AARNet's experience with IPv6

Glen Turner

2008-11-19

Australian 2008 IPv6 Summit

aarnet
Australia's Academic and Research Network

Where is AARNet?

- Native IPv6 service to our customers
  - Not-for-profit education and research, health, cultural institutions
- IPv6 broker
  - A best effort service to the greater community, especially developers
- Low deployment by customers
  - Didn't used to matter: by definition research has low initial usage
  - Slowly becoming a strategic issue, and we're trying various approaches to see what will fix that

Network address translaton

aarnet
Australia's Academic and Research Network

The core issue

- IPv6 deployment has failed
  - This summit should be a “wrap party”
  - Failure a result of a vicious circle involving ISPs, customers, vendors plus a notorious historical regulatory failure inhibiting a regulatory response
- So now IPv4 NAT by the ISPs is required for ISPs to provide internet service to new customers

“Carrier-class NAT”

- NAT in the ISP as well as the customer premises equipment

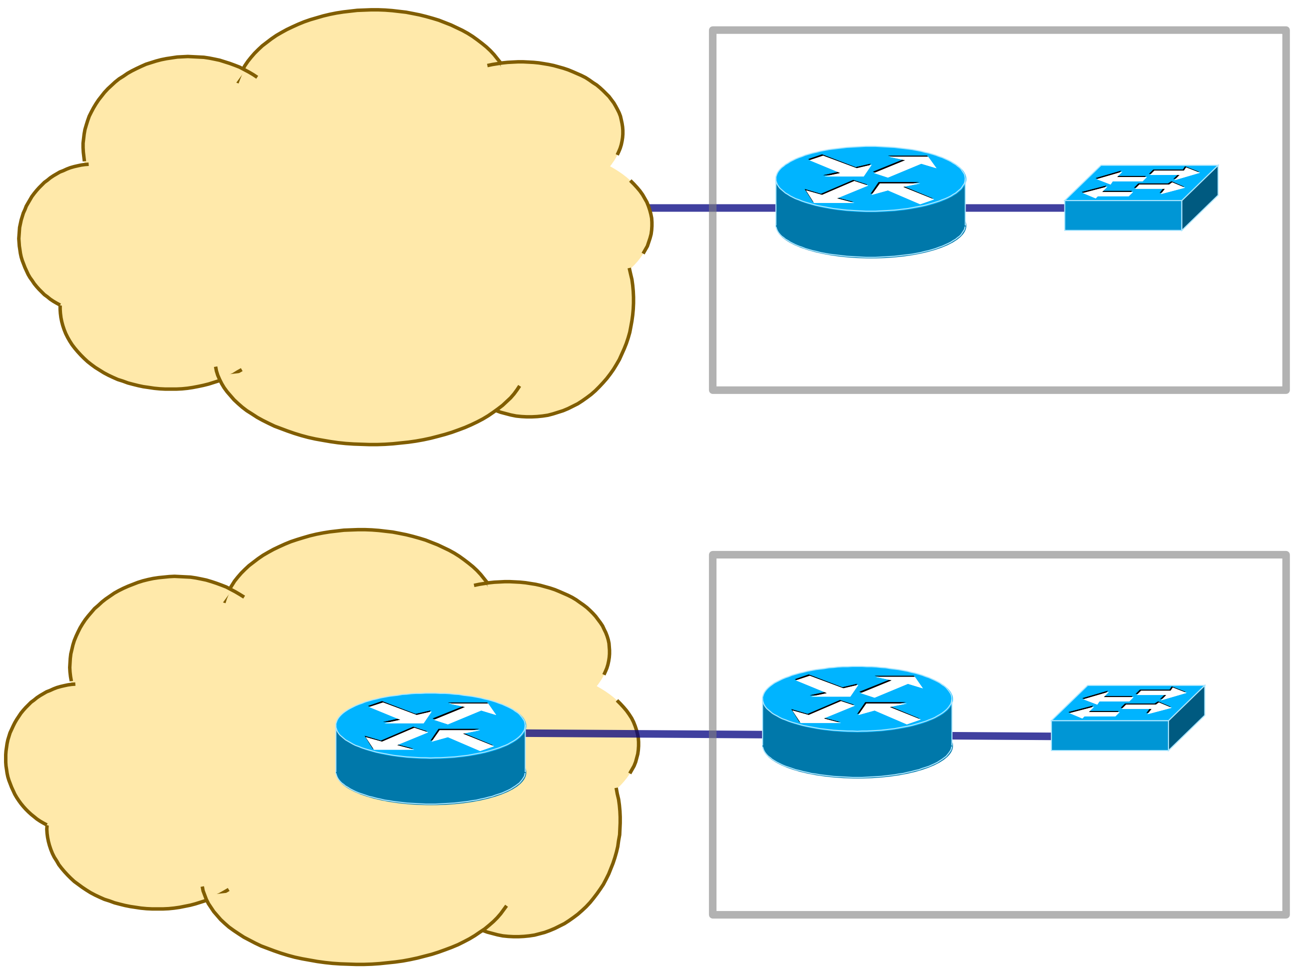

How does NAT work?

