Migrating a Campus Network from Flat to Routed

Campus Network Design & Operations Workshop

UNIVERSITY OF OREGON

Last updated 2nd October 2024

NSRC
Network Startup Resource Center

Migrating a Campus Network from Flat to Routed

- Topics to be considered for any migration:
  - DHCP
  - Broadcasts
  - Planning Migration
  - Other Hints

Migrating a Campus Network: Flat to Routed
Campus Network Design & Operations Workshop

SECTION 1: DHCP

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

Ideal routed campus network

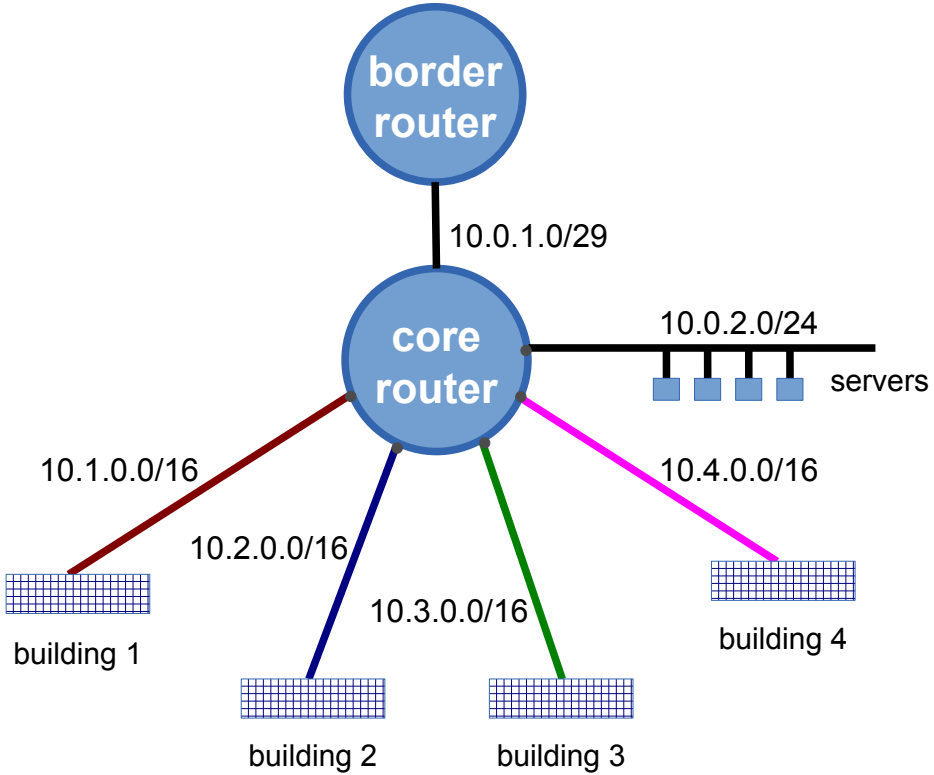

Changing from flat network implies:

- Nearly everything needs renumbering!
  - Well, you can keep one subnet on its old addresses
  - What's hardest to renumber – servers perhaps?
- So, first get as much as possible onto DHCP
- This lets you renumber centrally
- Each subnet will need a separate DHCP pool of addresses

Quick refresher: DHCP (RFC2131)

- A DHCP exchange is 4 UDP messages:
  - Client sends “Discover” (broadcast)
  - One or more servers replies with “Offer”
  - Client picks one offer and sends “Request”
  - Server responds with “Ack” to confirm
- Address is granted for a finite “lease time”
  - When this is nearly over, client must request again to continue using the address

DHCP options (RFC2132)

- DHCP response can also contain other settings to configure the client
  - Netmask, default gateway
  - DNS servers, default domain
  - SIP server (IP phones)
  - TFTP boot server (PXEboot / diskless clients)
- Centralises all client network configuration

Managing Devices

- Highly recommended to use DHCP to configure even devices with “static” IP addresses like printers, phones, admin workstations
  - DHCP servers can be configured with a mapping of MAC address to fixed IP address
- DHCP logs are a useful source of address pool availability information

DHCP Broadcasts

- You need to respond to the DHCP Discover broadcasts on every subnet
- Option 1: run DHCP service on the core router itself
  - Can be awkward to manage if you have a lot of static MAC address mappings or custom options
  - Not all routers include a DHCP server (e.g. NX-OS does not)
- Option 2: use a feature on the router called “DHCP relay” or “DHCP helper”
  - Relays requests to one or more remote DHCP servers

