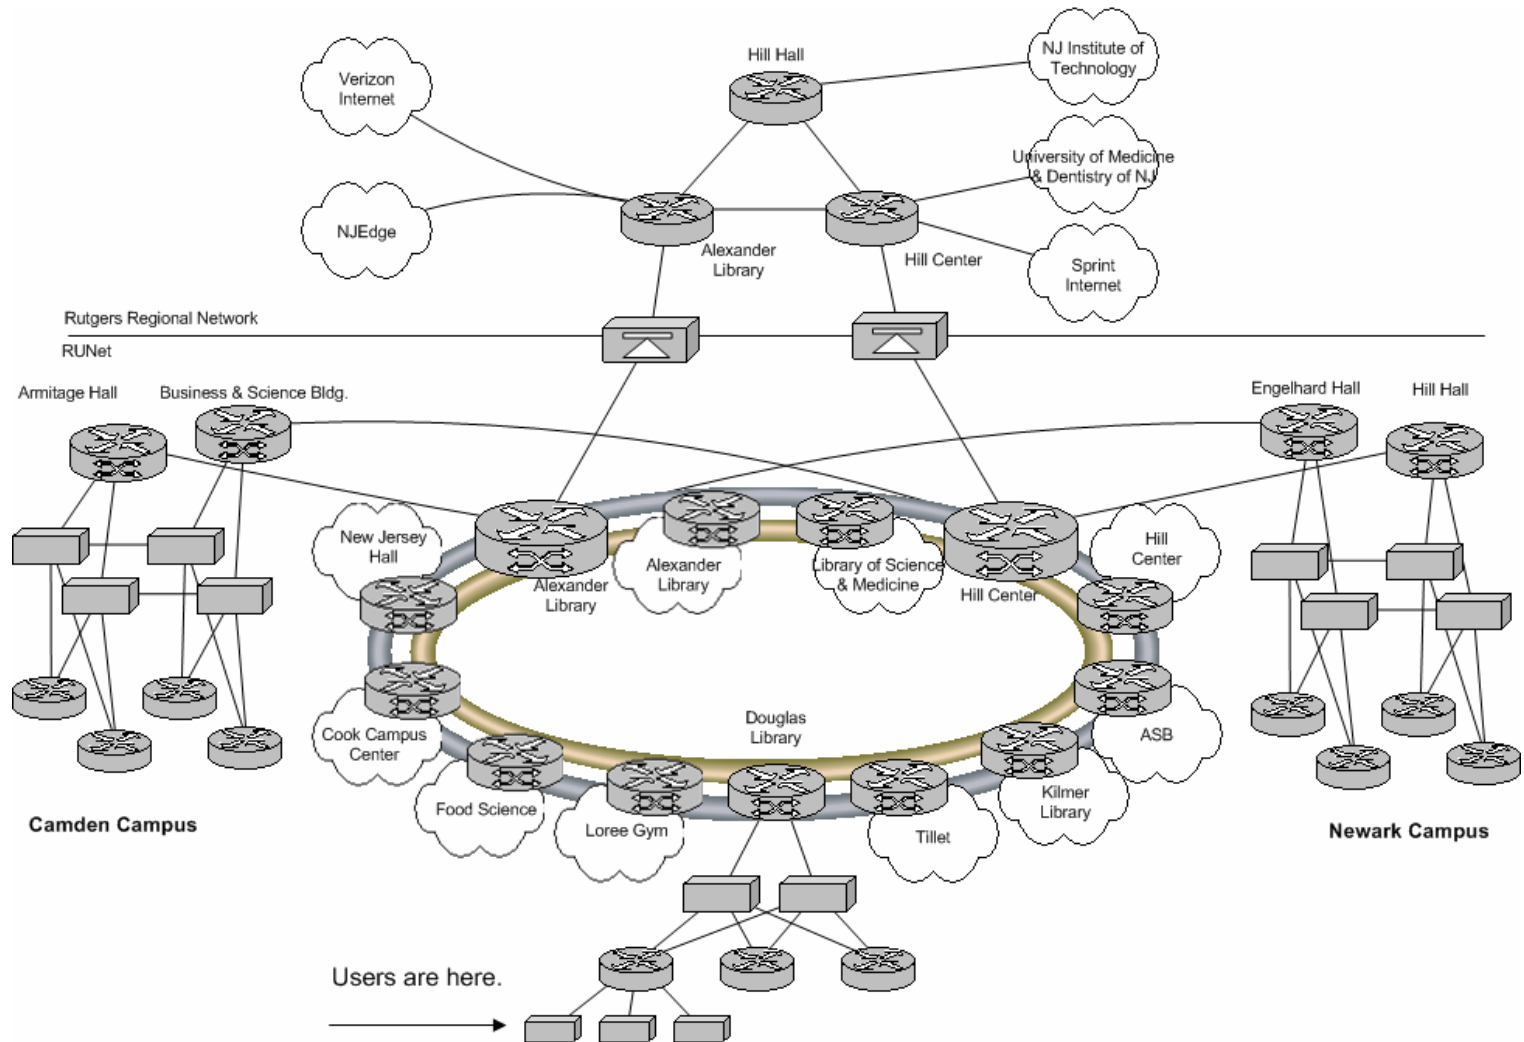
# Today's discussion

- Quick review of RUNet structures and services
- Review of bridging and switching
- New Layer 2 services
- Concepts for future RUNet development

