



IPv6 Implementation fundamentals for ISPs

NANOG47 / ARIN XXIV
Dearborn, MI

Aaron Hughes, 6connect, Inc.

Unifying Internet Infrastructure

www.6connect.net

AGENDA

- Obtain an IPv6 allocation
- My first IPv6 packets
- Dual stacking the backbone
- OSPFv3 implementation (IS-IS if you want)
- IPv6 iBGP implementation
- IPv6 eBGP implementation
- IPv6 peering
- Extending dual stack from the backbone
- IPv6 for your customers

AUDIENCE

- Successful IPv6 implementation requires:
 - + ARIN
 - × Supporting policies
 - × Allocations / Assignments
 - × Decision influencers

 - + NANOG
 - × Designers / Architects
 - × Implementors
 - × Operators

REMEMBER YOUR ROLE!

- Attending and participating is not enough. We must influence our respective companies to make good decisions about future viability. This includes the decision to implement IPv6.
- Typically Strategy / Revenue generation / P&L / BisDev / Marketing / Customer demand drives product.
- IPv6 decisions based on survivability + costs \$'s
- Business justification -> another presentation.

IPV6 OBJECTIONS....

- Dispel the myths!
 - + Obtaining IPv6 space from my RIR is hard
 - + My transit providers don't support IPv6 yet
 - + No BGP multihoming
 - + Lack of general support
 - + IPv6 is hard to implement
 - + Existing infrastructure doesn't support IPv6
 - × This one is harder to address, but better plan now!
 - × Question of `_when_` not `_if_`

OBTAINING AN ALLOCATION IS HARD!

- Lots of coverage on this topic..
- <https://www.arin.net/resources/templates/v6-isp.txt>

- Dear RIR,
- I am *planning* on assigning IPv6 space to 200 customers in the coming five years.

- Dear LIR,
- **[ARIN-20090507.2451] IPV6 REQUEST --APPROVED**

From: hostmaster@arin.net
Subject: **Re: [Ops] [ARIN-20090507.2451] IPV6 REQUEST --APPROVED**
Date: June 5, 2009 2:30:38 PM PDT
To: Aaron Hughes
Cc: ops@6connect.net

-----BEGIN PGP SIGNED MESSAGE-----

Hash: SHA1

Hello-

Your request for network addresses has been completed.

Below is the network which will be visible in ARIN's WHOIS tomorrow :

OrgName: 6connect, Inc.
OrgID: CONNE-81
Address: 1743 Park Ave.
Address: Suite 130
City: San Jose
StateProv: CA
PostalCode: 95126
Country: US

NetRange: 2607:FAE0:0000:0000:0000:0000:0000 - 2607:FAE0:FFFF:FFFF:FFFF:FFFF:FFFF
CIDR: 2607:FAE0:0000:0000:0000:0000:0000/32
NetName: 6CONNECT-6-1
NetHandle: NET6-2607-FAE0-1
Parent: NET6-2600-1
NetType: Direct Allocation
Comment:
RegDate: 2009-06-05
Updated: 2009-06-05

OrgTechHandle: 6CONN-ARIN
OrgTechName: 6connect Operations
OrgTechPhone: +1-408-329-6901
OrgTechEmail: ops@6connect.net

WHERE ARE WE?

- ✓ Obtain an IPv6 allocation
- My first IPv6 packets
- Dual stacking the backbone
- OSPFv3 implementation
- IPv6 iBGP implementation
- IPv6 eBGP implementation
- IPv6 peering
- Extending dual stack from the backbone
- IPv6 for your customers

MY PROVIDER DOESN'T SUPPORT IPV6

- No need to cry, let's just route around them.



IPV6 TRANSIT IS FREE! (right now...)



**HURRICANE ELECTRIC
INTERNET SERVICES**

Exchange Name	IP Address	Speed
AMS-IX	2001:7f8:1::a500:6939:1	10000
BigApe	2001:458:26:2::500	100
CoreSite - Any2 California	2001:504:13::1a	10000
DE-CIX	2001:7f8::1b1b:0:1	10000
Equinix Ashburn	2001:504:0:2::6939:1	10000
Equinix Chicago	2001:504:0:4::6939:1	10000
Equinix Dallas	2001:504:0:5::6939:1	10000
Equinix Hong Kong	2001:de8:7::6939:1	10000
Equinix Los Angeles	2001:504:0:3::6939:1	10000
Equinix Newark	2001:504:0:6::6939:1	10000
Equinix San Jose	2001:504:0:1::6939:2	20000
Equinix San Jose	2001:504:0:1::6939:1	20000
Equinix Tokyo	2001:de8:5::6939:1	10000
Equinix Zurich	2001:7f8:c:8235:194:42:48:80	10000
HKIX	2001:7fa:0:1::ca28:a19e	1000
KleyReX	2001:7f8:33::A100:6939:1	1000
LAIX	2001:504:a::a500:6939:1	1000
LINX Brocade LAN	2001:7f8:4:0::1b1b:1	10000

LONAP	2001:7f8:17::1b1b:1	1000
NetNod Stockholm	2001:7f8:d:fe::187	20000
NetNod Stockholm	2001:7f8:d:fb::187	20000
NetNod Stockholm	2001:7f8:d:ff::187	20000
NL-IX	2001:7f8:13::a500:6939:1	1000
NOTA	2001:478:124::176	10000
NYIIX	2001:504:1::a500:6939:1	10000
PAIX New York	2001:504:f:39	10000
PAIX Palo Alto	2001:504:d::10	10000
PaNAP	2001:860:0:6::6939:1	10000
SIX	2001:478:180::40	10000
SOL-IX	2001:7f8:21:10::101	1000
SOL-IX	2001:7f8:21:9::101	1000
STHIX	2001:7f8:3e::a500:0:6939:1	1000
Telx Atlanta	2001:478:132::75	10000
TorIX	2001:478:245:1::112	10000

30 IPv6 locations as of Sept 25th, 2009 10

Exchange Name	IP Address	Speed
AMS-IX	2001:7F8:1::A501:9151:1	10000
ChIX	2001:504:14::a501:9151:1	1000
CIIX (formerly LAAP)	2001:504:A::A501:9151:1	10000
CoreSite - Any2 California	2001:504:13:0:0:0:0:B	10000
DE-CIX	2001:7F8::4ACF:0:1	10000
Equinix Ashburn	2001:504:0:2:0:1:9151:1	20000
Equinix Chicago	2001:504:0:4:0:1:9151:1	10000
Equinix Dallas	2001:504:0:5:0:1:9151:1	10000
Equinix Los Angeles	2001:504:0:3:0:1:9151:1	10000
Equinix Newark	2001:504:0:6:0:1:9151:1	10000
Equinix San Jose	2001:504:0:1:0:1:9151:1	10000
LINX Brocade LAN	2001:7F8:4::4ACF:1	10000
LONAP	2001:7F8:17::4ACF:1	1000
MadIX	2607:F388:0:2200::3	1000
NOTA	2001:478:124::167	10000
NYIIX	2001:504:1::a501:9151:1	10000
PAIX Atlanta	2001:504:10::15	10000
PAIX New York	2001:504:F::9151:1	10000
PAIX Palo Alto	2001:504:D::9151:1	10000
PAIX Seattle	2001:504:12::25	10000
SIX	2001:504:16::4ACF	2000
Telx Atlanta	2001:478:132::12	10000
TorIX	2001:478:245:1::110	1000

Others??