- Inspect outgoing traffic
  - Collect (*src_addr*, *src_port*, *dst_addr*, *dst_port*)
- Re-write *src_addr* to my exterior interface, find an unused source port on my exterior interface and re-write *src_port* to that
- Record these addresses and ports

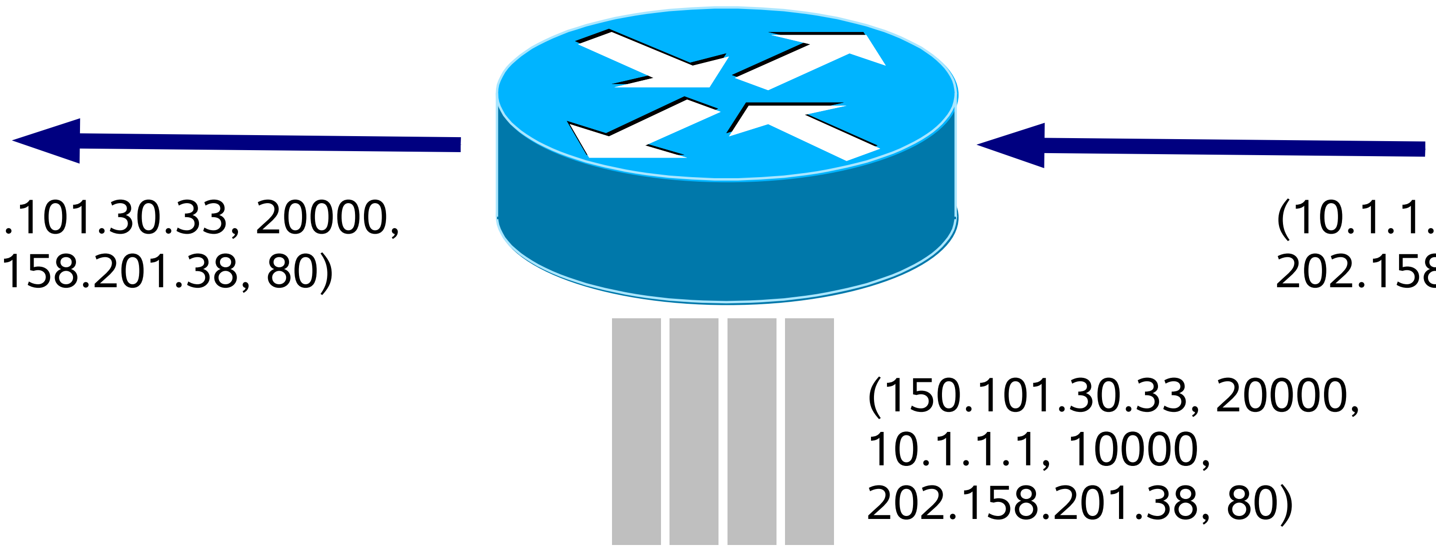

How does NAT work?

- Inspect incoming traffic
- Is the incoming (*src_addr*, *src_port*, *dst_addr*, *dst_port*) in the NAT table?
- Re-write the *dst_addr* and *dst_port* to the original values in the table