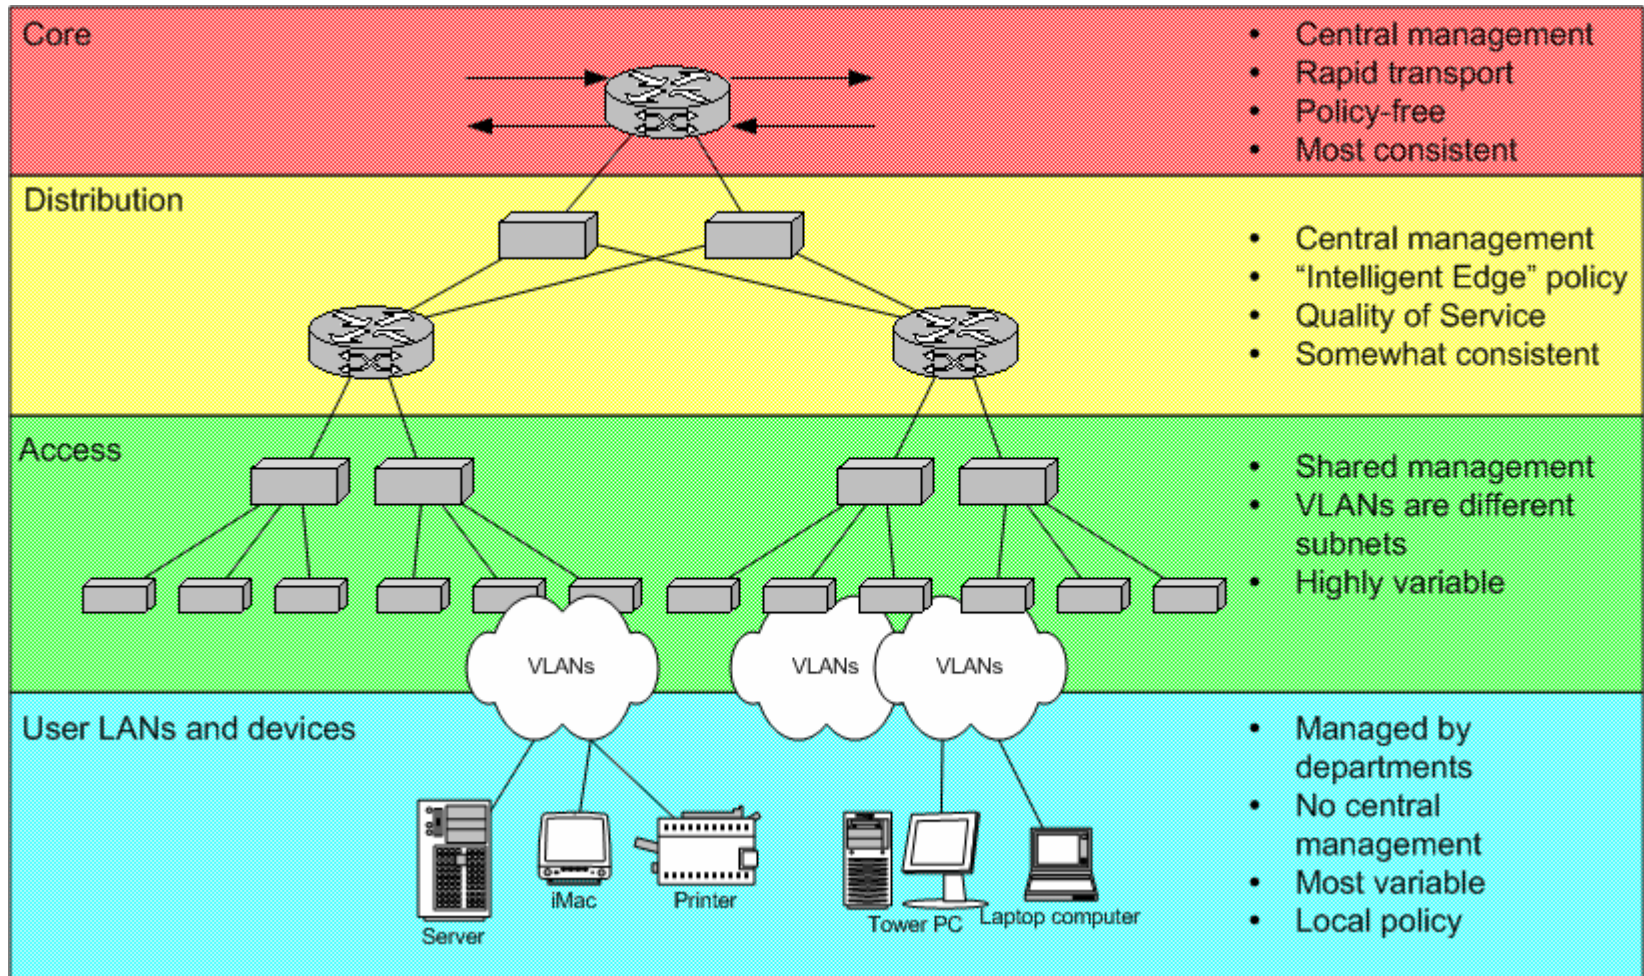
# Why is RUNet successful?

- Design standards and best practices
- Common, modular solutions replicated many times
- Core – Distribution – Access
- Planned aggregation practices
- Designed-in fault management
- Automated monitoring and management
- Our staff!

# RUNet Structure



# RUNet Layers



# Introducing the Catalyst 4500

- New access layer switch selection
- Modular chassis
- High density port capacity
- New capabilities in the closet
  - 28 or 64Gbps; 21 or 48Mpps
  - L3/L4 capability
  - Supports current 802.1 standards
  - Support for anticipated demands
- Reduces access layer complexity
- Projected 7 year life cycle
- Cost effective