23 IPv6 locations as of Sept 25th, 2009

WHERE TO START?

- Existing IX locations
 - + IX provider,
 - × I am an existing customer at location X, Y, and Z.
 - × My existing IPv4 addresses are X.X.X.X, X.X.X.X, etc.
 - × What are my IPv6 addresses for those respective locations?

- You likely have already been assigned IPv6

PEERING INFORMATION LIST

- Make a list of all relevant peering information:
 - + \$My_Company info:
 - + <http://asXXXX.peeringdb.com/>
 - + AS: XXXX
 - + AS-SET: ALTDB AS-COMPANY
 - + Equinix Ashburn IPv4: 206.223.115.X
 - + Equinix Ashburn IPv6 : 2001:504:0:X:XXXX:1
 - + Equinix San Jose IPv4 : 206.223.116.X etc..
 - + Equinix San Jose IPv6 : 2001:504:0:1:0:X:XXXX:1
 - + PAIX Palo Alto IPv4 : 198.32.176.X
 - + PAIX Palo Alto IPv6 : 2001:504:D::XXXX
 - + Peering Contact : peering@domain.com
 - + NOC Contact : support@domain.com

UPDATE PEERINGDB

- Add your IPv6 records and check the v6 box.

Navigation	Company Information				Public Peering Locations					
Home Page	Company Name	<input type="text" value="Company Inc"/>			Public Exchange Point	ASN	IP Address	Mbit/sec	Delete	
Logout	Primary ASN / IRR Record	<input type="text" value="8038"/>	<input type="text" value="AS-6CONNECT"/>		CoreSite - Any2 Calif	8038	<input type="text" value="2001:504:13::"/>	<input type="text" value="1000"/>	<input type="checkbox"/>	
Your Records	Also Known As (Aliases)	<input type="text"/>			CoreSite - Any2 Calif	8038	<input type="text" value="206.223.143."/>	<input type="text" value="1000"/>	<input type="checkbox"/>	
Peering Record	Company Website URL	<input type="text" value="http://www.6connect.net/"/>			Equinix Ashburn	8038	<input type="text" value="2001:504:0:2::2"/>	<input type="text" value="1000"/>	<input type="checkbox"/>	
User Account	Approx Prefixes Announced	<input type="text" value="500"/>			Equinix Ashburn	8038	<input type="text" value="206.223.115."/>	<input type="text" value="1000"/>	<input type="checkbox"/>	
Search Records	Network Type	NSP (Network Service Provider) / Backbone			Equinix Los Angeles	8038	<input type="text" value="206.223.123."/>	<input type="text" value="1000"/>	<input type="checkbox"/>	
Networks	Approx Traffic Levels	5-10 Gbps			Equinix Los Angeles	8038	<input type="text" value="2001:504:0:3::2"/>	<input type="text" value="1000"/>	<input type="checkbox"/>	
Exchange Points	Approx Traffic Ratios	Balanced			Equinix San Jose	8038	<input type="text" value="2001:504:0:1::2"/>	<input type="text" value="1000"/>	<input type="checkbox"/>	
Facilities	Geographic Scope	North America			Equinix San Jose	8038	<input type="text" value="206.223.116."/>	<input type="text" value="1000"/>	<input type="checkbox"/>	
Common Points	Looking Glass URL (http)	<input type="text"/>			PAIX Palo Alto	8038	<input type="text" value="198.32.176."/>	<input type="text" value="1000"/>	<input type="checkbox"/>	
Suggestions	Route Server URL (belnet)	<input type="text"/>			PAIX Palo Alto	8038	<input type="text" value="2001:504:Dr::"/>	<input type="text" value="1000"/>	<input type="checkbox"/>	
Comments	Public Notes	<input type="text"/>			SFMIX	8038	<input type="text" value="206.197.187."/>	<input type="text" value="1000"/>	<input type="checkbox"/>	
New Exchange	Private Notes	<input type="text"/>			Select Value		<input type="text"/>	<input type="text"/>	<input type="checkbox"/>	
New Facility	Protocols Supported	Unicast IPv4 <input checked="" type="checkbox"/> Multicast <input checked="" type="checkbox"/> IPv6 <input checked="" type="checkbox"/>			<input type="button" value="Update"/>					
Help	Peering Policy Information				Private Peering Locations					
FAQ	Peering Policy URL	<input type="text" value="http://www.6connect.net/peering.html"/>			Facility Name	ASN	SONET	Ethr	ATM	Delete
Statistics	General Peering Policy	Open			CoreSite Los Angeles (Dne Wishine)	23342	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
	Multiple Locations	Preferred			Equinix Ashburn (DC1-DC4)	23342	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
	Ratio Requirements	No			Equinix Los Angeles (SA1)	23342	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
	Contract Requirements	Not Required			Equinix San Jose (SV1)	23342	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
	<input type="button" value="Update"/>				eXchange, 200 Paul	23342	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
	Peering Contacts				PAIX Palo Alto	23342	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
	Role	Contact Name	Telephone	E-Mail Address	The Main Exchange 365 Main San Francisco	23342	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
	Policy	Aaron Hughes	831-824-4863	peering@6connect.net	Select Value		<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
	Ops	Support	888-xxx-xxxx	support@6connect.net	<input type="button" value="Update"/>					
	Tech	Support	415-xxx-xxxx	network@6connect.net						
	Sales	Sales	888-xxx-xxxx	sales@6connect.net						
	Select									
	<input type="button" value="Update"/>									

WHAT NEXT?

- We've a direct allocation
- We've IPv6 addresses for each of our IX locations
- We've made a list of info we will need to start
- We've updated peeringdb.com

DISCLAIMER

- Follow your own company change process!



CONFIGURING IPV6

- Locate existing IPv4 peering interfaces
- Enable IPv6 (Cisco)
- Configure the IPv6 address on the peering int
- Test

- First Cisco, then Juniper

CONFIGURING IPV6 CONT..

```
br01-1w-lax#conf t
Enter configuration commands, one per line. End with CNTL/Z.
br01-1w-lax(config)#ipv6 unicast-routing
br01-1w-lax(config)#^Z
```

- Enable IPv6 (Cisco)

```
br01-1w-lax#sh ip route 206.223.123.108
Routing entry for 206.223.123.0/25
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Redistributing via ospf 23342, bgp 23342
  Advertised by ospf 23342 metric-type 1 subnets
  Routing Descriptor Blocks:
  * directly connected, via Vlan204
    Route metric is 0, traffic share count is 1
```

- Find the v4 int

```
!
interface Vlan204
  description [UL:PEERING:EQUINIX LAX] To EQX Peering
  ip address 206.223.123.108 255.255.255.128
  ip flow ingress
  ip flow egress
  load-interval 30
  no mop enabled
  no mop sysid
!
```

- Verify existing config

CONFIGURING IPV6 CONT..

```
br01-1w-lax#conf t
Enter configuration commands, one per line. End with CNTL/Z.
br01-1w-lax(config)#int vlan 204
br01-1w-lax(config-if)#ipv6 enable
br01-1w-lax(config-if)#ipv6 address 2001:504:0:3:0:2:3342:1/64
br01-1w-lax(config-if)#^Z

br01-1w-lax#ping 2001:504:0:3:0:2:3342:1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:504:0:3:0:2:3342:1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/0 ms
br01-1w-lax#wr
Building configuration...
Compressed configuration from 60064 bytes to 21962 bytes
[OK]
```

- Configure
- Test

- We are passing IPv6 packets!