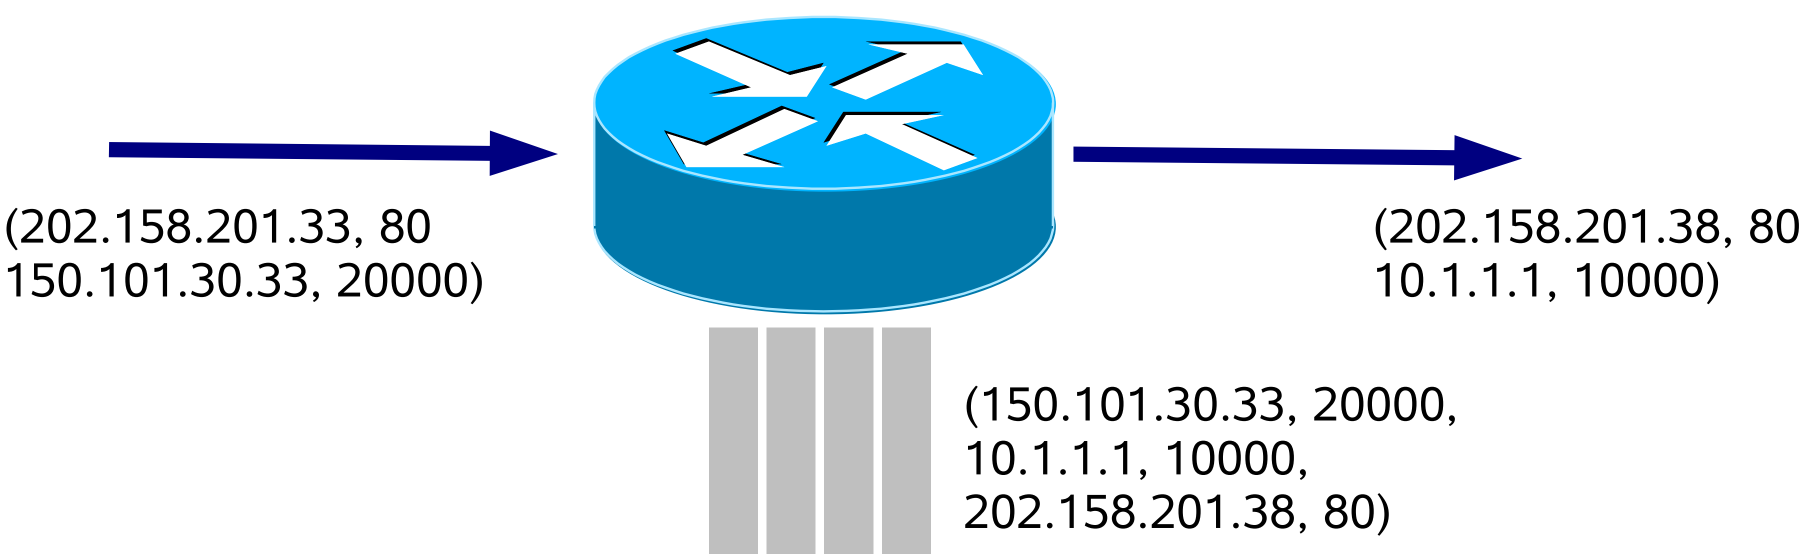

Wrinkles with NAT

- Some protocols embed IPv4 addresses
  - These need to be rewritten too
  - May be complex and thus dangerous to do in the forwarding plane
    - eg: SNMP uses ASN.1 encoding
- Some protocols embed forthcoming connection information
    - FTP
- These are typically handled by “NAT modules” which do deeper inspection of the traffic to add entries to the expectation table

NAT is deep packet inspection

- Complex
  - Forwarding plane moves from ASIC to CPU
- Jitter and complexity attacks
  - Some packets need a lot more work than others
- Exploits of code with errors
  - Complex code, so errors certain
- Huge amounts of state
  - Abundant opportunity for resource exhaustion
- Timeouts
  - Some traffic simply isn't suitable

Implications of carrier-class NAT

- The pain of deep packet inspection is exploitable
  - Contrary to the typical IETF practice of soft state protocols
- Latency will increase
  - These will be expensive boxes, so there will be only a few in a ISP's network
  - Gamers will love IPv6
- There is no end-to-end visibility

No end-to-end visibility

- We're sort of used to that: sharing photos on Flickr rather than on a home router
- Real IPv4 addresses are already special
  - Skype supernode
  - Who wants to volunteer to run a real IPv4 address in a NAT world?
- Potential for evil ISPs to move the Internet from a low-rent transport to a “walled garden” where the only services available are those selected by the ISP

Customers and the walled garden

- No research
  - Especially research which disrupts ISP business plans
- NAT performs poorly
  - It's deep packet inspection
  - We've already got severe TCP performance problems with normal routers
- NAT is a poor fit to sensor networks
  - Timeouts and 30s keepalives
  - UDP blasting from big sensors

Our customers' customers

- Internet traffic is language-based
- Australia – a small English-speaking country on the far edge of Asia – is an exception
- So it is possible for some language groups to move to IPv6 but not others
  - If IPv4 addresses are priced, then that price will be beyond customers in developing countries
- Noting that our customers' customers come from greater Asia

Practicalities of staged deployment

aarnet
Australia's Academic and Research Network

1. Paperwork

- Allocate IPv6 prefix
- Develop addressing plan
  - Lay IPv6 design over IPv4 design
  - There are 16 bits for subnetting, use the top 4 or so for site aggregation, leaving about 12 for subnets per site
  - Allocate a /64 per leaf subnet