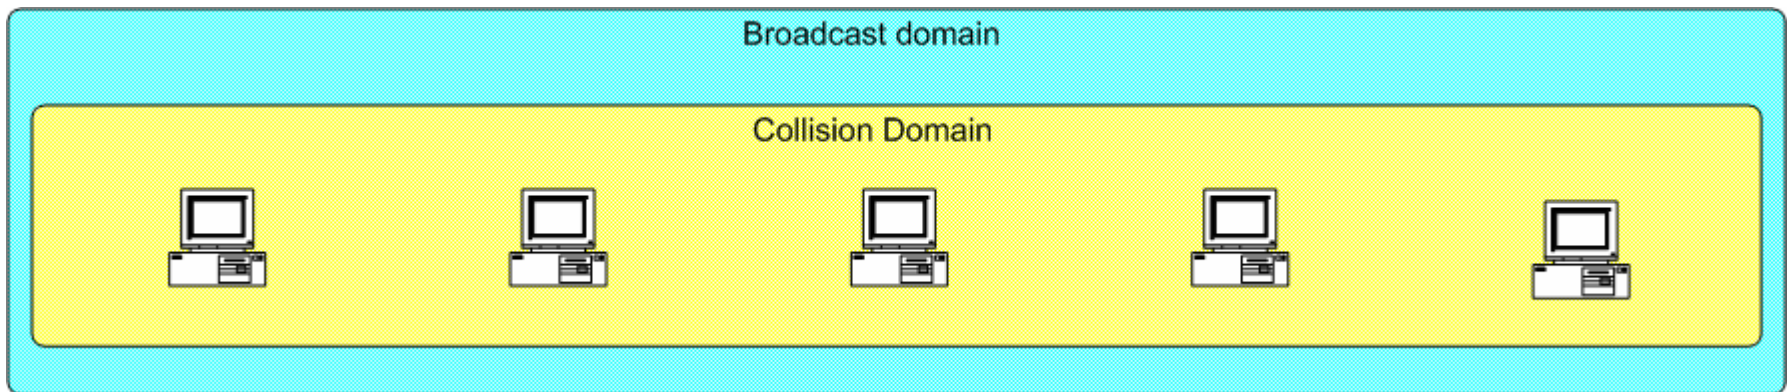


# Access Layer

- Shared access introduces management challenges
  - Many VLANs across several devices
  - Shared occupancy buildings mean several groups making changes
  - Need for strict configuration management
- Size and scale
  - Access layer is the most challenging layer to manage
  - Most devices on RUNet are in the access layer
  - Balance need for consistency with flexibility
- TD evolving its standards and practices
  - Access layer reference standard evolving with technology
  - New technologies introduce new service opportunities
  - New ways of addressing old challenges

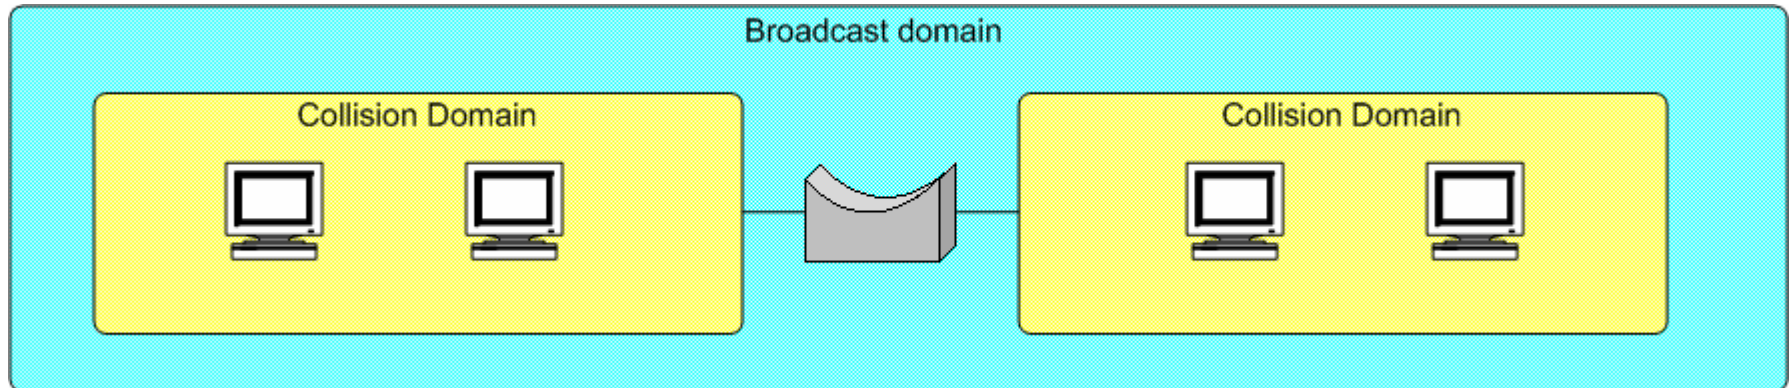
# Shared networks

- Shared networks place several hosts in a common broadcast/collision domain
  - Only one host can transmit data at a time
  - Collision domain limited by propagation delay, which limited the number of hosts on the segment



