Enterprise Networks

Jennifer Rexford
COS 461: Computer Networks
Lectures: MW 10-10:50am in Architecture N101

http://www.cs.princeton.edu/courses/archive/spr12/cos461/

Networking Case Studies

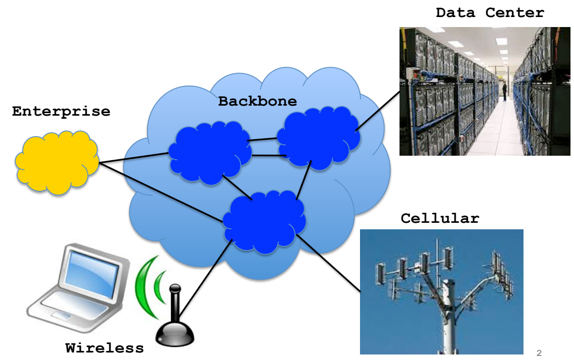

Simple Enterprise Design

- Single layer-two subnet
  - Hubs and switches
  - Gateway to the Internet
  - Single IP address block
- Local services
  - DHCP
  - DNS

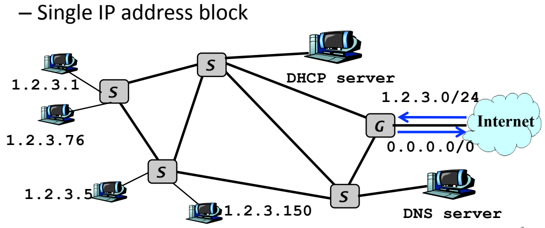

Limitations of Simple Design

- Ethernet scalability and performance
- Single ISP reliability and performance
- Limited IP address space
- Unwanted Internet traffic
- Privacy and isolation within the enterprise
- Detecting and preventing bad behavior from inside

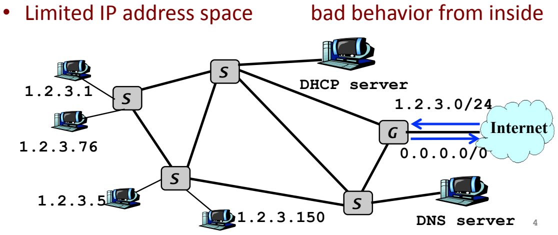

Beyond Ethernet Switching

Scalability Limitations of Ethernet

- Spanning tree
  - Paths that are longer than necessary
  - Bandwidth wasted for links not in the tree
- Forwarding tables
  - Bridge tables grow with number of hosts
- Broadcast traffic
  - ARP, DHCP, and broadcast applications
- Flooding
  - Frames sent to unknown destinations

Hybrid of Switches and Routers

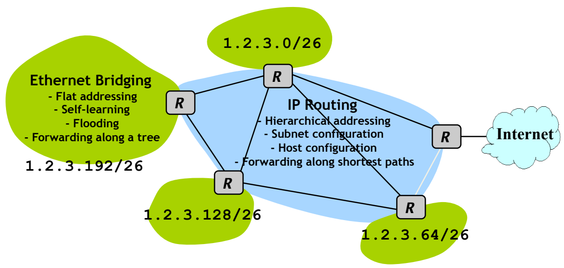

Limitations of Hybrid Design

- No plug-and-play and mobility between subnets
- Need consistency between IP addressing & routing

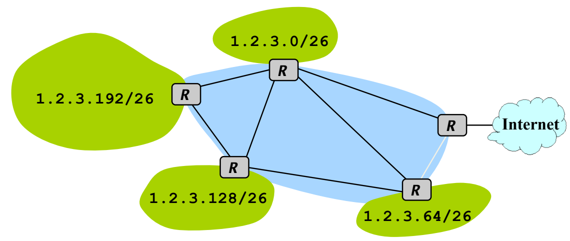

Virtual Local Area Networks

Early Days of Ethernet LANs

- Thick cables snaked through cable ducts
  - Every computer they passed was plugged in
- All people in adjacent offices on the same LAN
  - Whether they belonged together or not
- Users grouped based on physical layout
  - Rather than organizational structure
- Security, privacy, and scalability limitations…

Today's Ethernet LANs

- Changes introduced by hubs and switches
  - Every office connected to central wiring closets
  - Often multiple LANs (k hubs) connected by switches
  - Flexibility in mapping offices to different LANs
- Can group by organizational structure
  - Better privacy: snooping in promiscuous mode
  - Separate IP addresses: one IP subnet per LAN
  - Better security: access control at IP routers
  - Better load management: isolate broadcast/flooding

People Move, and Roles Change

- Organizational changes are frequent
  - E.g., faculty office becomes a grad-student office
  - E.g., graduate student becomes a faculty member
- Physical rewiring is a major pain
  - Requires unplugging the cable from one port
  - … and plugging it into another
  - … and hoping the cable is long enough to reach
- Would like to "rewire" the building in software
  - The resulting concept is a Virtual LAN (VLAN)