DHCP Relay

- DHCP relay agent

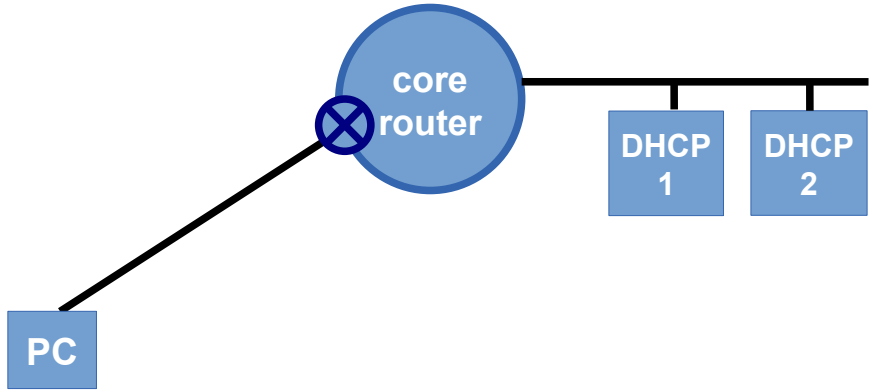

DHCP Relay

DHCP relay agent

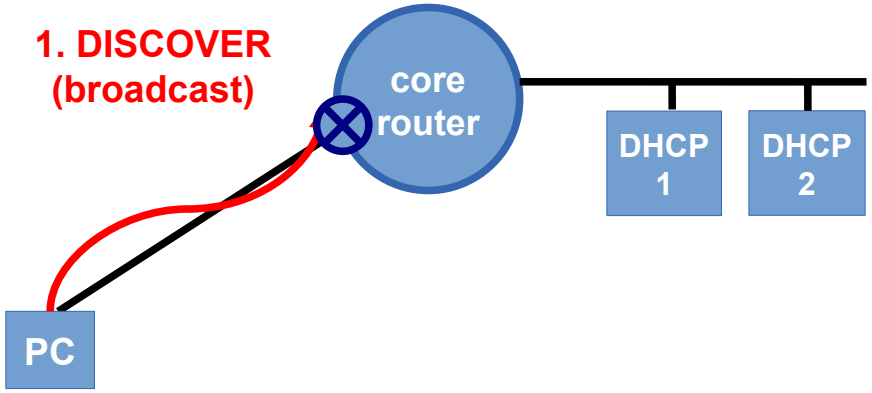

* Client can request broadcast response using the B flag

DHCP Relay

DHCP relay agent

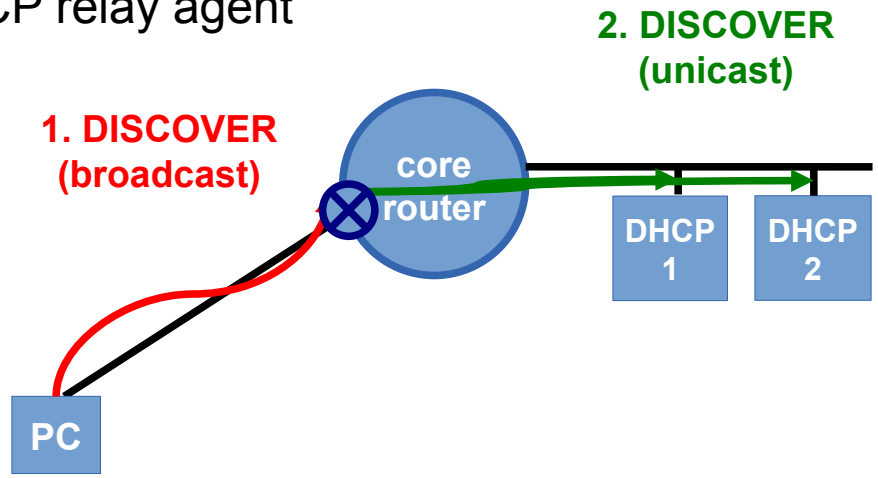

* Client can request broadcast response using the B flag

DHCP Relay

DHCP relay agent

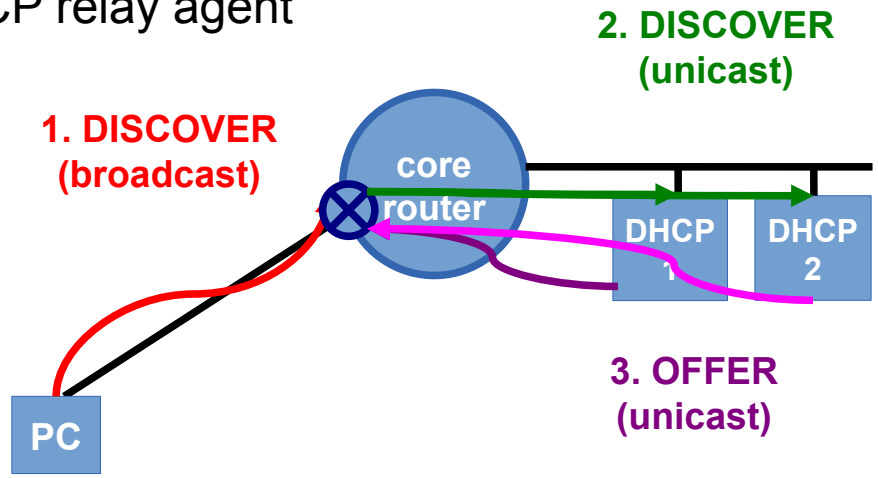

* Client can request broadcast response using the B flag

DHCP Relay

DHCP relay agent

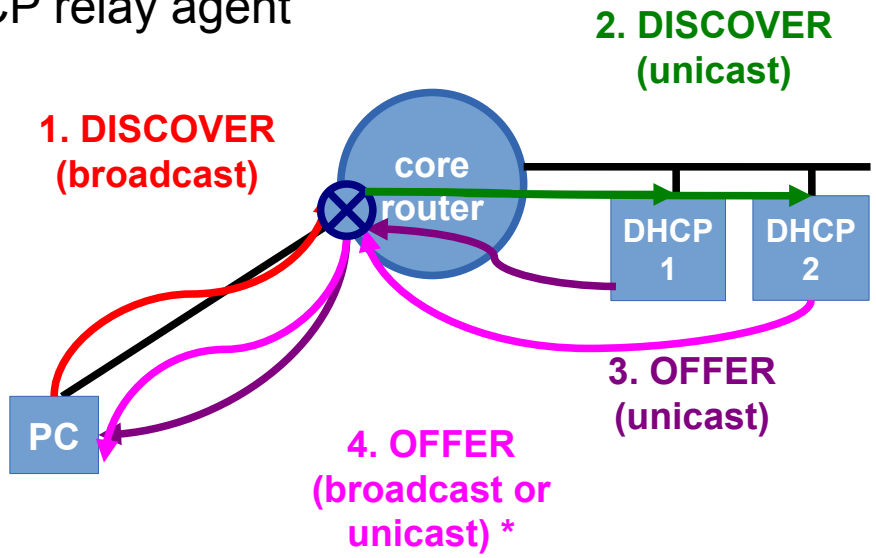

* Client can request broadcast response using the B flag

DHCP relay configuration

- Cisco: Repeat for every interface where DHCP service is required

```
interface Vlan100
  ip address 10.1.1.1 255.255.255.0
  ip helper-address 10.0.2.4
  ip helper-address 10.0.2.5
```

- Nexus (NX-OS) is similar

```
feature dhcp
ip dhcp relay
interface Vlan100
  no shutdown
  ip address 10.1.1.1/24
  ip dhcp relay address 10.0.2.4
  ip dhcp relay address 10.0.2.5
```

DHCP relay configuration

- Juniper: Repeat the last statement for every other interface that needs DHCP

```
set forwarding-options helpers bootp relay-agent-option
set forwarding-options helpers bootp server 10.0.2.4
set forwarding-options helpers bootp server 10.0.2.5
set forwarding-options helpers bootp interface irb.100
```

DHCP server configuration

- Define each subnet where service is required
  - (Windows DHCP server: "DHCP scope")
- Example for Linux isc-dhcp-server:

```
subnet 10.1.1.0 netmask 255.255.255.0 {
    option routers 10.1.1.1;
    option subnet-mask 255.255.255.0;
    range 10.1.1.100 10.1.1.199;
}
```

- Remember that every subnet has a *different* gateway

Questions?

Migrating a Campus Network: Flat to Routed
Campus Network Design & Operations Workshop

SECTION 2: BROADCASTS

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

Removing dependence on broadcasts

- Subnetting means that broadcasts don't propagate
- You may *think* that some services require broadcasts for devices to find each other, e.g.
  - Wireless access points to contact their controller
  - VOIP phones to contact their voice gateway
  - Windows clients to locate their servers, printers etc
- In reality, there is almost always another way to do this
- You just need to find (and test) the right solution for each one

Example 1: Unifi access points

- Unifi access points will find their controller on a different subnet if you provide a special DHCP response attribute (number 43)
  - The value is typically entered in hex as 0104 + IP address, e.g. 10.2.3.64 = 0x01040A020340
  - Other clients on the same subnet will ignore this attribute
- Alternatively, they can resolve "unifi" + default domain in the DNS
  - The default domain also comes from DHCP
  - This approach requires that you have managed DNS on your site
- For other wireless vendors, check their documentation

Example 2: VOIP phones, cameras etc

- Often these have similar mechanisms (DHCP or DNS)
- At worst, you can statically configure the hostname or IP address of the voice gateway on the phones themselves

Example 3: Windows clients/servers

- You might have an application where the client is a Windows .exe that talks to a local Windows server (e.g. a finance application)
- Modern Windows clients can use the DNS
- Older applications can use the WINS name resolution service
  - In DHCP, set the "WINS server" attribute to point to the IP address of one or two local servers which will maintain a name mapping table

Questions?