ZOOM IT A LITTLE

- Enabling IPv6 unicast on the router
 - + Cisco:
 - × 'ipv6 unicast-routing'
 - + Juniper
 - × enabled by default

INTERFACE CONFIG

■ Cisco

- + interface \$interface_name
- + ipv6 enable
- + ipv6 address 2001:1::1/64

■ Juniper

- + set interface \$interface_name unit \$unit family inet6 address 2001:1::1/64;

REACHING ACROSS THE INTERFACE

- Now that we have configured an interface and we know we can ping ourself, let's see if we can exchange some packets over IPv6 with the outside world.
- Finding other hosts in the subnet is no longer as simple as sending ICMP to the broadcast address. Instead IPv6 utilizes neighbor discovery.
- Let's take the easy way out and use peeringdb.

REACHING ACROSS THE INTERFACE

Navigation		Public Exchange Point Detailed View				List of Peers at this Exchange Point (Total: 63)				
Home Page	Common Name	Equinix Los Angeles				Peer Name	Local ASN	IP Address	IPs	Policy
Logout	Long Name	Equinix Los Angeles Exchange				Abovenet Communications Inc.	6461		1	Restrictive
	City	Los Angeles				Advanced Video Communications	46294	206.223.123.124	1	Open
Your Records	Country	US				Akamai Technologies	20940	206.223.123.102	1	Open
Peering Record	Continental Region	North America				AT&T US - AS7132	7132	206.223.123.79	1	Selective
User Account	Media Type	Ethernet				BandCon	26769	206.223.123.26	1	Selective
	Protocols Supported	Unicast IPv4 <input checked="" type="checkbox"/> Multicast <input type="checkbox"/> IPv6 <input type="checkbox"/>				BitGravity, Inc.	40009	2001:504:0:3::4:9:1	2	Open
Search Records	Contact Information				CENIC / CalREN	2152		1	Selective	
Networks	Company Website	http://peering.equinix.com/				Chunghwa Telecom	9505	206.223.123.51	1	Open
Exchange Points	Traffic Statistics Website					Cox Communications	22773	206.223.123.42	1	Selective
Facilities	Technical E-Mail	support@equinix.com				DALnet IRC Network	31800	206.223.123.88	1	Open
Common Points	Technical Phone					EarthLink	4355	206.223.123.32	1	Open
Suggestions	Policy E-Mail	support@equinix.com				Equinix Corp Network	14609	206.223.123.126	1	Open
Comments	Policy Phone					FAT Networks, LLC	14076	206.223.123.30	1	Open
New Exchange	IP Address Blocks				Giganews	30094	206.223.123.114	1	Open	
New Facility	Type	Address Block	Reverse DNS Scan		Global NAPS, GNAPS	1784	206.223.123.41	1	Open	
	IPv4 Unicast	206.223.123.0/25	Link		Guam Cablevision, LLC.	3605	206.223.123.3	1	Open	
	IPv6 Unicast	2001:504:0:3::/64	Unsupported		Highwinds Network Group, Inc	12989	2001:504:0:3::1:2989:1	2	Selective	
Help	Local Facilities				Hurricane Electric	6939	2001:504:0:3::6939:1	2	Open	
FAQ	Facility Name	City	Country	Participant Count	Hypermedia Systems	30212	206.223.123.9	1	Open	
Statistics	Equinix El Segundo (LA3)	El Segundo	US	10	ICANN	40528	2001:504:0:3:0:4:0528:1	2	Open	
	Equinix Los Angeles (LA1)	Los Angeles	US	77	imeem, Inc.	36119	206.223.123.53	1	Selective	
					Integra Telecom	7385	2001:504:0:3::7385:1	2	Open	
					Internap	22212	206.223.123.29	1	Selective	
					1 2 3 of 3 Next > Last >>					

```
br01-1w-lax#ping 2001:504:0:3::6939:1
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 2001:504:0:3::6939:1, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/2/4 ms
```


CONFIGURING IPV6 ON A JUNIPER

- Find the IPv4 interface

```
syntax error, expecting <command>.
aaronh@br01-eqx-ash> show route 206.223.115.154

inet.0: 281580 destinations, 632286 routes (281580 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

206.223.115.154/32 *[Local/0] 19w3d 08:10:51
                   Local via ge-1/3/0.3
```

- Verify the config

```
aaronh@br01-eqx-ash> show configuration interfaces ge-1/3/0 unit 3
description "[PEER:EQX:ASH] Equinix Peering Fabric";
vlan-id 200;
family inet {
    filter {
        input all;
        output all;
    }
    address 206.223.115.154/24;
}
```

- Configure the IPv6 address

```
aaronh@br01-eqx-ash# set interfaces ge-1/3/0 unit 4 family inet6 address address 2001:504:0:2::2:3342:1/64;
```