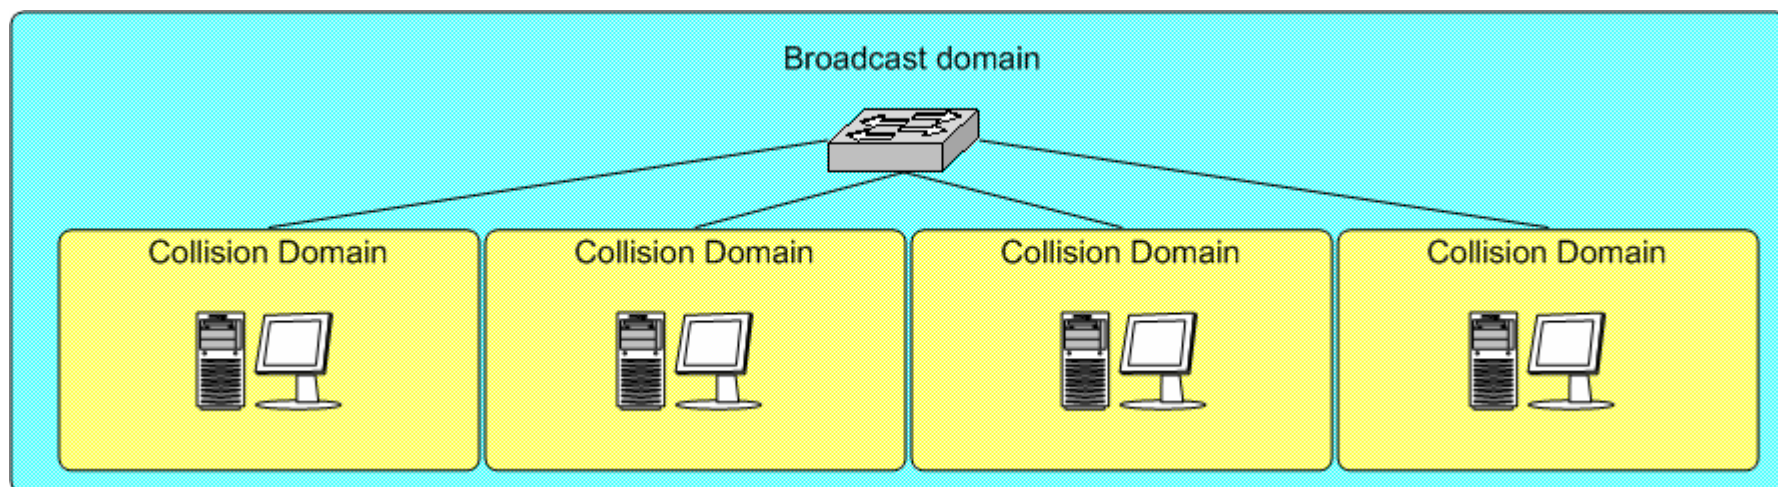
# Bridged networks

- Shared networks could be segmented into bridged domains to reduce the number of hosts in a collision domain
- Broadcasts still have to traverse the length of the network segment
- Increased performance
- Careful not to introduce loops into topology