2. Link to ISP

- Configure a IPv6 address and routing on existing ISP link
  - copying design from IPv4
- Static routing or BGP, depending upon site and ISP requirements
- Create or inject interior default route

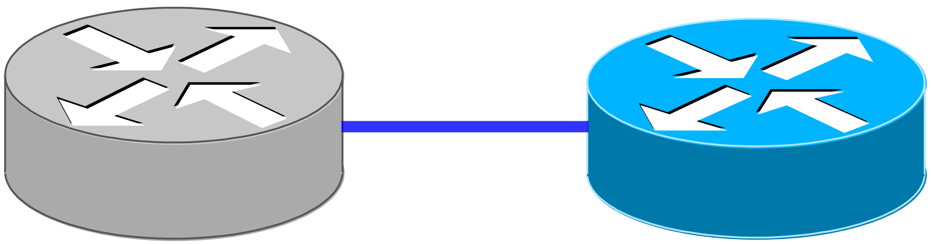

3. Activate IPv6 on backbone

- This brings the first problem: the poor quality of IPv6 support on some firewalls and other middleboxes
- Don't use EUI-64, but be compatible

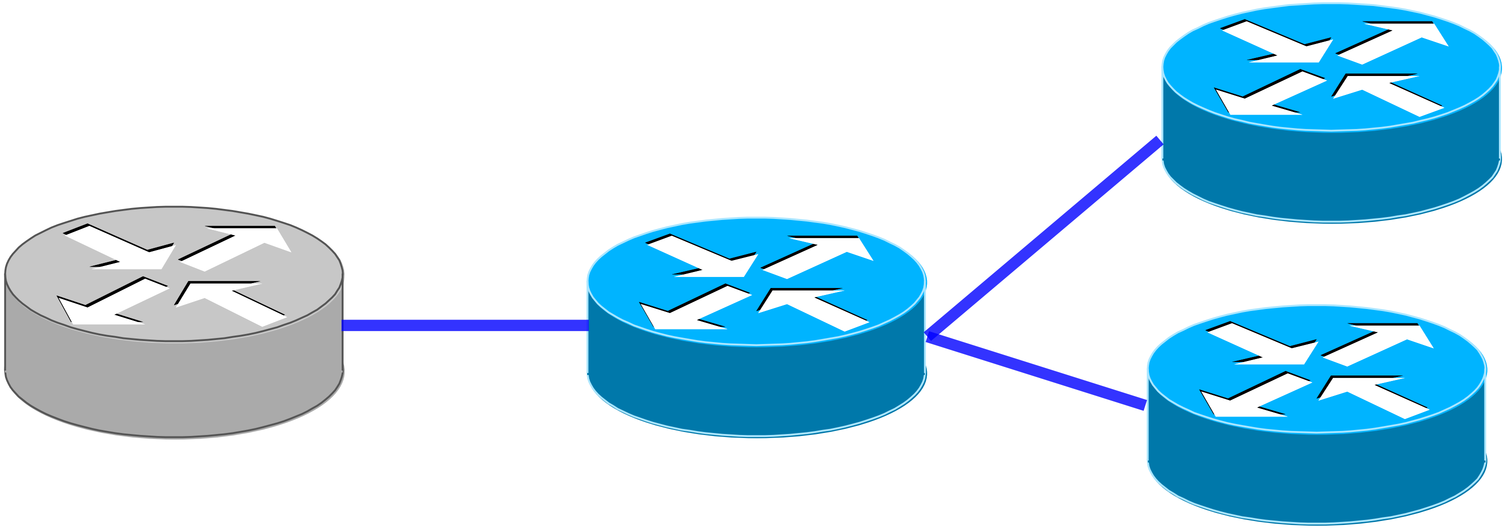

4. Establish networking servers

- Unless good reason otherwise use autoconfiguration (EUI-64 addressing) with stateless DHCP
- Stateless DHCP provides DNS and NTP server addresses
  - These will be IPv4 addresses, because of Windows Xp
- Use Dynamic DNS for the average host
- If you plan on IPv6-only devices then use an anycast IPv6 server on the well-known addresses

