Migrating a Campus Network: Flat to Routed

Brian Candler
Network Startup Resource Center
brian@nsrc.org

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

Ideal routed campus network

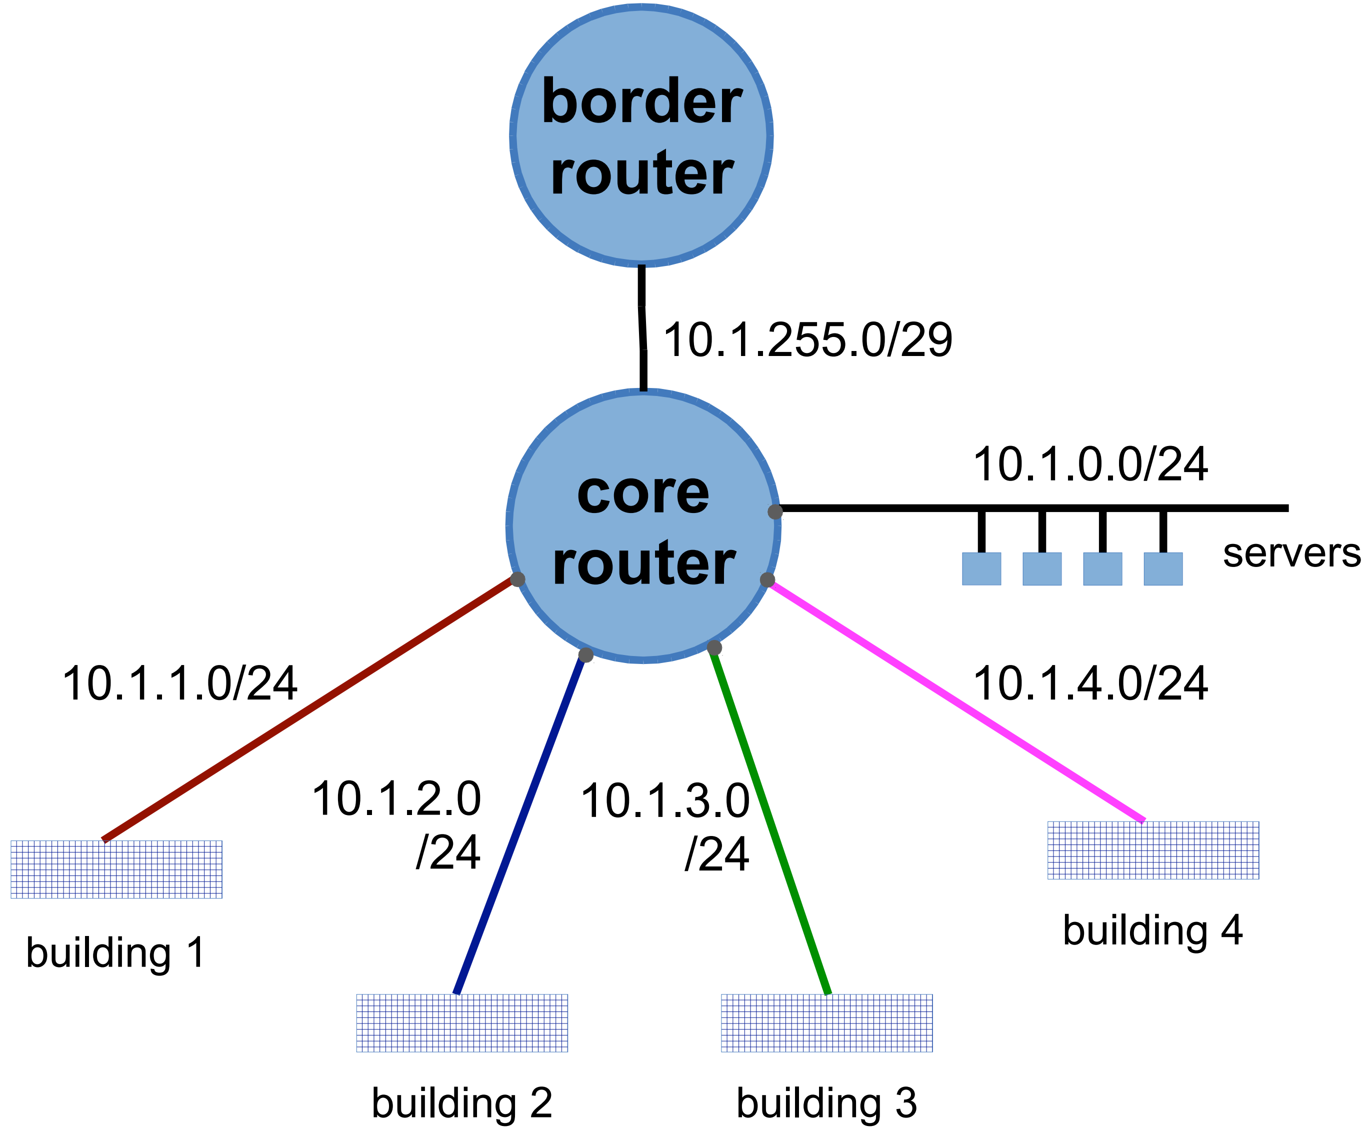

Changing from flat network implies:

- 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

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

Lease time

- It's a good idea to reduce the 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

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 availability information

DHCP broadcasts

- You need to respond to the DHCP Discover *broadcasts* on every subnet
- Option 1: run DHCP service on the router itself
  - can be awkward to manage if you have a lot of custom options or static MAC address mappings
- Option 2: use a feature on the router called "DHCP relay" or "DHCP helper"
  - relays requests to one or more DHCP servers

