Making NATs work for Online Gaming and VoIP

Dr. Cullen Jennings

fluffy@cisco.com

Topics

- Requirements that Gaming & VoIP place on NATs
- Solutions with NATs
- Types of NATs
- Protocols to work with NATs
- NAT Market
- How to Build Good NATs
- IETF Work

Gaming, VoIP, and Collaboration

- **Real time response is needed to serve these applications**
  - **Need low latency**
  - **Applications use significant bandwidth**

- **Data flows between 2 or more end points**
  - **Client to Client not Client to Server**

- **UDP is usually used for Real Time data**
  - **TCP Retransmission increases latency too much**

What NATs Do

- Allow many computers with private IP addresses to sit behind a single public IP address
- Send packets that arrive at the public IP address to the correct computer behind the NAT
- Reduce number of public IP addresses needed
- Allow partitioning of who manages the IP address space

The NAT Problem

- Alice wants to call Bob, whose phone is behind a NAT
- Bob needs to tell Alice where to send her IP packets to let them traverse his NAT
- STUN (RFC 3489) solves this for most NATs

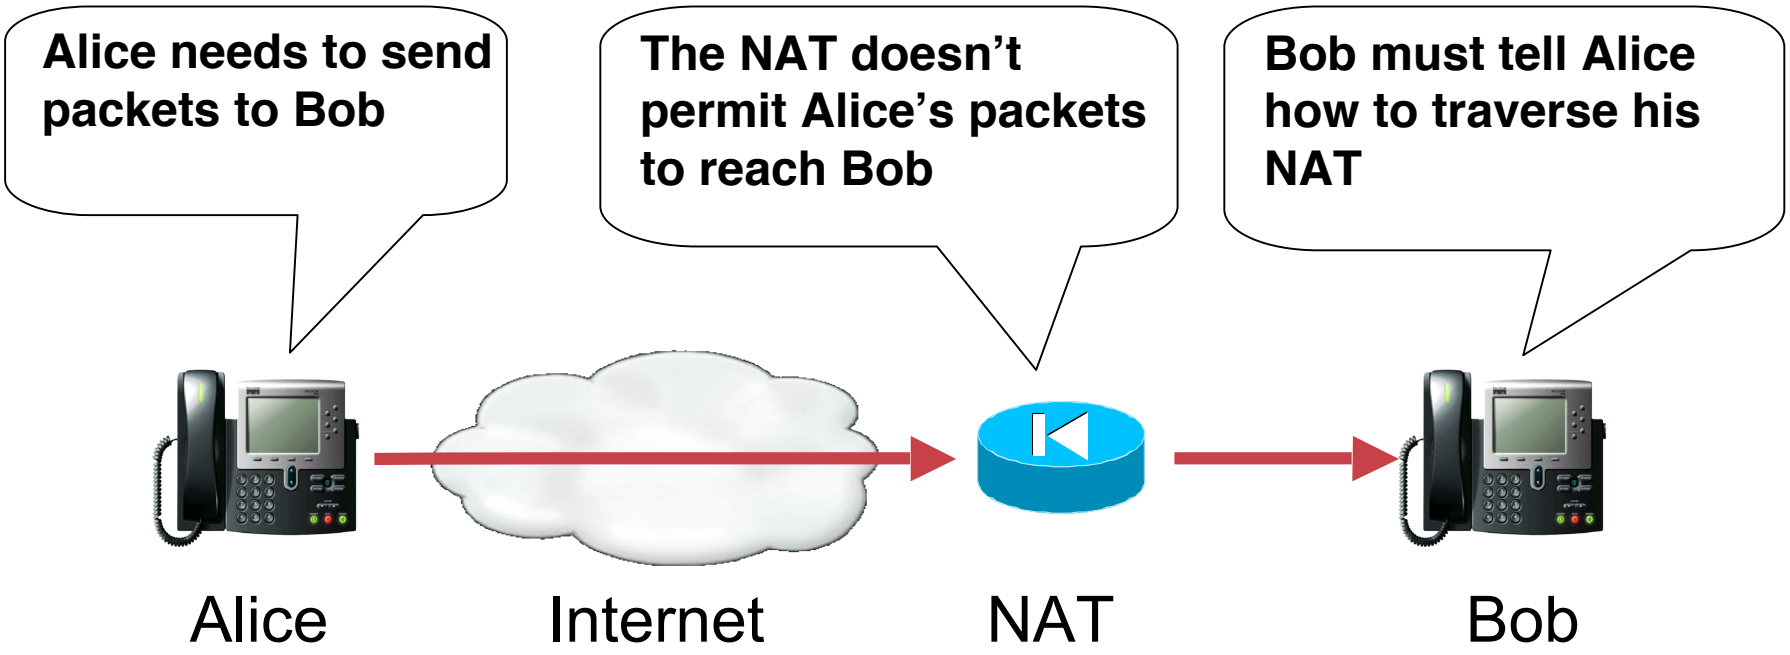

Relay Solution
(Not appropriate for Real Time data)