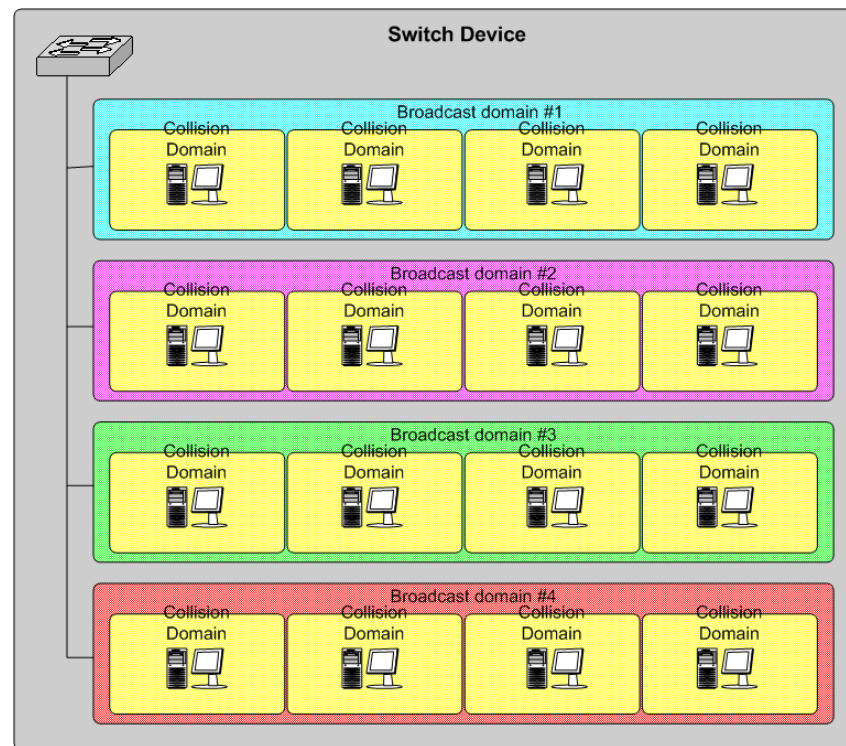
# Switched networks

- Multi-port bridges evolved into switches over time
- More ports per bridge, more memory for store & forward
- Costs came down, making single-host collision domains practical
- Increased efficiency – many hosts can transmit at the same time
- Networks can grow arbitrarily large – store & forward eliminates propagation delay
- Spanning tree helps reduce risks of loops



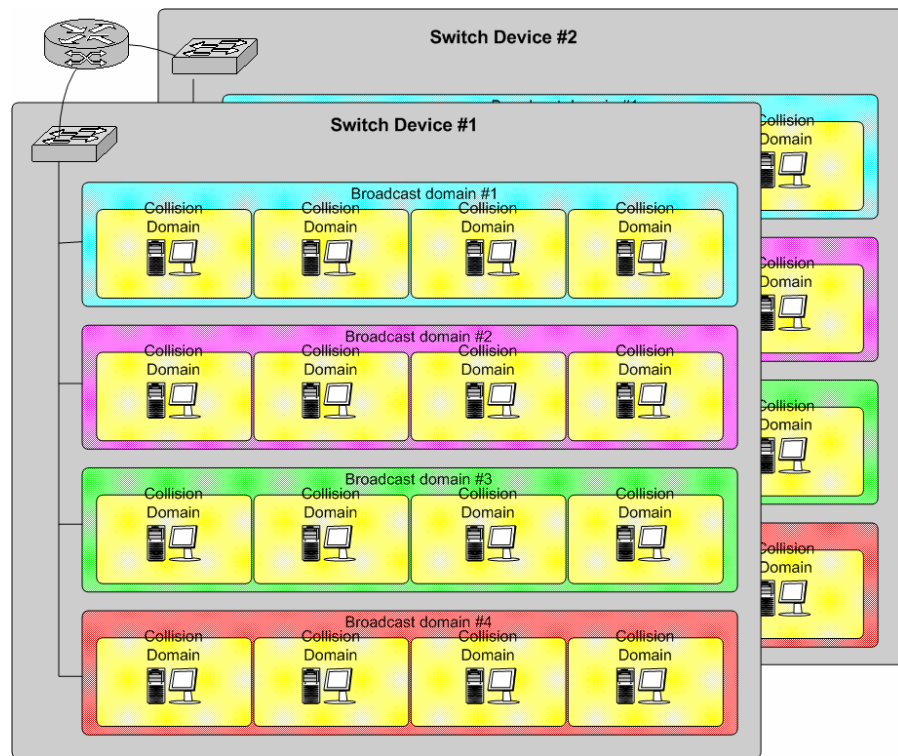
# Virtual LANs

- Switches can be partitioned into several “Virtual LANs” to segment broadcast domains
- Each VLAN has its own forwarding and address tables
- Packets cannot move between VLANs without a gateway, such as a router