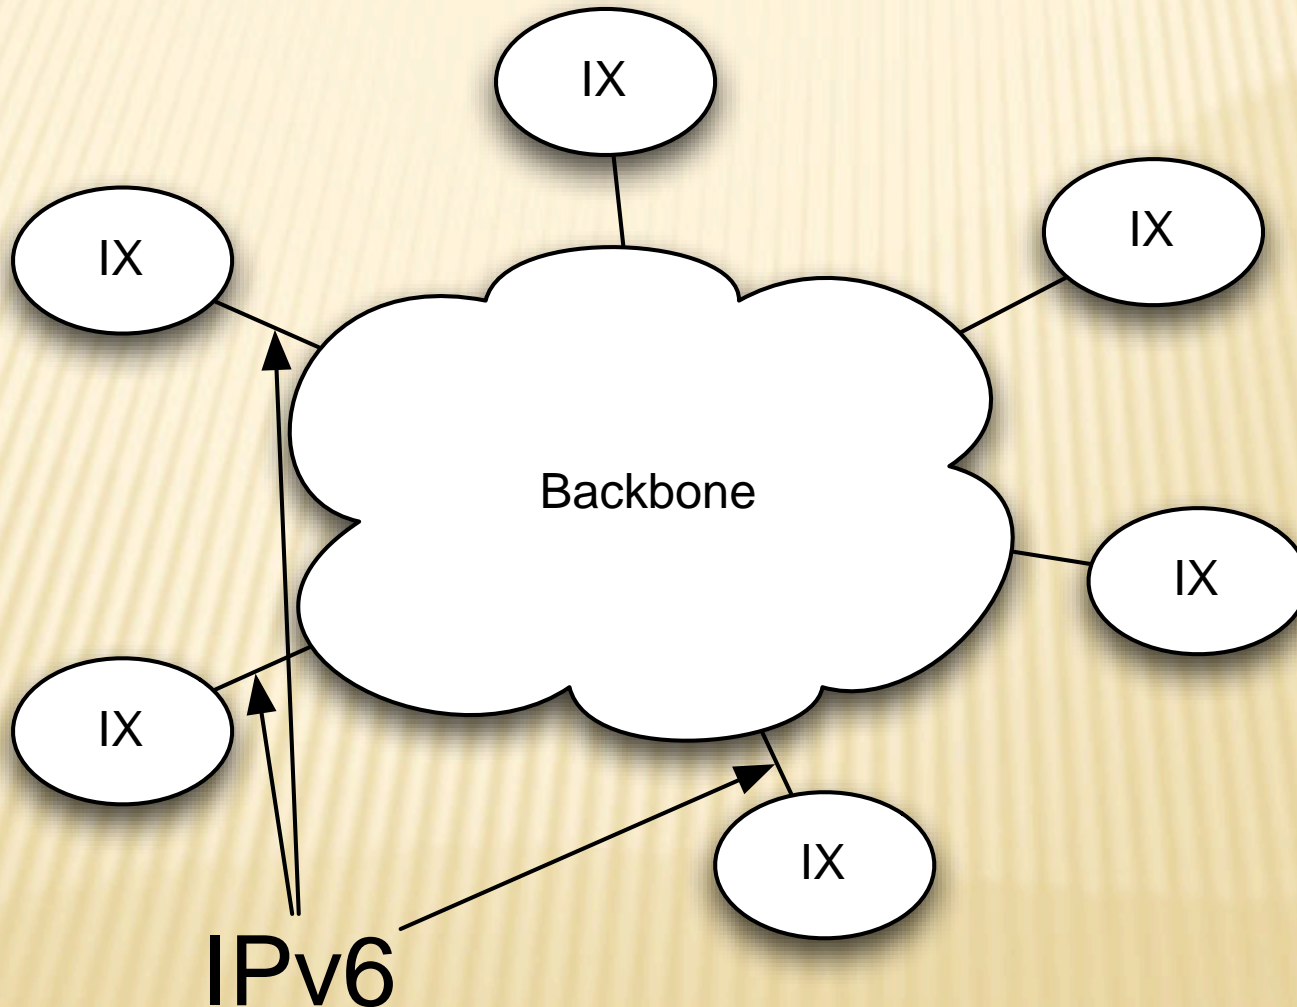
CONFIGURING ON A JUNIPER CONT.

```
aaronh@br01-eqx-ash> ping 2001:504:0:2::2:3342:1
PING6(56=40+8+8 bytes) 2001:504:0:2:0:2:3342:1 --> 2001:504:0:2:0:2:3342:1
16 bytes from 2001:504:0:2:0:2:3342:1, icmp_seq=0 hlim=64 time=0.265 ms
16 bytes from 2001:504:0:2:0:2:3342:1, icmp_seq=1 hlim=64 time=0.144 ms
16 bytes from 2001:504:0:2:0:2:3342:1, icmp_seq=2 hlim=64 time=0.161 ms
16 bytes from 2001:504:0:2:0:2:3342:1, icmp_seq=3 hlim=64 time=0.171 ms
16 bytes from 2001:504:0:2:0:2:3342:1, icmp_seq=4 hlim=64 time=0.173 ms
16 bytes from 2001:504:0:2:0:2:3342:1, icmp_seq=5 hlim=64 time=0.172 ms
16 bytes from 2001:504:0:2:0:2:3342:1, icmp_seq=6 hlim=64 time=0.156 ms
^C
--- 2001:504:0:2::2:3342:1 ping6 statistics ---
7 packets transmitted, 7 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 0.144/0.177/0.265/0.037 ms
```

- IPv6 packets are passing!

WHERE ARE WE?

- IPv6 addresses are configured on the IX edges



WHERE ARE WE?

- ✓ Obtain an IPv6 allocation
- ✓ My first IPv6 packets
 - Dual stacking the backbone
 - OSPFv3 implementation
 - IPv6 iBGP implementation
 - IPv6 eBGP implementation
 - IPv6 peering
 - Extending dual stack from the backbone
 - IPv6 for your customers

IPV6 INTERNAL ASSIGNMENTS

- Keeping track of your peering interface address is one thing, however, keeping track of your internal assignments is an entirely different thing.
- If you have the resources to do so, write a tool to manage IPv6 DNS and assignments.
- If not?
 - + Spreadsheet?
 - + Database?
 - + DNS zone files?

ALLOCATION AND ASSIGNMENT TOOLS

- If you have the resources, write a tool

6connect

Welcome Aaron Hughes of 6connect, Inc. 1743 Park Ave.
Suite 130
San Jose, CA, US, 95126

View as:

Home Leads Opportunities Accounts Contacts Assets Support Monitoring Ticketing Cabinets Locations IP Logout

IPv4 - IPv6	Hosts	Is SWIPed	Is Aggregate	CustID	Description	Company Name	Action Filter:(Available Assigned) ARIN Report
2607:fae0:0:0::/64	2^64	0	0	1		6connect, Inc.	Description Reclaim whois viewSWIP sendSWIP
2607:fae0:0:1::/64	2^64	0	0	7		6connect Available	Assign
2607:fae0:0:2::/64	2^64	1	0	8		InfinIT Consulting	Description whois viewSWIP sendDESWIP
2607:fae0:0:3::/64	2^64	0	0	7		6connect Available	Assign
2607:fae0:0:4::/64	2^64	0	0	7		6connect Available	Assign
2607:fae0:0:5::/64	2^64	0	0	7		6connect Available	Assign
2607:fae0:0:6::/64	2^64	0	0	7		6connect Available	Assign
2607:fae0:0:7::/64	2^64	0	0	7		6connect Available	Assign
2607:fae0:0:8::/64	2^64	0	0	7		6connect Available	Assign
2607:fae0:0:9::/64	2^64	0	0	7		6connect Available	Assign
2607:fae0:0:a::/64	2^64	0	0	7		6connect Available	Assign
2607:fae0:0:b::/64	2^64	0	0	7		6connect Available	Assign
2607:fae0:0:c::/64	2^64	0	0	7		6connect Available	Assign
2607:fae0:0:d::/64	2^64	0	0	7		6connect Available	Assign
2607:fae0:0:e::/64	2^64	0	0	7		6connect Available	Assign
2607:fae0:0:f::/64	2^64	0	0	7		6connect Available	Assign
2607:fae0:0:10::/64	2^64	0	0	7		6connect Available	Assign
2607:fae0:0:1000::/64	2^64	0	0	7		6connect Available	Assign
2607:fae0:0:1001::/64	2^64	0	0	7		6connect Available	Assign
2607:fae0:0:1002::/64	2^64	0	0	7		6connect Available	Assign
2607:fae0:1:0::/48	2^48	0	0	7		6connect Available	Assign
2607:fae0:2:0::/48	2^48	1	0	8		InfinIT Consulting	Description whois viewSWIP sendDESWIP
2607:fae0:3:0::/48	2^48	0	0	7		6connect Available	Assign
2607:fae0:4:0::/48	2^48	0	0	7		6connect Available	Assign
2607:fae0:5:0::/48	2^48	0	0	7		6connect Available	Assign
2607:fae0:6:0::/48	2^48	0	0	7		6connect Available	Assign
2607:fae0:7:0::/48	2^48	0	0	7		6connect Available	Assign

