NAT Tutorial

Dan Wing, dwing@cisco.com

IETF77, Anaheim

March 21, 2010

Agenda

- NAT and NAPT
  - Types of NATs
- Application Impact
  - Application Layer Gateway (ALG)
  - STUN, ICE, TURN
- Large-Scale NATs (LSN, CGN, SP NAT)
- IPv6/IPv4 Translation (“NAT64”)
- NAT66

Agenda

- NAT and NAPT
  - Types of NATs
- Application Impact
  - Application Layer Gateway (ALG)
  - STUN, ICE, TURN
- Large-Scale NATs (LSN, CGN, SP NAT)
- IPv6/IPv4 Translation (“NAT64”)
- NAT66

NAT

- First described in 1991
- 1:1 translation
  - Does not conserve IPv4 addresses
- Per-flow stateless
- Today's primary use is inside of enterprise networks
  - Connect overlapping RFC1918 address space

NAT Diagram

- Hosts seem to have multiple IPv4 addresses – almost like “ghosts”

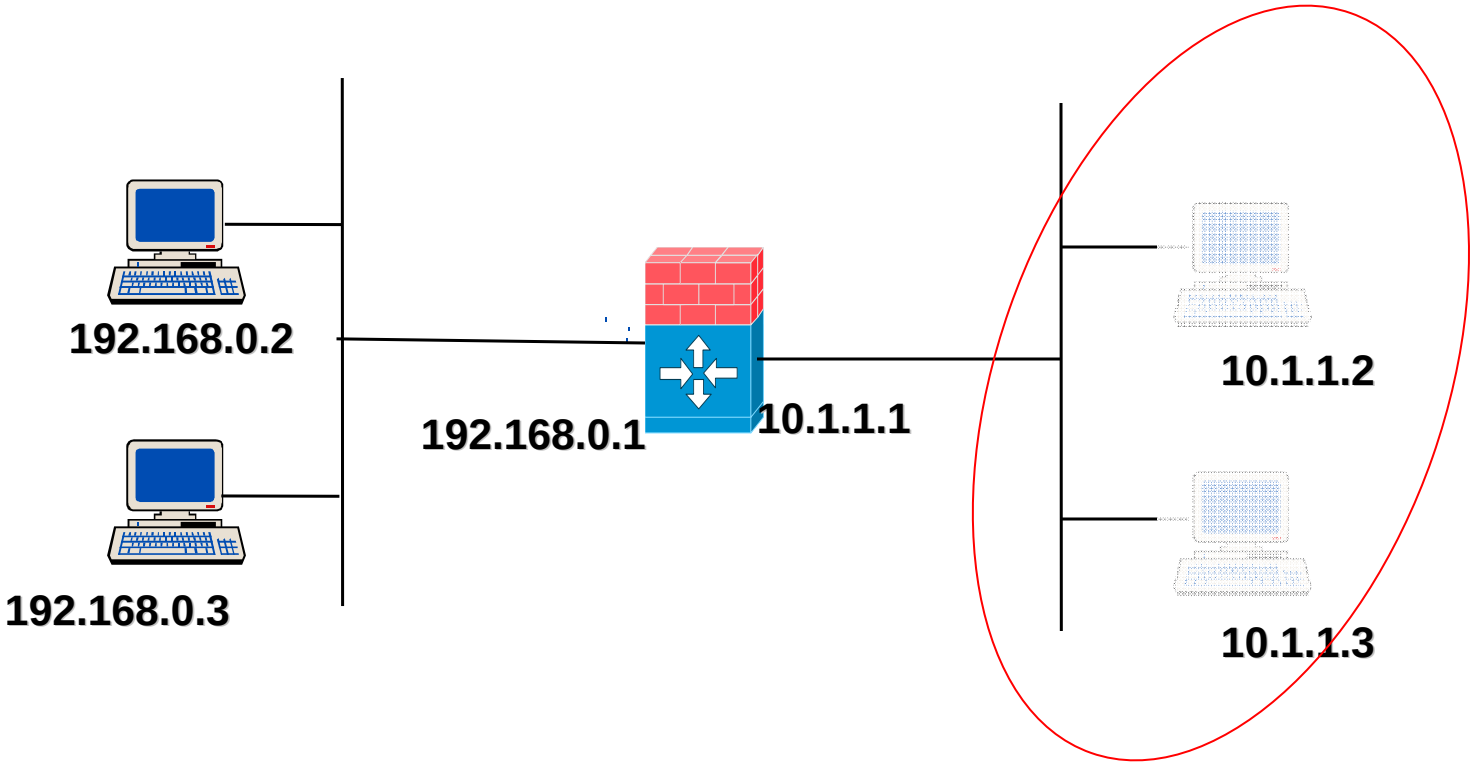

NAPT

- Described in 2001 (RFC3022)
- 1:N translation
  - Conserves IPv4 addresses
  - Allows multiple hosts to share one IPv4 address
  - Only TCP, UDP, and ICMP
  - Connection has to be initiated from ‘inside’
- Per-flow stateful
- Commonly used in home gateways and enterprise NAT