5. Find a ~~sucker~~ early adopter

- Computer science, engineering, ourselves
- System administration team

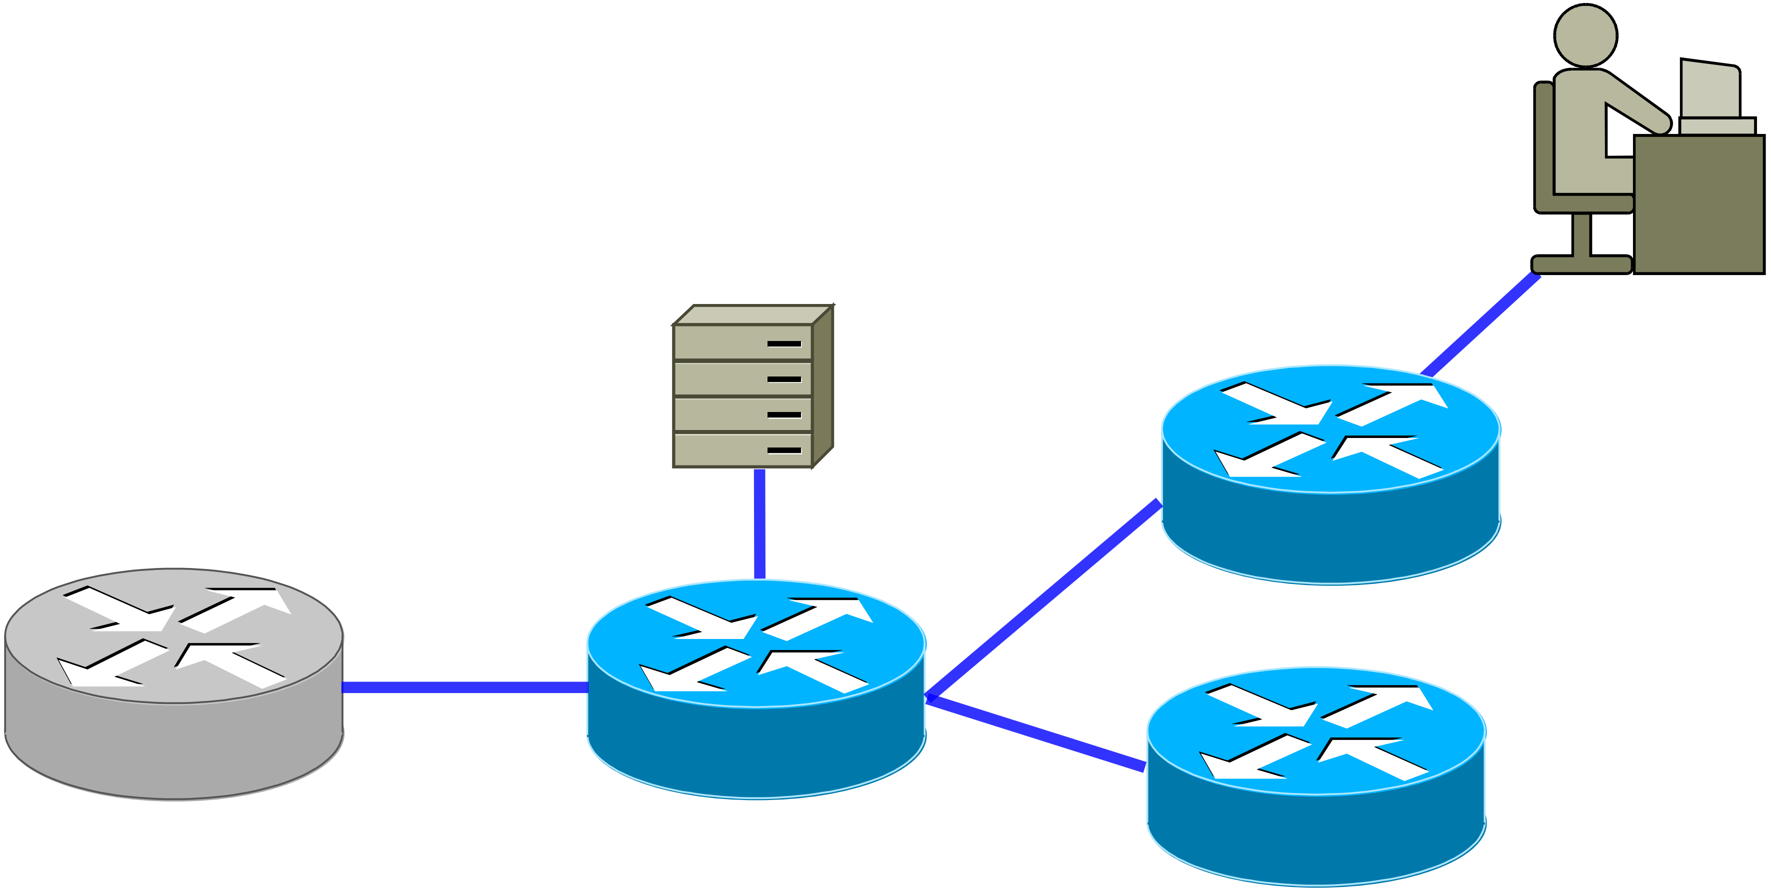

6. Transition public-facing services

- Web, e-mail, …
- Issue: Microsoft Exchange 2003
- Decision: EUI-64 or fixed address in the /64