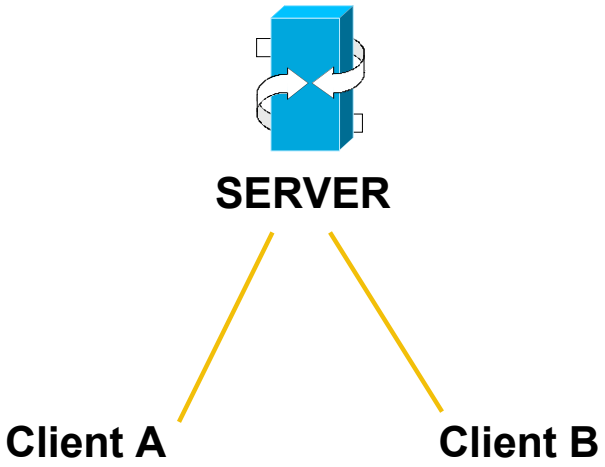

- A communicates with B through a relay
- Server hosting must have bandwidth for all traffic from A → B
- Resulting latency is higher
- Relay needs bandwidth for all the data among all clients

The Latency Problem

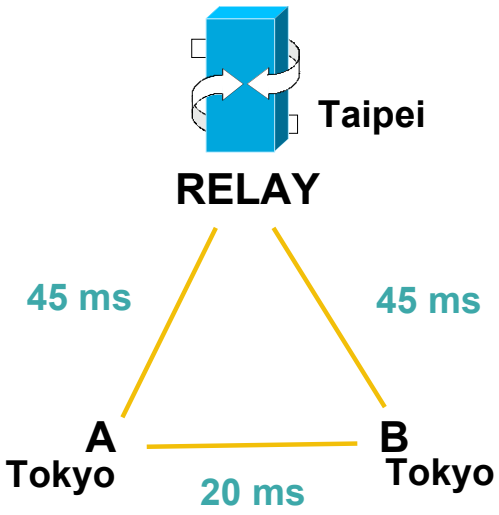

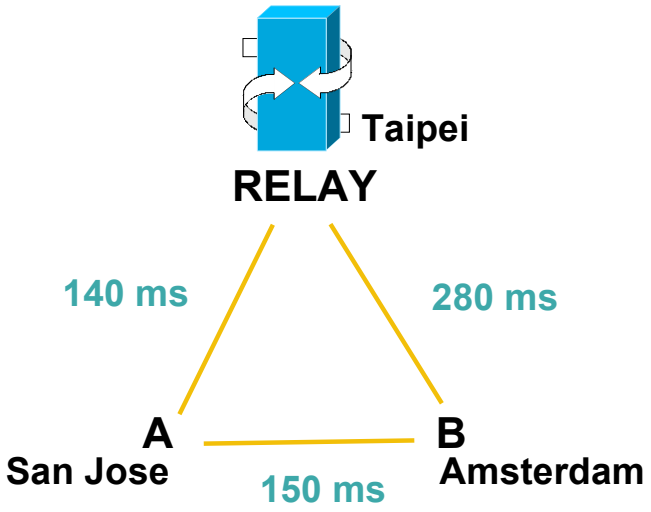

- Communication is often between parties in same geography
- When parties are separated, relay is often off path

- Human communications work best at < 150ms latency
  - Arcade games require even less latency

Echo Server Solutions

- STUN (RFC 3489) is an example of this class of solution
- Used for online gaming & VoIP for many years

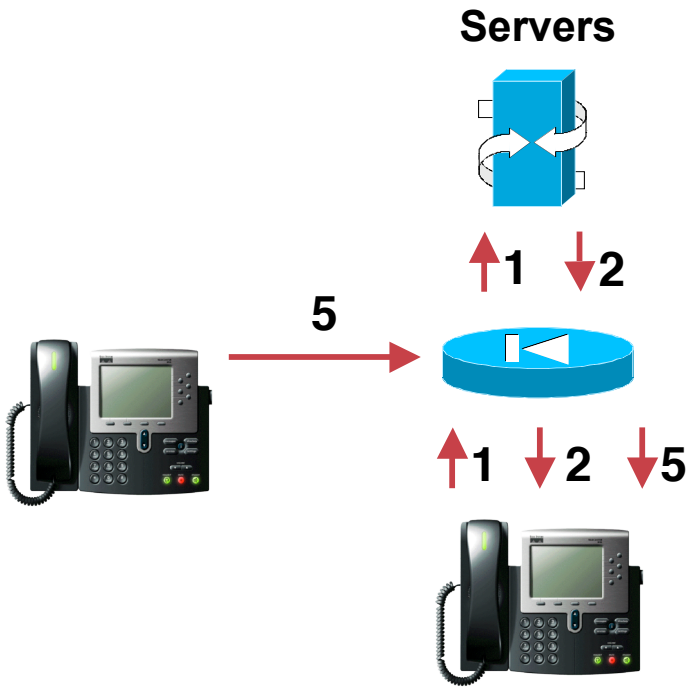

1. What's my public IP address?
2. It is a.b.c.d
3. Tell server when client can receive data
4. Server tells client where to send data
5. Client sends data directly to other client

