

Enhanced IP

Dr. William Chimiak, Stephen Janansky,
Sam Patton

Design Choices

- Evolve the Internet with additions to IPv4
- Evolve the Internet through deployment of a new protocol (IPv6)

Talk Outline

- 1992-1994: A Time Period of Significance to Today
- IPv6 Challenges
- Enhanced IP

Some Important R&D 1992-1994

- Nov 1992 - publication of EIP as RFC 1385
- Jan 1993 - original ACM SIGCOMM paper describing NAT
- 1992 - 1994 - Robert Ullman's IPv7, TP/IX, and CATNIP
- Dec 1993 - RFP by IPng working group

Some Important R&D

1992-1994

- Mar 1994 - Brian Carpenter's AEIOU
- Jul 1994 - SIPP is chosen by the IPng Directorate to become IPv6 after changing address size from 64-bit to 128-bit

What Happened with:

- EIP
- IPv7, TP/IX
- AEIOU

Changes in router architecture since 1992

- What obstacles would IPv7, AEIOU, and EIP face today?
- Introduction of the fast path and slow path concept in router design around the time of the Cisco 7000 in the late 1990's and early 2000's.
- IP options are kicked up to the slow path for processing and routing.

What Obstacles are there today?

- Layering of NAT (Nat444, Nat64, DS-Lite)
- UC Berkeley paper on IP option usage
 - 160 machines used to traceroute to each other. 139 machines usable. $160 \times 139 = 22240$ paths
 - Traceroutes for 21,051 of the possible pairs or 94.65%
 - Sent standard ICMP traceroutes, and traceroutes containing NOP, Timestamp, and Record Route options.
 - Their conclusion: between 85% and 92% of drops occurred in edge ASs with a majority of the drops in a minority of ASs

IPv6 Challenges

The IPv6 Peering Challenges

Route Table Source	IPv4 RIB	IPv6 RIB
bgpmon.net [1]	477101	10654
potaroo.net APNIC R&D [2]	409381	8683
potaroo.net Route Views Oregon [3]	421241	8759

[1] bgpmon.net IPv4 and IPv6 weathermap prefixes on 4/20/2012

[2] Fri, 20 Apr 2012 18:13:38 GMT according to <http://bgp.potaroo.net/index-bgp.html>

[3] Fri, 20 Apr 2012 18:13:38 GMT according to <http://bgp.potaroo.net/index-bgp.html>

[4] We would like to get he.net numbers since they work towards IPv6 peering everywhere. Current measurements we found for RIB on he.net focus on AS's that have IPv4 and IPv6 support versus routing table entry counts.

IPv6 Implementation Challenges

- 2) The CPE Problem
 - According to Brian Carpenter's IPv6 Task Force talk: Billing Systems, Handsets, management interfaces and systems, DSLAMs, Routers, Traffic mgmt boxes, load balancers, VPN boxes, SIP boxes, firewalls
- 3) End to End principle violated to support LTE needs
 - NAT64 to support IPv6 LTE subscribers reachability to legacy IPv4
 - RFC 6586

Enhanced IP

Enhanced IP Outline

- Protocol Details
- Working code between University of Maryland and University of Delaware over Internet 2 link (over 7 router hops)
- Demonstrated Enhanced IP running in an LTE system (details later)
- Demonstrated Enhanced IP running on an Android (Alpha state)
- The transition plan
- Positive and Negative viewpoints

EnIP Protocol Details

- EnIP demonstrates an alternate way to solve the address depletion problem
- Creates a 64-bit IP network as an overlay on the IPv4 Internet


EnIP Protocol Details

- Routing of Enhanced IP packets are forwarded using the IPv4 routing table
- EnIP uses IP options to carry more address bits
- EnIP continues to use ARP, DHCPv4, and IPv4 routing protocols. Packets are routed as IPv4 datagrams so less disruption in these areas


EnIP Protocol Details

- **65.127.221.1.10.3.3.2**
- All Enhanced IP addresses have a site address and a host address
- **Site address: 65.127.221.1**, is used to route packets over the public Internet to a router/NAT that is aware of Enhanced IP packet format. This would generally be a public IPv4 address.
- **Host address: 10.3.3.2**, used to route packets to a node behind the router/NAT that has the outside address of 65.127.221.1