Assigned Hosts: 3.68937696224E+19, Total Hosts: 3.68936851799E+20 Percentage used: 10.000022881

[Generate next /48](#)
[Generate next /64 for Customer](#)
[Generate next /64 Give for Infrastructure](#)

NUMBERING PLAN

```

/16  /32  /48  /64  /80  /96  /112
  v   v   v   v   v   v   v
2001:aaaa:bbbb:cccc:dddd:eeee:ffff:1111

```

Here is a little shortcut to hopefully help you subnet IPv6.

```

/16 /32 /48 /64 /80 /96 /112
  v  v  v  v  v  v  v
2001:aaaa:bbbb:cccc:dddd:eeee:ffff:1111

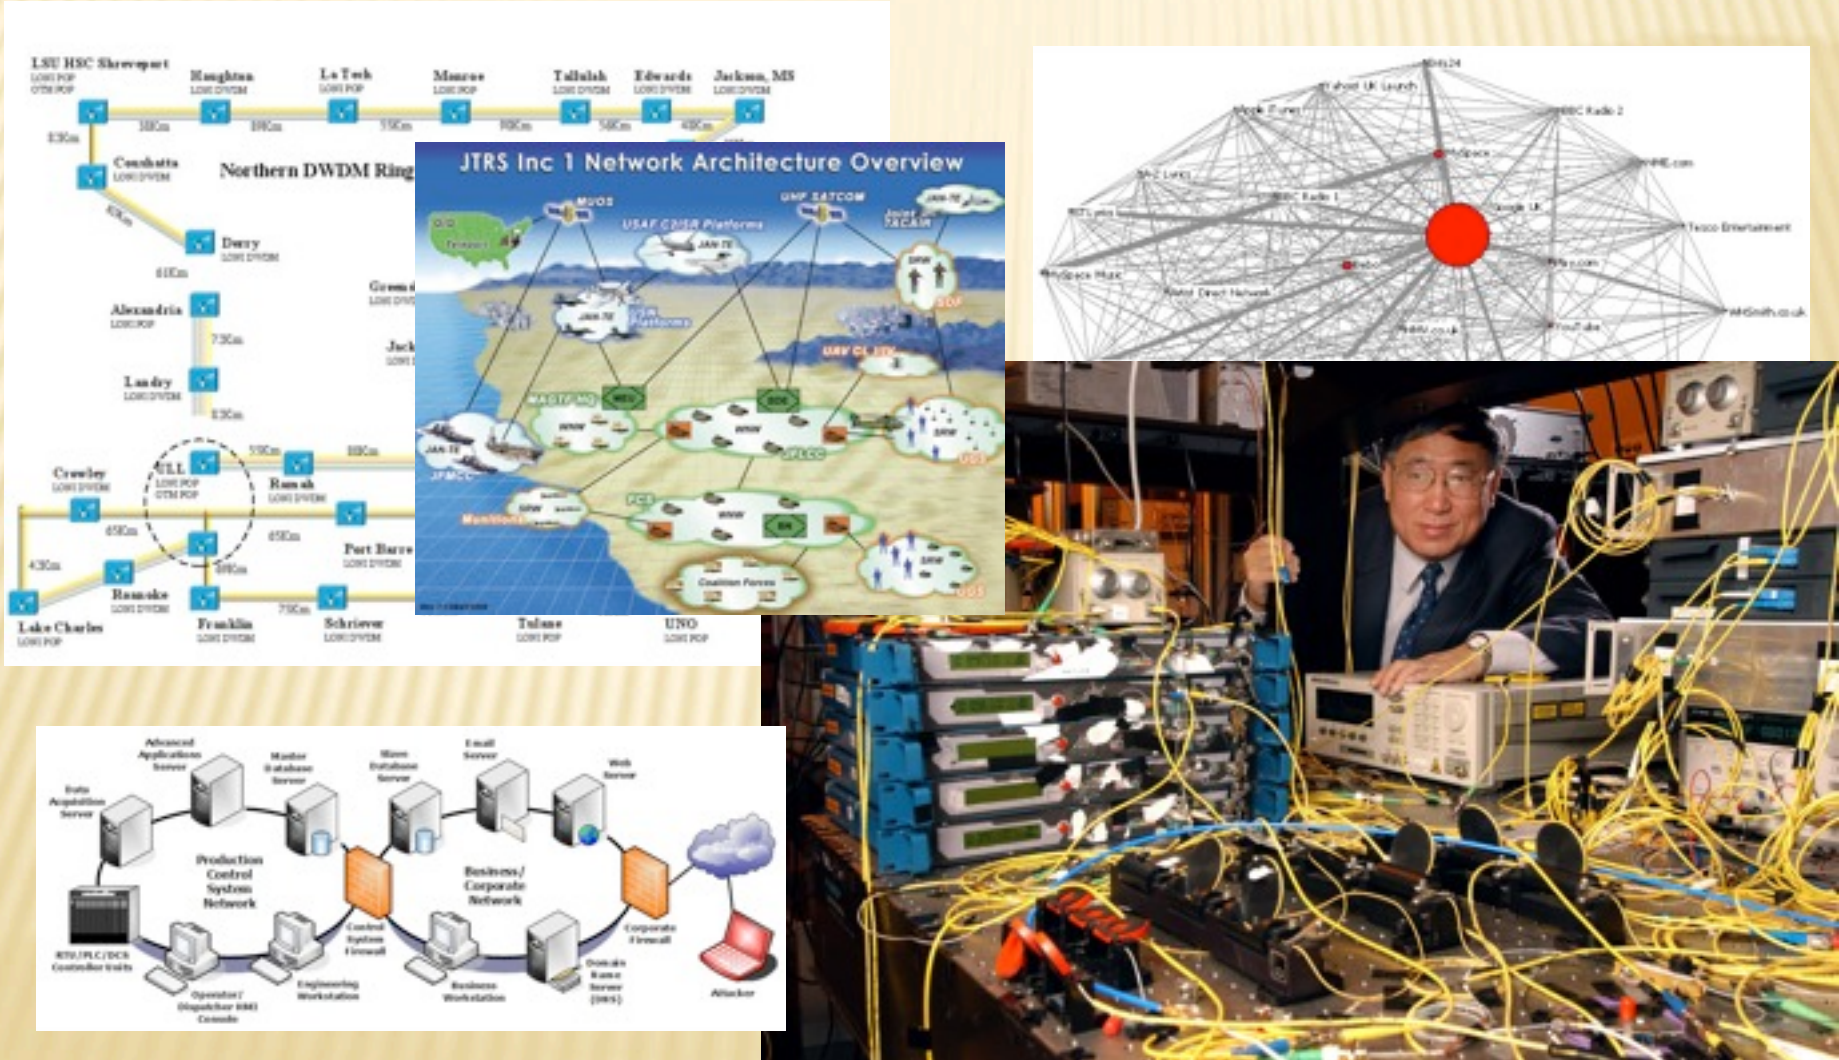
```

