ISP Systems Design

ISP Workshops

Agenda

- DNS Server placement
- Mail Server placement
- News Server placement
- Services network design
- Services Network Security

ISP Services

- Most ISP services such as DNS, Mail, etc are easily deliverable on low budget hardware platforms
  - Single Rack Unit in height (1RU)
  - Dual processor is “default” now
  - RAM is very cheap (may as well use 2Gbytes or more)
  - Hard drives are very cheap (SCSI more reliable)
  - Unix like operating systems (FreeBSD, Debian, Ubuntu, CentOS) are very common
    - In addition to commercial operating systems such as Solaris, RedHat Enterprise Linux &c
  - Minimal overhead, minimal OS install, plus the service required

ISP Services: DNS

- Domain Name System
  - Provides name and address resolution
  - Servers need to be differentiated, properly located and specified
    - Primary nameserver
    - Secondary nameserver
    - Caching nameserver – resolver

ISP Services: DNS

- Primary nameserver
  - Holds ISP zone files
    - Forward zone (list of name to address mappings) for all ISP's and any customer zones
    - Reverse zone (list of address to name mappings) for all ISP's address space
  - Hardware & OS: easily satisfied by simple specification
  - Located in secure part of net, e.g. NOC LAN
  - Usually run as "hidden master" – secondary nameservers are the official listed nameservers

ISP Services: DNS

- Secondary nameserver
  - Holds copies of ISP zone files
  - At least two are required, more is better
  - Hardware & OS: easily satisfied by simple specification
  - Strongly recommended to be geographically separate from each other and the primary DNS
    - At different PoPs
    - On a different continent e.g. via services offered by ISC, PCH and others
    - At another ISP

ISP Services: Secondary DNS Example

```
$ dig apnic.net ns
;; ANSWER SECTION:
apnic.net.              10800   NS      ns1.apnic.net.
apnic.net.              10800   NS      ns3.apnic.net.
apnic.net.              10800   NS      ns4.apnic.net.
apnic.net.              10800   NS      ns5.apnic.com.
apnic.net.              10800   NS      cumin.apnic.net.
apnic.net.              10800   NS      ns-sec.ripe.net.
apnic.net.              10800   NS      tinnie.arin.net.
apnic.net.              10800   NS      tinnie.apnic.net.

;; ADDITIONAL SECTION:
ns1.apnic.net.          3600    A       202.12.29.25
ns3.apnic.net.          3600    A       202.12.28.131
ns4.apnic.net.          3600    A       202.12.31.140
ns5.apnic.com.          10800   A       203.119.43.200
cumin.apnic.net.        3600    A       202.12.29.59
tinnie.apnic.net.       3600    A       202.12.29.60
ns-sec.ripe.net.        113685  A       193.0.0.196
tinnie.arin.net.        10800   A       199.212.0.53
```

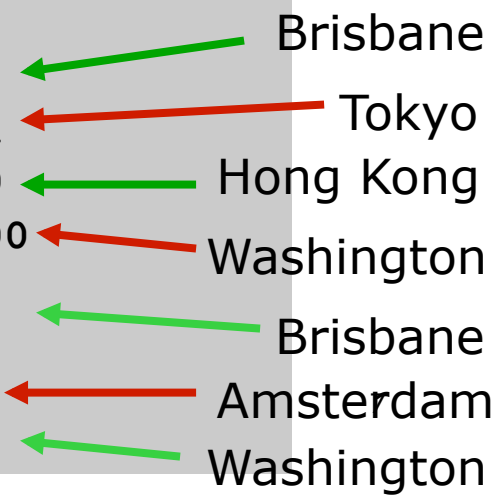

ISP Services: Secondary DNS Example

- apnic.net zone
  - Primary DNS in Brisbane (ns1.apnic.net)
  - Secondary DNS run all over the world by APNIC:
    - Brisbane
    - Hong Kong
    - Tokyo
    - Washington
  - Zone secondaried by
    - RIPE NCC in Amsterdam
    - ARIN in Washington
  - Geographical and service provider redundancy – this is the perfect example!

ISP Services: DNS

- Caching nameserver
  - This is the resolver – it is the DNS cache
  - Your customers use this as resolver, NEVER your primary or secondary DNS
  - Provides very fast lookups
  - Does NOT secondary any zones
  - One, or preferably two per PoP (redundancy)
  - Hardware & OS: easily satisfied by simple specification