How STUN (RFC 3489) works

- **Bob pings the STUN server to discover the NAT's public IP address and create a forwarding in the NAT.**
- **Bob then tells this address to Alice.**

1. Bob sends packet to stun server
2. NAT maps packet to be from 1.2.3.4:5555
3. STUN replies and says address packet came from is 1.2.3.4:5555
4. NAT forwards to Bob
5. Bob tells Alice to send to 1.2.3.4:5555 and sends a packet to where Alice will send from
6. Alice sends to 1.2.3.4:5555
7. NAT forwards to Bob

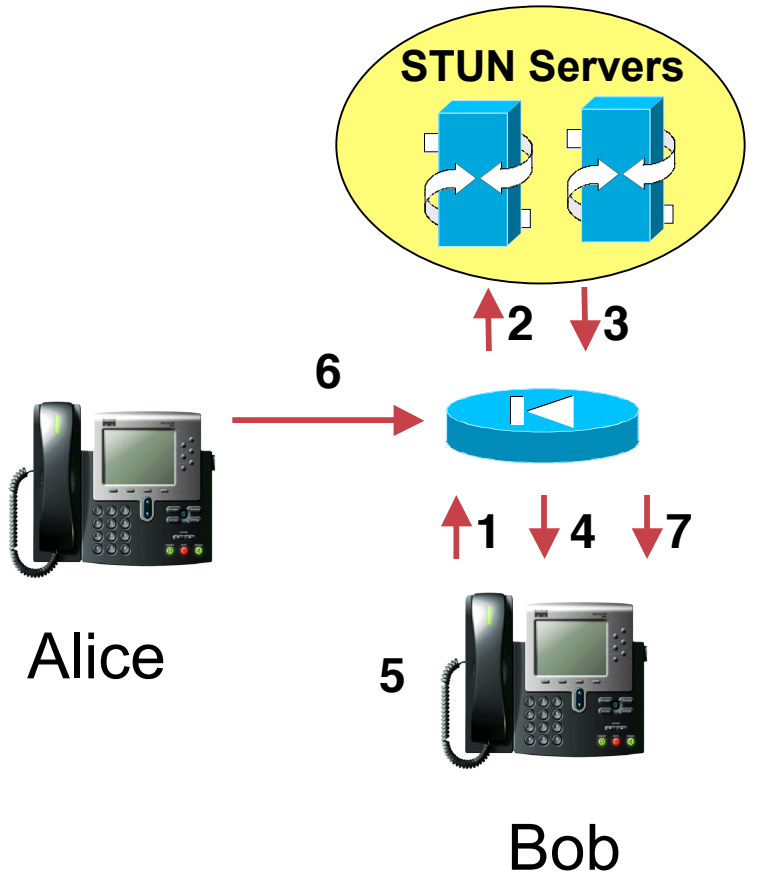

When does STUN work?