NAPT Diagram

- Hosts share an IPv4 address

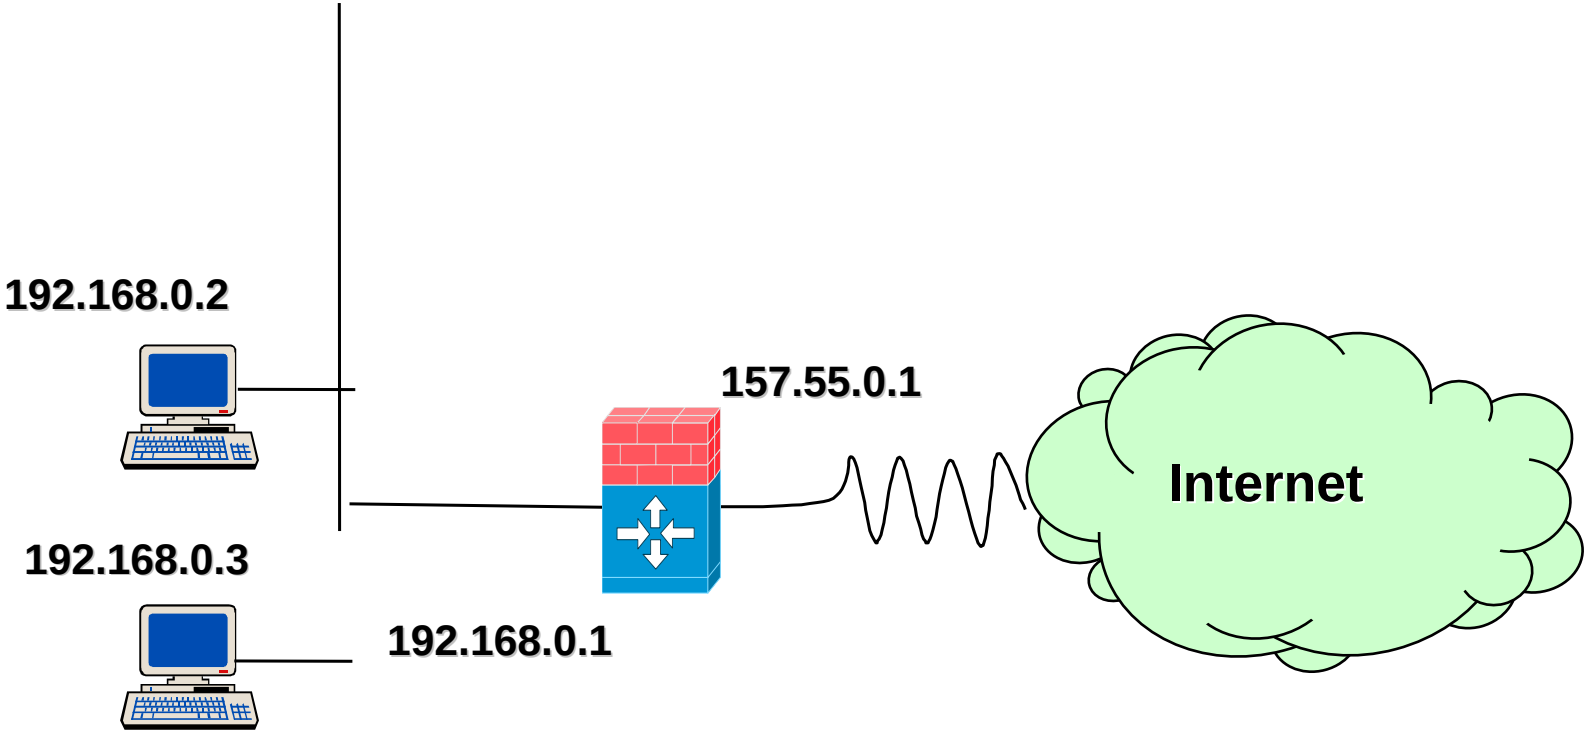

NAPT complications

- NAPT requires connections initiated from ‘inside’
- Creates state in the network (in the NAPT)
  - This is bad
  - NAPT crashes -> connections break
- When to discard state?
  - TCP RST? Spoofed RSTs?
  - Timeout?

Terminology

- “NAT” is spoken/written instead of “NAPT”
  - Even though NAPT is often more accurate
  - The more accurate “PAT” never caught on
- So, it’s “NAT”
- Now, often called “NAT44” to differentiate from NAT64 and NAT46

Types of NAT (old terms)

- Full Cone
- Restricted Cone
- Port Restricted Cone
- Symmetric

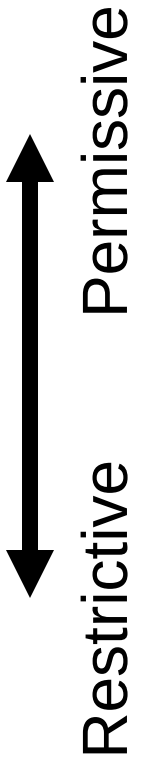

RFC3489 (obsoleted)

Types of NAT (new terms)

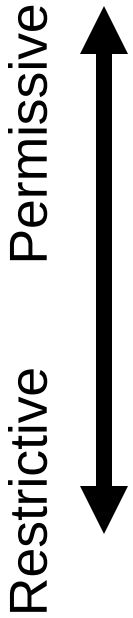

Mapping

- Endpoint-Independent
- Address-Dependent
- Address and Port-Dependent

Filtering

- Endpoint-Independent
- Address-Dependent
- Address and Port-Dependent

RFC4787

Agenda

- NAT and NAPT
  - Types of NATs
- Application Impact
  - Application Layer Gateway (ALG)
  - STUN, ICE, TURN
- Large-Scale NATs (LSN, CGN, SP NAT)
- IPv6/IPv4 Translation (“NAT64”)
- NAT66

NAT Philosophy

- “Be transparent”
- This means it isn’t a proxy
  - Applications are generally unaware of a NAT
- Problem with IP addresses inside the application
  - Generally called a “referral”
  - Example: SIP

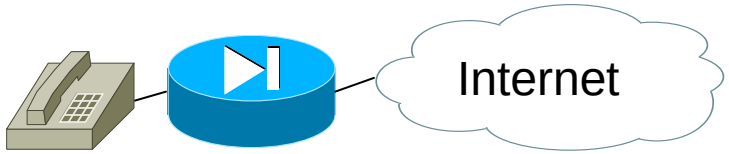

“my address is 10.1.1.1” Internet sees 161.44.1.1

NAPT and servers

- NAPT: connection initiated from inside
- Incoming connections are difficult
- Significant problem for servers
  - Webcam, VoIP, RTSP receivers, etc.
- Port forwarding (“pinholing”, etc.)
  - web or CLI configuration
  - UPnP IGD, NAT-PMP
  - All have drawbacks

Application Layer Gateway (ALG)

- Application awareness inside the NAT
- ALG modifies IP addresses and ports in application payload, and creates NAT mapping
- Each application requires a separate ALG
  - FTP, SIP, RTSP, RealAudio, …

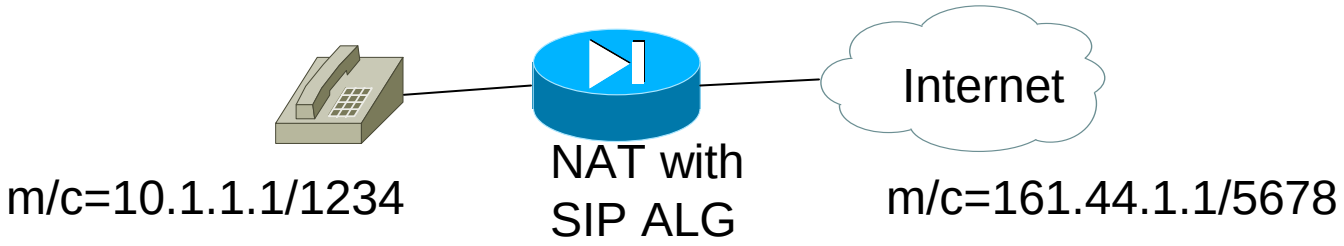

Problems with ALGs

- Requires ALG for each application
- Requires ALG that understands *this particular* application's nuance
  - Proprietary extensions / deviations
  - New standard extensions
- ALG requires:
  - Un-encrypted signaling (!)
  - Seeing application's signaling and media/data
    - easy with stub network; harder with mesh network