Enhanced IP Network



Enhanced IP Network



IPv4 Header

Version	IHL	ToS	Total Length	
Identification			Flag	Fragment Offset
Time to Live	Protocol		Header Checksum	
Source Address				
Destination Address				

Enhanced IPv4 Header

Version	IHL	ToS	Total Length		
Identification		Flag	Fragment Offset		
Time to Live	Protocol		Header Checksum		
Source Address					
Destination Address					
Option ID	Option Length	ESP	EDP	Reserved	
Extended Source Address					
Extended Destination Address					

Enhanced IPv4 Header

Version	IHL	ToS	Total Length		
Identification		Flag	Fragment Offset		
Time to Live	Protocol		Header Checksum		
Source Address					
Destination Address					
0x9a	Option Length	ESP	EDP	Reserved	
Extended Source Address					
Extended Destination Address					

Enhanced IPv4 Header


Version	IHL	ToS	Total Length		
Identification		Flag	Fragment Offset		
Time to Live	Protocol		Header Checksum		
Source Address					
Destination Address					
0x9a	12	ESP	EDP	Reserved	
Extended Source Address					
Extended Destination Address					

Enhanced IPv4 Header

Version	IHL	ToS	Total Length		
Identification		Flag	Fragment Offset		
Time to Live	Protocol		Header Checksum		
Source Address					
Destination Address					
0x9a	12	ESP	EDP	0	
Extended Source Address					
Extended Destination Address					

Enhanced IPv4 Header


Version	IHL	ToS	Total Length		
Identification		Flag	Fragment Offset		
Time to Live	Protocol		Header Checksum		
Source Address					
Destination Address					
0x9a	12	ESP	EDP	0	
Extended Source Address					
Extended Destination Address					




Version	IHL	ToS	Total Length		
Identification		Flag	Fragment Offset		
Time to Live	Protocol	Header Checksum			
10.1.1.2					
65.127.221.2					
0x9a	12	0	1	0	
255.255.255.255					
10.3.3.2					



Version	IHL	ToS	Total Length		
Identification		Flag	Fragment Offset		
Time to Live	Protocol	Header Checksum			
10.1.1.2					
65.127.221.2					
0x9a	12	0	1	0	
255.255.255.255					
10.3.3.2					