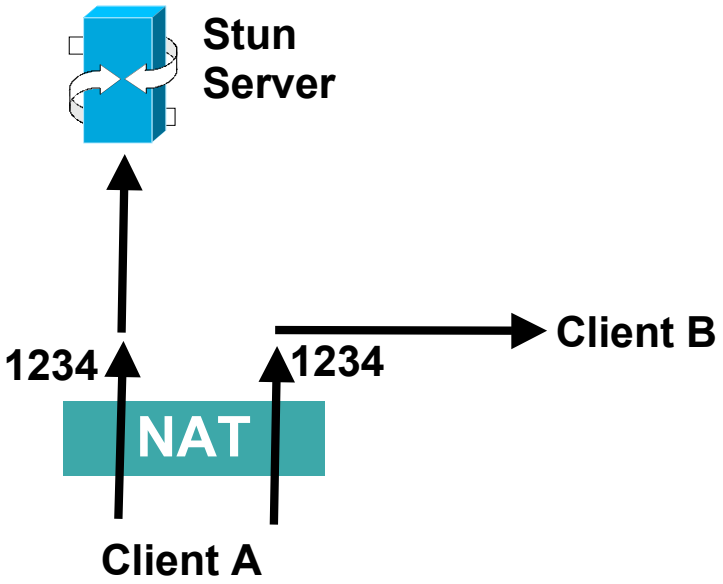

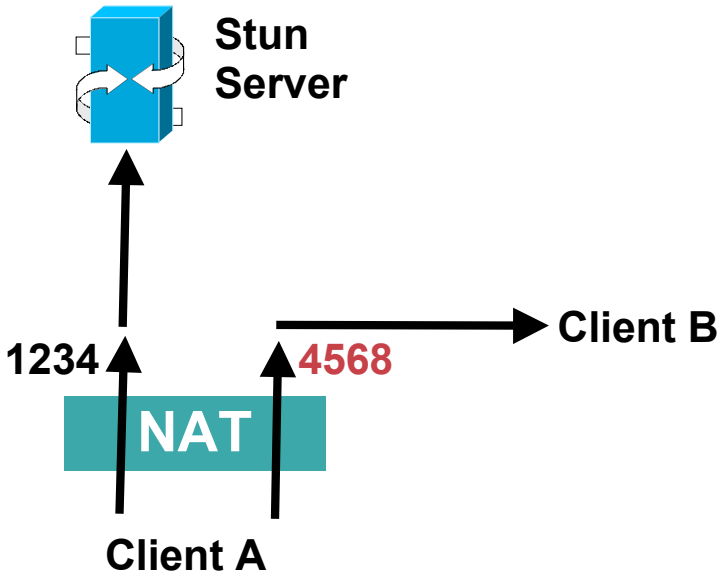

- Echo server works when NAT binding is endpoint independent

- Echo server does NOT work when ports change
- This is bad

Types of NATs: Full Cone

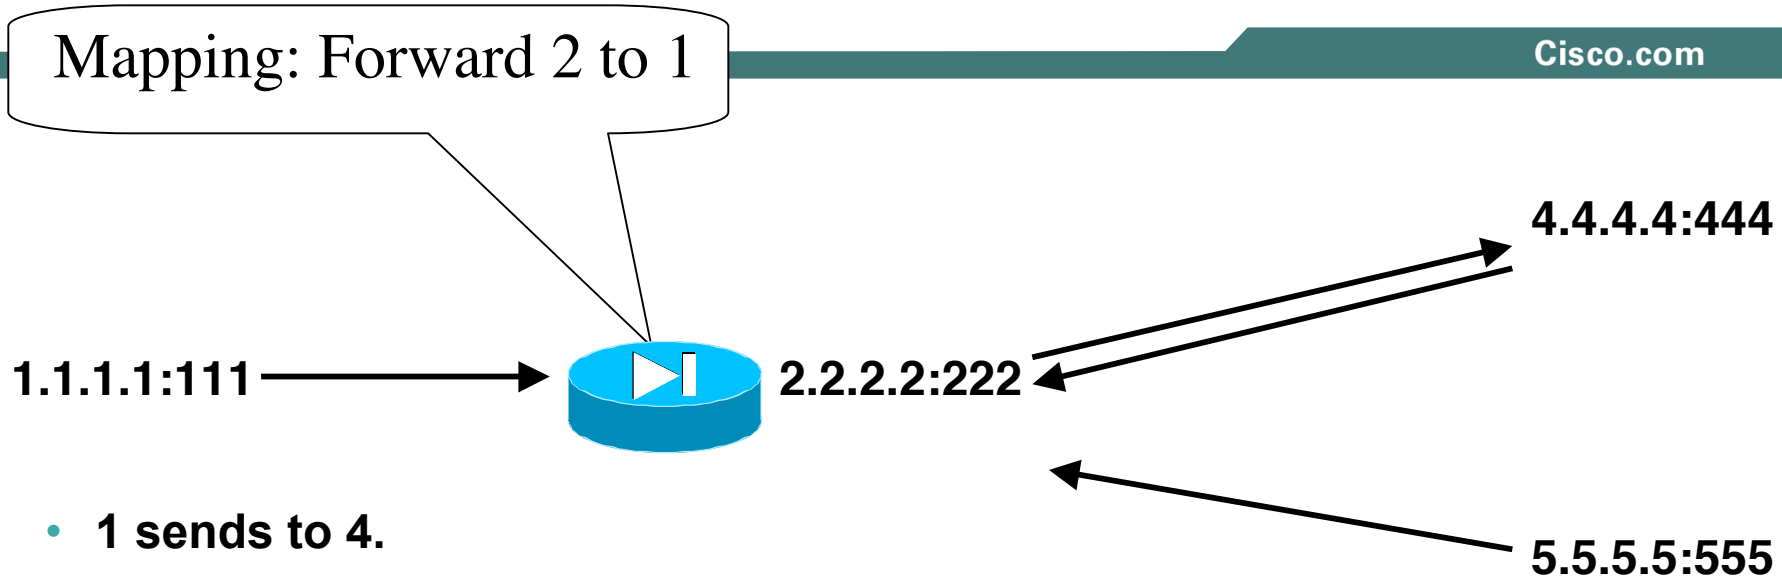

- 1 sends to 4.
- The NAT creates a mapping and forward from 2 to 4 and sends the packet to 4 from 2
- Now any packets that arrive at 2 are forwarded to 1
- Both 5 and 4 can send a packet to 2 and have it forwarded to 1
- Works with STUN