Migrating a Campus Network: Flat to Routed
Campus Network Design & Operations Workshop

SECTION 3: PLANNING MIGRATION

Planning Migration

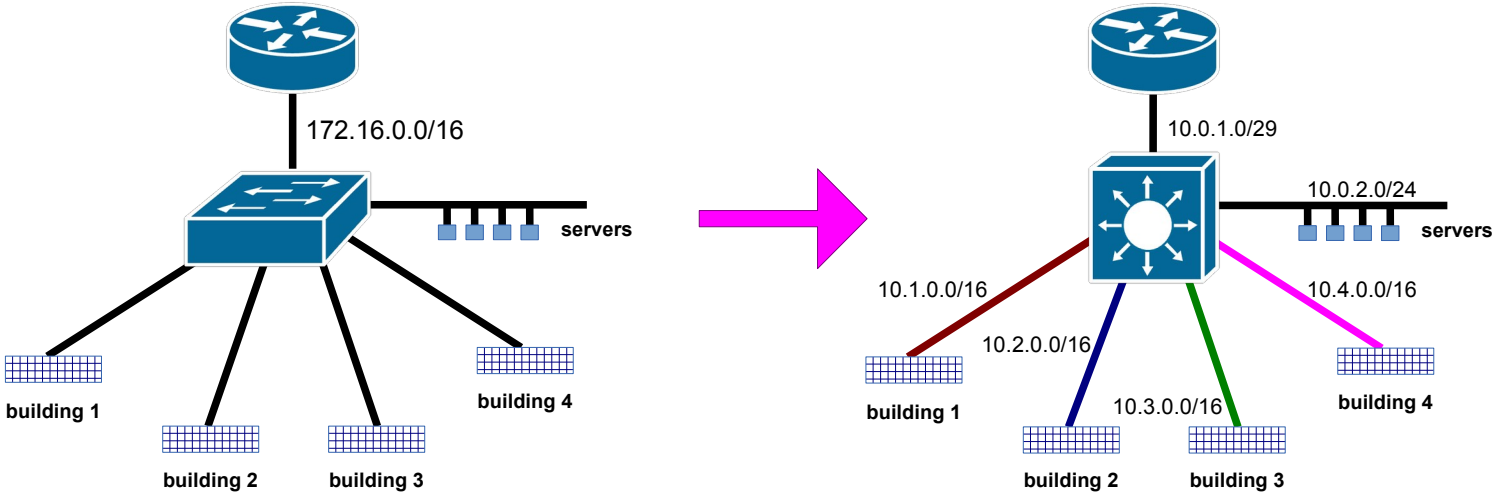

General Principles

- No "big bang"!
- Series of small, incremental changes
- Test at each stage
- Plan to rollback at each stage
  - You will discover things that break
  - Understand the problem, correct and try again
- Localize outages and give advance warning

Managing Complexity

- Incremental steps means you will be running parts of old and new configuration in parallel
- Remember to strip out old configuration when it is no longer needed
  - So it's understandable
  - So you are not left with any configuration which might be important but actually isn't
- It all gets easier with experience

Quick Example

- You want to replace an old switch with a new one
  - How would you go about it?

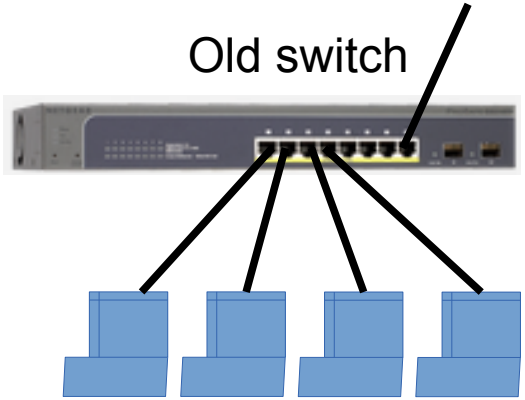

For discussion!

Longer Example

- Migrate one building from the flat network onto three new subnets (*e.g wired, wireless, guest*)

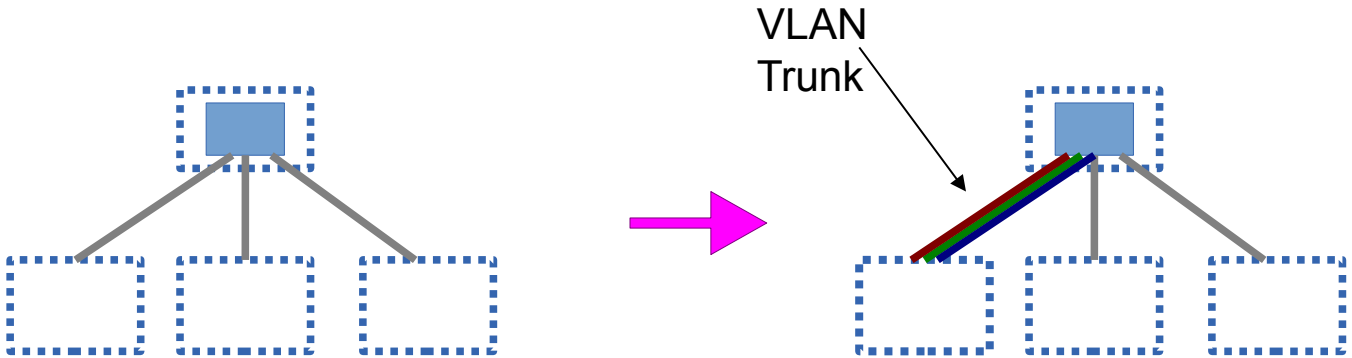

Before (detail)