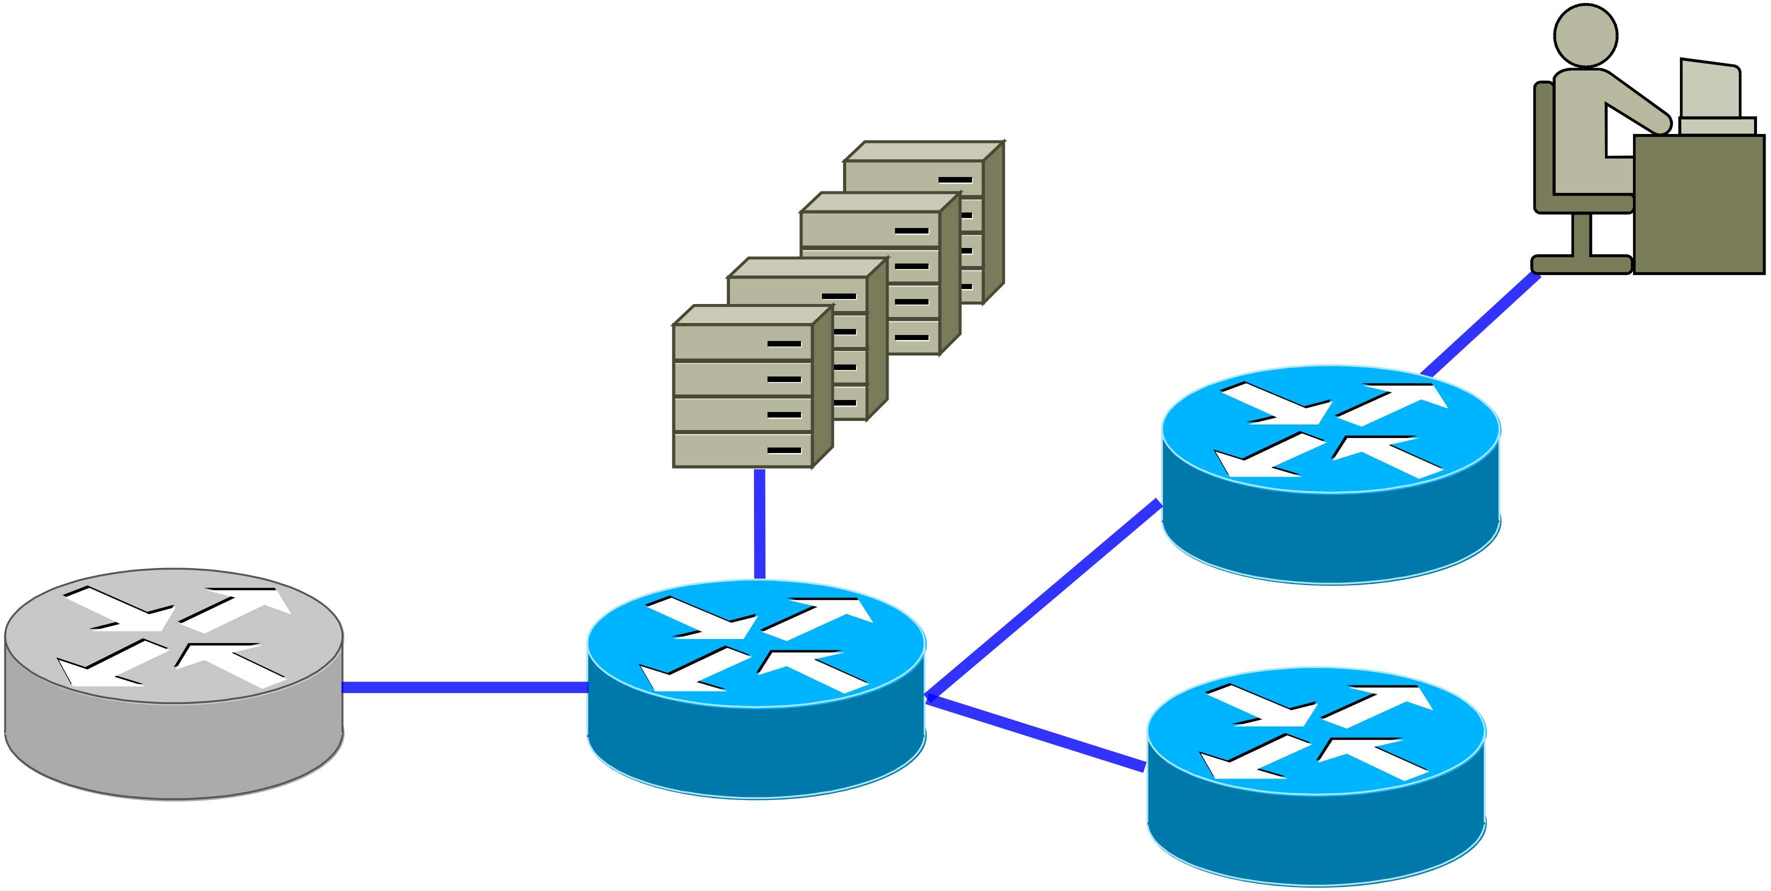

7. Transition the masses

- Issue: people how travel to other sites which have IPv6 configured but no connectivity
- Issue: another round of fighting with middlerubbish such as VPN servers and clients
- Issue: accounting