Application Solutions

- Applications cannot successfully rely on ALGs
- So, Applications have developed their own solutions
- FTP PASV
  - Data connection always to server. Has security side-effects.
- ICE, STUN, TURN
  - Intelligence in endpoint
  - Useful for offer/answer protocols (SIP, XMPP, probably more)
  - Standardized in MMUSIC and BEHAVE
  - (more on next slides)
- RTSP supports ‘interleaved data’ (RFC2326)
  - Streaming over RTSP’s TCP control channel
- RTSPv2 with ICE-like NAT traversal
- HTTP delivery
  - Flash (e.g., YouTube)

STUN, ICE, TURN

- Request/response protocol, used by:
  - STUN itself (to learn IP address)
  - ICE (for connectivity checks)
  - TURN (to configure TURN server)
- The response contains IP address and port of request
  - Runs over UDP (typical) or TCP, port 3478
- Think http://whatismyip.com

STUN, **ICE**, TURN

- Procedure for Optimizing Media Flows
- Defines SDP syntax to indicate ‘candidate addresses’
- Uses STUN messages for connectivity checks
  - Sent to RTP peer, using **same ports** as RTP
- First best path wins

- Think: gather all my IP addresses, send them to my peer, and do connectivity checks

STUN, ICE, **TURN**

- Media Relay Protocol and Media Relay Server
- Only used when:
  - **both** endpoints are behind ‘Address and Port-Dependent Filtering’ NATs (rare, about 25% of NATs), or
  - one endpoint doesn’t implement ICE, and is behind a ‘Address and Port-Dependent Filtering’ NAT

ICE Deployments

- Google chat (XMPP)
- Microsoft MSN
- Yahoo
- Counterpath softphone

Agenda

- NAT and NAPT
  - Types of NATs
- Application Impact
  - Application Layer Gateway (ALG)
  - STUN, ICE, TURN
- Large-Scale NATs (LSN, CGN, SP NAT)
- IPv6/IPv4 Translation (“NAT64”)
- NAT66