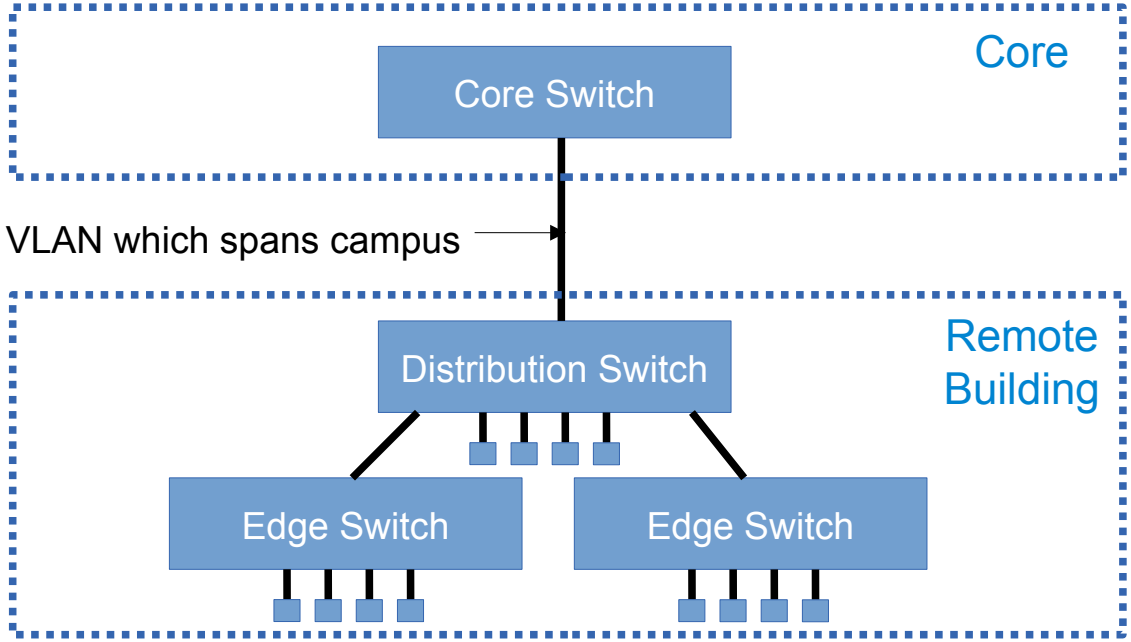

After (detail)

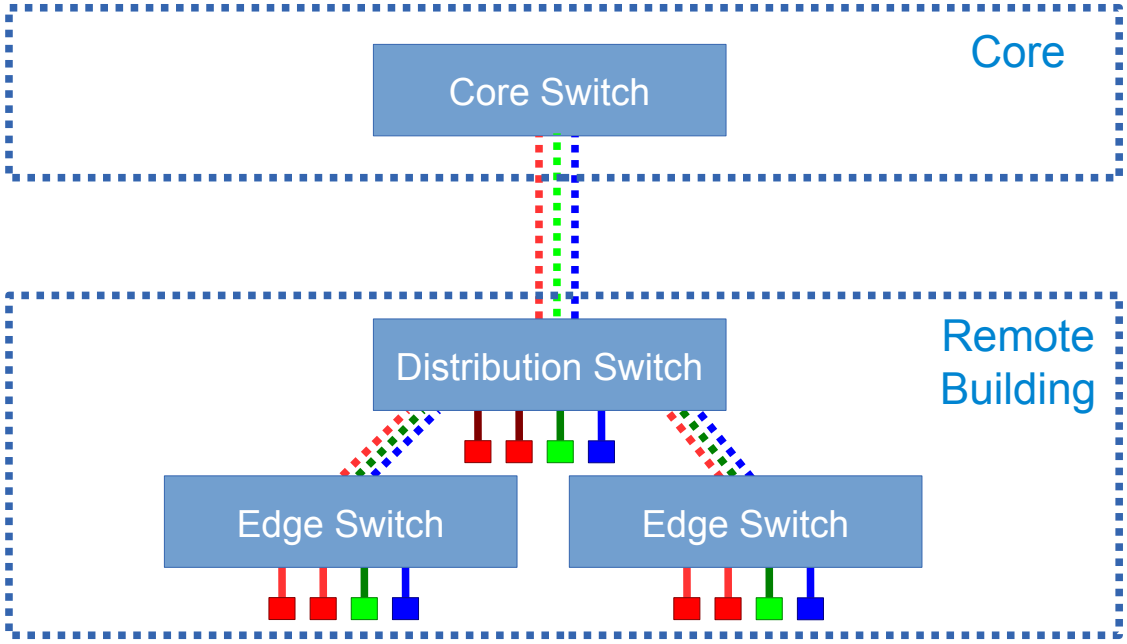

How do we get from there to here?

1. Create new VLANs in core

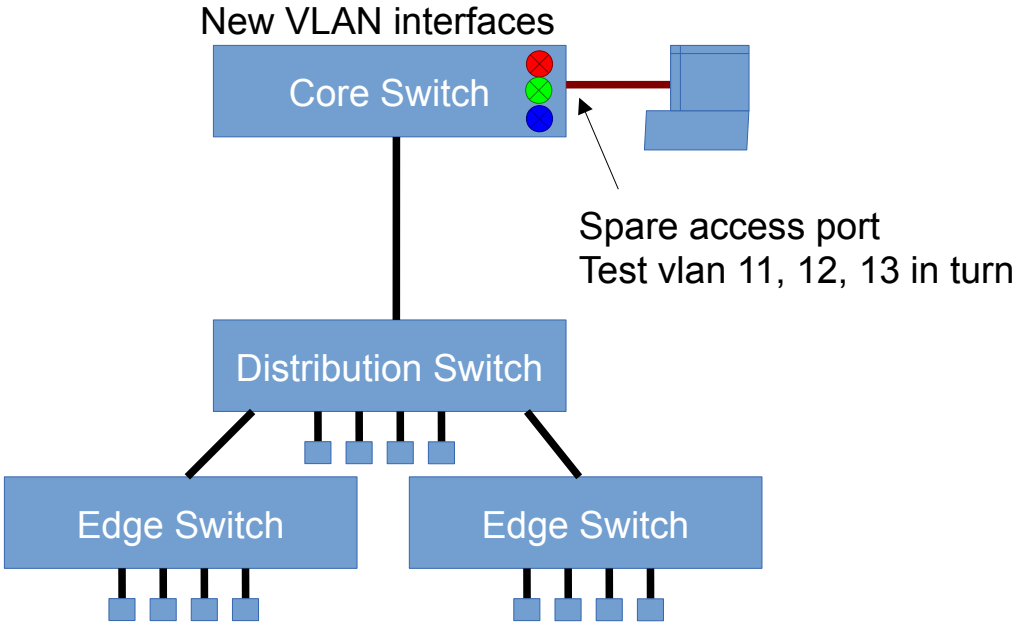

Test all client functionality, e.g. DHCP, routing

Rollback plan

- Undo changes to core switch
- Take a copy of the configuration before you start making any changes, so you have a reliable reference