DHCP relay

⊗ DHCP relay agent

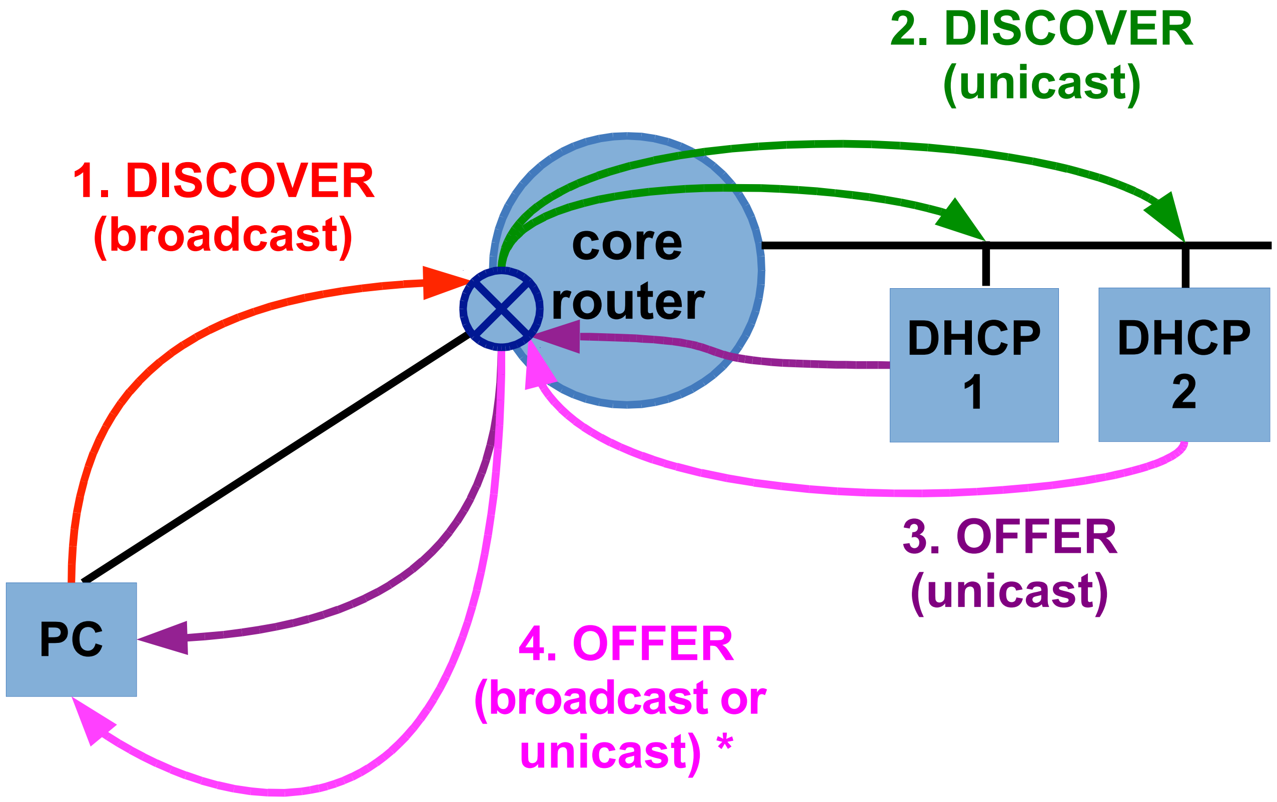

* Client can request broadcast response using the B flag

DHCP relay configuration

- Repeat for every interface where DHCP service required

```
interface Vlan100
  ip address 10.1.1.1 255.255.255.0
  ip helper-address 10.1.0.4
  ip helper-address 10.1.0.5
```

DHCP server configuration

- Define each subnet where service is required
  - (Windows DHCP server: "DHCP scope")

```
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;
}
```

Questions?

Planning Migration

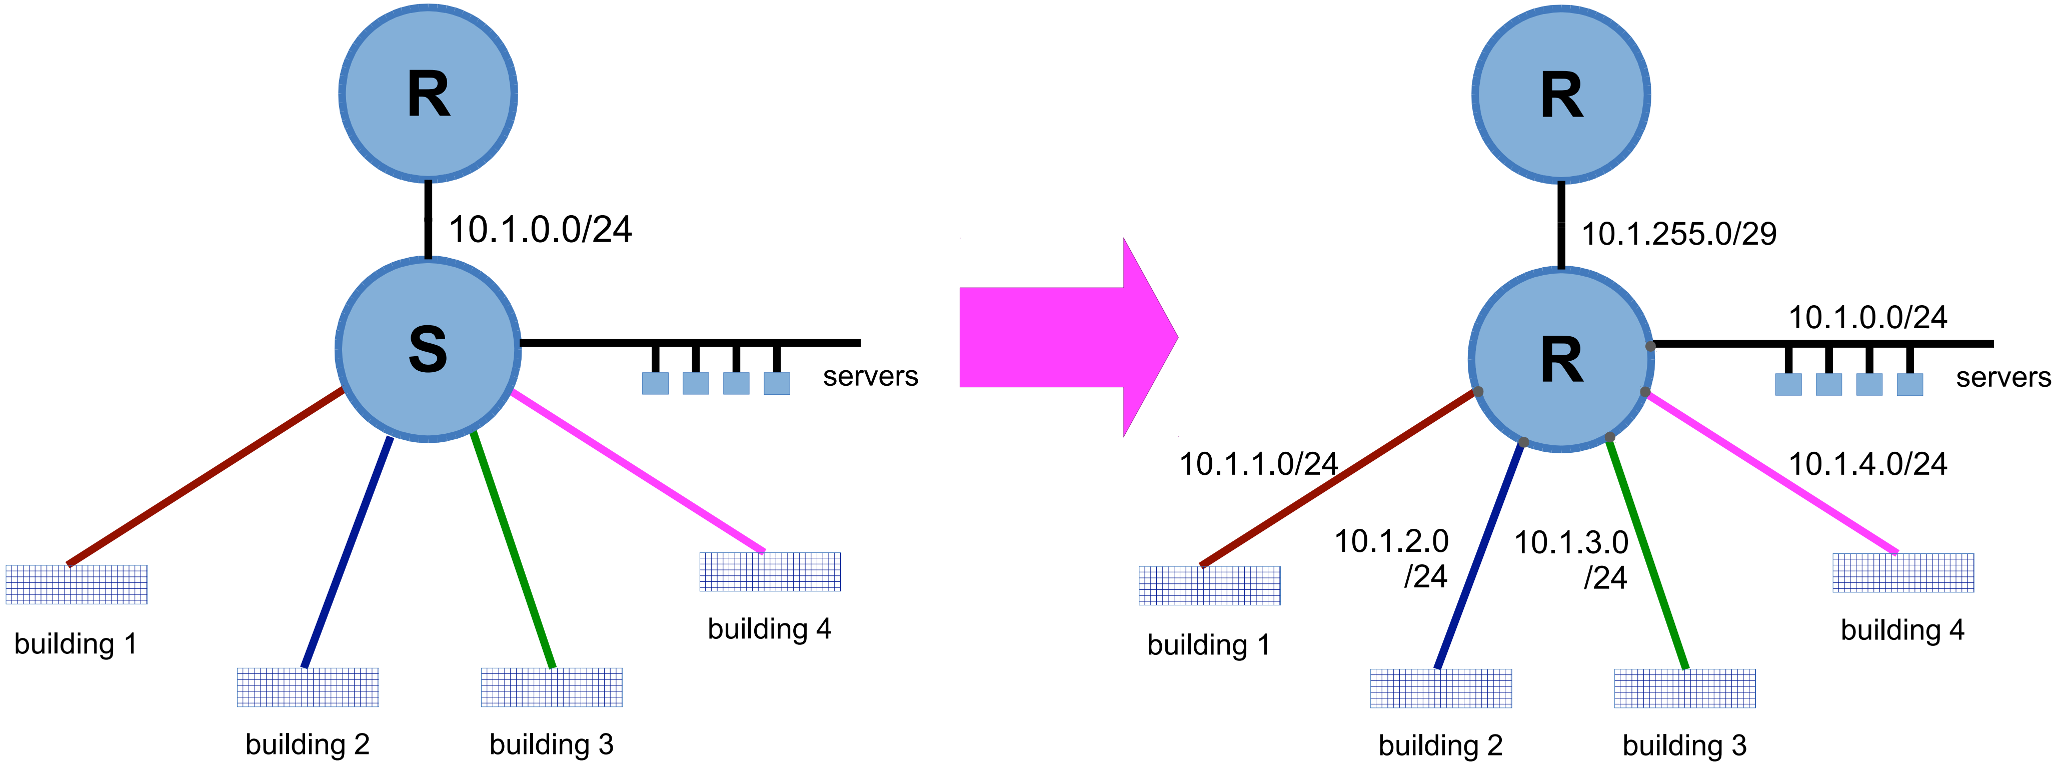

General principles

- 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 config 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?

