A Distributed Waypoint Service Approach to Connect Heterogeneous Internet Address Spaces

T.S. Eugene Ng, Ion Stoica, Hui Zhang
Carnegie Mellon University
{eugeneng,istoica,hzhang}@cs.cmu.edu

Evolution of Internet Addressing

- Size of IPv4 address space is a fundamental limitation
  - Only 32-bit and poor utilization
- Solutions
  - Upgrade and deploy networks using IPv6
  - Deploy networks using reusable-IPv4 (or private-IPv4) addresses
- IPv4, IPv6, reusable-IPv4 address spaces will coexist in the Internet for the foreseeable future
- **Important goal**: Maintain universal connectivity in this environment with heterogeneous address spaces

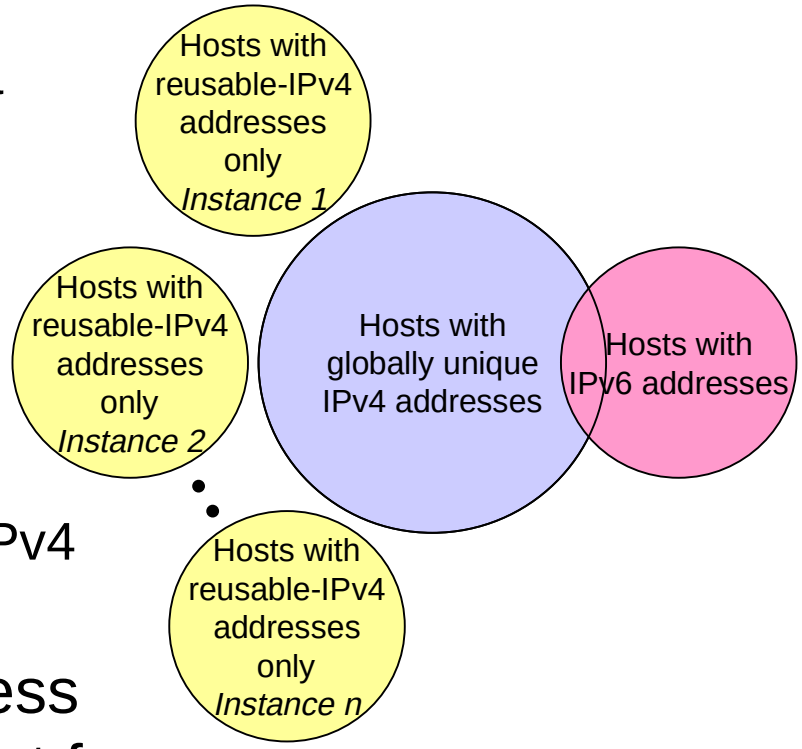

Wait, Isn't IPv6 Going to Prevent This?

- "Dual-stacking" makes IPv6 hosts simultaneously act as IPv4 hosts
  - No visible heterogeneity
  - Every IPv6 host must consume an IPv4 address
- May not be feasible
  - IPv4 addresses are difficult and expensive to obtain
  - IPv4 address space may be exhausted before transition to IPv6 is complete
  - Additional implementation complexity