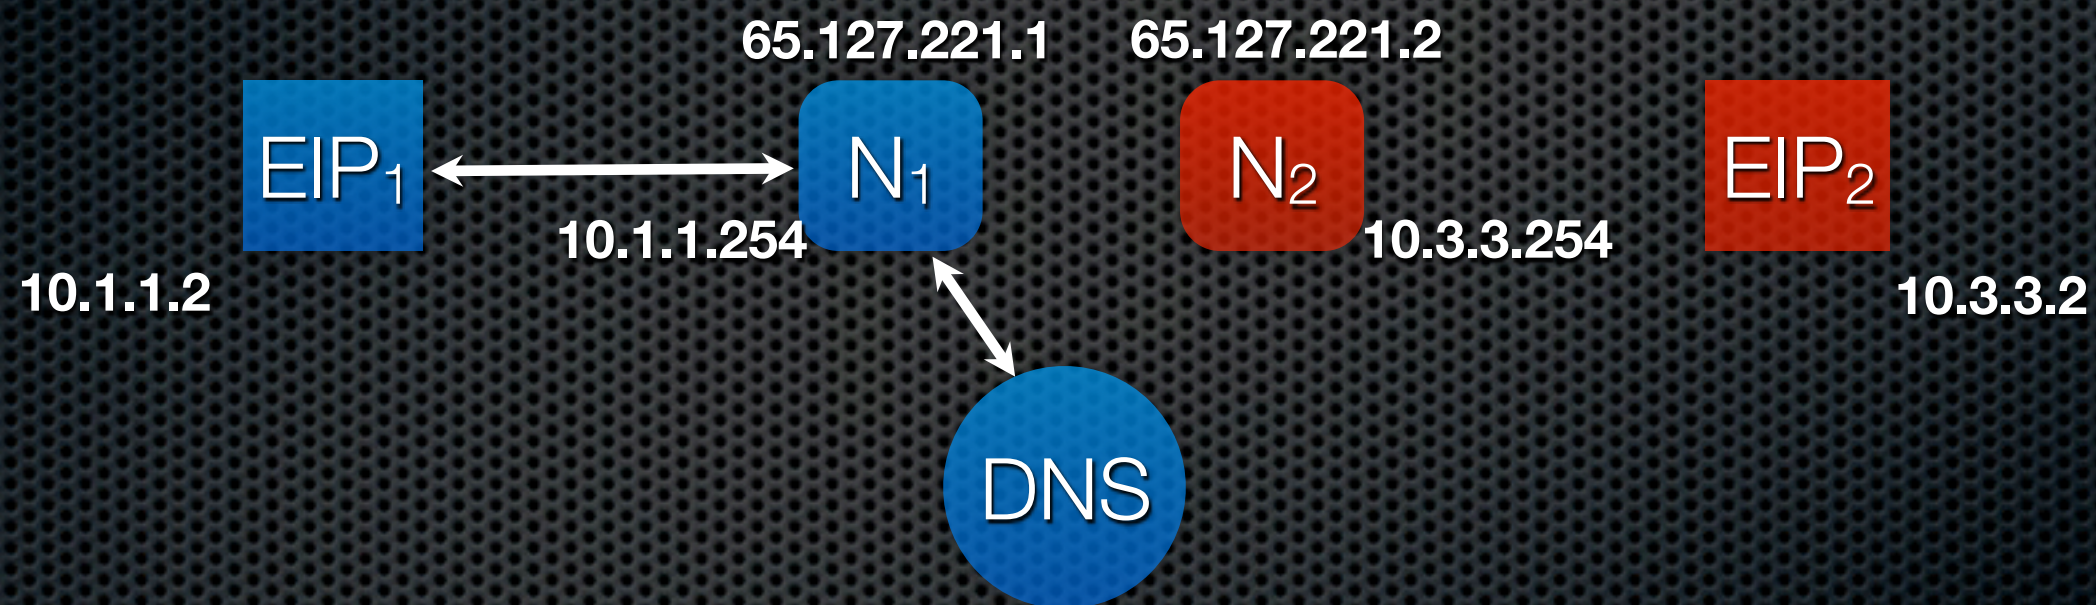


Version	IHL	ToS	Total Length		
Identification		Flag	Fragment Offset		
Time to Live	Protocol	Header Checksum			
65.127.221.1					
65.127.222.2					
0x9a	12	1	1	0	
10.1.1.2					
10.3.3.2					