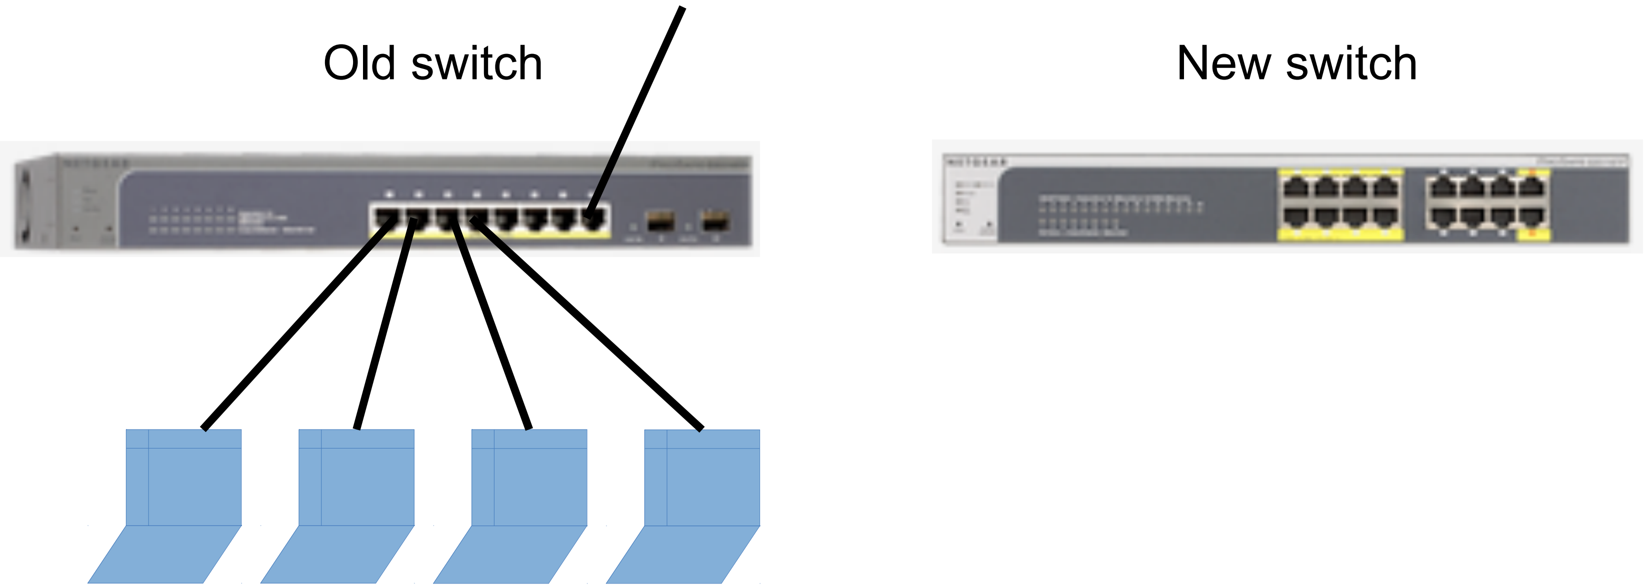

For discussion!

Longer example

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

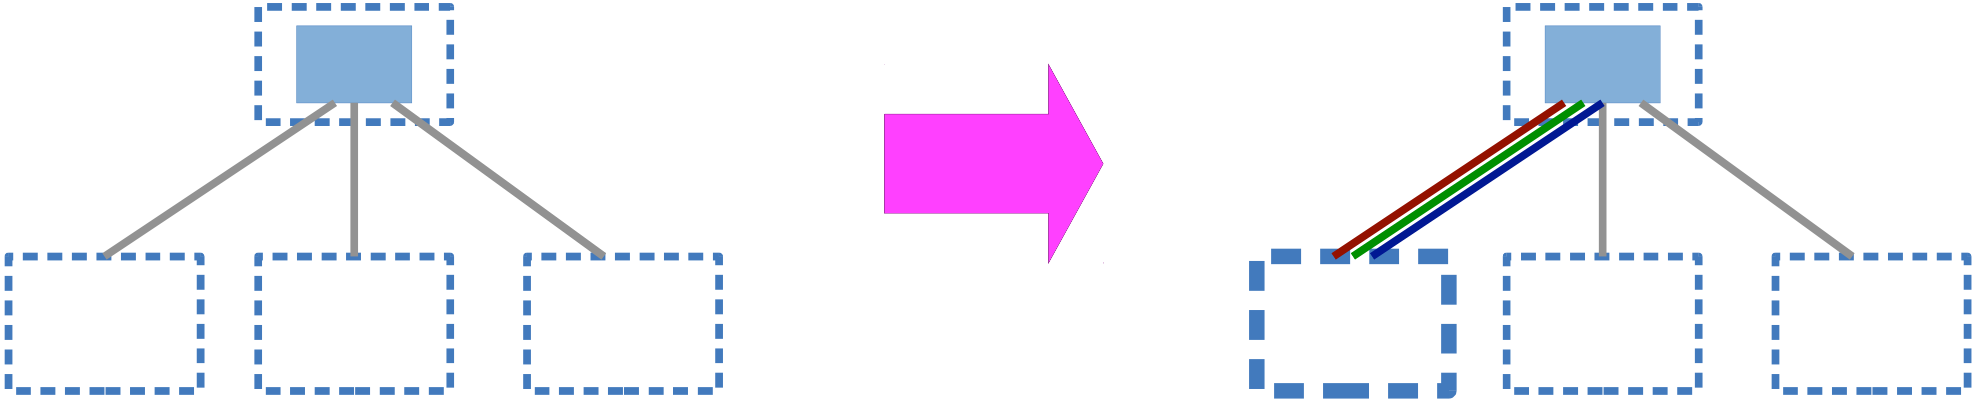

Before (detail)