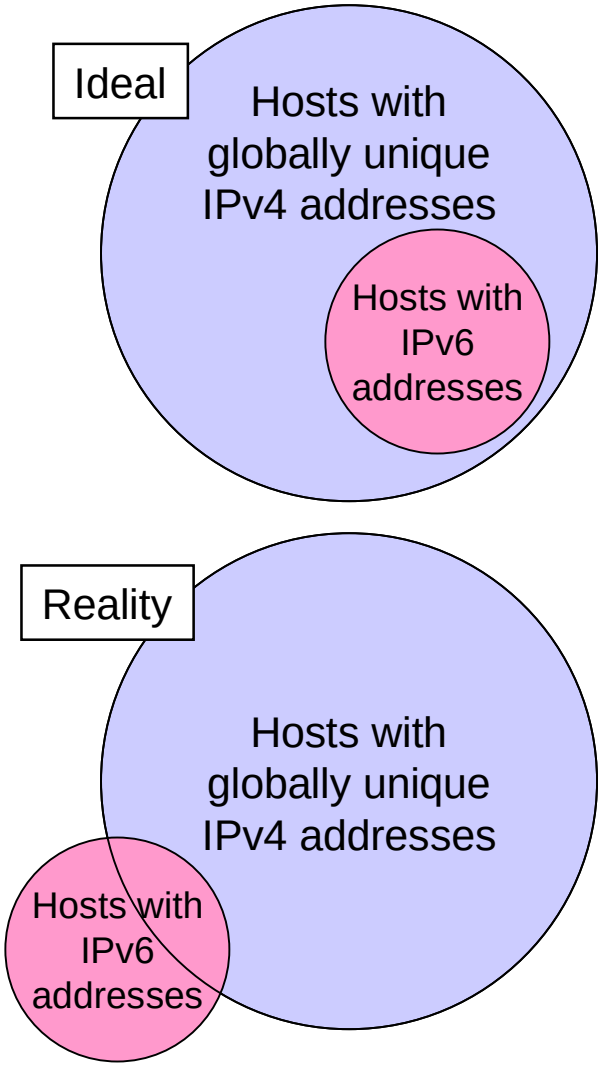

Maintaining Universal Connectivity

		Responder		
		IPv4	IPv6	Reusable-IPv4
Initiator	IPv4	Trivial	(a) Hard	(b) Hard
	IPv6	NAT-PT	Trivial	Reduces to (b)
	Reusable-IPv4	NAT	Reduces to (a)	Reduces to (b)

- An initiator of any address space type must be able to initiate a connection to a responder of any address space type
- Case (a) and (b) are hard because **the responder fundamentally cannot be addressed by the IPv4 initiator**
  - Consider case (b) only for simplicity; results apply directly to case (a)