NAT444 = NAT44 + NAT44

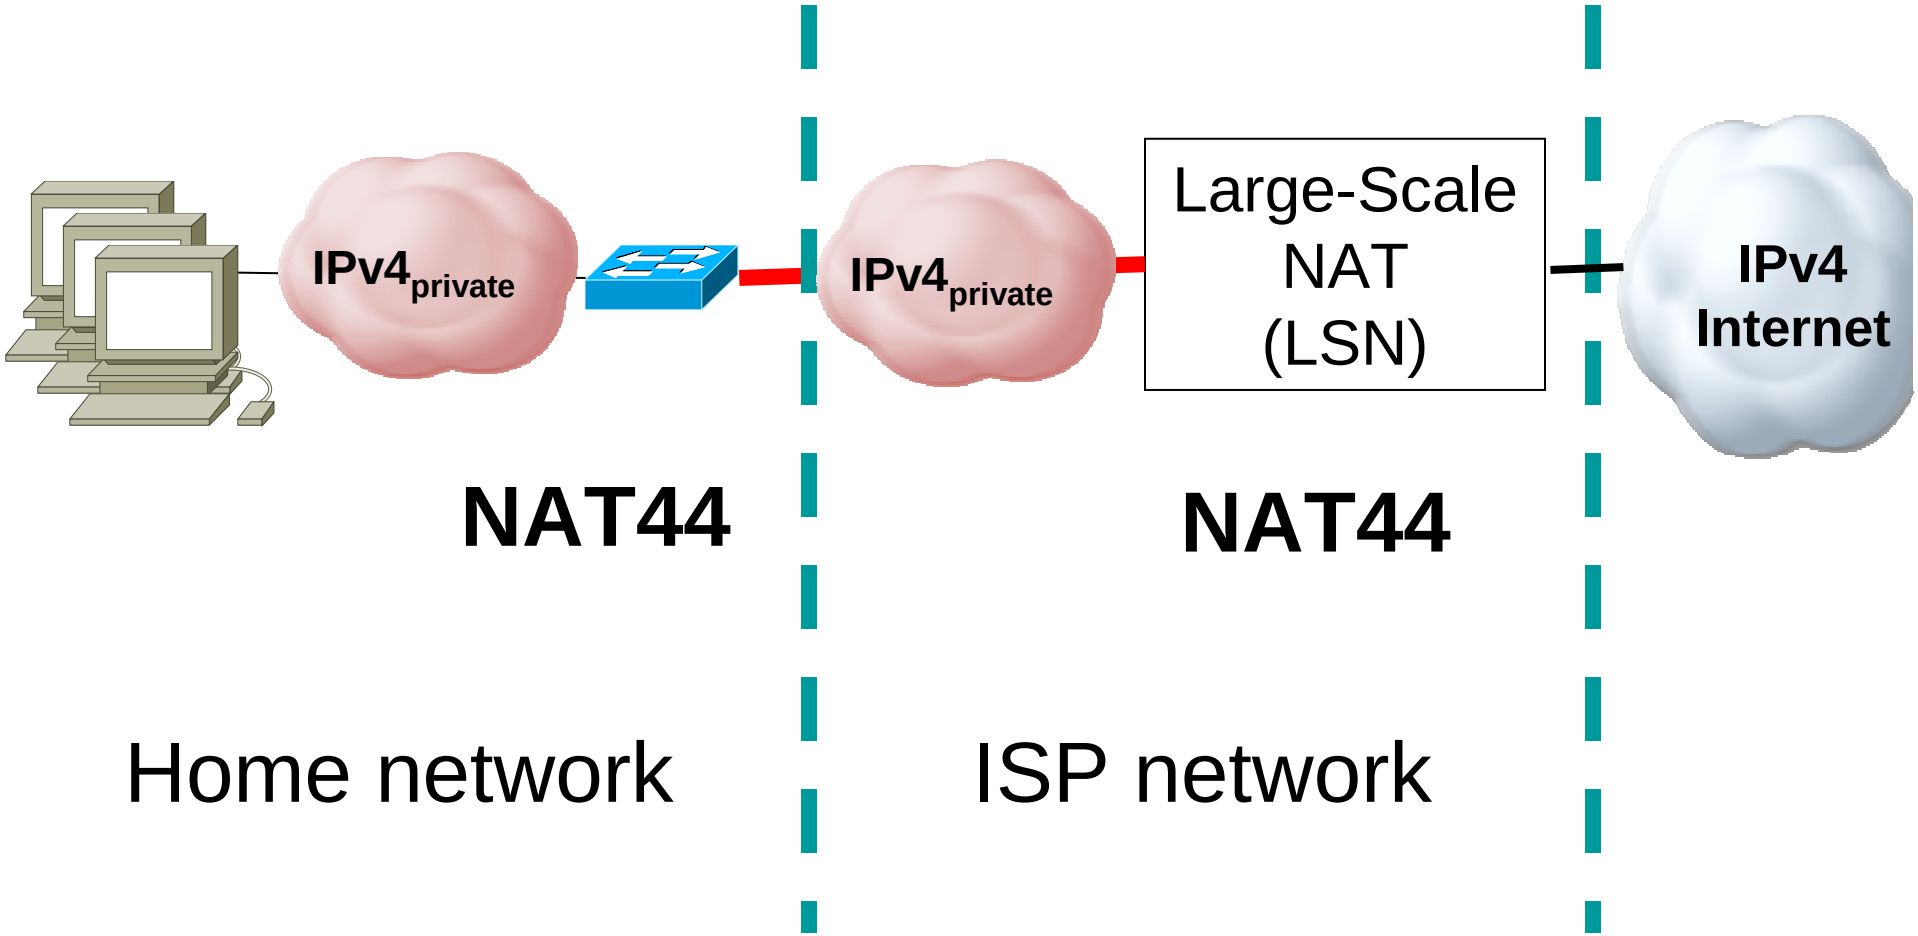

Large Scale NAT (LSN)

- Essentially, just a big NAPT44
- Needs per-subscriber TCP/UDP port limits
  - Prevent DoS
  - If too low, can interfere with applications
    - Classic example: Google maps
- How to number network between subscriber and LSN?
  - RFC1918 conflicts with user's space, breaks some NATs
  - Using routable IPv4 addresses is … wasteful

LSN and ALG

- Operationally complex in a LSN
  - Application X works but Application Y breaks. Upgrade ALG??
  - How long is vendor turn-around for patches?
- Interfering with competitor’s over-the-top application (e.g., SIP, streaming video)

IPv4 Address Sharing

- Problem most noticed with LSN
- Reputation and abuse reporting are based on IPv4 address
  - Shared IP address = shared suffering
  - Law Enforcement
  - “Which subscriber posted on www.example.com at 8:23pm?”
  - Requires LSN log source port numbers
  - Requires web servers log source port numbers
- Everybody can’t get port 80
- Geo-location breaks

Agenda

- NAT and NAPT
  - Types of NATs
- Application Impact
  - Application Layer Gateway (ALG)
  - STUN, ICE, TURN
- Large-Scale NATs (LSN, CGN, SP NAT)
- IPv6/IPv4 Translation (“NAT64”)
- NAT66

The Ideal IPv6/IPv4 Translation

Translation versus Tunneling

- If you have a choice, tunnel
  - 6rd (IPv6 over IPv4)
  - Dual-Stack Lite (IPv4 over IPv6)
- Translate only when crossing between address families
  - IPv4-only host to IPv6-only host
  - IPv6-only host to IPv4-only host

Then, Why Translate?

- Will exhaust IPv4 addresses in 2011-2012
- Need to access IPv4-only content from IPv6-only clients
- Long tail of IPv4-only content
  - Children's soccer practice schedule

- Longer term: need to access IPv6-only servers from IPv4-only clients

NAT-PT

- NAT-PT combined all scenarios
  - IPv4 to IPv6 is problematic; IPv6 space is bigger
  - Broke DNSSEC
- RFC4966 said IPv6/IPv4 translation causes other side effects
  - And some are not solvable

- But:
- IPv4 addresses running out
- Effectively no IPv6 Internet access and no IPv6 content anywhere in the world
- We can’t tunnel everywhere