ISP Services: Caching Nameserver

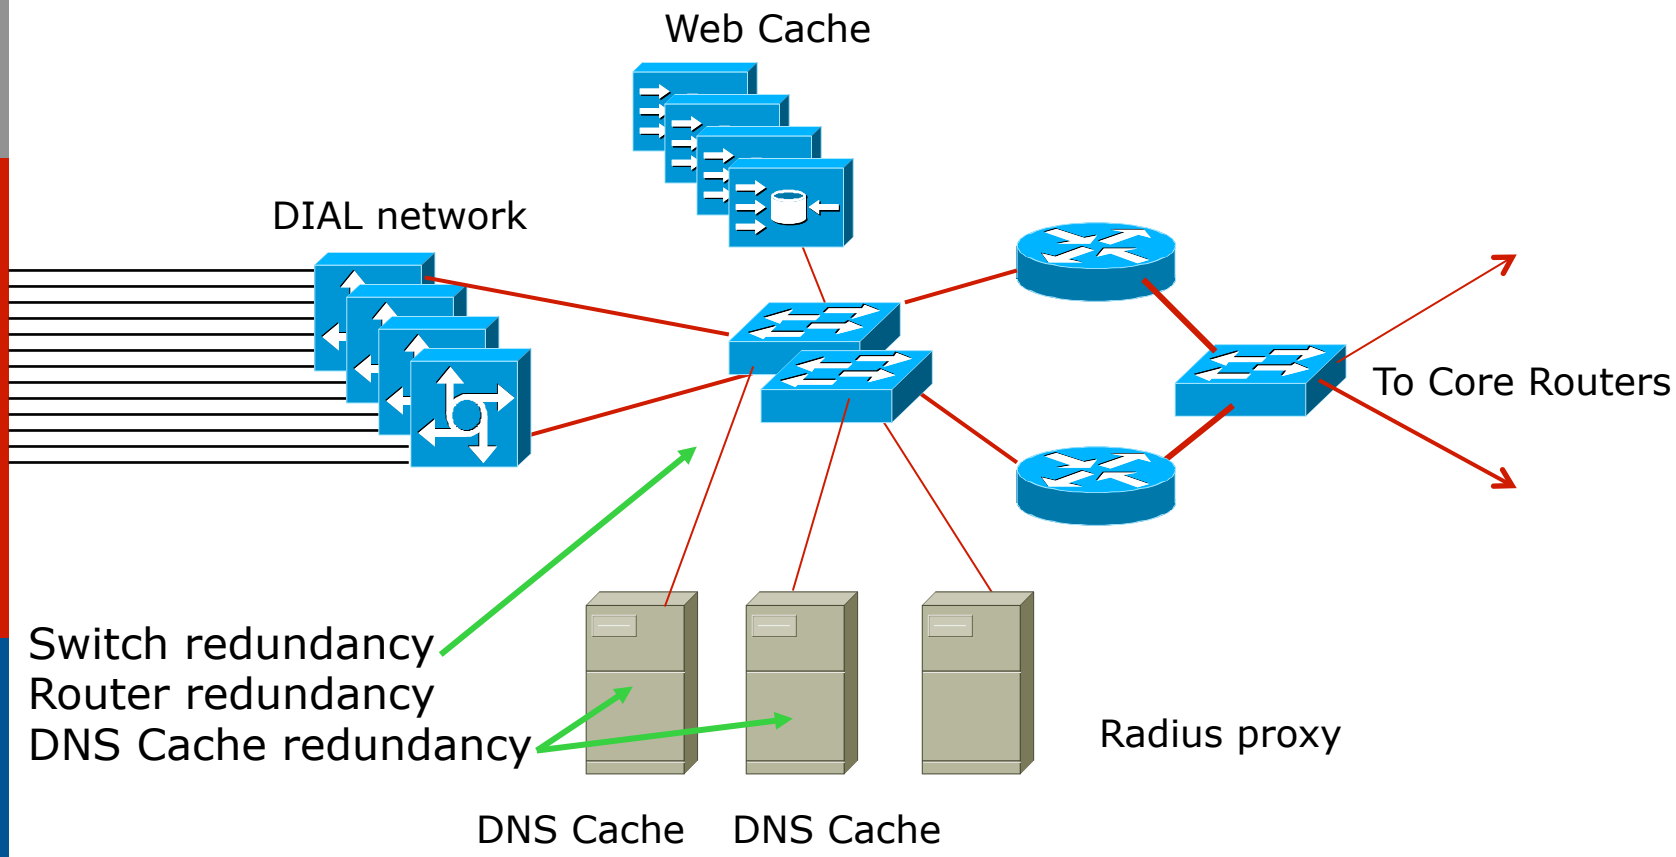

- DIAL users automatically given the IP addresses of DNS caches when they dial in

ISP Services: Anycasting the Caching Nameserver

- One trick of the trade
  - assign two unique IP addresses to be used for the two DNS resolver systems
  - use these two IP addresses in every PoP
  - route the two /32s across your backbone
  - even if the two resolver systems in the local PoP are down, the IGP will ensure that the next nearest resolvers will be reachable
  - Known as IP Anycast

ISP Services: DNS

- Efficient and resilient design
  - Primary DNS – keep it secure
  - Secondary DNS – geographical and provider redundancy
    - Don't ever put them on the same LAN, switched or otherwise
    - Don't put them in the same PoP
  - Caching DNS – one or two per PoP
  - Reduces DNS traffic across backbone
  - More efficient, spreads the load