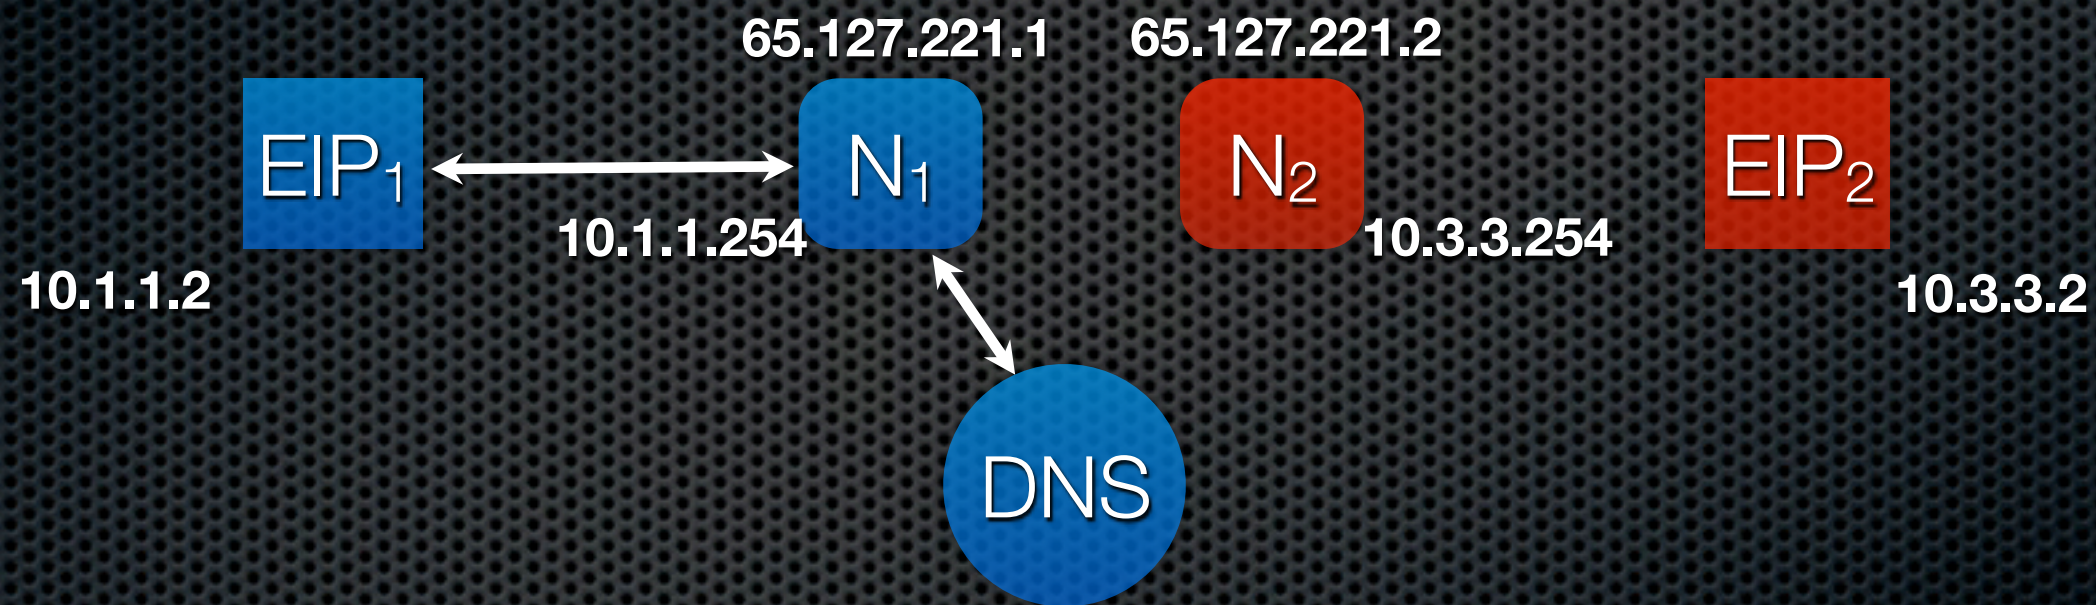
Version	IHL	ToS	Total Length		
Identification		Flag	Fragment Offset		
Time to Live	Protocol		Header Checksum		
65.127.221.1					
10.3.3.2					
0x9a	12	1	0	0	
10.1.1.2					
0.0.0.0					

DNS: Using AAAA to lookup 64 bits




- EIP_1 sends a AAAA request for `eip2.somesite.com` and receives back `2001:0101:417F:DD02:0a03:0302::0`

DNS: Using AAAA to lookup 64 bits



- `2001:0101:417F:DD02:0a03:0302::0` is really `65.127.221.2.10.3.3.2`

EnIP Proposed Architecture



Enhanced IP: Transition

Transition Plan

- Phase 1: Router vendors start pushing fast path upgrades for selected IP options, propose that mobile device and network industry are the first to upgrade.
- Phase 2: Content Providers may selectively upgrade servers.
- Phase 3: Home and enterprise users are upgraded as needed or not at all. For example, a Mobile provider may wish to offload voice calls from cell towers to a wired connection within a home via a home's wifi network.

Making selected options fast

- Some router designs have moved towards flexibility in the fast path (e.g. BBN 50 Gbps router design used microprocessor and used 85 instructions to forward a packet in the common case) This was 1998.
- Would passing a specific IP option in the fast path bring this case to ~90 instructions?
- ASICs, FPGAs, NPUs, CPUs. Can we upgrade the fast path?
- Looking at another hardware refresh cycle.