Translation Evolution S-Curve

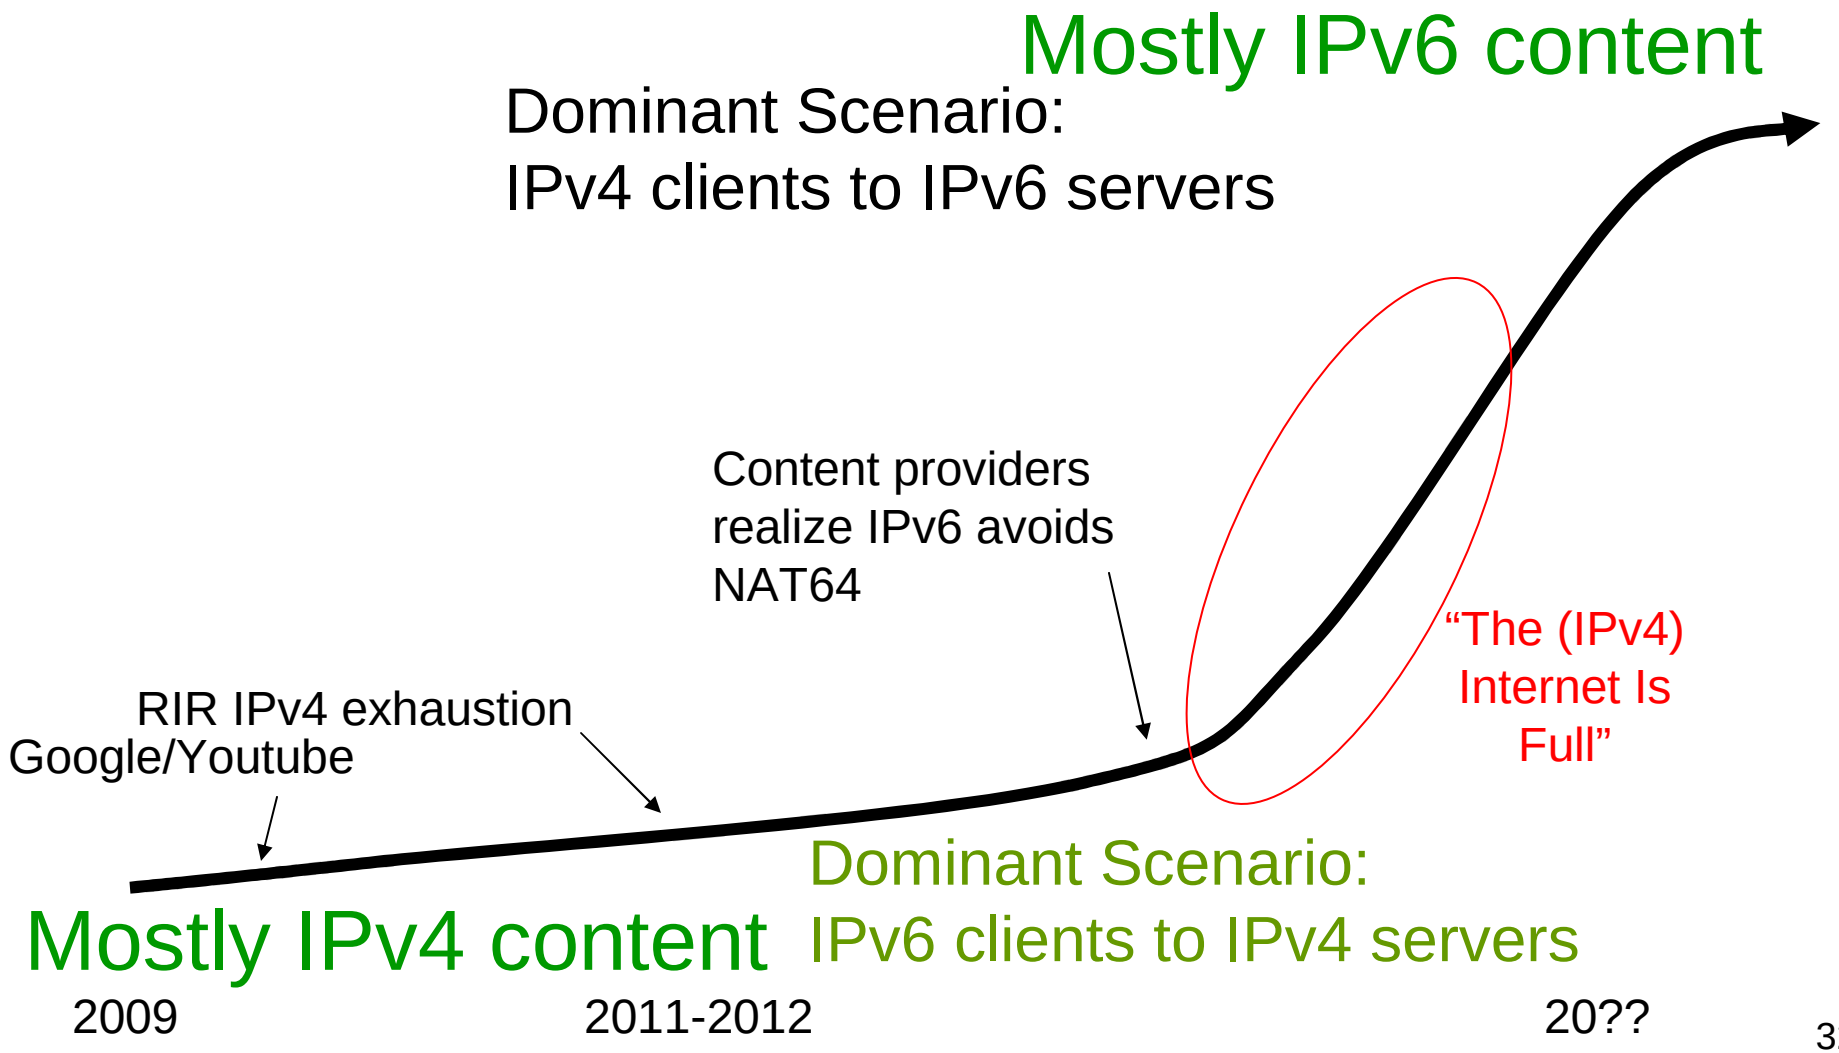

BEHAVE's Approach

- Do first part of S-Curve first
- Split problem into separate documents
  - Framework
    - Lists all 8 scenarios
  - Address format
  - 6/4 translation (1:1), including fragmentation
  - Stateful translation (1:N)
  - DNS64
  - FTP64 ALG
- Later scenarios in S-Curve done later

IPv6/IPv4 Translation: some detail

- Connecting an IPv6 network to the IPv4 Internet
  - You built an IPv6-only network, and want to access servers on the IPv4 Internet
- Connecting the IPv6 Internet to an IPv4 network
  - You have IPv4 servers, and want them available to the IPv6 Internet
- Connecting the IPv4 Internet to an IPv6 network
  - You built an IPv6-only network, and want its servers available to the IPv4 Internet