ISP Services: DNS

- Software
  - Make sure that the BIND distribution on the Unix system is up to date
    - The vendor's distribution is rarely current
  - Pay attention to bug reports, security issues
  - Reboot the DNS cache on a regular (e.g. monthly) basis
    - Clears out the cache
    - Releases any lost RAM
    - Accepted good practice by system administrators

ISP Services: DNS

- Implementation
  - Put all your hosts, point-to-point links and loopbacks into the DNS
    - Under your ISP's domain name
    - Use sensible/meaningful names
  - Put all your hosts, point-to-point links and loopbacks into the REVERSE DNS also
    - Don't forget about in-addr.arpa and ip6.arpa – many ISPs do
    - Some systems demand forward/reverse DNS mapping before allowing access

ISP Services: Mail

- Must have at least two mail hosts (MX records) for all supported domains
  - Geographical separation helps
- Dedicated POP3 server
  - Consumers/mobile users get mail from here
- SMTP gateway dedicated to that function
  - Consumers/mobile users send mail via here
- Mail relay open to CUSTOMERS only!
  - Don’t let outside world use your mail relay
- Block port 25 outbound for all customers
  - Insist that outbound e-mail goes through SMTP relay
  - SMTP relay does virus (ClamAV) and spam (Spamassassin) filtering