# Virtual LANs

- Protocols like VTP make it convenient to extend VLANs between multiple switches, creating extended switched domains
- Each device in the infrastructure maintains forwarding and address tables for every device in each VLAN
- Lots of memory



# Virtual Trunk Protocol

- Cisco protocol designed to manage and distribute VLAN ID tags
- Designates a switch to be a SERVER, and others to be CLIENTS
- VLANs created on the SERVER are propagated to the CLIENTS
- VTP version 2
  - Introduces domain passwords
  - Clients still listen to the device with the highest database version
  - Huge security holes
- VTP version 3 introduces server authentication, redundant servers and failover
  - More secure than version 2
  - Not fully supported across all RUNet devices
- RUNet currently standardized on VTPv2

# IEEE 802.1q

- IEEE 802.1q introduced a standard way of carrying frames for multiple VLANs over a single bridge port
- 802.1q inserts a VLAN-ID field in the Ethernet frame header to identify the what VLAN originated the frame
- Ports participating in 802.1q framing are called “trunk” ports, and use different forwarding rules
- Trunked ports participate in a VLAN called the “native VLAN” and transmit to it using standard frames


# IEEE 802.1s

- IEEE 802.1d, Spanning Tree Protocol (STP), transmitted discovery frames on the physical topology, and pruned any loops
  - One topology for all VLANs
  - Topology changes intrusive across the whole infrastructure
- IEEE 802.1s moves STP into the VLANs, transmits discovery frames within each VLAN
  - Allows different VLANs to be configured on different physical links
  - Allows traffic engineering for VLANs
  - Topology changes only affect VLANs on those physical links
  - Loops are detected and pruned on a per-VLAN basis
- Improves security – topology problems in one VLAN only affect that VLAN

# HOWTO: Run your own switching domain

- TD Side
  - TD and UCS agree what VLANs to exchange
  - TD sets VTP domain name and password
  - TD creates VLANs in our domain
  - TD assigns a port to trunk, disables all dynamic protocols and sets native VLAN
  - TD assigns your VLANs to our trunk port for forwarding
  - TD forwards/receives frames to your VLANs using agreed VLAN IDs
- Your side
  - Set your own VTP domain and password, or use VTP TRANSPARENT
  - Create VLANs on your switch
  - Assign a port to trunk with TD, disable dynamic protocols and set native VLAN
  - Assign VLANs you want to receive from TD to your port
  - Forward/receive frames on your infrastructure using the agreed tags

# HOWTO: Run your own switching domain




# HOWTO: Trunk a PIX firewall with 802.1q

- TD Side
  - TD and UCS agree what VLANs to exchange
  - TD sets VTP domain name and password
  - TD creates VLANs in our domain
  - TD assigns a port to trunk, disables all dynamic protocols and sets native VLAN
  - TD assigns your VLANs to our trunk port for forwarding
  - TD forwards/receives frames to your VLANs using agreed VLAN IDs

- PIX Firewall One-arm example

```
interface GigabitEthernet 1
  no ip address
  no security-level
  no nameif
interface GigabitEthernet 1.100
  vlan100
  nameif outside
  security 0
  ip address 172.16.28.202 255.255.255.252
interface GigabitEthernet 1.99
  vlan99
  nameif inside
  security 100
  ip address 192.168.1.1 255.255.255.0
```