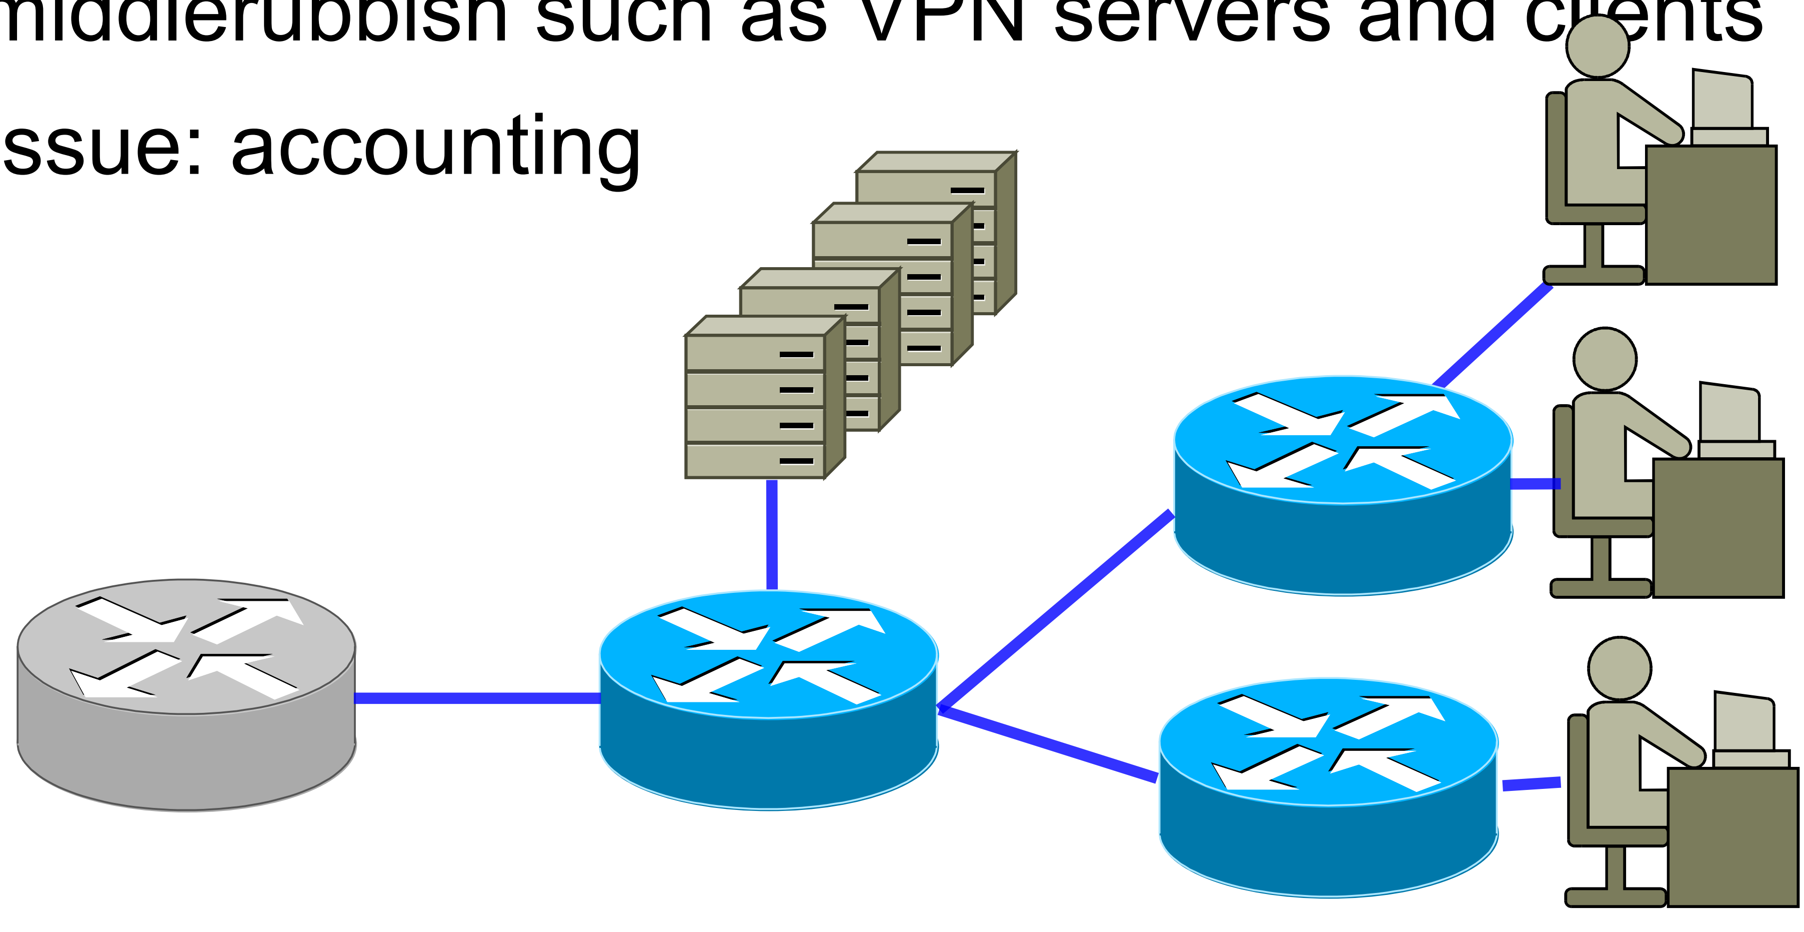

8. Transition inward-facing services

- Problem: disconnect between network engineering and applications programmers
  - "You want us to upgrade PeopleSoft so you can get IPv6 support?"
  - "You want deployment prior to the annual production line shutdown?"

9. Finish the job

- Delegation using IPv6 to DNS servers
  - Not available to edu.au
- Activate equivalent IPv6 features on switches as used on IPv4
  - To prevent address spoffing and so on
- Be careful not to deploy services which really only make sense for IPv4
  - VRRP
- Monitoring systems

Applications: get them running

- Even a trivial task such as finding a IP address needs more work than expected
  - IPv4: [0-9]+\.[0-9]+\.[0-9]+\.[0-9]+
  - IPv6: First network addresses which are not regular
  - IPv6: Uses different characters, “:” was a error
- Applications' deployment timelines are a lot longer than network engineering
- Not fair to only get them running after network engineering and systems administration have finished
  - You can use a tunnel broker to get them IPv6 for testing

A few things we've learned

aarnet
Australia's Academic
and Research Network

IPv6 applications

- “Finding each other” applications
  - Peer-to-peer networks
  - Videoconferencing
- Simple old-fashioned Internet
  - Why does the web server on my laptop stop working when I use the home network?
  - Why can't I directly ssh to my laptop when on my home network?