Example: Two Virtual LANs

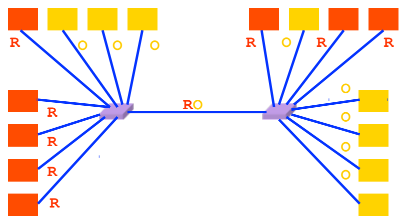

13

Making VLANs Work

- Changing the Ethernet header
  - Adding a field for a VLAN tag
  - Implemented on the bridges/switches
  - … but can still interoperate with old Ethernet cards
- Bridges/switches trunk links
  - Saying which VLANs are accessible via which interfaces
- Approaches to mapping access links to VLANs
  - Each interface has a VLAN color
  - Each MAC address has a VLAN color

14

Uses of VLANs (See the Survey Paper)

- Scoping broadcast traffic
- Simplifying access control policies
- Decentralizing network management
- Enabling host mobility

15

Problem: Limited Granularity

- Limited number of VLANs
  - Placing multiple groups in the same VLAN
  - Reusing limited VLAN
- Limited number of hosts per VLAN
  - Divide a large group into multiple VLANs
- One VLAN per access port
  - Supporting VLANs on the end host
  - Supporting multiple groups at the router

16

Problem: Complex Configuration

- Host address assignment
  - Wasting IP addresses
  - Complex host address assignment
- Spanning tree computation
  - Limitation of automated trunk configuration
  - Enabling extra links to survive failures
  - Distributing load over the root bridges

Open question: can we do better than VLANs?

17

Multiple Internet Connections

18

Motivation for Multi-Homing

- Benefits of multi-homing
  - Extra reliability, e.g., survive single ISP failure
  - Financial leverage through competition
  - Better performance by selecting better path
  - Gaming the 95^{th}-percentile billing model

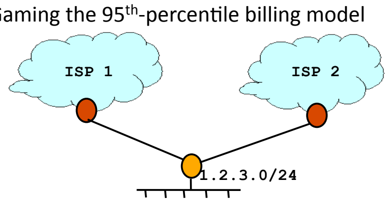

Multi-Homing Without BGP

Inbound Traffic

- Ask each ISP to originate the IP prefix
- … to rest of the Internet

Outbound Traffic

- One ISP as a primary, the other as a backup
- Or simple load balancing of all traffic

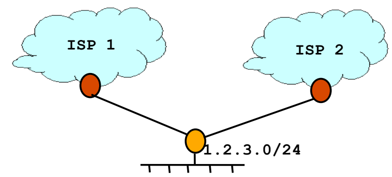

Multi-Homing With BGP

- Inbound traffic
  - Originate the prefix to both providers
  - Do *not* allow traffic from one ISP to another
- Outbound traffic
  - Select the “best” route for each remote prefix
  - Define BGP policies based on load, performance, cost

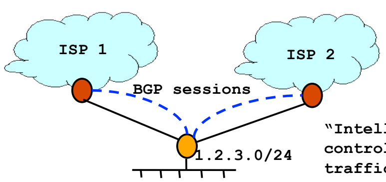

“Intelligent route control” or “multi-homed traffic engineering”.

Interconnecting Multiple Enterprise Sites

Challenges

- Challenges of interconnecting multiple sites
  - Performance
  - Reliability
  - Security
  - Privacy
- Solutions
  - Connecting via the Internet using secure tunnels
  - Virtual Private Network (VPN) service
  - Dedicated backbone between sites

Connecting Via the Internet

- Each site connects to the Internet
  - Encrypted tunnel between each pair of sites
  - Packet filtering to block unwanted traffic
  - But, no performance or reliability guarantees

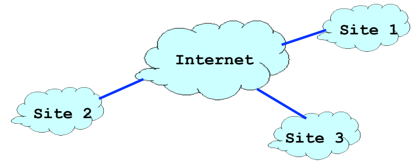

Virtual Private Network (VPN)

- Each site connects to a common VPN provider
  - Provider allows each site to announce IP prefixes
  - Separate routing/forwarding table for each customer
  - Performance guarantees

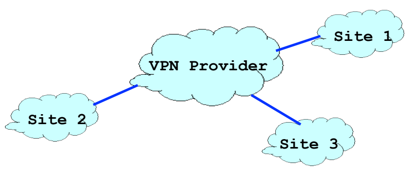

Middleboxes

Enterprise Internet Connection

- Multiple middleboxes
  - Intrusion prevention system
  - Network address translator
  - Firewall
  - Traffic shaper
- Handling bad internal users
  - Filtering IP packets with spoofed source IP addresses
  - Logging which MAC address has each IP address

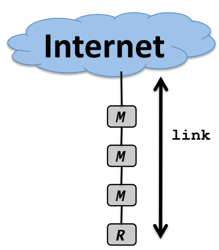

Internal Middleboxes

- Network divided into regions
  - E.g., departments within a campus
  - E.g., public computers (servers, WiFi) vs. private
- Network divided by roles
  - E.g., human resources vs. engineering
  - E.g., faculty vs. students
- Sometimes physically separate networks
  - E.g., ATM machines, campus safety, media streaming

Princeton Campus Network

http://www.net.princeton.edu/index.html
http://www.net.princeton.edu/statistics/
http://www.net.princeton.edu/whatsnew.html