The Future....

- NAT continues to be used to expand the Internet
- Layers of NAT become ubiquitous because of IP scarcity (DS-Lite, Nat444, etc)
- The IPv6 designers all retire
- IPv6 continues to have CPE and peering issues

Or.....

- Multiple levels of NAT are deployed and deployment of IPv6 and EnIP in these networks becomes impractical.
- A workable solution to IPv6 peering and CPE problems is found.
- Someone comes up with IPTg. The argument of extending IPv4 versus new IP datagram structure takes us back to issues being considered during 1992-1994 except this time the focus would be towards extension research.

Security Considerations

Security Considerations

- Similarly to IPv6, Enhanced IP means increased connectivity for IP addresses that were previously unreachable because of NATs.
- Firewall software is not presently aware of Enhanced IP addresses. IPv4 functions still work just not filtering of specific EnIP addresses. This is probably workable in many situations.
- Enhanced IP NAT devices must only forward packets to networks that are directly connected (e.g. 10.0.0.0/8)

EnIP: Positive and Negative

What people might like about EnIP

- The 8 byte address scheme is familiar
- Each 4 byte part of the address is still addressable as an IPv4 node
- Enhanced IP addresses can talk to legacy IPv4 addresses

What people might like about EnIP

- With a single public IP address it's possible to address a 10.0.0.0/8 network
- Logical and simple extension to NAT design
- Avoids the peering problem
- Security models are similar to IPv4
- Implementation requires smaller amount of code

What people might not like about EnIP

- Use of IP options means wide-scale deployment can't start today.
- Legacy IPv4 nodes can still reach an Enhanced IP node through use of existing NAT port forwarding mechanisms. Beauty is in the eye of the beholder.
- Enhanced IP is not 2^{128} . Use of RFC 1918 addresses behind public IPs you already have.
- Will not work through multiple layers of NAT

Enhanced IP: Alpha Implementation


Enhanced IP: Alpha Implementation

- Supports TCP and UDP sockets over 64-bit Enhanced IP
- ICMP is not implemented. More research here.
- Forward lookups for DNS are implemented. (reverse not)
- Enhanced IP NAT extensions are implemented to demonstrate interoperation with IPv4 NAT
- Demonstrated ssh, samba file sharing, and firefox+apache web server over Enhanced IP sockets. No modifications to ssh, scp, samba, firefox, or apache software.

Development and Test Environment

- Gentoo Linux running as Virtualbox guest OS
- 2.6.38 Kernel
- 4 virtual machines: **EIP1, N1, N2, EIP2**

Enhanced IP Network



Enhanced IP code

- Linux OS changes
 - ~200 lines in the kernel, ~500 lines in user-space
- edge device (SOHO router)
 - ~450 line driver, 8 line patch to NAT code
- Linux Utilities
 - ~3500 lines: ping, traceroute, netcat-like program, measurement programs

Conclusion

- In our opinion certain conditions were in place to enable the Enhanced IP research. These conditions were not in place in the 1992-1994 period when IP options were last considered. IPng was the focus.
- Single layer NAT became pervasive
- In 2012, Peering and CPEs are two problem areas in IPv6 deployment. Multiple layers of NAT are one competing solution.

Conclusion

- We came up with the 64-bit format, the IP swapping technique used in the EnIP NAT and found a way to overload the AAAA record to avoid addition of a new DNS record. The EnIP focus is on mobile devices and infrastructure.
- IP options are processed in the slow path. If we want a fall-back plan passing a few IP options in the fast path can buy some design options. Engineering cost would be low. We would probably have to wait.

Conclusion

- Current EnIP R&D focuses heavily on mobile space. In the EnIP model home and enterprise adoption could come later or not at all.
- Peer review is requested. Community interest will determine whether the research is continued.
- Is there community interest in issuing RFP for IPTg focused on extensions to IP? EnIP would be one entry.

Project Info.

- enhancedip at enhancedip.org
- <http://www.enhancedip.org/>

Questions?