Network Address Translator Is Insufficient

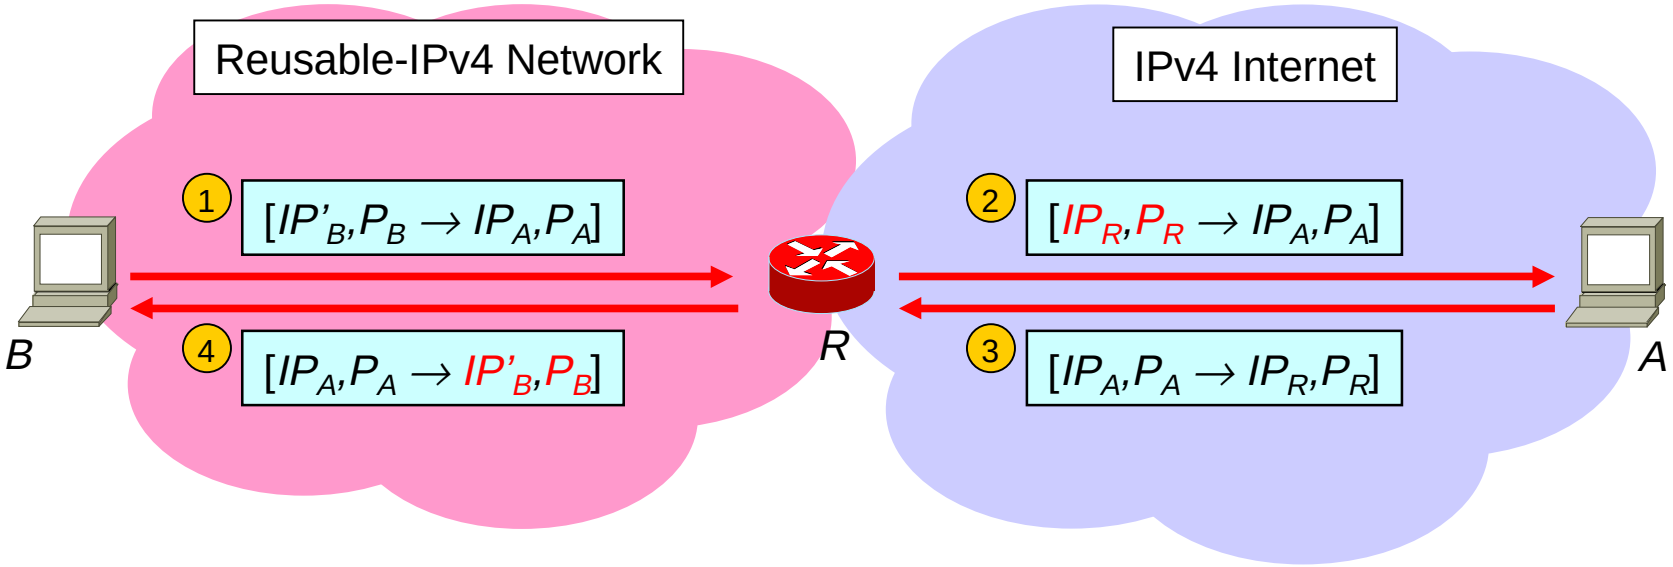

- Works well when reusable-IPv4 hosts initiate connections to IPv4 hosts (out-bound)
  - Bind reusable-IPv4 hosts to arbitrary client port numbers
- But only **one** reusable-IPv4 host is reachable by IPv4 hosts (in-bound)
  - IP_R can be bound to only one reusable-IPv4 host

Potential Solutions for IPv4 to Reusable-IPv4 Connectivity

- Static port binding
  - Bind service ports on NAT gateway to reusable-IPv4 hosts
  - Poor connectivity
- Encapsulation
  - IPv4 initiator generates encapsulated packets
  - Needs sophisticated extensions to DNS and all IPv4 hosts
- Naming layer (e.g. HIP [Moskowitz ‘00])
  - The real destination can be encoded in a new naming layer
  - Needs sophisticated extensions to DNS and all IPv4 hosts
- Application layer naming (e.g. HTTP)
  - The real destination can be encoded in higher level protocols
  - All applications need to be rewritten

The Case for Waypoints

- All existing solutions either provide poor connectivity or are difficult to deploy
- Fundamental question: How should the Internet be evolved to adapt to new requirements?
  - Change all routers simultaneously?
  - Change all end systems simultaneously?
- Waypoints provide a perfect platform to adapt the Internet
  - Shared 3rd-party network agents
  - Independent non-intrusive deployment
  - Functionality immediately benefits global Internet