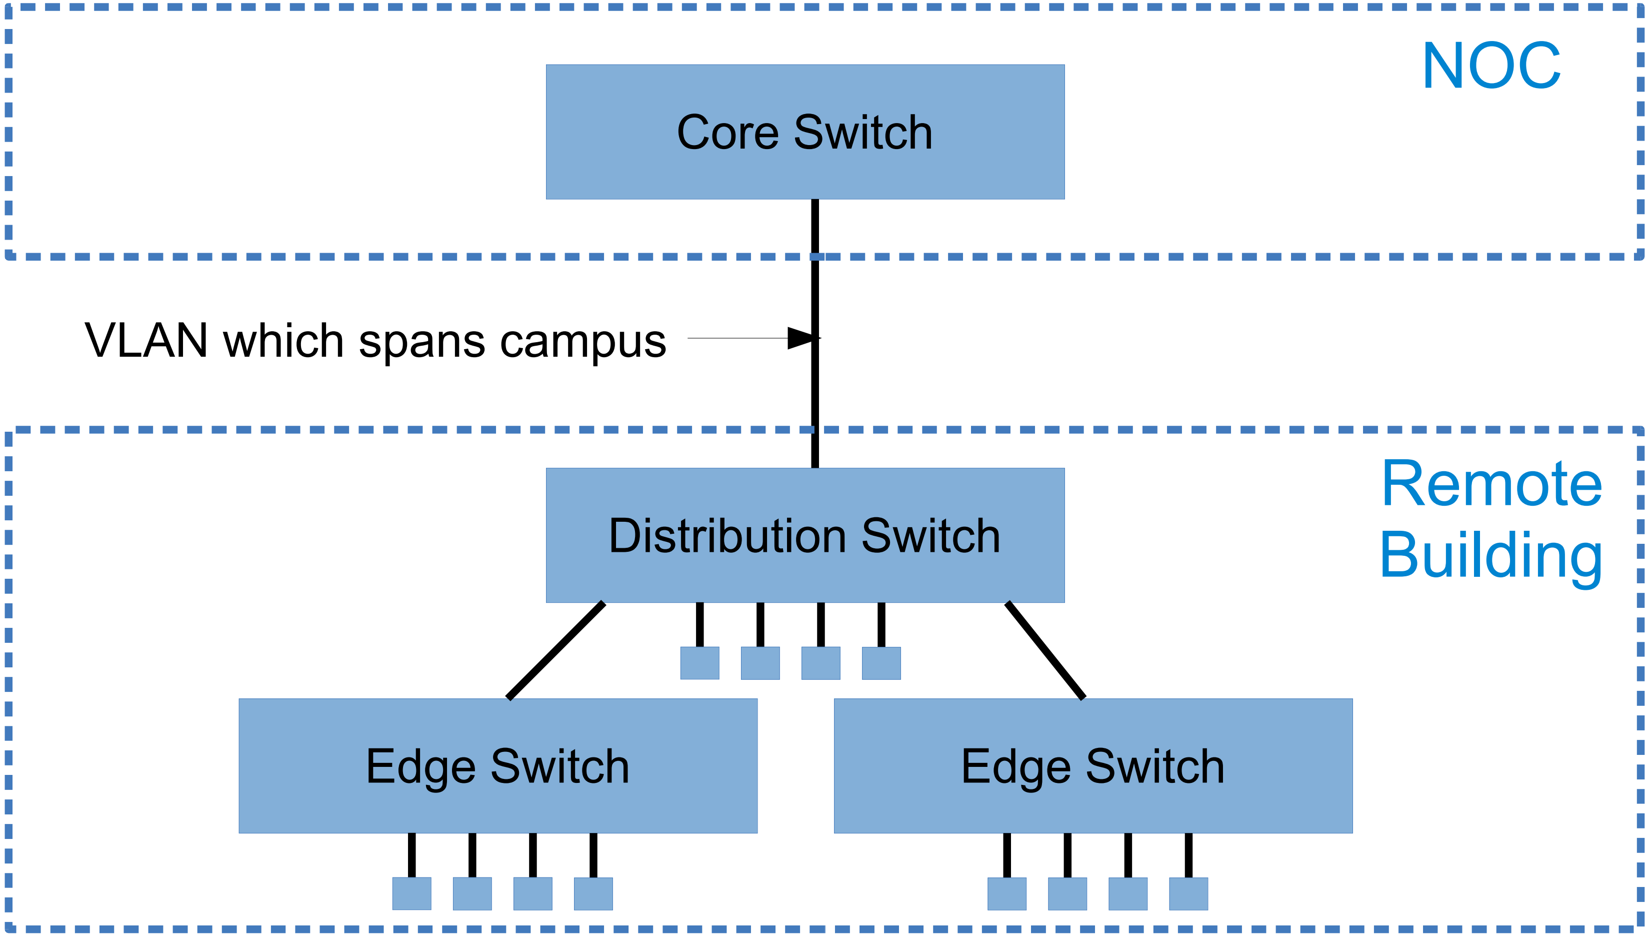

1. Create new VLANs in core

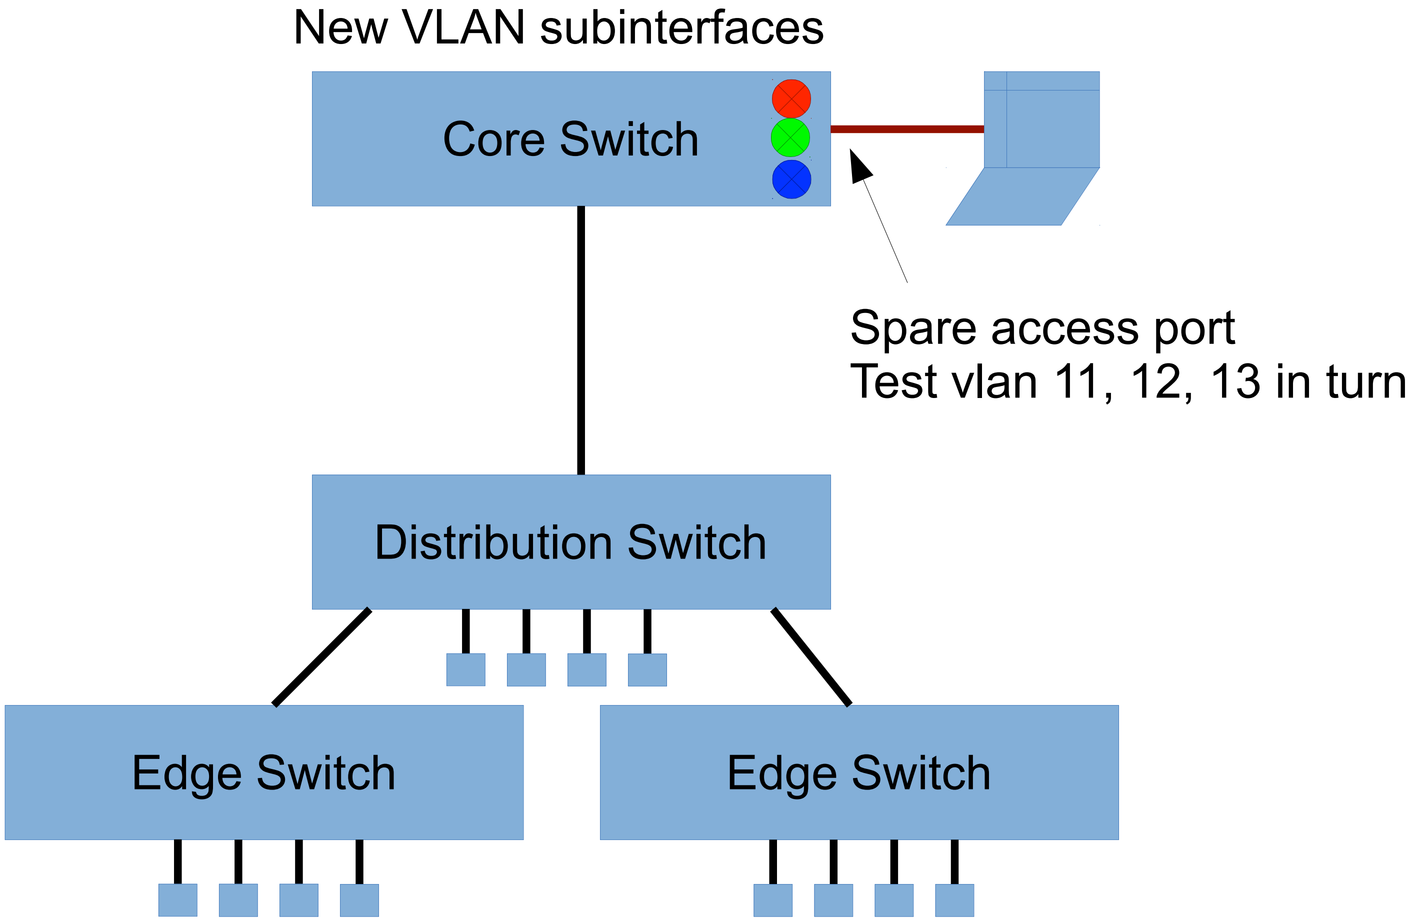

- Test all client functionality, e.g. DHCP, routing

Rollback plan

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