Prefix	/48 count	/56 count	Number of /64 Subnets	Number of Hosts
/64			1	18,446,744,073,709,551,616 (2 ⁶⁴) (quintillion)
/63			2	36,893,488,147,419,103,232
/62			4	73,786,976,294,838,206,464
/61			8	147,573,952,589,676,412,928
/60			16	295,147,905,179,352,825,856
/59			32	590,295,810,358,705,651,712
/58			64	1,180,591,620,717,411,303,424 (sextillion)
/57			128	2,361,183,241,434,822,606,848
/56		1	256	4,722,366,482,869,645,213,696 (2 ⁷²)
/55		2	512	9,444,732,965,739,290,427,392
/54		4	1,024	18,889,465,931,478,580,854,784
/53		8	2,048	37,778,931,862,957,161,709,568
/52		16	4,096	75,557,863,725,914,323,419,136
/51		32	8,192	151,115,727,451,828,646,838,272
/50		64	16,384	302,231,454,903,657,293,676,544
/49		128	32,768	604,462,909,807,314,587,353,088
/48	1	256	65,536	1,208,925,819,614,629,174,706,176 (2 ⁸⁰) (septillion)

65,536 1,208,925,819,614,629,174,706,176 (2⁸⁰) (septillion)

ARCHITECTURE OF YOUR NETWORK

- Opportunity to change vs. keep the same as v4

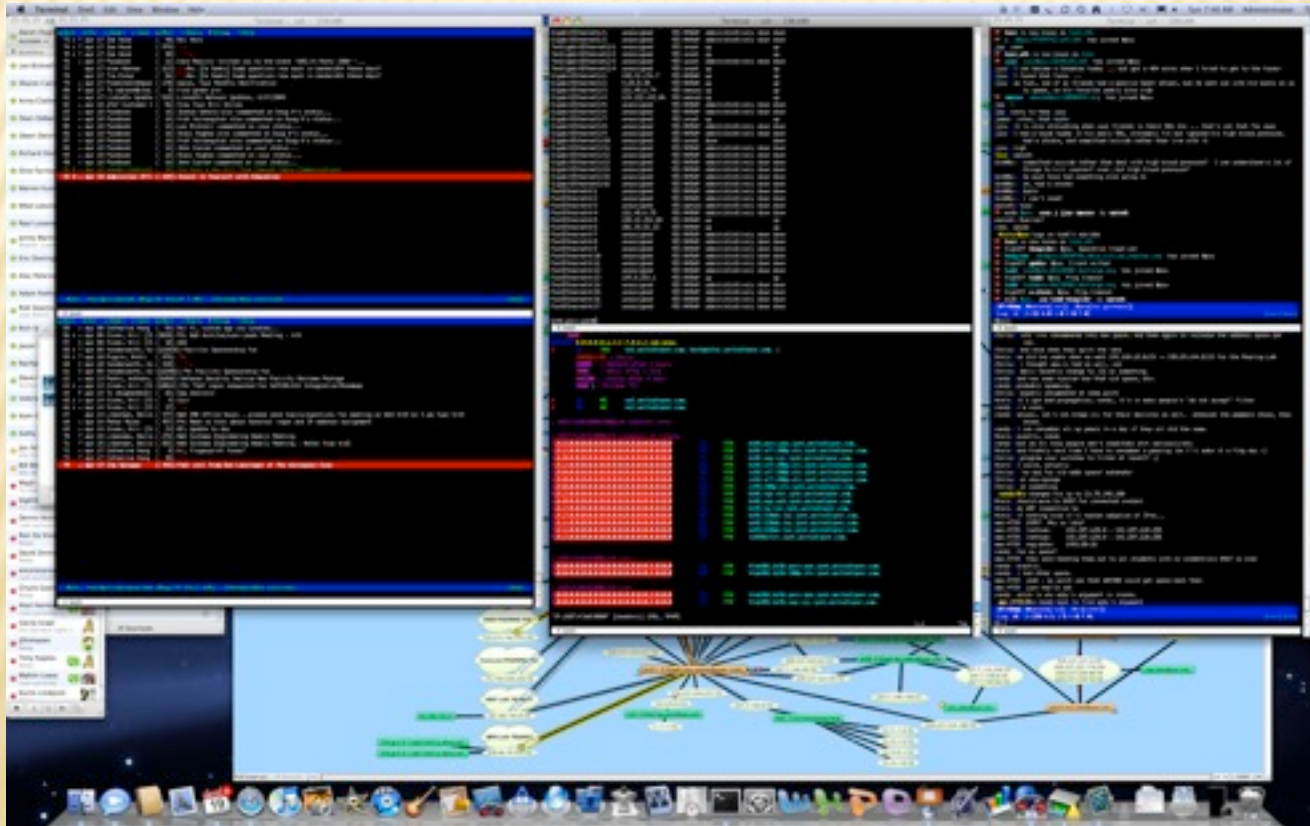


ARCHITECTURE CHOICES

- There are a ton of opinions about how to architect a network. For the purpose of this presentation, I've just picked one.
- Basic Network Architecture
 - + Loopbacks and connected infrastructure into OSPF
 - + iBGP full mesh sourced off loopbacks
 - + iBGP next-hop-self
 - + All connected except loopbacks into iBGP
 - + eBGP distribution via route-maps and communities

CONFIGURING YOUR BACKBONE

- Open your mail client, DNS zone editor and whatever application you use to access routers.



STARTING ROUTER

- Pick one router connected to an IX to start with.

```
ipv6 router ospf 23342
```

- Some versions of IOS require this.

```
!  
interface Loopback0  
description [UL:LOOPBACK]  
ip address 209.237.224.247 255.255.255.255  
ipv6 address 2607:F3A0::247/128  
ipv6 enable  
ipv6 ospf 23342 area 0  
!
```

```
br01-paix-pao(config)#ipv6 ospf name-lookup  
br01-paix-pao(config)#
```

```
interface Vlan903  
description [UL:VLAN] ->br01-eqx-sjc-v903  
mtu 9216  
ip address 207.7.159.53 255.255.255.252  
ip ospf cost 2  
ipv6 address 2607:F3A0:0:2::1/64  
ipv6 enable  
ipv6 ospf 23342 area 0  
!
```

- Enable IPv6 on the interface
- Add the IPv6 Address
- Enable IPv6 OSPF name-lookup
- Extend to edge interfaces facing backbone routers.