Connecting an IPv6 network to the IPv4 Internet

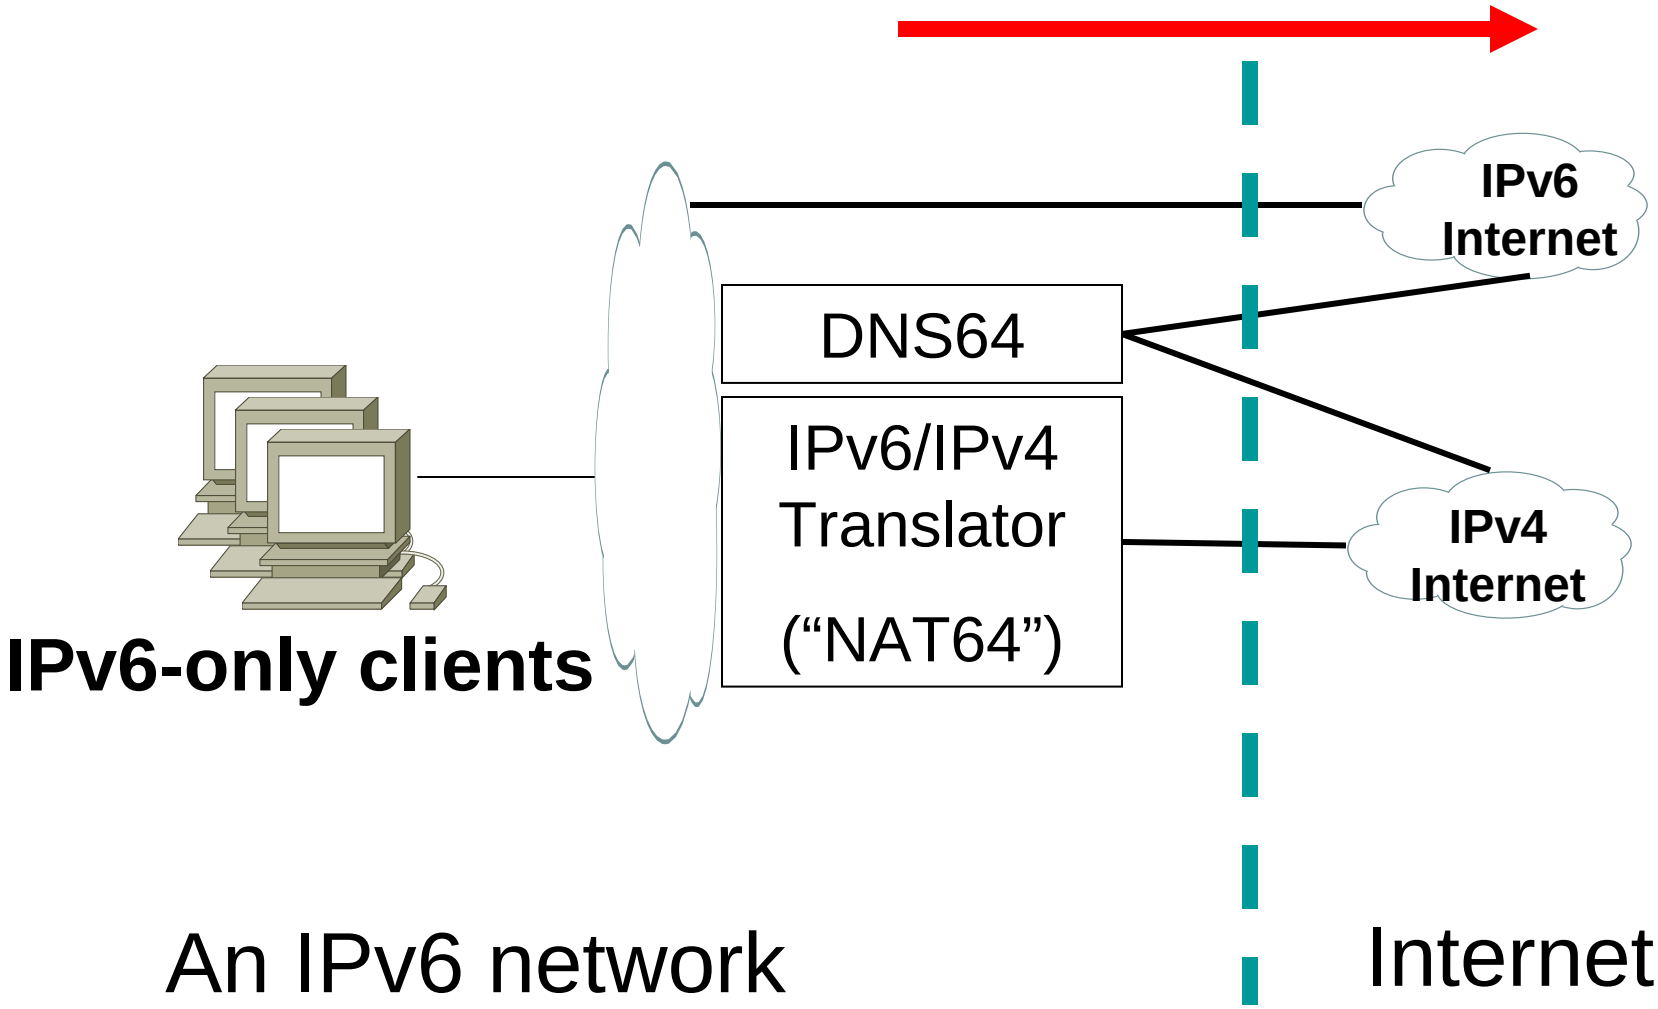

DNS64

- Synthesizes AAAA records when not present
  - With IPv6 prefix of NAT64 translator
- Works for applications that do DNS queries
- Breaks for applications that don’t

IPv6/IPv4 Translation

Stateless

- 1:1 translation
- “NAT”
- Any protocol
- No IPv4 address savings
  - Just like dual-stack

Stateful

- 1:N translation
- “NAPT”
- TCP, UDP, ICMP
- Saves IPv4 addresses

IPv6/IPv4 translation issues

- IPv4 address literals
  - http://1.2.3.4
  - SIP, RTSP, SAP
- IP Family sensitive protocols
  - FTP (EPSV, PASV)

- How to resolve?
  - Application proxies, make application smarter, ALG (FTP64)