ISP Services: Mail Configuration

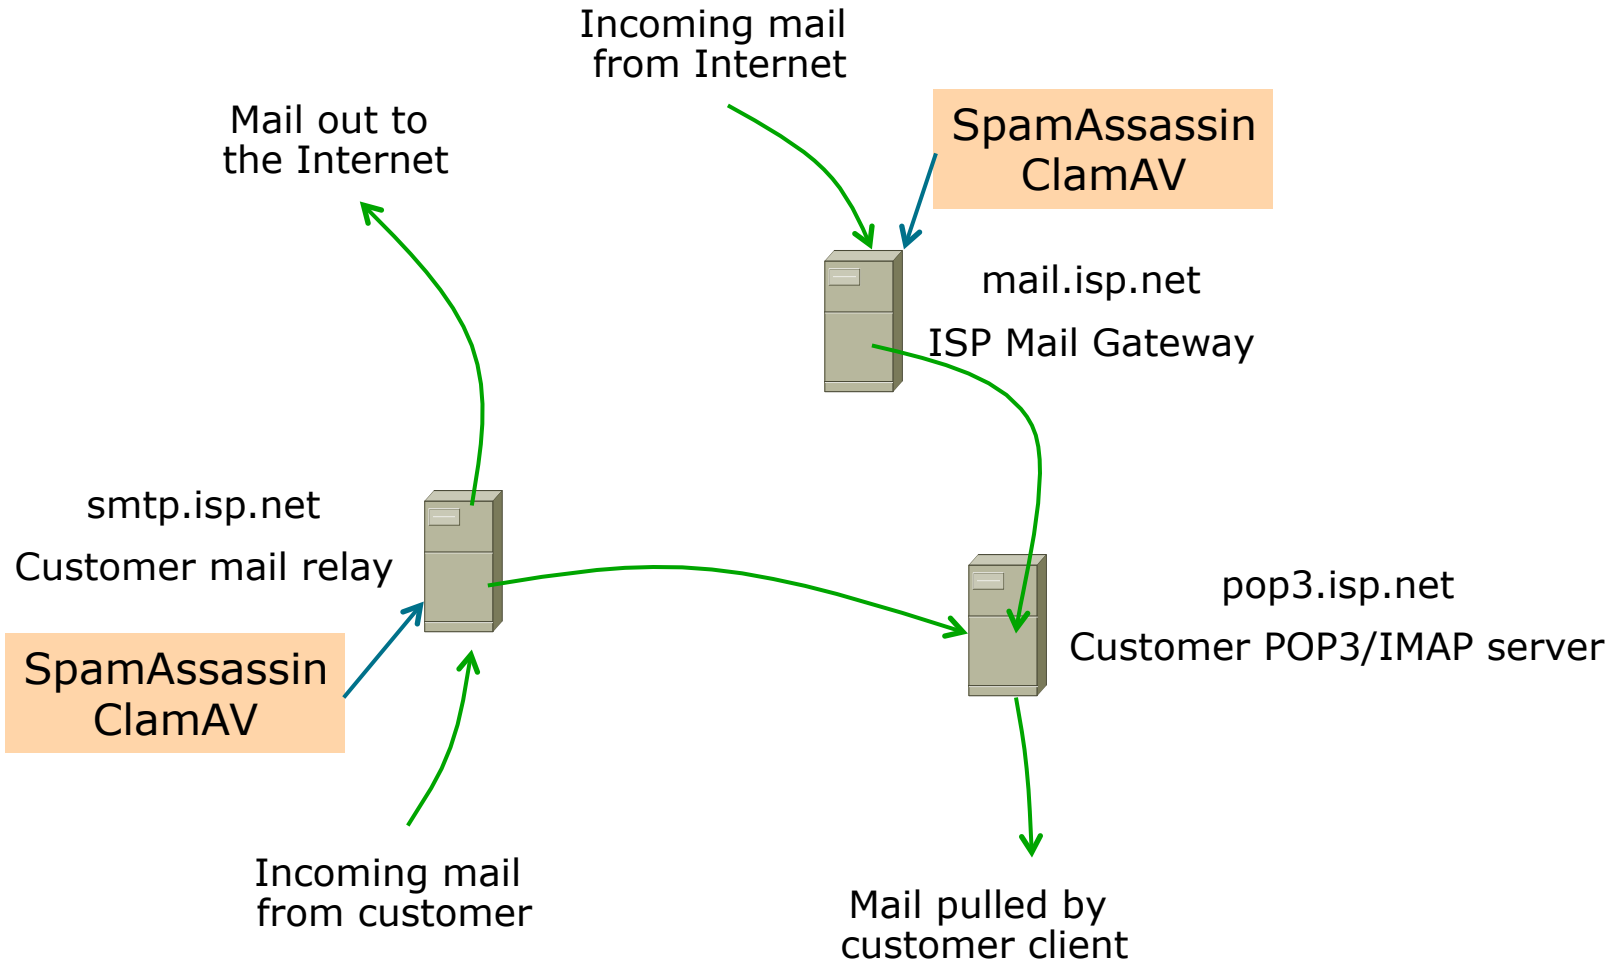

ISP Services: Mail Example

- cisco.com mail (MX records)
  - primary MX are 6 systems in San Jose
  - Three backup MXes in RTP, Amsterdam and Sydney
  - backup MX only used if primary unavailable

```
$ dig cisco.com mx

;; ANSWER SECTION:
cisco.com.          86400   MX      10 sj-inbound-a.cisco.com.
cisco.com.          86400   MX      10 sj-inbound-b.cisco.com.
cisco.com.          86400   MX      10 sj-inbound-c.cisco.com.
cisco.com.          86400   MX      10 sj-inbound-d.cisco.com.
cisco.com.          86400   MX      10 sj-inbound-e.cisco.com.
cisco.com.          86400   MX      10 sj-inbound-f.cisco.com.
cisco.com.          86400   MX      15 rtp-mx-01.cisco.com.
cisco.com.          86400   MX      20 ams-inbound-a.cisco.com.
cisco.com.          86400   MX      25 syd-inbound-a.cisco.com.
```

ISP Services: Mail

- Software
  - Make sure that the MAIL and POP3 distributions on the Unix system are up to date
    - The vendor distributions are rarely current
  - Pay attention to bug reports, security issues, unsolicited junk mail complaints

IMPORTANT: Do NOT allow non-customers to use your mail system as a relay

ISP Services: News

- News servers provide a Usenet news feed to customers
- Distributed design required
  - Incoming newsfeed to one large server
  - Distributed to feed servers in each PoP
  - Feed servers provide news feed to customers
  - Outgoing news goes to another server
  - Separate reading news system
  - Separate posting news system