Types of NATs: Address Restricted

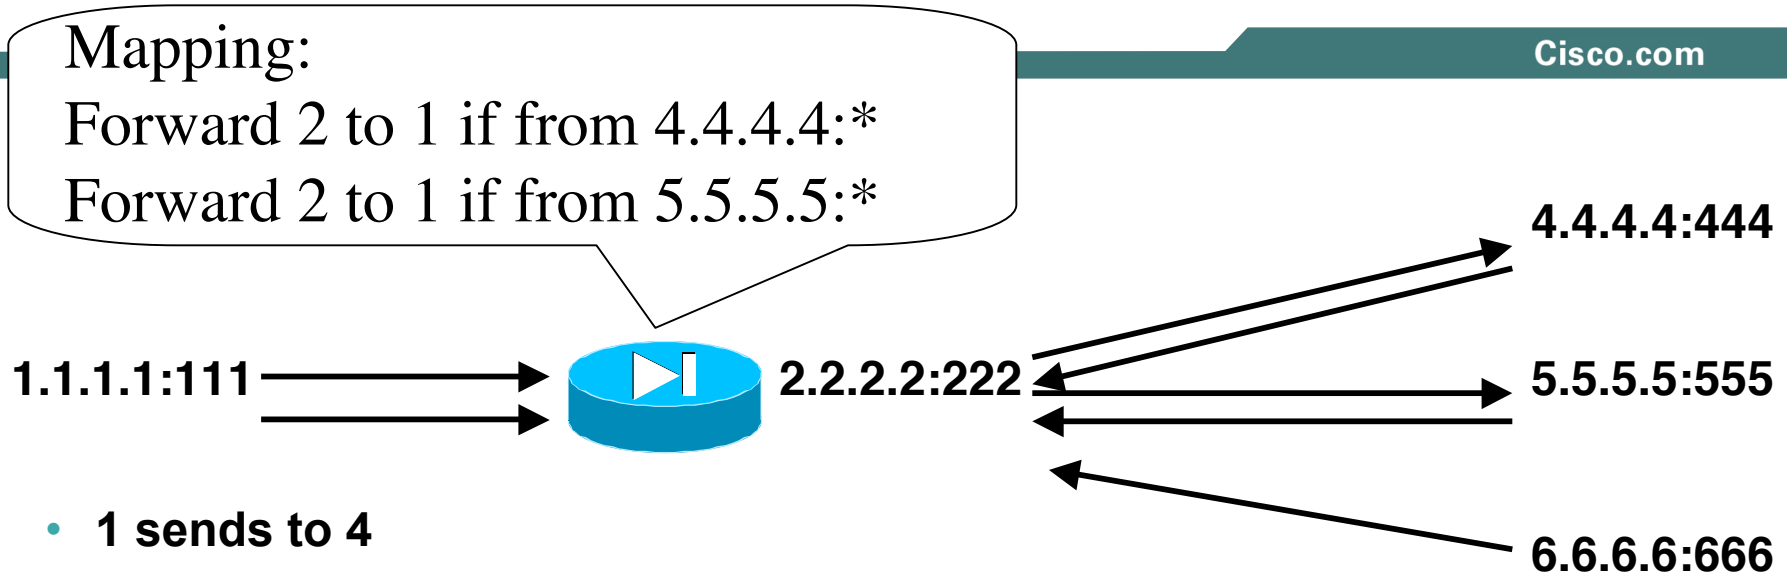

- 1 sends to 4
- The NAT creates a mapping and forward from 2 to 4 and sends the packet to 4 from 2
- 1 sends to 5 and NAT creates similar binding
- Now any packets that arrive at 2 from 4 or 5 are forwarded to 1.
- Packets from 6 get dropped because 1 never sends to 6
- Most Restricted NATs are port restricted, not address restricted
- Works with STUN - can send RTP from any port