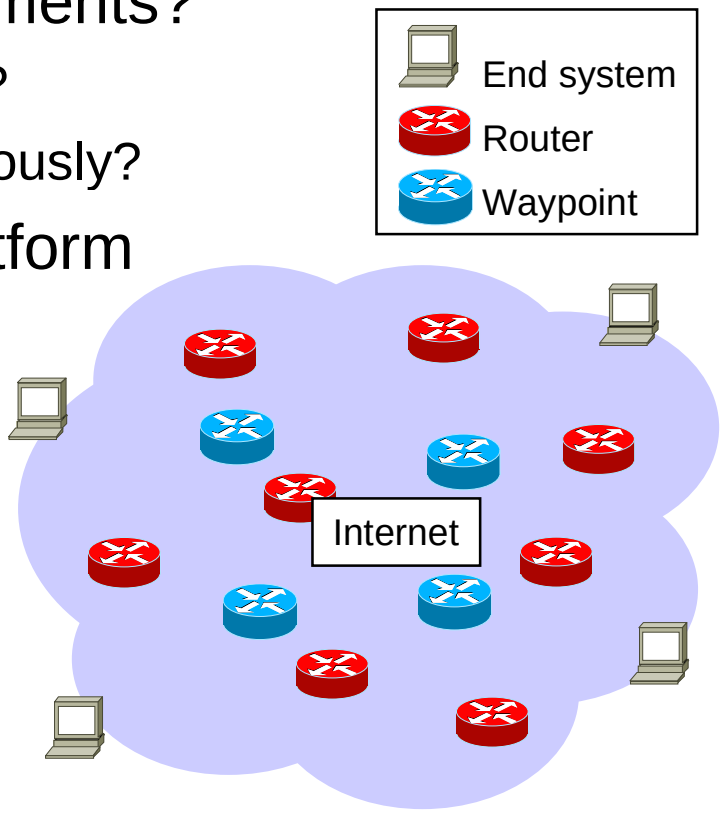

AVES: A Distributed Waypoint Service

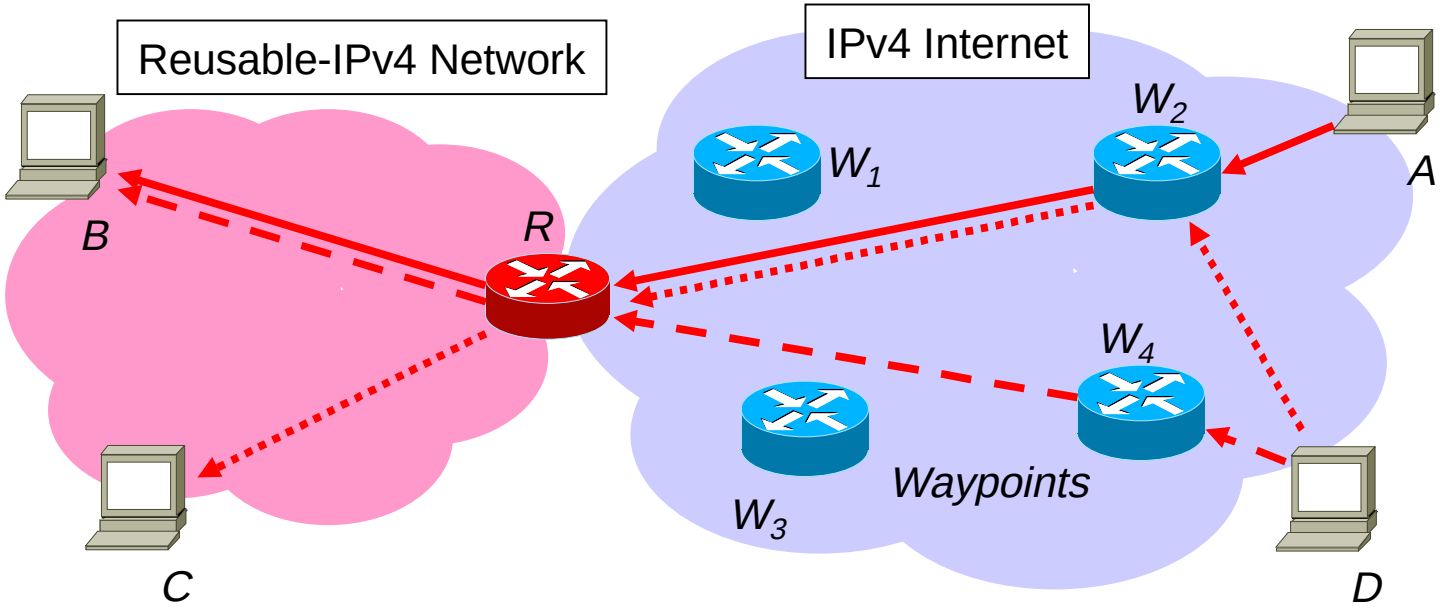

- Virtual expansion of IPv4 address space
  - Small number of IPv4 waypoints virtually represent large number of reusable-IPv4 hosts (e.g. W_2 represents *B* and *C*)
  - Heterogeneity is hidden from IPv4 hosts
- Transparent AVES access via DNS name resolution