ISP Services:
News System Placement

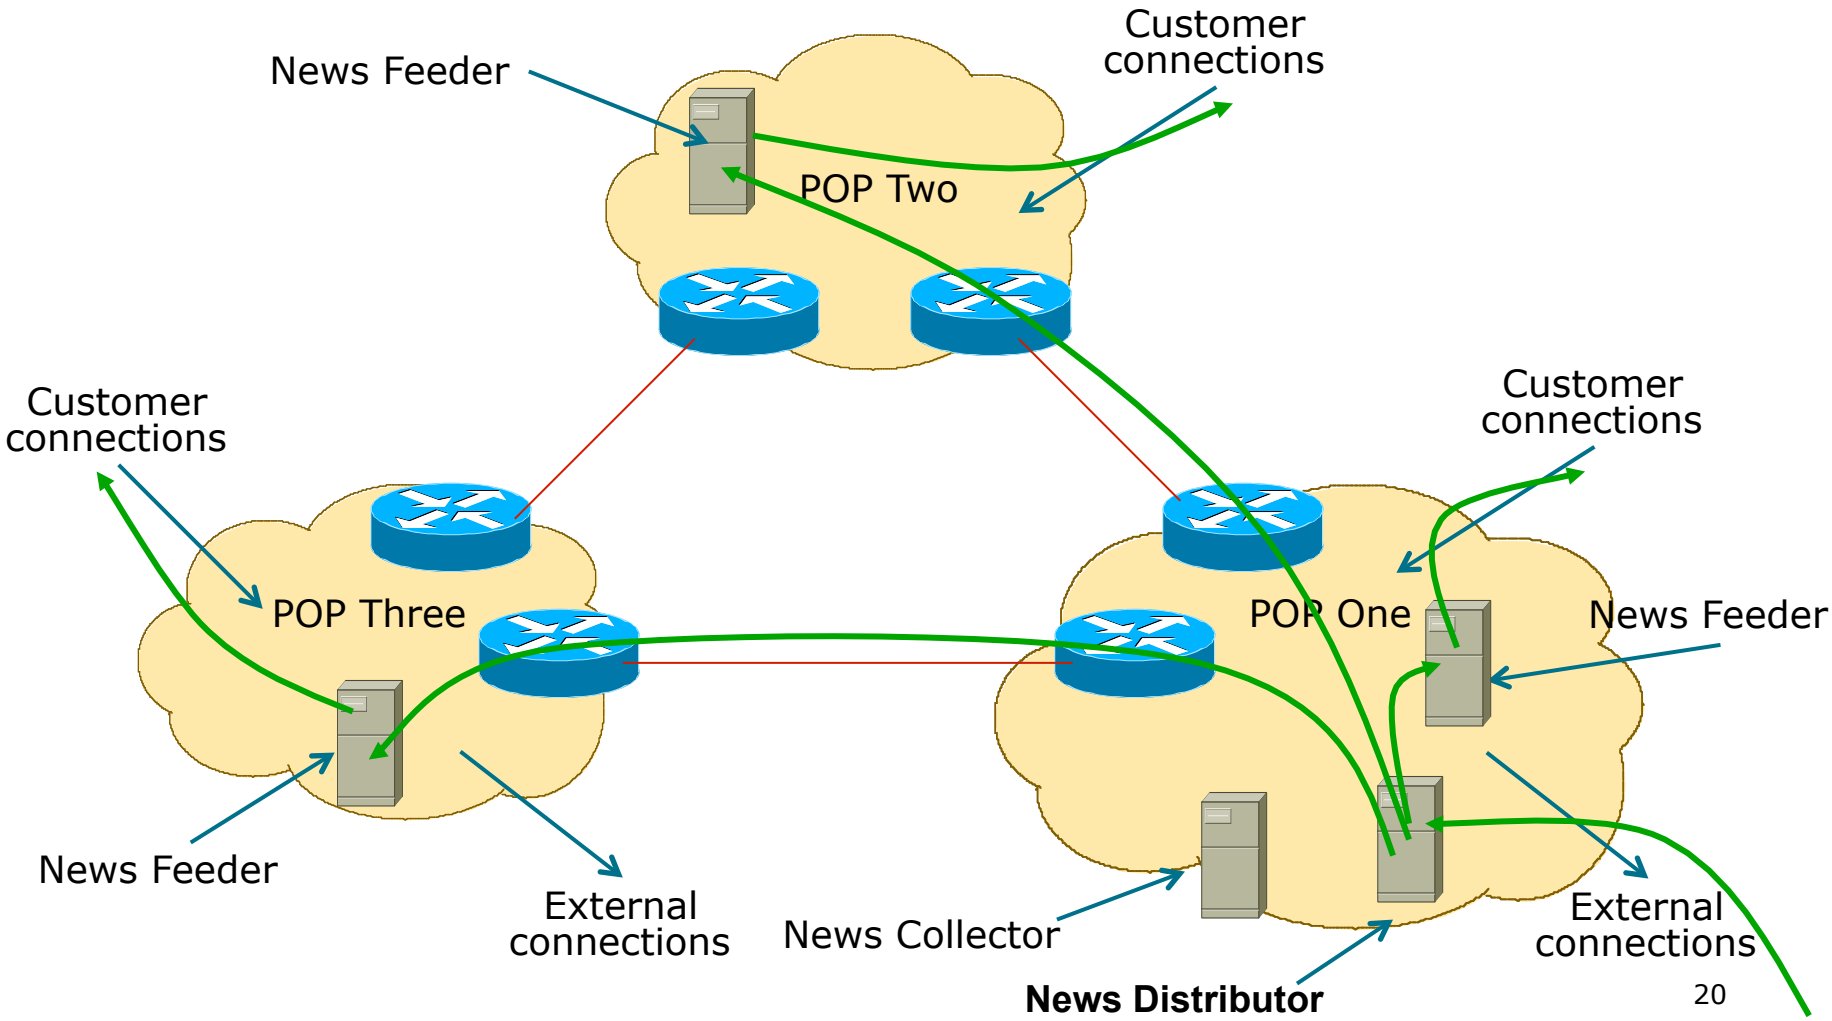

ISP Services:
News System Placement

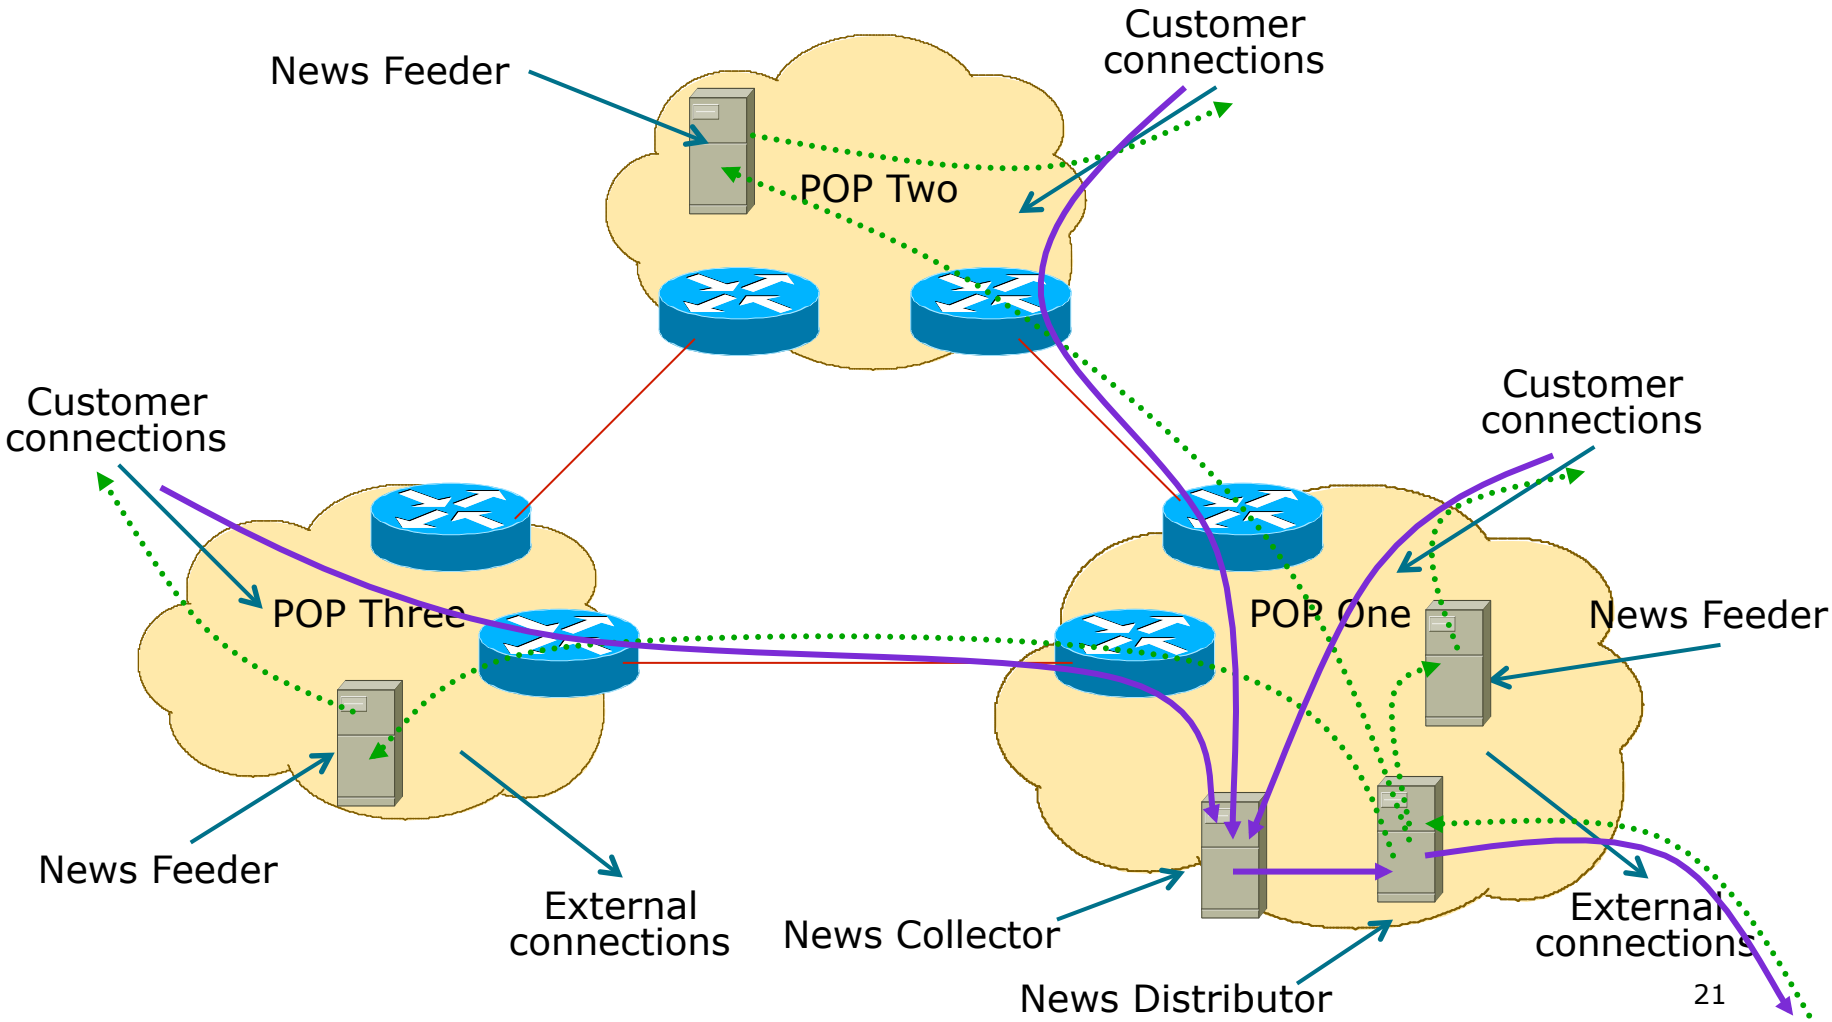

ISP Services: News

- Software
  - Make sure that the Internet News distribution on the Unix system is up to date
    - The vendor distributions are rarely current
  - Pay attention to bug reports, security issues, unsolicited junk posting complaints

IMPORTANT: Do NOT allow non-customers to use your news system for posting messages

Services Security

Security

- ISP Infrastructure security