2. Add new VLANs to trunk

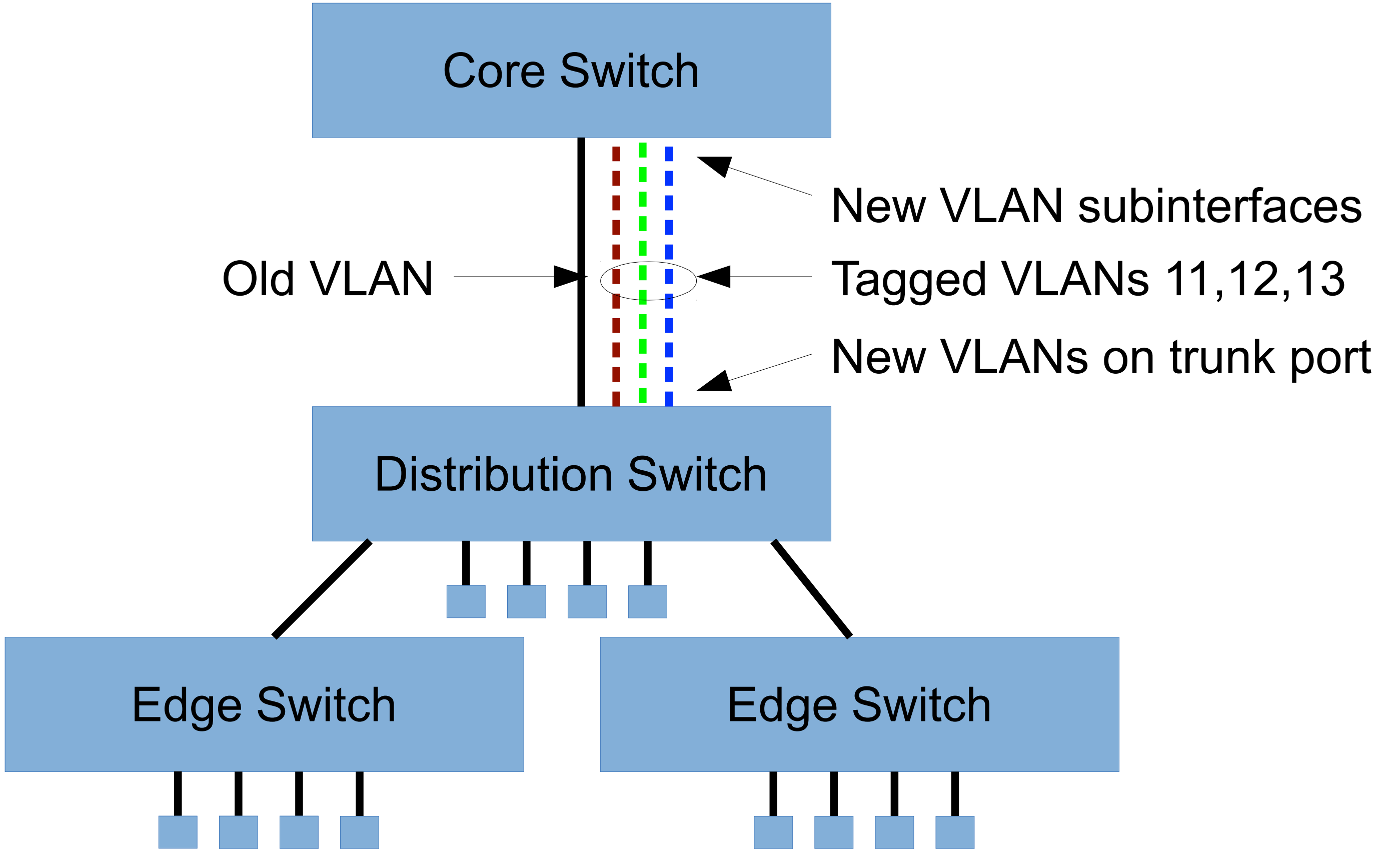

- Should not break anything! (But check anyway)

Choice to make

- Run the old VLAN untagged, 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