Types of NATs: Port Restricted

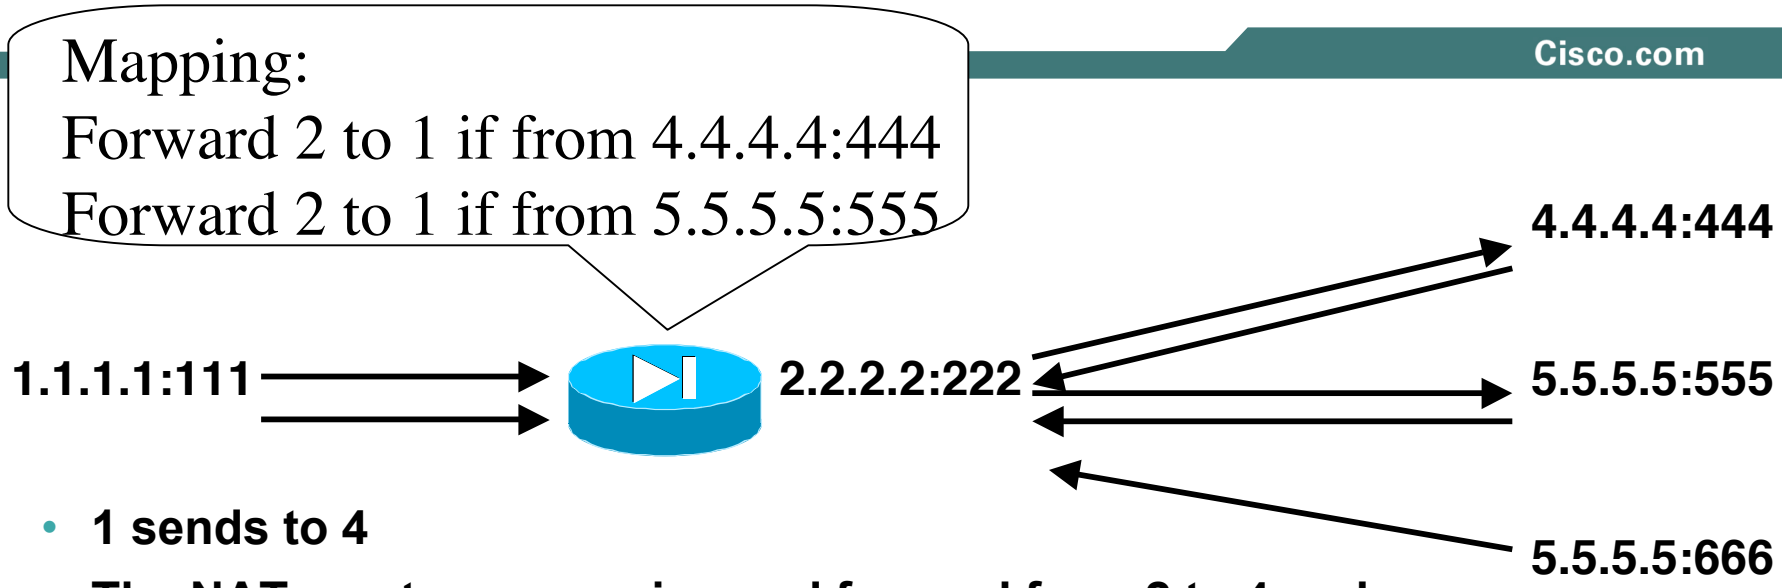

- 1 sends to 4
- The NAT creates a mapping and forward from 2 to 4 and sends the packet to 4 from 2
- 1 sends to 5 and NAT creates similar binding
- Now any packets that arrive at 2 from 4 or 5.5.5.5:555 are forwarded to 1
- Packets from 5.5.5.5:666 get dropped because 1 never sends to 5.5.5.5:666
- Works with STUN - must send & receive RTP from same port