Connecting the IPv6 Internet to an IPv4 network

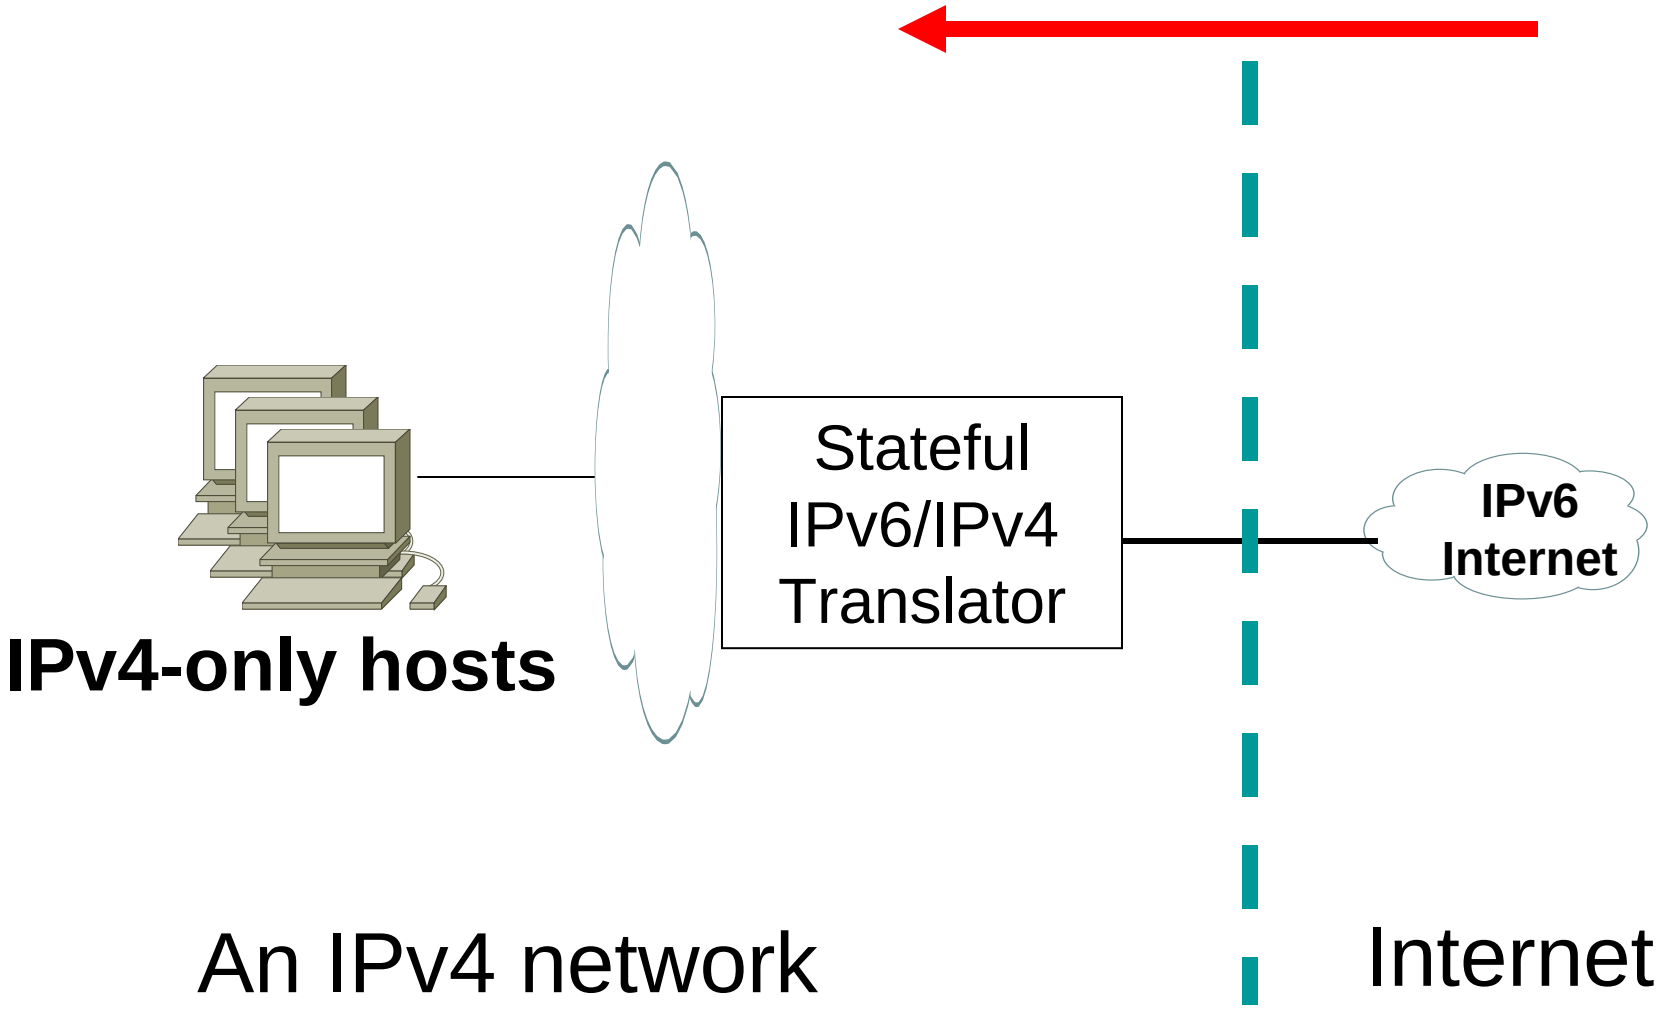

Connecting the IPv6 Internet to an IPv4 network

- Makes IPv4-only servers accessible on the IPv6 Internet
- Requires stateful translation
  - Because IPv6 Internet is bigger than IPv4
  - (can’t represent every address in IPv4)
- All connections come from translator’s IPv4 address
  - Problem for abuse logging
  - Lack of X-Forwarded-For: header
- Maybe application proxy is superior?
  - E.g., lighthttpd
  - But has poor TLS interaction