Internet Connections

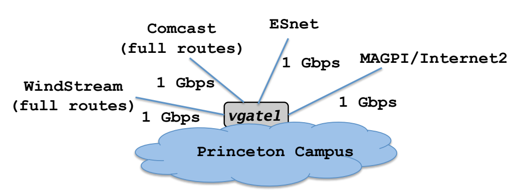

- Two commercial ISPs: Comcast and WindStream
- Two research networks: ESnet and Internet2
- Non-profits: McCarter Theater, Princeton Public Library, and Princeton Regional Schools

Princeton Public Internet Traffic

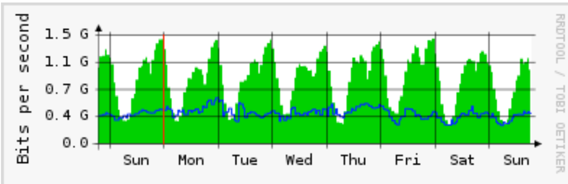

- Traffic volumes over the past week
  - Green: traffic *from* the Internet
  - Blue: traffic *to* the Internet

Three Internal Networks

- Campus Data Network
  - Connects dorms, academic and administrative buildings, campus WiFi, etc.
- Princeton Private Network
  - Environmental systems, power, security cameras, building locks
- VoIP Network
  - VoIP phones in data center, chemistry, neuroscience, Forrestal campus, and all new construction
  - Separate for disaster recovery & traffic management

Campus Data Network

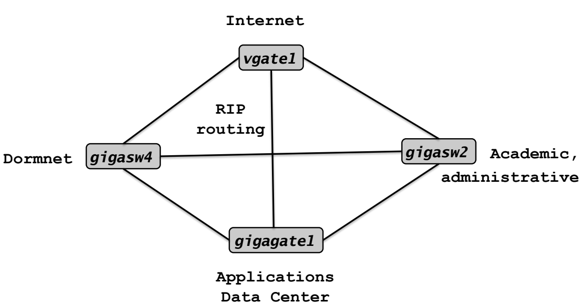

Data Center (Forrestal Campus)

- 40,000 square feet with 1800 computers
- Multiple tiers of backup power
- Minimizes energy for cooling and power

Virtual Private Network (VPN)

- Online campus resources
  - E.g., some Princeton University library resources
  - Not available from outside of campus
- External resources with Princeton subscription
  - E.g., digital libraries from ACM and IEEE
  - Accessible from a Princeton IP address
- Princeton VPN service (vpn.princeton.edu)
  - Secure network connection layered over IP network
  - … connects you to an internal Princeton machine

Aruba WiFi Access Points

- Adaptive radio management
  - Automatically assigns channel and power settings
  - Channel load balancing to distribute clients
  - Coverage hole detection

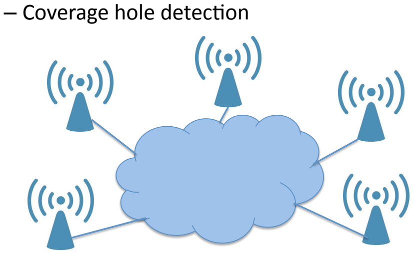

WiFi Anecdote (puwireless)

- Single large VLAN
  - Enabling seamless mobility on campus
- Limited address space
  - 16K or 32K IP addresses
  - 3 hour DHCP leases
- Frequently a large number of users
  - Several thousand to up to 10,000
  - … may soon run low on IP addresses

37

WiFi Anecdote (puwireless), Continued

- Bug in Android and IOS smart phones
  - Don't release DHCP lease on IP addresses
  - Offloads ARP processing to the chipset, to avoid waking up sleeping device on ARP requests
  - … but DHCP timeout is handled by the processor
- So, can have IP address collisions
  - DHCP lease expires, but the phone doesn't know
  - DHCP server gives the IP address to someone else
  - … and both devices respond to ARP requests!

http://www.net.princeton.edu/android/android-stops-renewing-lease-keeps-using-IP-address-11236.html

38

WiFi Anecdote (puwireless), Continued

- Working with Google and Apple on the problem
- Longer-term solution
  - Move to larger, private address block (10.0.0.0/8)
  - Use network address translation (NAT) to communicate with the public Internet
- Benefits
  - Avoids running out of IP addresses
  - Introduces long delay before reusing an address
  - Seems like a good solution, right?

39

WiFi Anecdote (puwireless), Continued

- Solution makes troubleshooting harder
  - Public IP addresses shared by many users
  - … due to network address translation
- Example: DMCA violations
  - Student downloads copyrighted material on WiFi
  - Company comes to Princeton to complain
- Given IP address, can OIT identify the student?
  - With NAT, cannot pinpoint a unique MAC address
  - … without much more detailed (flow-level) logs

40

Conclusions

- Enterprise networks
  - Campuses and companies
  - Access to local services and the Internet
- Challenges
  - IP address limitations
  - Hybrid switch and routed network
  - Load balancing over upstream ISPs
  - Protecting users and the Internet from each other
- Next time: data-center networks

41