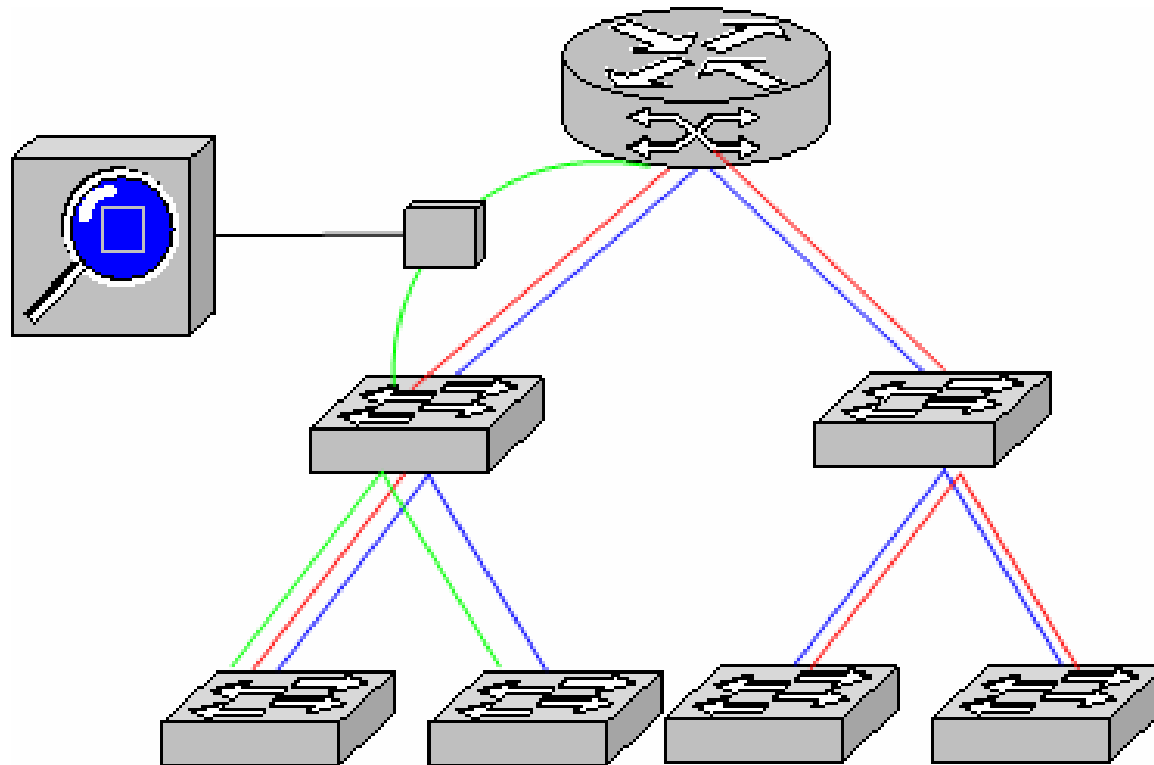
# HOWTO: Trunk a PIX firewall with 802.1q



# HOWTO: IDS on TD managed infrastructure

- First, we recommend the Cisco ASA platform, a routed solution with integrated firewall, IDS and Anomaly detection
  - No mucking about with the Layer 2.
- TD has discouraged IDS because we do in-band management, and will not allow other devices to be inserted between our own
- IEEE 802.1s allows us to reassign your VLANs to a different trunk port, bypassing our management trunks
- You must purchase additional ports
- Once you have purchased the port, you may install an optical or (if applicable) copper tap to feed your IDS or packet sniffer device

# HOWTO: IDS on TD managed infrastructure



# RUNet Tomorrow

## Technologies being evaluated for use on RUNet

- Wave Division Multiplexing (CWDM and DWDM)
- MPLS Virtual Routers and Traffic Engineering
- 10 Gigabit Ethernet
- Firewall Service Module service
- IPv6
- Wireless backbone connections
- Central VoIP services

# Questions?

## Websites referenced during the presentation

- TD presentations and papers (<http://www.td.rutgers.edu/papers>)
- TD Tools (<http://www.td.rutgers.edu/tools/>)
- For any follow-up questions about this presentation, please contact the NOC at 5-7541, or by email at [noc @ rutgers.edu](mailto:noc@rutgers.edu)