AVES Data Path Operations

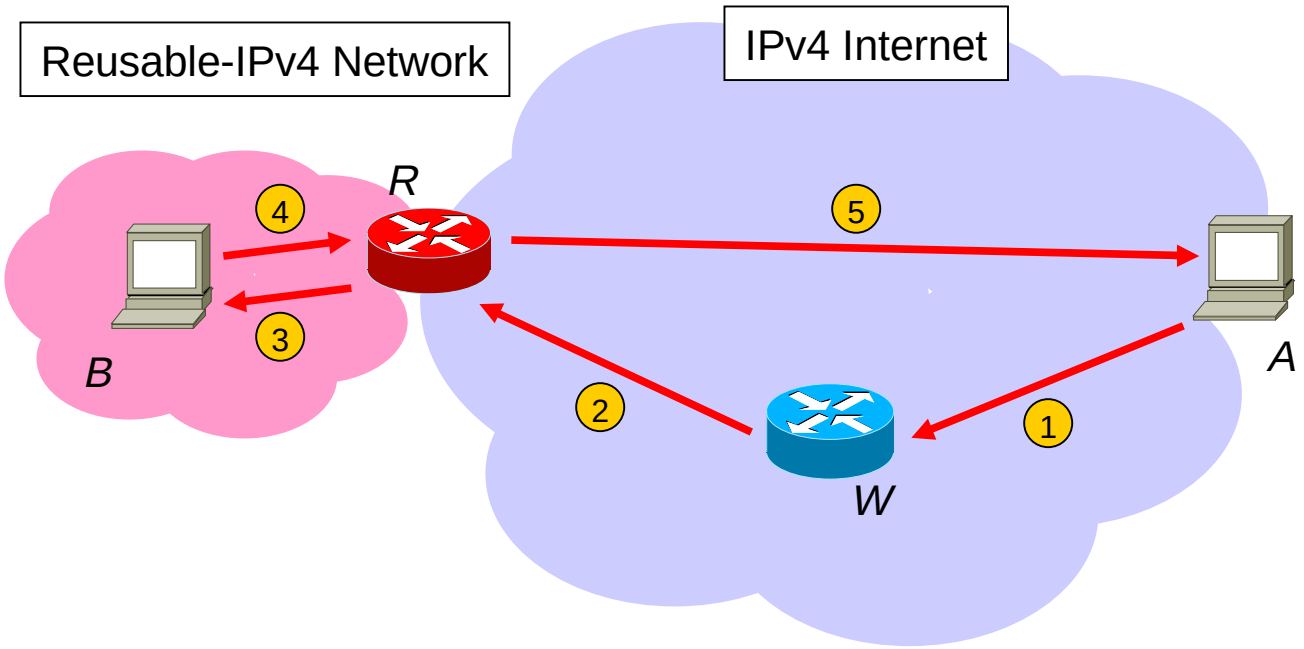

Step	Packet sent
1	[$IP_A \rightarrow IP_W$]
2	[$IP_W \rightarrow IP_R$ [$IP_A \rightarrow IP'_B$]]
3	[$IP_A \rightarrow IP'_B$]
4	[$IP'_B \rightarrow IP_A$]
5	[$IP_W \rightarrow IP_A$]