EXTENDING IPV6 INTO YOUR CORE

- Rinse. Repeat.
- Extend from the outside in and watch as the OSPFv3 IPv6 sessions come up

```
br01-paix-pao#sh ip os ne
```

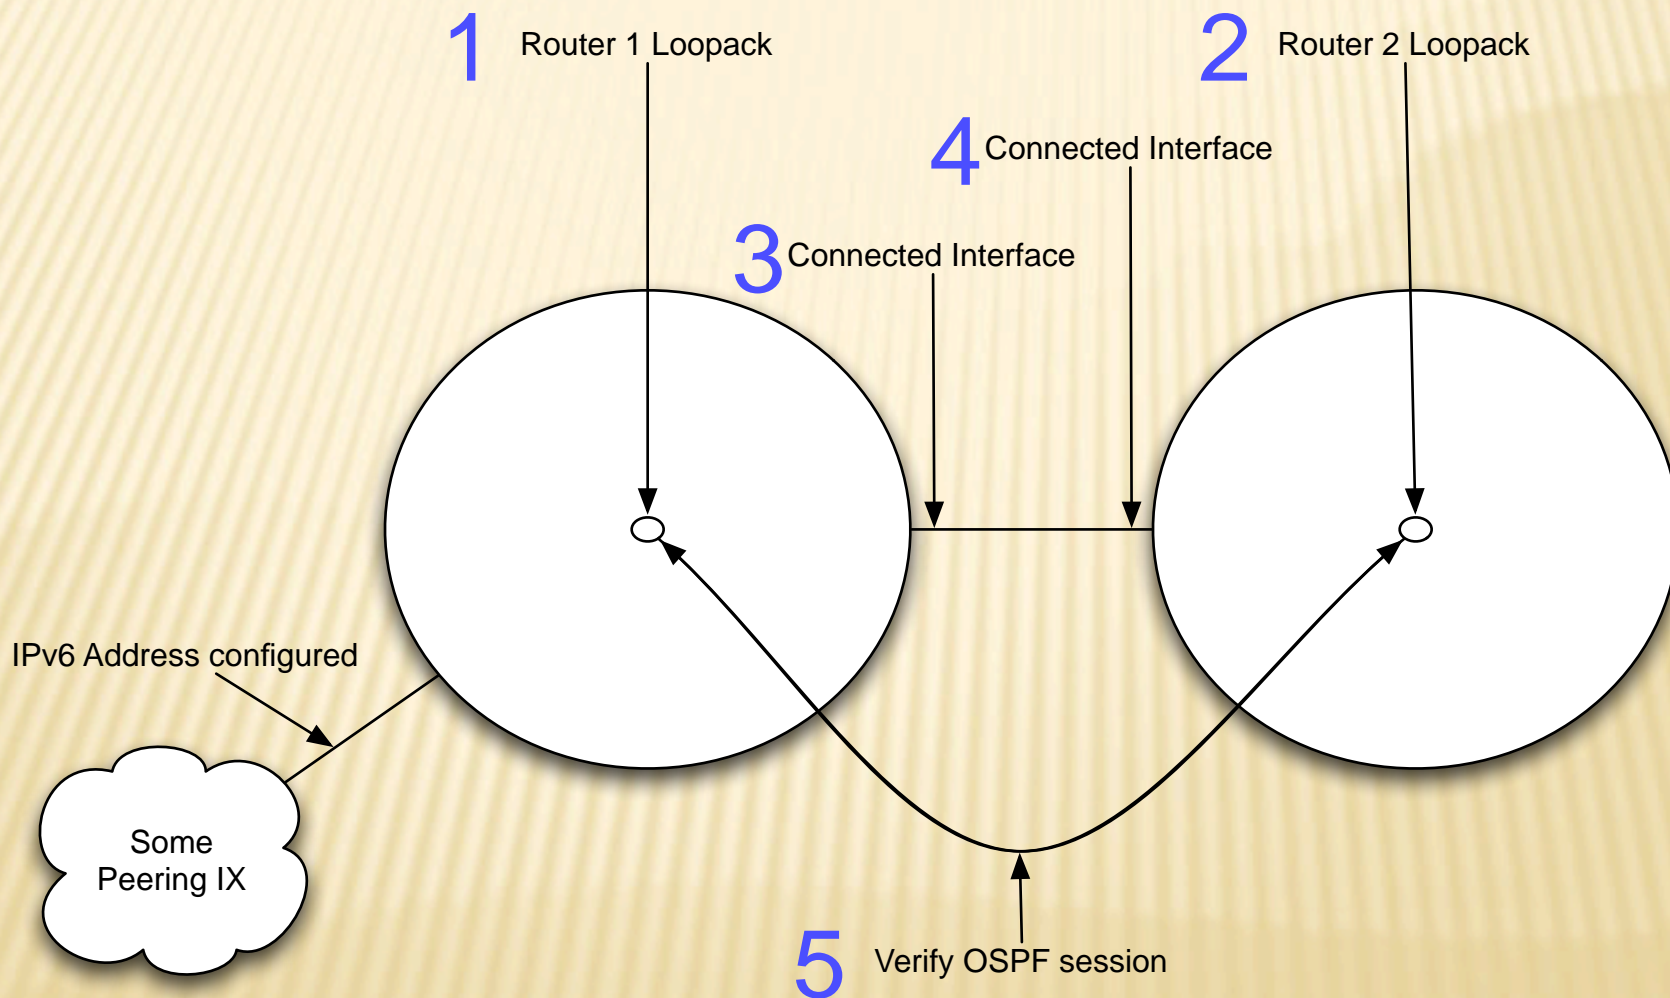
Neighbor ID	Pri	State	Dead Time	Address	Interface
cr01-55sm-sjc.u	1	FULL/DR	00:00:31	207.7.159.110	Vlan905
br01-eqx-sjc.un	1	FULL/DR	00:00:34	207.7.159.54	Vlan903
br02-sf7-200p-s	1	FULL/DR	00:00:39	207.7.129.74	Vlan902

```
br01-paix-pao#sh ipv6 os ne
```