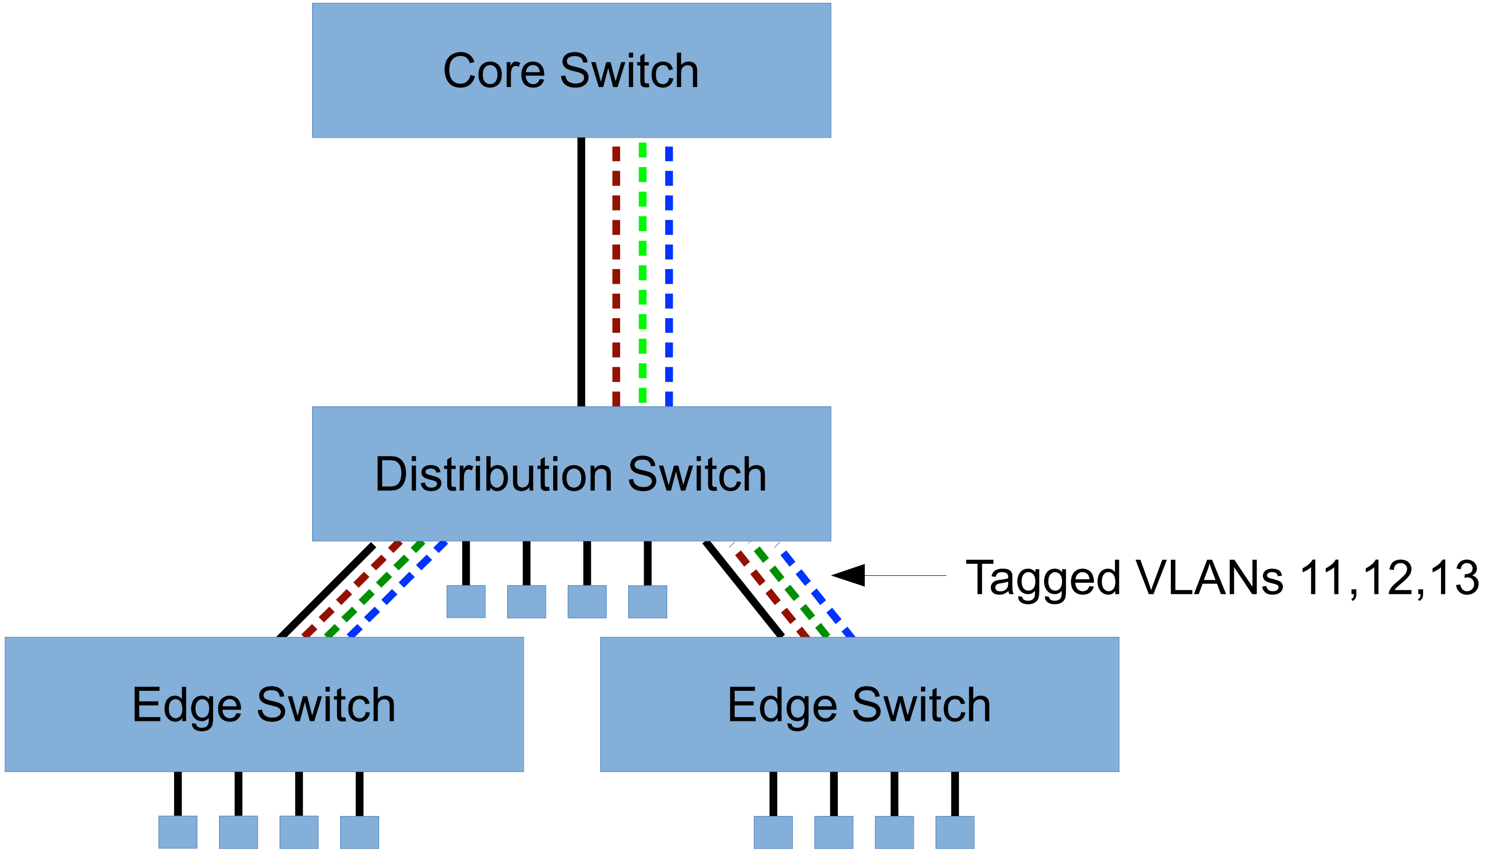

- Again, nothing should break

4. Test with spare access ports

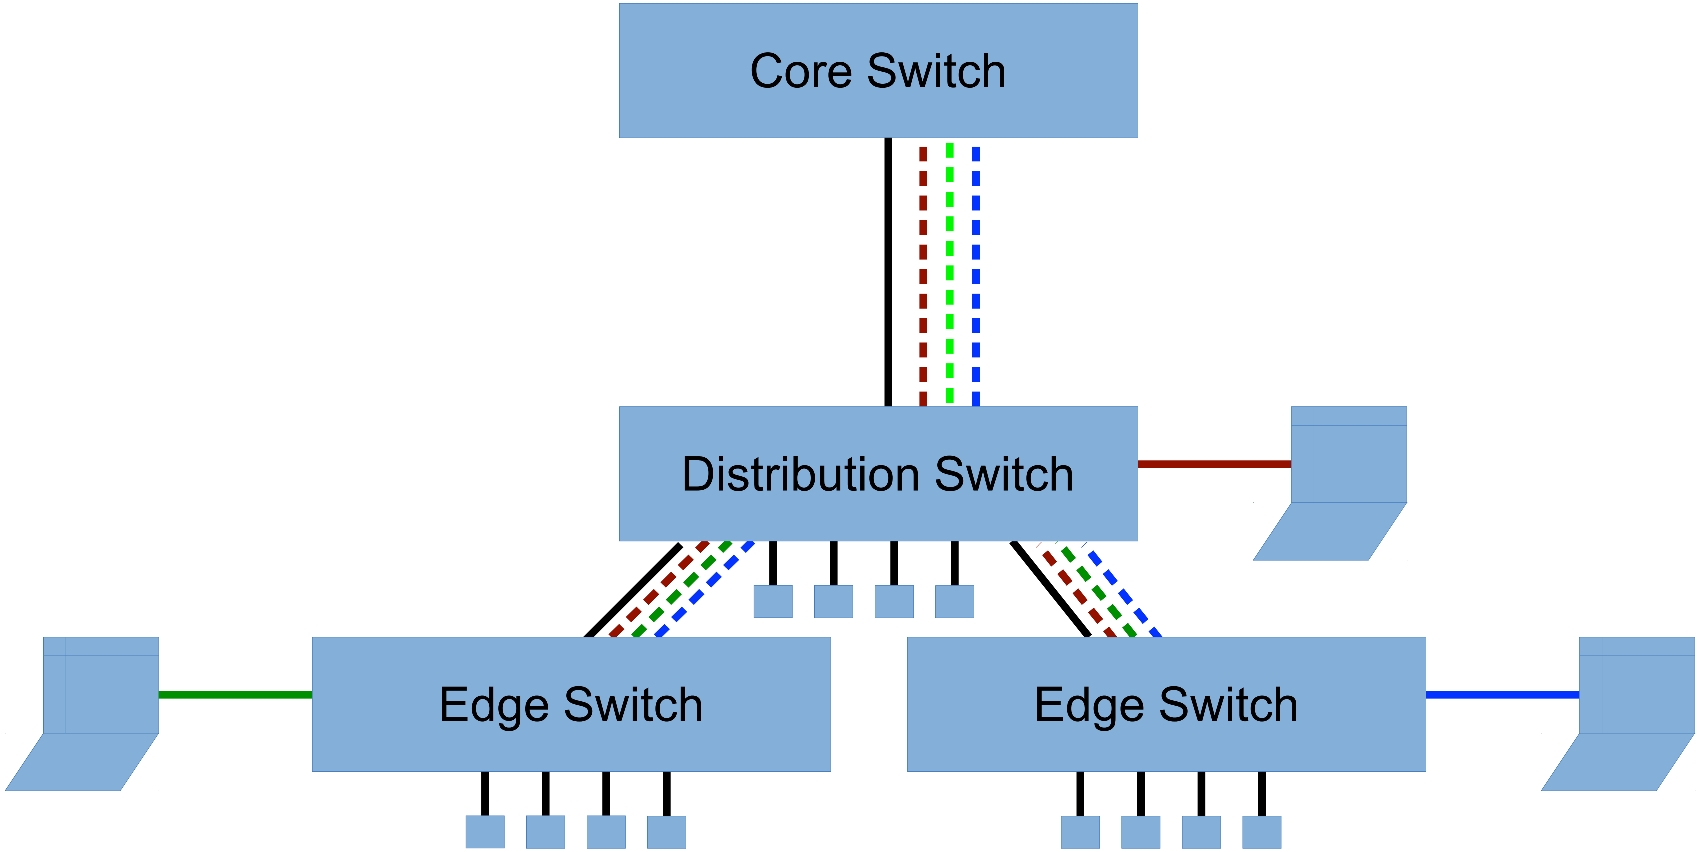

- Re-test all client functionality, DHCP, routing

5. Re-assign edge ports individually

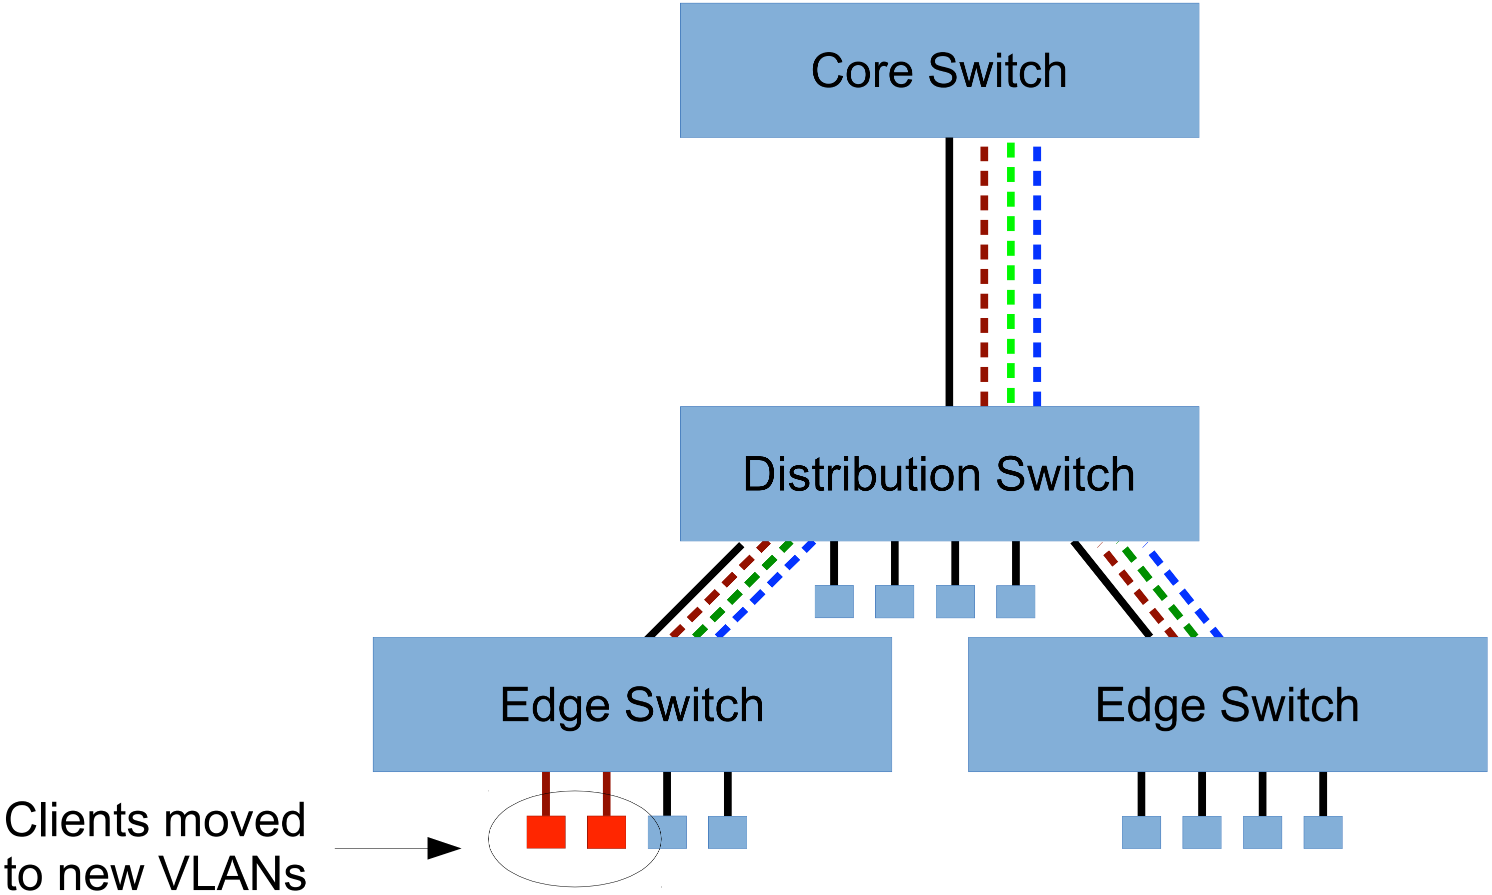

- Controlled interruption to service

6. Move all the remaining clients

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

- 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 VLAN