AVES Control Path Operations

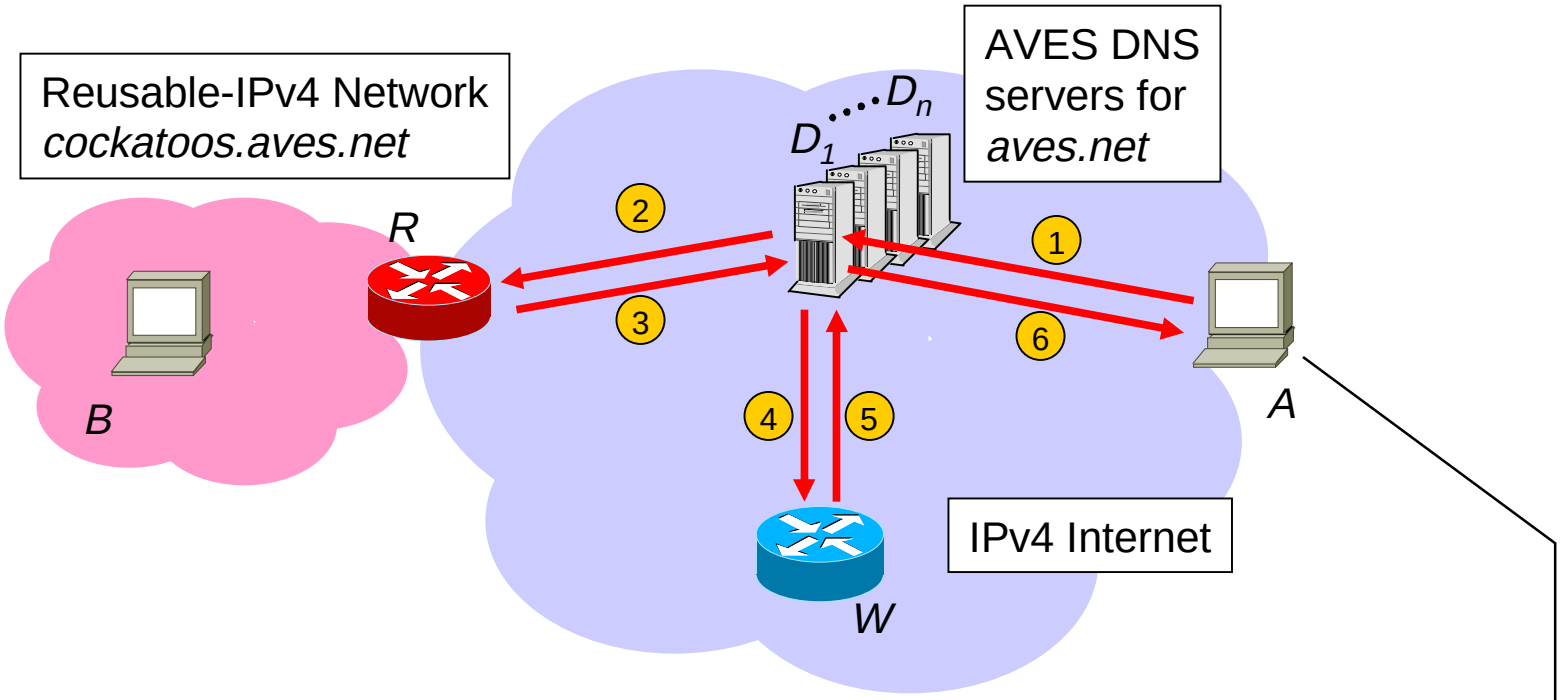

Step	Action
1	DNS query for *B*
2	DNS query for *B*
3	DNS reply for *B* (IP'_B)
4	SETUP message (IP_A, IP_R, IP'_B)
5	ACCEPT message
6	DNS reply for *B* (IP_W)

Note that *A* is configured to send *aves.net* queries directly to an AVES DNS server

AVES Connectivity Properties

- Using *N* IPv4 addresses, every IPv4 host can simultaneously reach up to *N* reusable-IPv4 hosts
- Every reusable-IPv4 host can be reached by an unlimited number of IPv4 hosts