Types of NATs: Symmetric

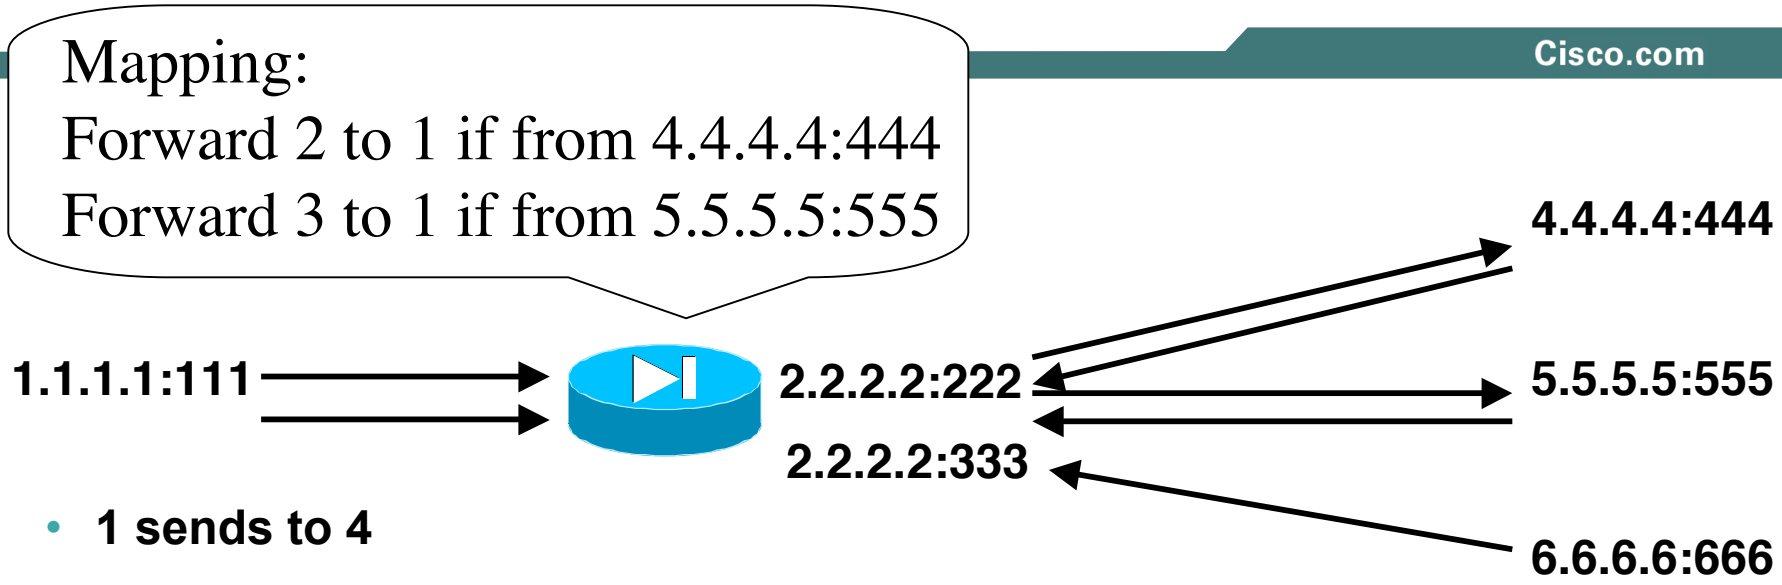

- 1 sends to 4
- The NAT creates a mapping and forward from 2 to 4 and sends the packet to 4 from 2
- 1 sends to 5 and NAT creates mapping from new port 3 to 5
- Now any packets that arrive at 2 from 4 or at 3 from 5 are forwarded
- Packets from 6 get dropped because 1 never sends to 6
- Does NOT work with STUN - needs TURN or other media relay

Security Implications

- Endpoint independent bindings do not change the security properties of NATs
- NATs can accept packets from anyone, or they can decide to only accept packets only from computers to which they have sent a packet to (reciprocal)
- Either way, NATs should have endpoint independent binding

Hairpin Media

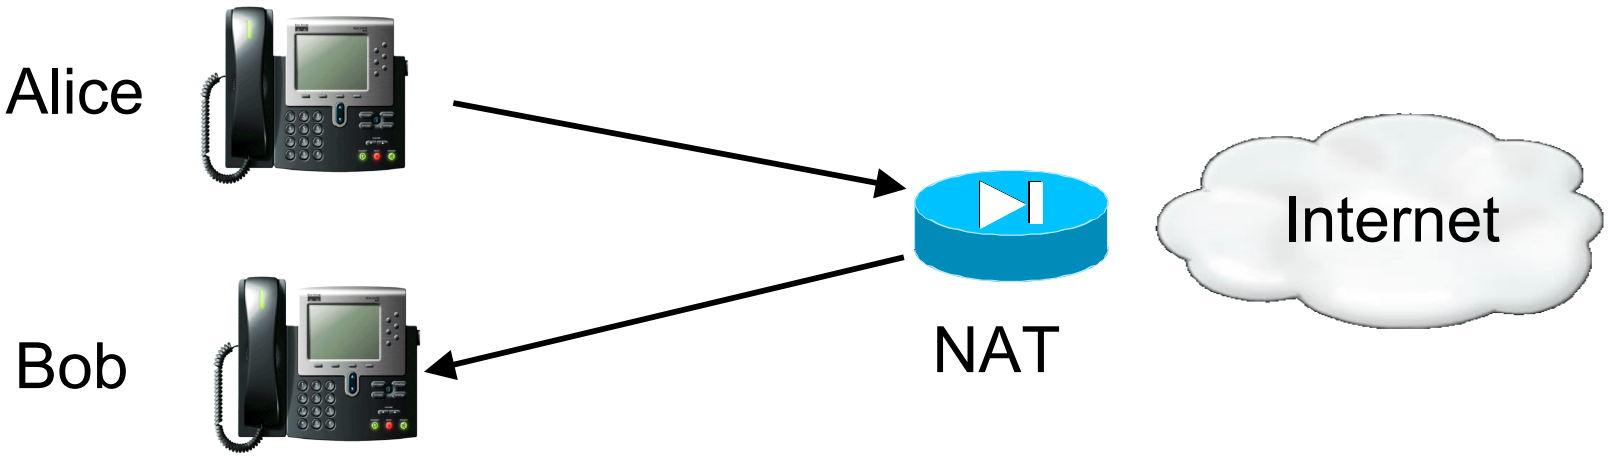

- **Happens when both clients are behind the same NAT**
- **NAT must send data from client A to the NAT's public IP where it loops back to client B**