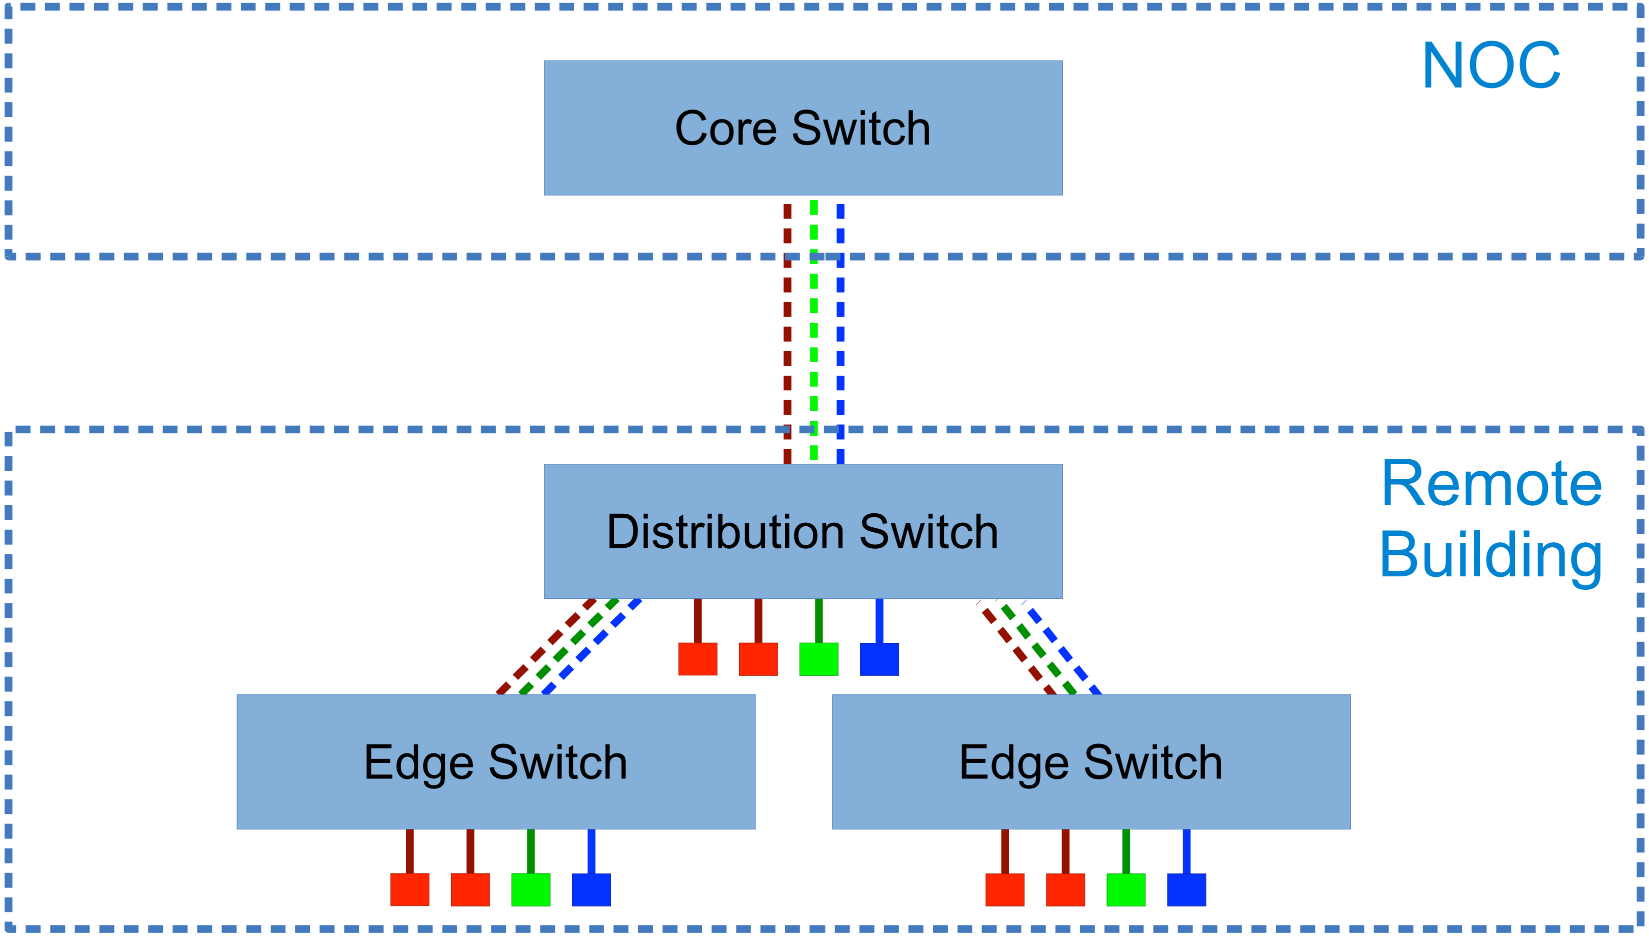

- Final test to sign-off

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 can give you extra confidence nothing has broken

Plan within your constraints

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

Other hints and tips

- If your core switch has only SFP ports, a copper gigabit SFP is useful for testing
- 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 ip arp-cache" may be required