2. Add new VLANs to trunk

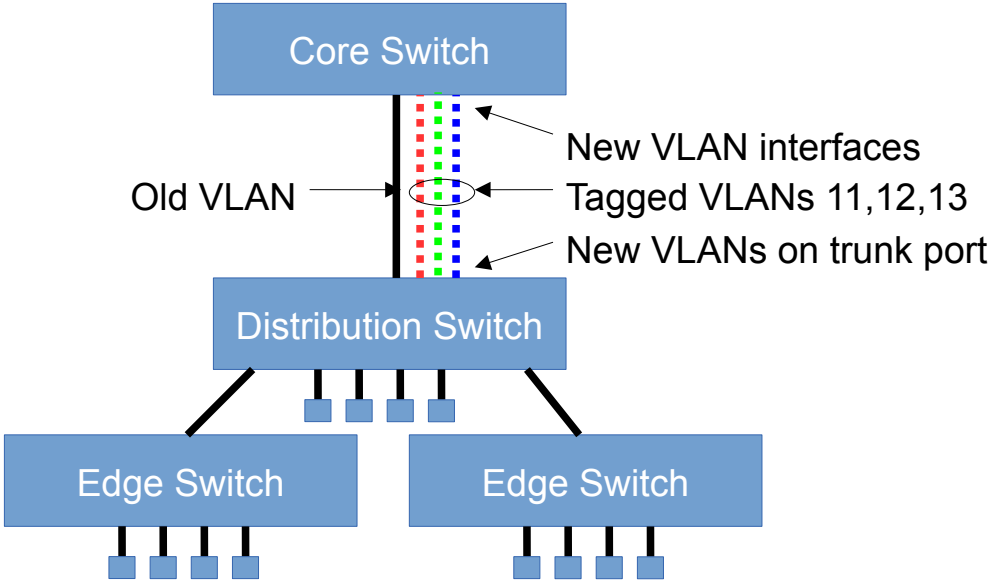

Should not break anything! (But check anyway)

Choice to make

- Run the old VLAN untagged (a.k.a. "native"), together with the new VLANs tagged; OR
- Change the old VLAN to tagged at both ends
  - Bigger change, but may be easier to understand
- Whichever you are most comfortable with
- No clients should be affected yet
- Rollback plan: revert these small config changes