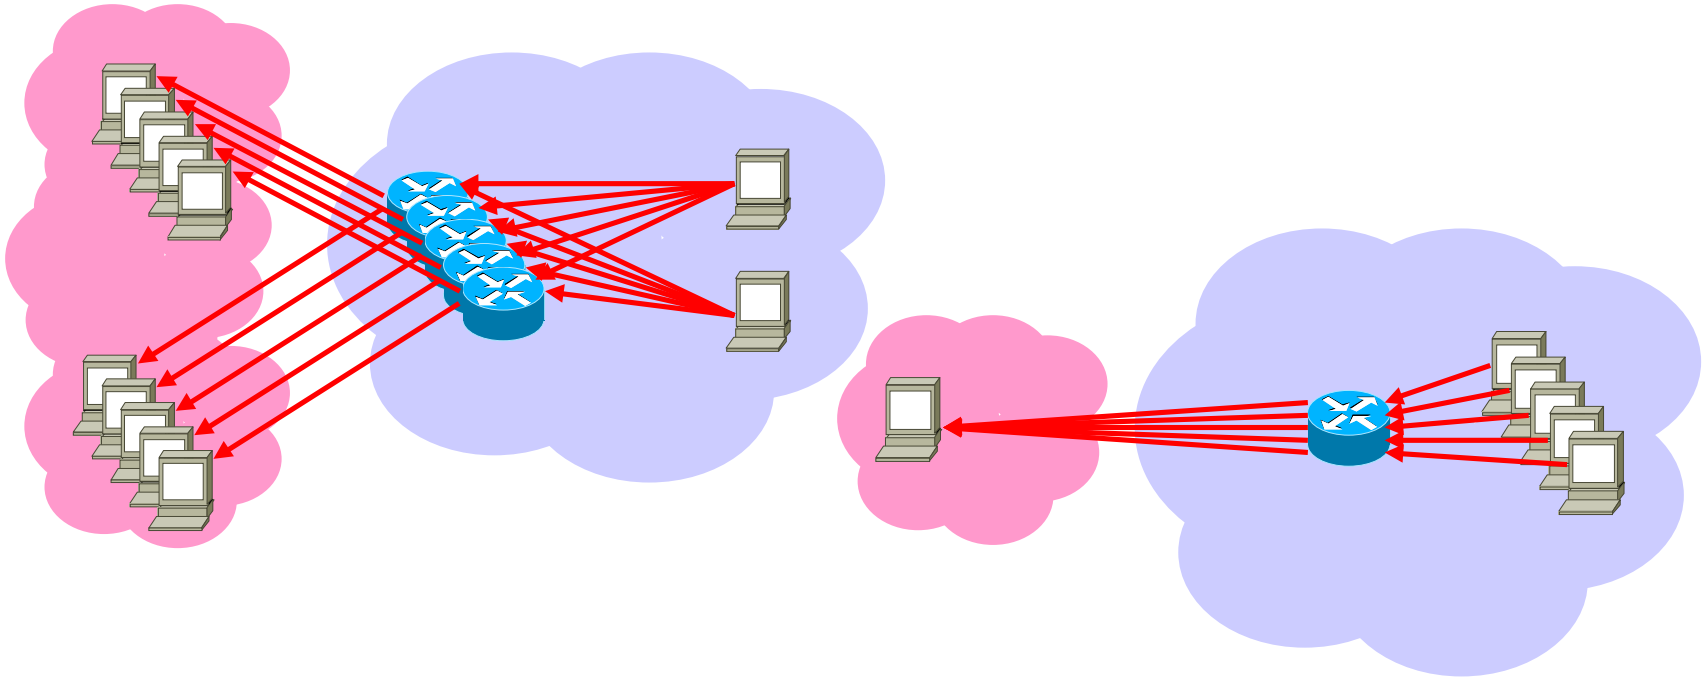

Summary

		NAT	NAT + Static port binding	NAT + Encap-sulation	NAT + Naming layer	NAT + App layer naming	NAT + AVES
Out-bound connectivity	**For each initiating reusable-IPv4 network**	65000 TCP connections to each responding (IP, port)	→	→	→	→	→
	For each responding IPv4 host	No additional restriction	→	→	→	→	→
In-bound connectivity	**For each responding reusable-IPv4 network**	One host reachable at a time	One host per port reachable at a time	All hosts reachable	All hosts reachable	All hosts reachable	All hosts reachable
	For each initiating IPv4 host	No additional restriction	→	→	→	→	Up to N reusable-IPv4 hosts
Deployability (changes required)		Base case	NATs	NATs + DNS + All end system software	NATs + DNS + All end system software	NATs + All apps	NATs + DNS config

Future Work

- Intelligent waypoint selection
  - Minimize latency
  - Maximize throughput
- Automatic waypoint discovery
  - Eliminate needs for manual configuration
  - Increase robustness
- Application of AVES to virtual private networking
  - When site-to-site tunneling cannot be deployed
  - When address assignment conflicts exist among sites