- ISP Services security
- Security is not optional!
- ISPs need to:
  - Protect themselves
  - Help protect their customers from the Internet
  - Protect the Internet from their customers
- The following slides are general recommendations
  - Do more research on security before deploying any network

ISP Infrastructure Security

- ISP server security
  - Usernames, passwords, TCP wrappers, IPTABLES
  - Protect **all** servers using routers with strong filters applied
- Hosted services security
  - Protect network from hosted servers using routers with strong filters
  - Protect hosted servers from Internet using routers with strong filters

ISP Infrastructure Security
ISP Server Protection

Access-list examples:

Allow tcp/established to all servers
ICMP
DNS 2ary: udp/53 and tcp/53
POP3: tcp/110
Mail Relay: tcp/25 and ISP address range only
News: tcp/119 and ISP address range only
DNS Cache: udp/53
Web server: tcp/80

Other necessary filters:

All servers: SSH (tcp/22) from NOC LAN only

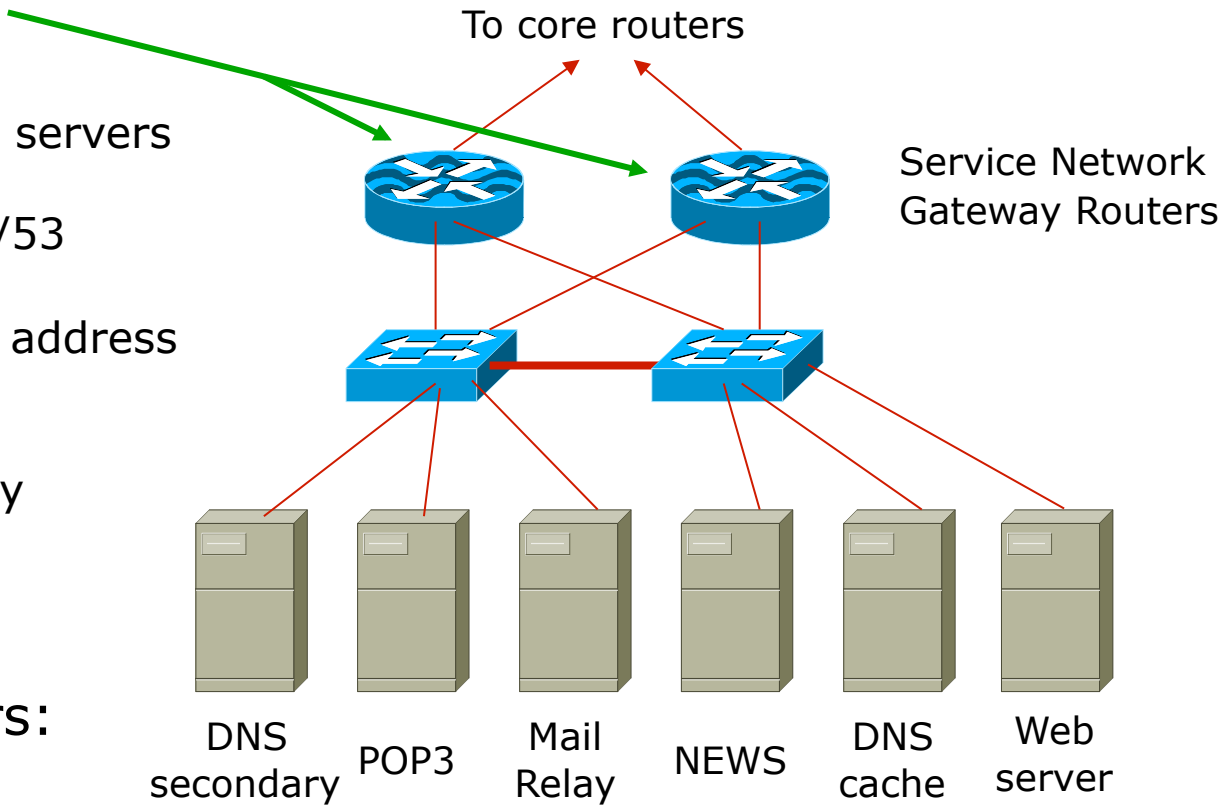

ISP Infrastructure Security
Hosted Server Protection

Access-list examples:

Inbound

Allow tcp/established to all servers
ICMP
Web server: tcp/80
SSH for customer access
Any other ports for services
sold to customers

Outbound

ICMP
Allow DNS udp/53 and
tcp/53
Block all access to ISP
address range

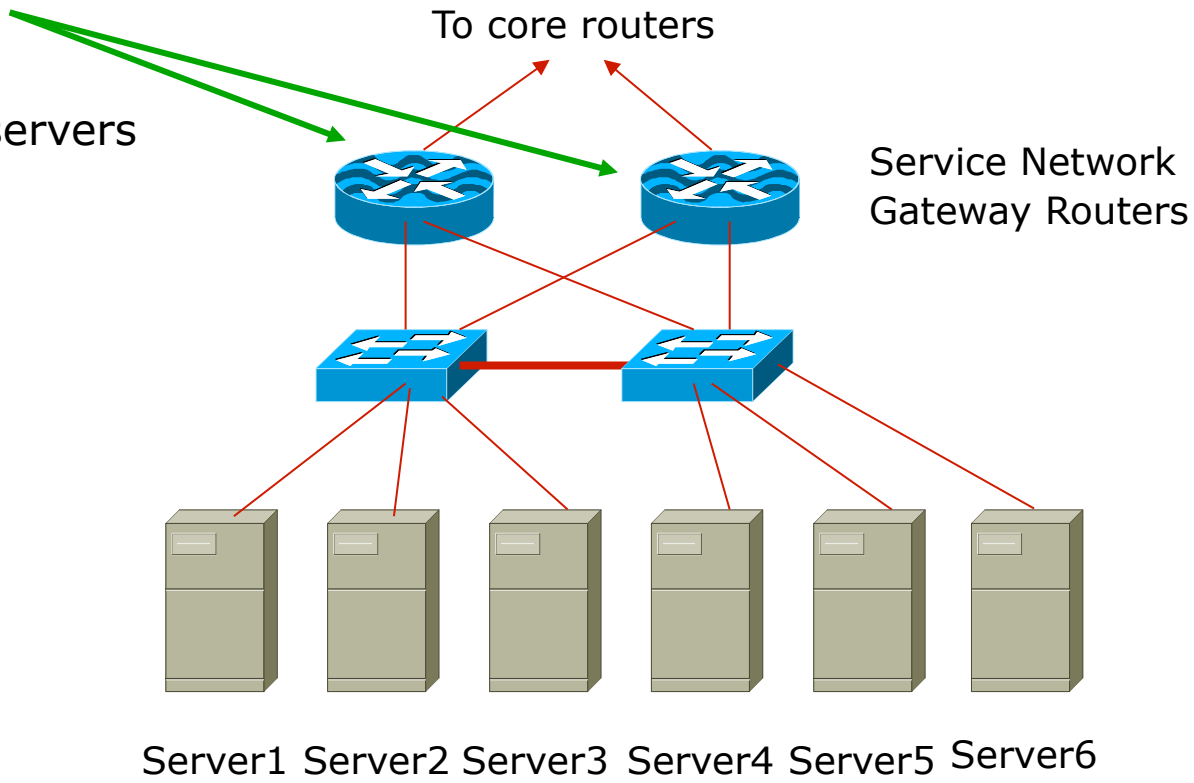

ISP Infrastructure Security

- Premises security
  - Locks – electronic/card key preferred
  - Secure access – 24x7 security arrangements
  - Environment control – good aircon
- Staff responsibility
  - Password policy, strangers, temp staff
  - Employee exit procedures
- RFC2196
  - (Site Security Handbook)
- RFC3871
  - (Operational Security Requirements for Large ISP IP Network Infrastructure)

ISP Network Security
Secure external access

- How to provide staff access from outside
  - Set up ssh gateway (Unix system with ssh daemon and nothing else configured)
  - Provide ssh client on all staff laptops
  - ssh available on Unix and Windows
  - ssh is Secure Shell – encrypted link
- How not to provide access from outside
  - telnet, rsh, rlogin – these are all insecure
  - Open host – insecure, can be compromised

ISP Systems Design

Summary

ISP Design Summary

- KEEP IT SIMPLE & STUPID ! (KISS)
- Simple is elegant is scalable
- Use Redundancy, Security, and Technology to make life easier for yourself
- Above all, ensure quality of service for your customers

ISP Systems Design

ISP Workshops