3. Extend VLANs to edge

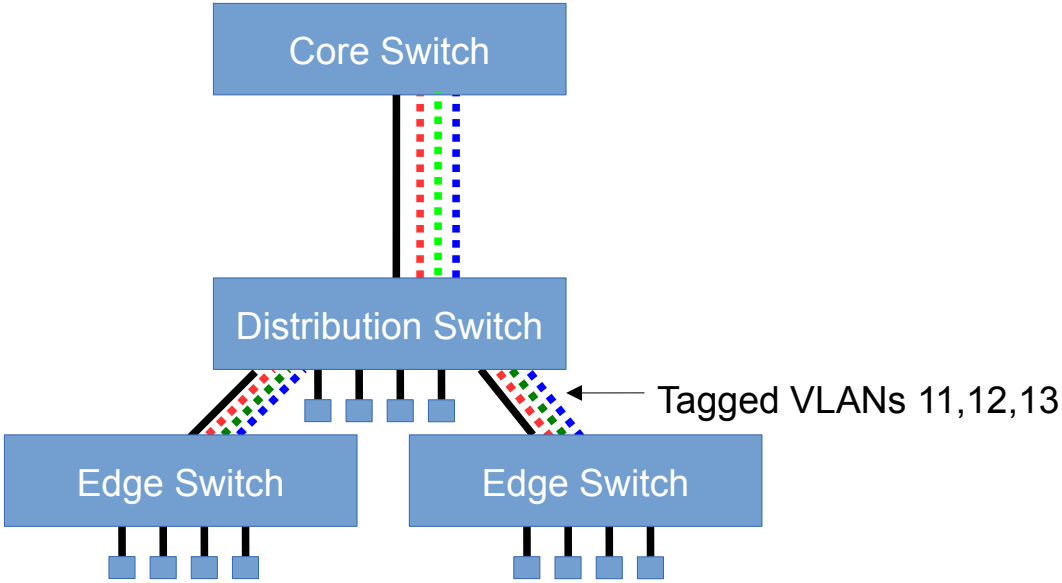

Again, nothing should break

4. Test with spare access ports

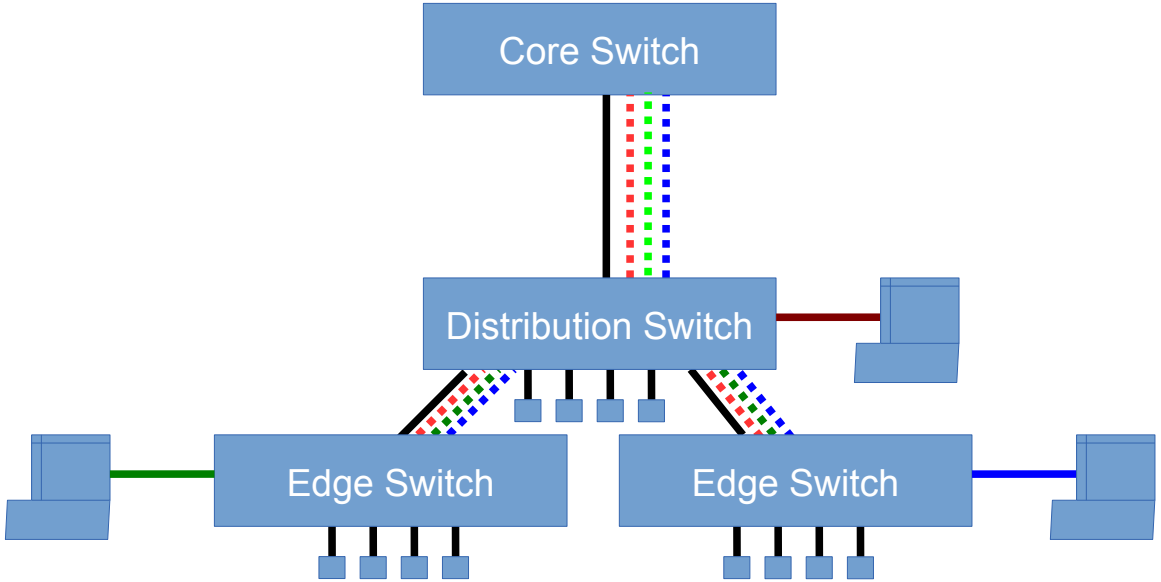

Re-test all client functionality, DHCP, routing

5. Re-assign edge ports individually

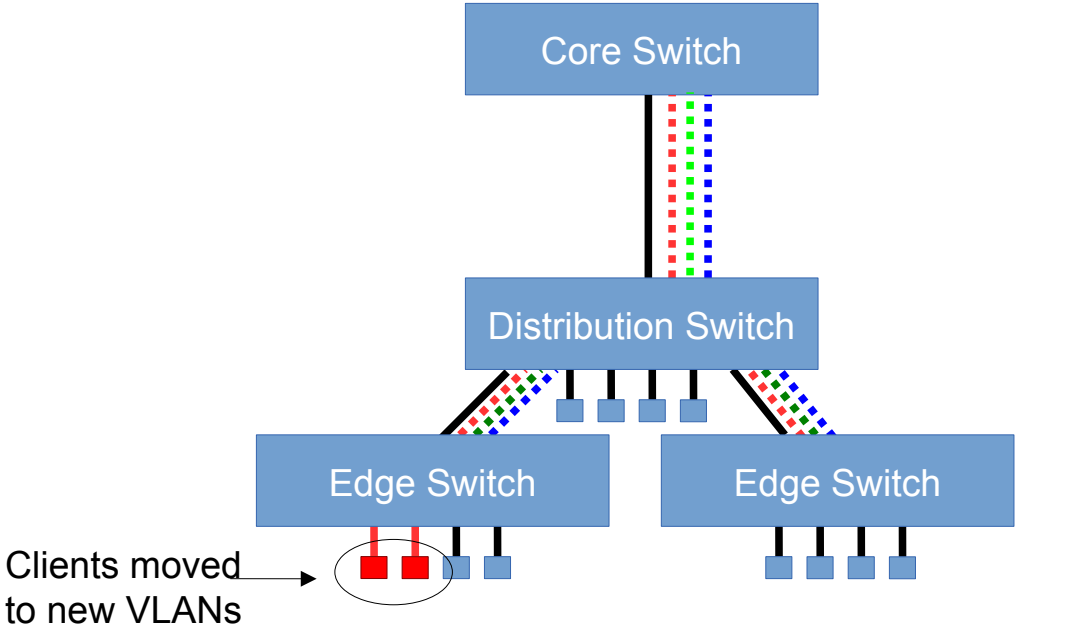

Controlled interruption to service

6. Move all the remaining clients

- Hint: a 5-second shutdown on the port can help force clients to re-DHCP

```
interface GigabitEthernet 0/3
 shutdown
 no shutdown
```

- Problematic clients can be rolled back to the old VLAN while you work out how to fix them
- For important devices, check in DHCP logs that they have come back

7. Renumber the switches

- Give the switches new management IP addresses on the appropriate new VLAN
  - Remember the default gateway will change
  - Try not to lock yourself out!
  - Serial console is safest way to do this
- Might choose to do this earlier (before moving clients)

8. Check nothing on old VLAN IPs

- wireshark / nmap / angry IP scanner are useful tools for this
  - Connect a laptop to each new VLAN, but configured statically with an IP address on the old VLAN range

```
nmap -sP -n x.x.x.x/x  # old range
```

  - You will discover any devices which are still statically configured with old IP addresses
  - Find them and correct them

9. Strip out the old (flat) VLAN

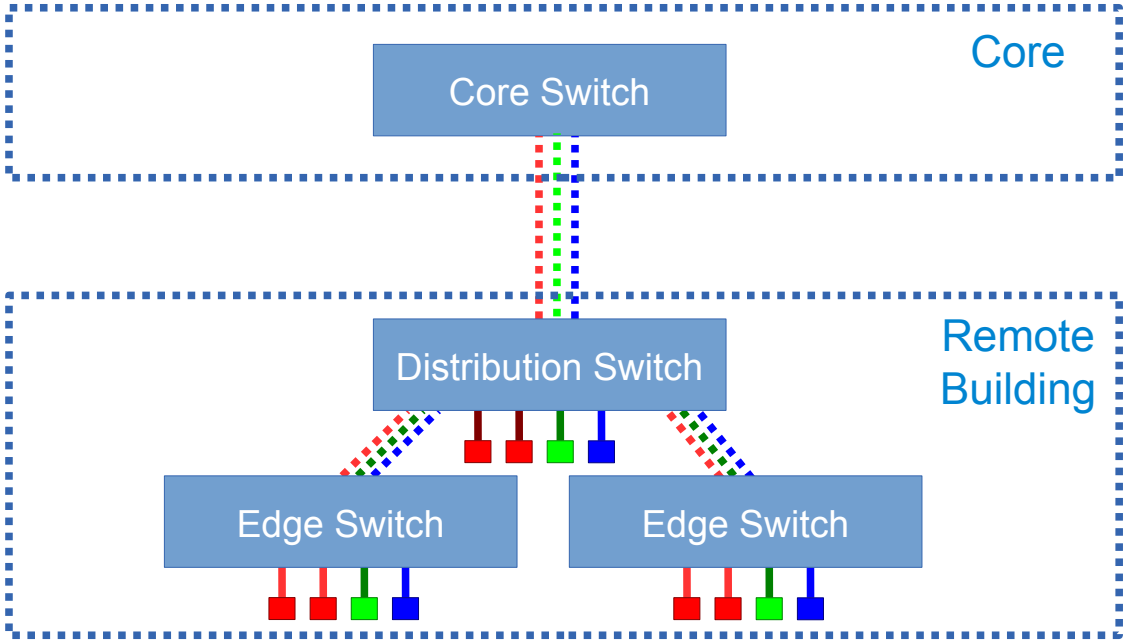

Final test to sign-off

Summary

- Lots of steps, but each one is easy to rollback
- Plan in advance what the final configuration will look like, and the steps to get there
- Make sure you know how to rollback any step
- Test **before** and **after** each change
  - Monitoring key devices with e.g. Nagios or LibreNMS can give you extra confidence nothing has broken

Migrating a Campus Network: Flat to Routed
Campus Network Design & Operations Workshop

SECTION 4: OTHER HINTS

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

Plan within your constraints

- Maybe some of your switches are dumb?
- Maybe some parts of your network must be in service at particular times?
- Make a plan which best fits your situation

Testing

- Make a "test building" using a spare distribution switch, and test it in the lab before cutting over real buildings
- Turn your tested configs for distribution and edge switches into templates that you can easily clone and modify

Lease time (1)

- It’s a good idea to reduce the DHCP lease time in advance of renumbering
  - e.g. say current lease time is 24 hours
  - Reduce this to 10 minutes then wait 24 hours
  - By this time you'll know every device is refreshing its address every 10 minutes
  - Minimises time for new addresses to be picked up
- Put back up after change tested and successful

Lease time (2)

- Choice is a balance between DHCP server load, and being able to reuse pool addresses quickly after clients have gone away
- Suggested final lease times:
  - between 4 hours and 24 hours for wired
  - between 5 minutes and 30 minutes for wireless
- Windows Server default of 7 days is way too long!

Other hints and tips

- If your core switch has only SFP ports, a copper SFP is useful for testing
  - Some are gigabit only, some auto-negotiate 10/100/1000
  - Some are SFP+ 10G or 10G/1G

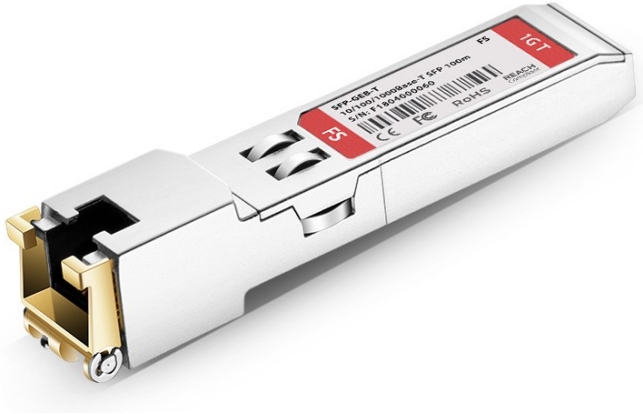

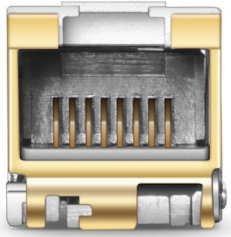

Other hints and tips

- If you move an IP address from one device to another, other devices may have the old MAC address cached in their ARP table for a while
  - Cisco routers are worst: 4 hour ARP timeout!
  - “clear arp-cache” or “clear ip arp x.x.x.x” may be required
- “write memory” as each change completed and tested