- "write mem" 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 IPs" can be useful when renumbering servers on the same VLAN
  - both old and new IPs active at the same time

Don't forget (static) routes

10.1.0.0/24 via 10.1.255.2
10.1.1.0/24 via 10.1.255.2
10.1.2.0/24 via 10.1.255.2
10.1.3.0/24 via 10.1.255.2
10.1.4.0/24 via 10.1.255.2
OR:
10.1.0.0/16 via 10.1.255.2

0.0.0.0/0 via 10.1.255.1
10.1.0.0/16 via Null0 (why?)

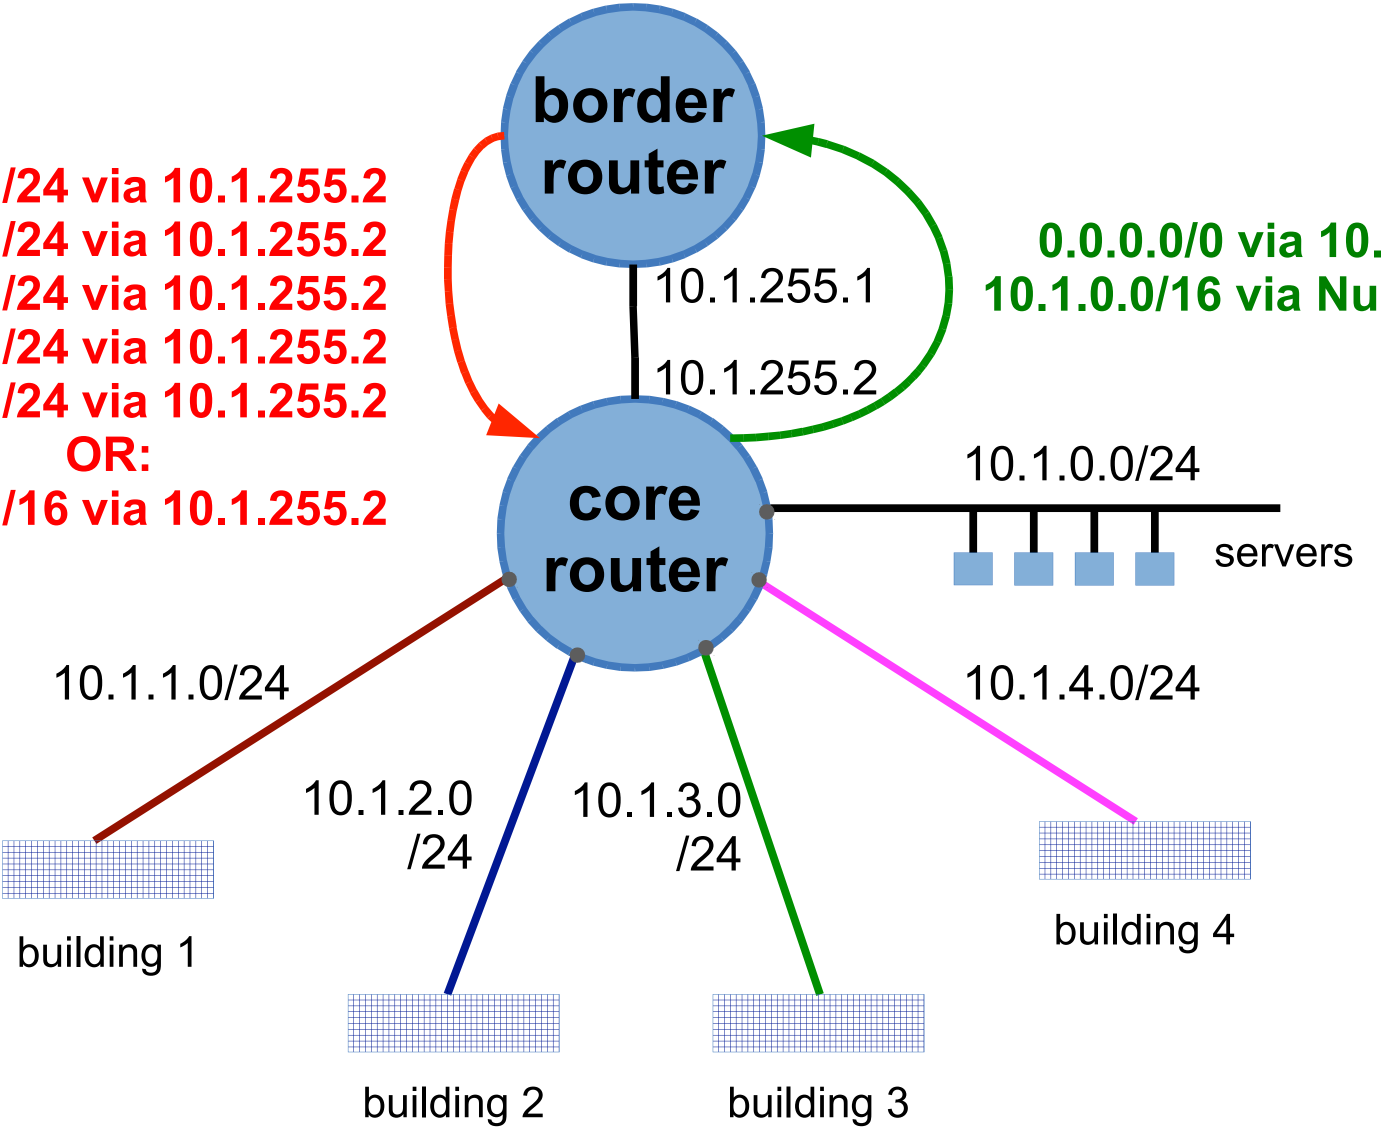

The End!