Later IPv6/IPv4 Scenarios

Connecting the IPv4 Internet to an IPv6 network

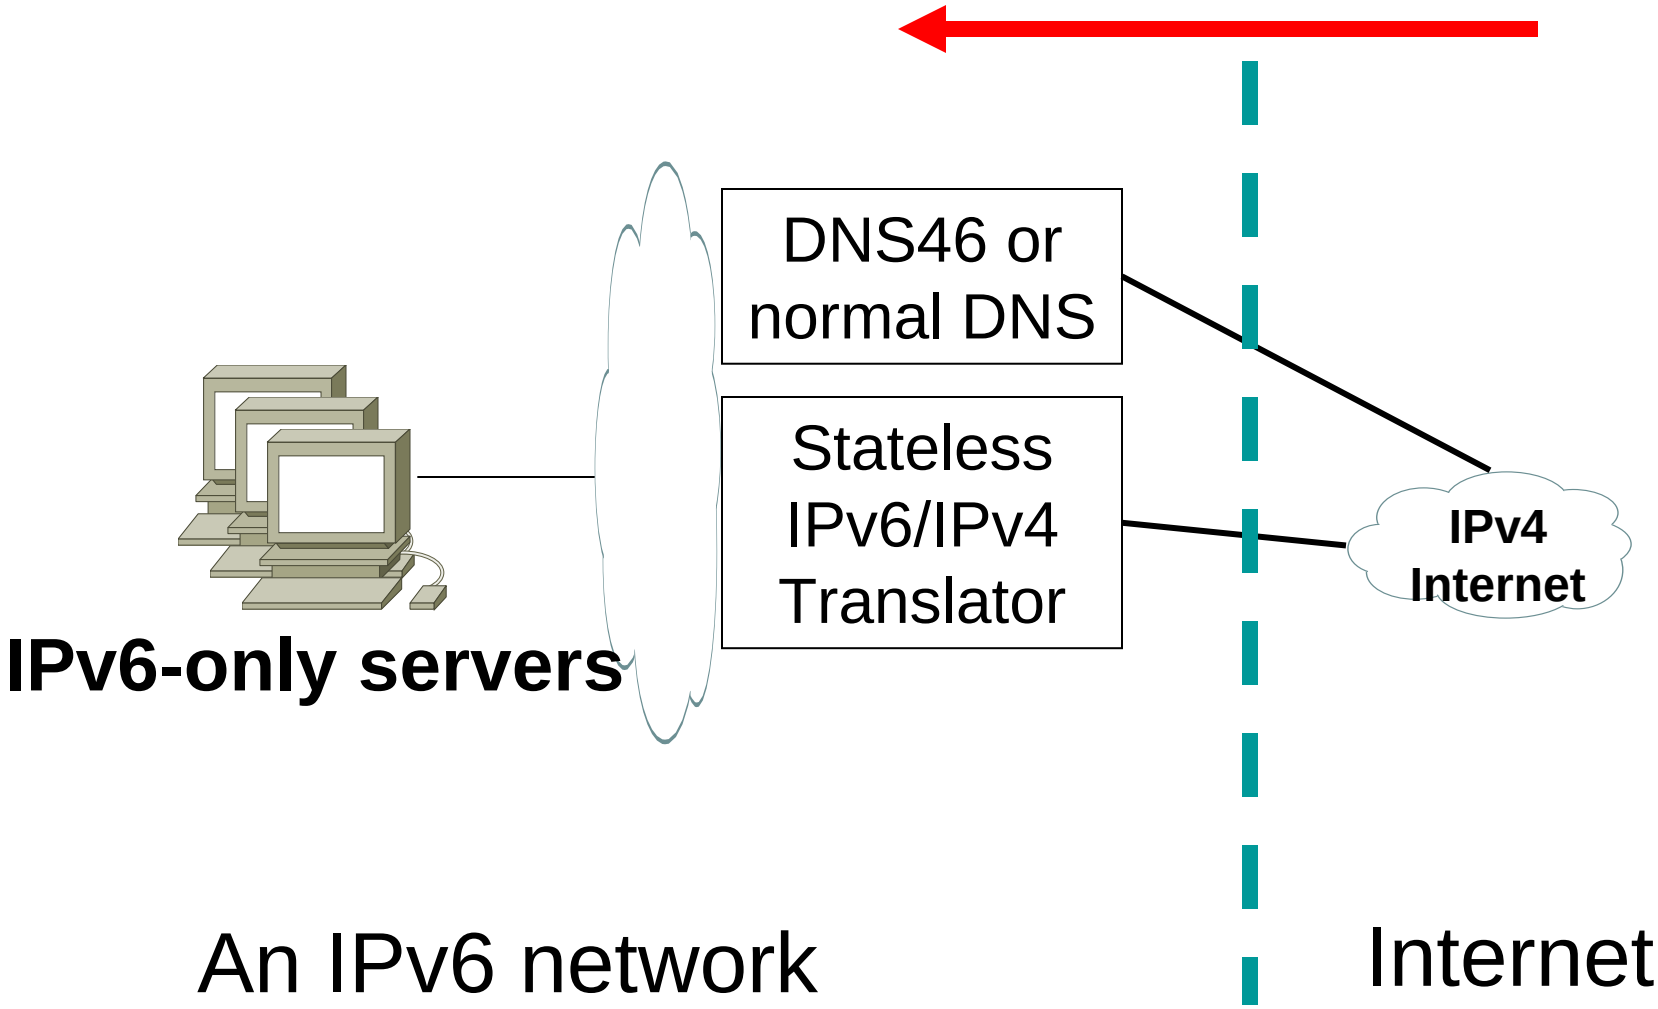

Connecting the IPv4 Internet to an IPv6 network

- Stateless works well, one IPv4 address for each IPv6 server
  - Same IPv4 consumption as dual-stack
- Just like with NAT64 case, don’t use IPv6 address literals
  - IPv4-only client can’t understand them!

Agenda

- NAT and NAPT
  - Types of NATs
- Application Impact
  - Application Layer Gateway (ALG)
  - STUN, ICE, TURN
- Large-Scale NATs (LSN, CGN, SP NAT)
- IPv6/IPv4 Translation (“NAT64”)
- NAT66

NAT66

- This is not NA**P**T66
- NAT66 is 1:1 translation
  - "IPv6 prefix rewriting"
- No per-flow state in the NAT66 device
- No manipulation of TCP or UDP headers

- "But, we don't need NAT with IPv6"

Multi-homed with Provider-Dependent (PD) addresses

- Can’t get PI space
- Don’t want Provider-Dependent space internally
  - renumbering, ACLs
- Lack of RFC4191 support in hosts
- 3GPP and work-at-home VPN tunnels

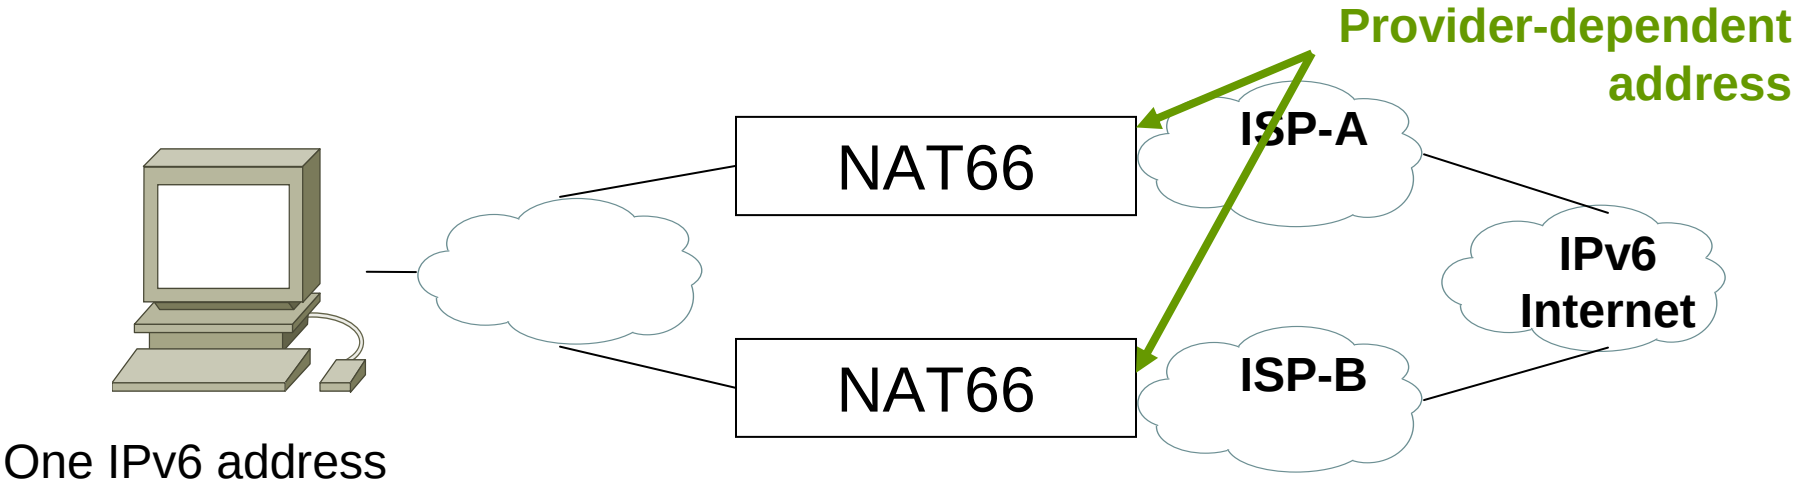

BEHAVE Status

BEHAVE Finished Work

- RFC
  - NAT44 behaviors: TCP, UDP, ICMP
- RFC Editor's queue
  - STUN, TURN, ICE (MMUSIC)

BEHAVE Nearly Finished Work

- IPv6/IPv4 Translation Scenarios
  - √ 1: an IPv6 network to the IPv4 Internet
  - – 2: the IPv4 Internet to an IPv6 network
  - √ 3: the IPv6 Internet to an IPv4 network
  - – 4: an IPv4 network to the IPv6 Internet
  - √ 5: an IPv6 network to an IPv4 network
  - – 6: an IPv4 network to an IPv6 network

BEHAVE Finished 6/4 Translation Documents

- draft-ietf-behave-address-format
- draft-ietf-behave-dns64
- draft-ietf-behave-v6v4-framework
- draft-ietf-behave-v6v4-xlate-stateful
- draft-ietf-behave-v6v4-xlate

BEHAVE Outstanding NAT Work

- draft-ietf-behave-ftp64
- draft-ietf-behave-sctpnat

Summary

- NAT and NAPT
  - Types of NATs
- Application Impact
  - Application Layer Gateway (ALG)
  - STUN, ICE, TURN
- Large-Scale NATs (LSN, CGN, SP NAT)
- IPv6/IPv4 Translation (“NAT64”)
- NAT66

Questions

Dan Wing, dwing@cisco.com