Renumbering servers

- If you are renumbering servers, remember to reduce the DNS TTL in advance of changes
  - Allow enough time for all caches to expire records with the old TTL
  - Put it back up afterwards
- “Secondary IP addresses” can be useful when renumbering servers on the same VLAN
  - Both old and new IP addresses are active at the same time

Don’t forget (static) routes

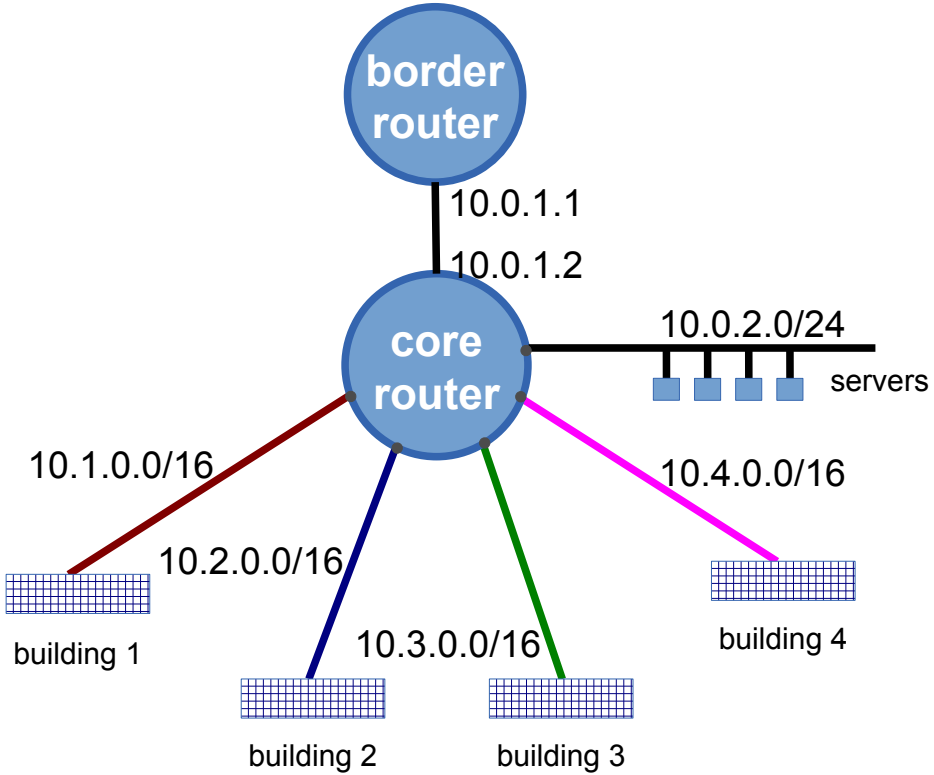

Don’t forget (static) routes

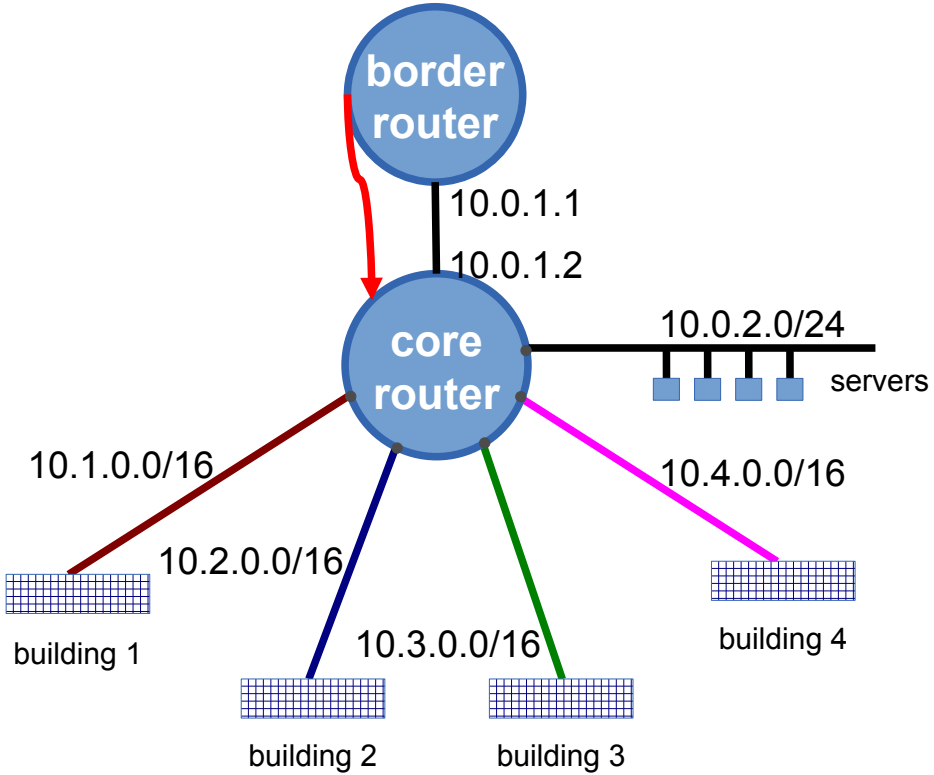

Don’t forget (static) routes

10.0.2.0/24 via 10.0.1.2
10.1.0.0/16 via 10.0.1.2
10.2.0.0/16 via 10.0.1.2
10.3.0.0/16 via 10.0.1.2
10.4.0.0/16 via 10.0.1.2
OR:
10.0.0.0/8 via 10.0.1.2