- Avoiding latency of NAT gateways
  - Gamers

Security

- Hosts
  - Not all firewall products understand IPv6, even when the host is running IPv6. You can guess the OS.
- Routers
  - It's a second protocol
    - ipv6 routing
      line vty 0 4
       ip access-group VTY-LIST
       ip access-group VTY-LIST6
- The real problem is support in corporate firewalls
  - And upgrade plans for those firewalls

Monitoring

- How a connection works:
  - Do I have a global address on default route interface?
  - Yes, look up DNS name using AAAA
    - Present, use that IPv6 address
    - Absent, try to look up the A record
  - No, try to look up the A record
  - Got a AAAA, try for IPv6 connection
Got a A, try for IPv4 connection
- What happens if we have a black hole on IPv6?
  - IPv6 traffic dies, IPv4-based monitoring system says all well

Reality of corporate networks

- Inadequate
  - Configuration control
  - Monitoring
  - Change control
  - Lab scenarios
- Firewalls are the new voodoo
  - Configuration changes induce fear
  - IPv6 changes the sense of firewall rules: match against lower /64
    - ::1 to ::ff Network
    - ::ff00 to ::ffff Servers
    - ::1234:1234:1243:1234 Autoconfed MAC

Training

- University computer science courses never show students an IPv6 address
- TAFE ditto
- Vendor training (MSCE, RHCE) ditto

AARNet's experience with IPv6

www.gdt.id.au/~gdt/presentations

Glen Turner
glen.turner@aarnet.edu.au

aarnet
Australia's Academic and Research Network