Survey of NATs (2004 Q1)

			Type	Hairpin	
Apple	Air Base Station	V5.2	C	Y	
DLink	704p	2.61 build 2	C	Y	OK
Dlink	DI-804	.30, Tue,Jun 24 2003	C	Y	
Netgear	RP614	4.00 April 2002	C	Y	
Belkin	F5D5321	V1.13	R	N	
DLink	DI 604	2.0 Jun 2002	C	N	
Linksys	BEFSR81	2.42.7.1 June 2002	R	N	
Linksys	WRV54G	2.03	R	N	OK but no hairpin without ICE
Microsoft	MN-700	02.00.07.0331	C	N	
Netgear	FVS318	V1.4 Jul. 15 2003	R	N	
SMC	7004ABR	V1.42.003	R	N	
US Robotics	USR8003	1.04 08	C	N	
Airlink	ASOHO4P	V1.01.0095	R-U	N	
Linksys	WRT54G	1.42.2	R-U	N	
SMC	2804WBRP-G	v1.00 (Oct 14 2003 18:20:25)	R-U	Y	
Toshiba	WRC-1000	1.07.03a-C024a	R-U	N	
ZOT	BR1014	Unknown	R-B	N	OK one phone
Hawkings	FR24	6.26.02h Build 0047 L:02	R-B	Y	
Network Everyw	NR041	Version 1.2 Release 03	R-B	Y	
Network Everyw	NR041	Version 1.0 Release 10	S	N	NO

Percentage Deployment of NAT in US

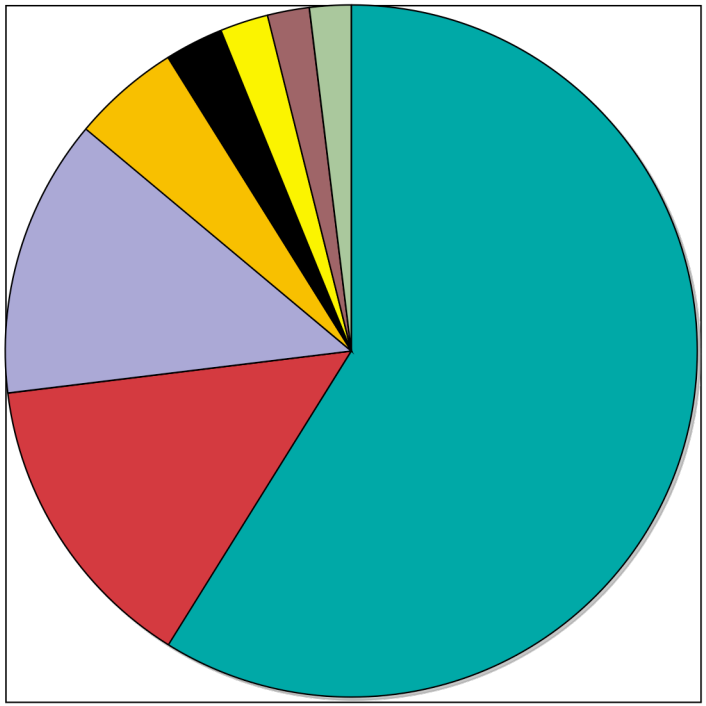

Linksys	59%
Netgear	14%
DLink	13%
Belkin	5%
Microsoft	3%
SMC	2%
Siemens	2%
Others	2%

- **Data from AOL study**

  Most data is this space is not public.
- **Consistent with other reports**
- **Fairly US Centric - not accurate for Asia**

(source http://www1.ietf.org/mail-archive/web/midcom/current/msg03507.html)

IETF Update

- **IETF has not encouraged NATs**
  - **IPv6 is a better solution**
  - **It will be many years before IPv6 is fully deployed**
- **At last IETF, a BOF on NAT Behavior was held**
  - **Plan to form working group to create formal RFC addressing best current practices around NAT behavior**
- **Read draft-audet-nat-behave-00.txt**

Key BEHAVE Draft Recommendations

- Bindings are endpoint independent
- UDP binding expiry time > 2 minutes
- Have SIP ALGs off by default
- Support Hairpin media
- Read draft-audet-nat-behave-00.txt

Recommendations & Predictions

- Online Gaming & VoIP will drive the NATs that service providers recommend and support
- Most vendors will build NATs that work this way
- Ensure that your NAT:
  - provides endpoint independent port translation
  - behaves consistently
  - can hairpin media
- This is no more work that doing the wrong thing
- Follow the advice of IETF drafts

  Applications will follow the advice of the drafts