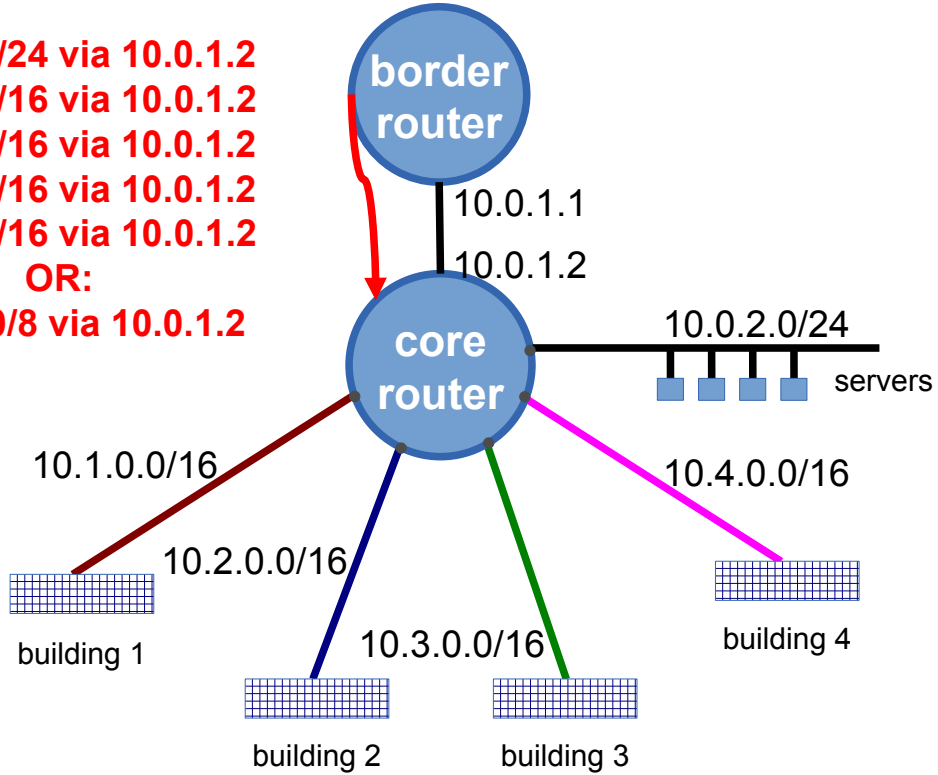

Don’t forget (static) routes

10.0.2.0/24 via 10.0.1.2
10.1.0.0/16 via 10.0.1.2
10.2.0.0/16 via 10.0.1.2
10.3.0.0/16 via 10.0.1.2
10.4.0.0/16 via 10.0.1.2
OR:
10.0.0.0/8 via 10.0.1.2

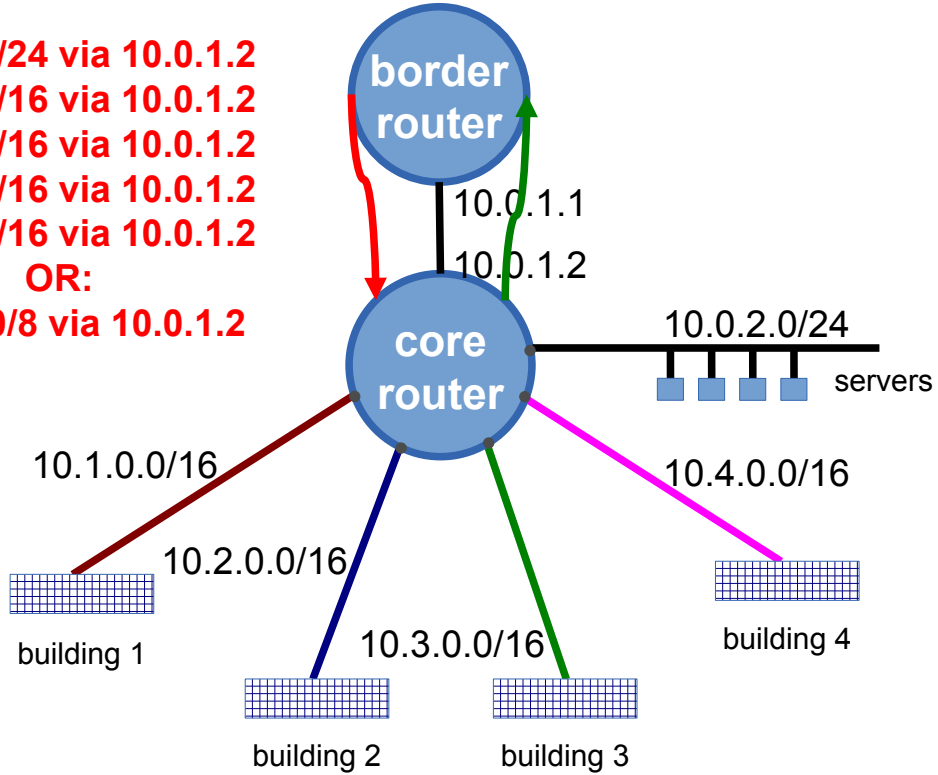

Don’t forget (static) routes

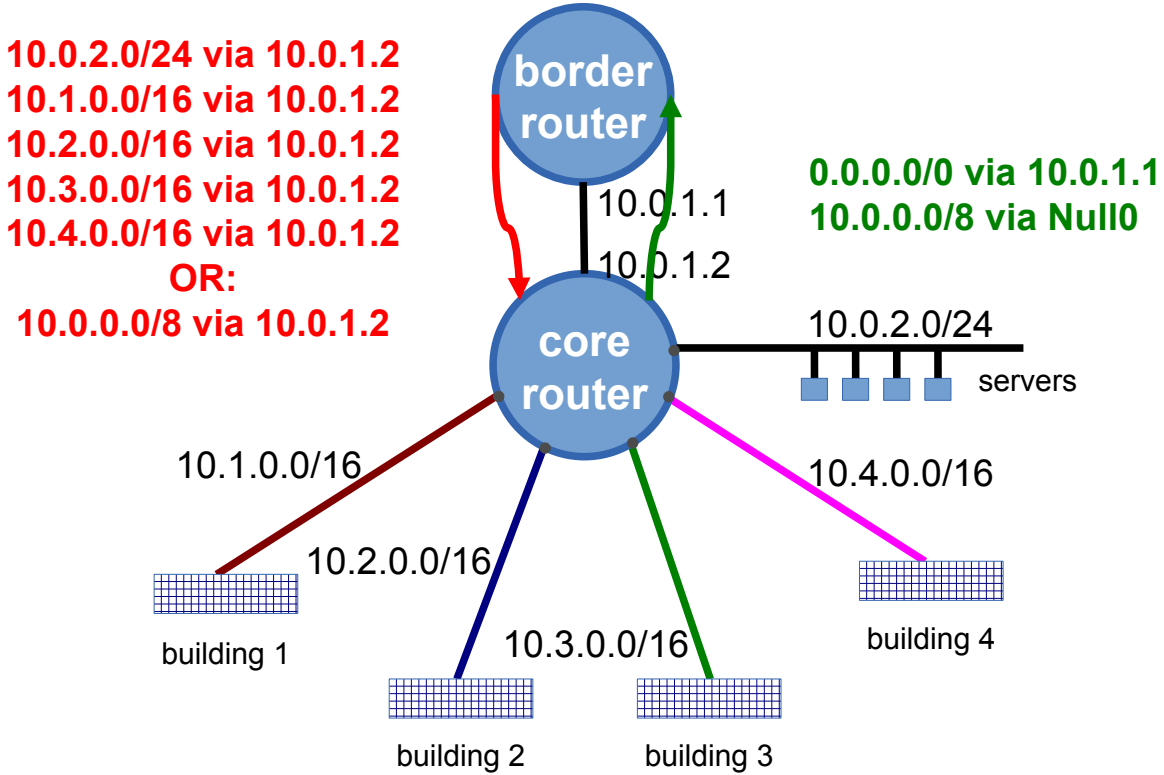

Questions?