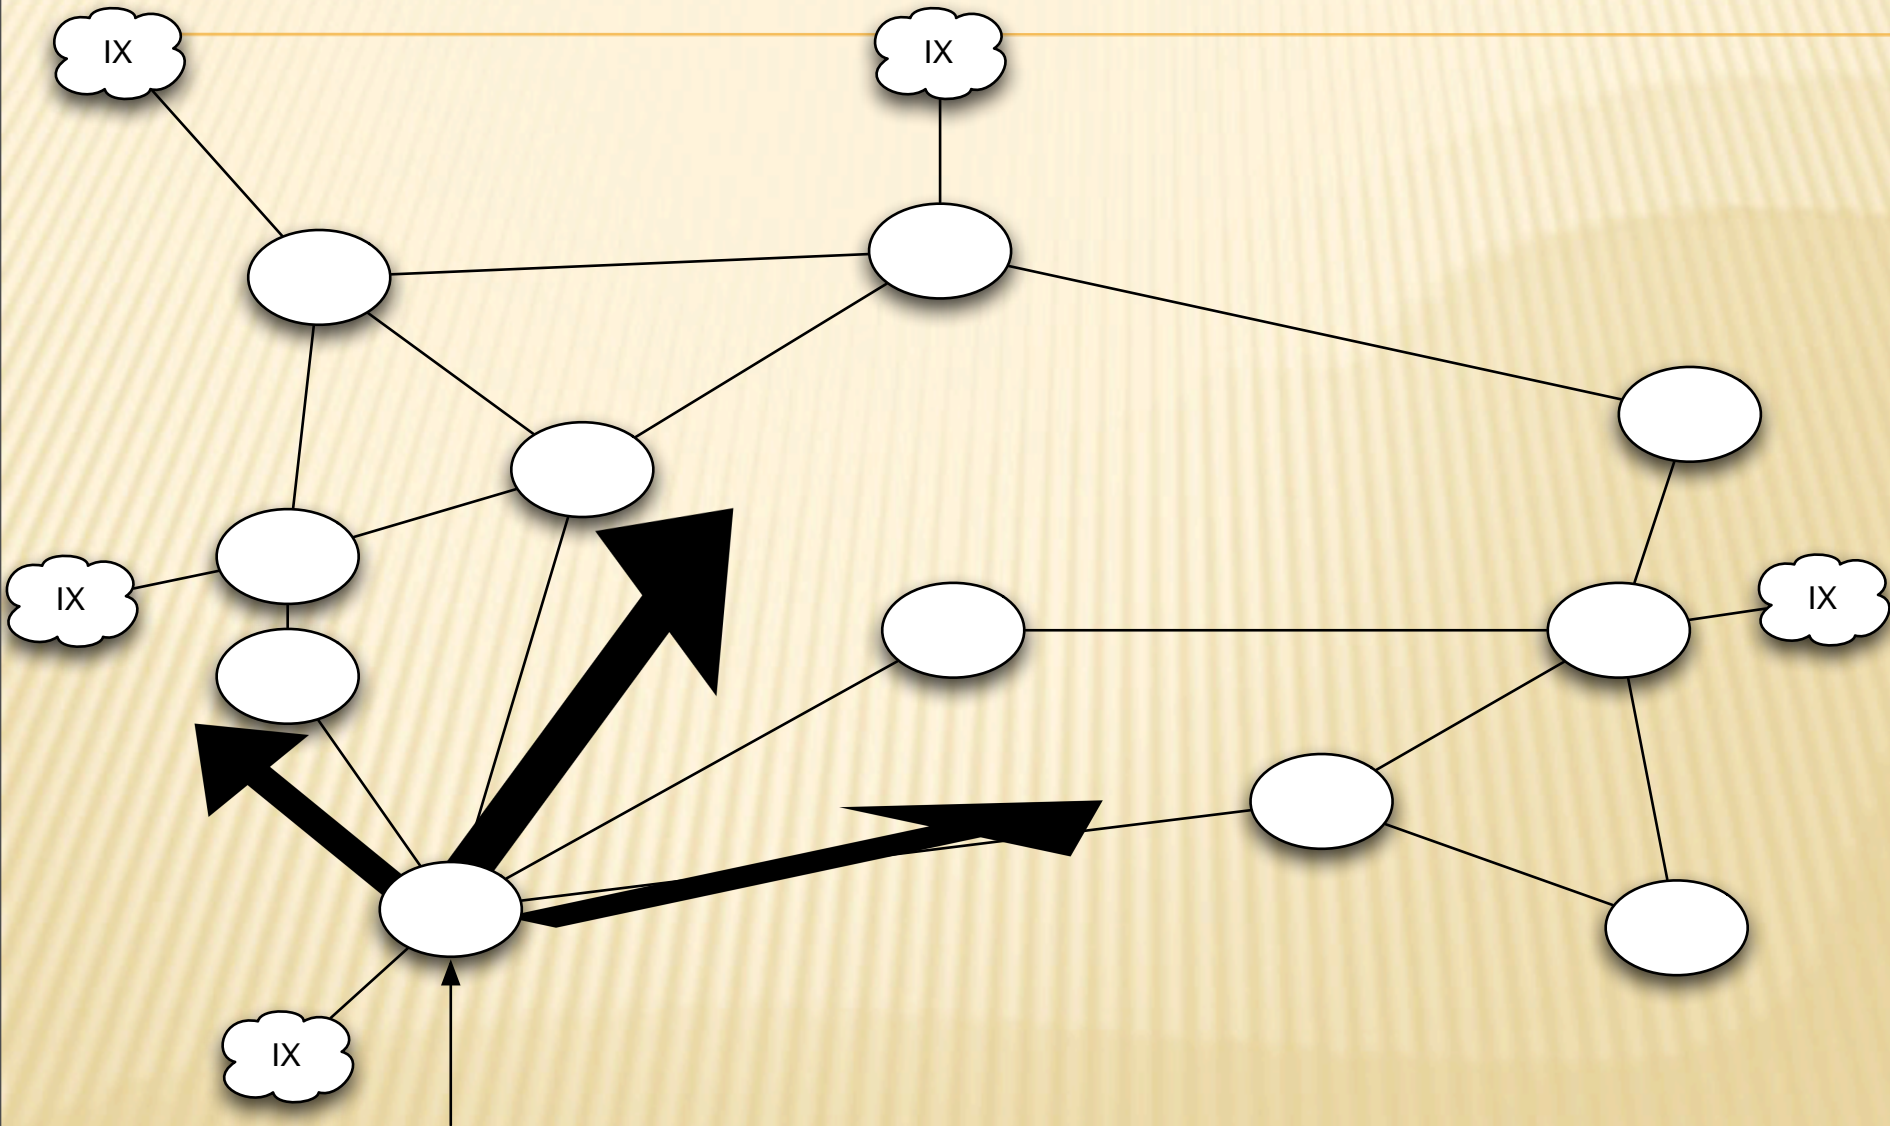
Neighbor ID	Pri	State	Dead Time	Interface ID	Interface
br01-eqx-sjc.un	1	FULL/DR	00:00:31	85	Vlan903
br02-sf7-200p-s	1	FULL/DR	00:00:36	79	Vlan902

```
br01-paix-pao#
```


EXTENDING IN FROM THE PEERING EDGE



10,000 FOOT VIEW



Pick a starting point

Slowly work your way out across the connected links.

MANAGING ASSIGNMENTS WITH DNS ZONE

- First /48 for all internal infrastructure
 - + Alternatively you can get a direct allocation for this.
- First /64 of the /48 for router loopbacks
- Second /64 is for your first connected interface

CUT, PASTE, INCREMENT

```
; 2607:f3a0:0:0001/64 Link  
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1.0.0.0      IN    PTR    Vlan902.br01-paix-pao.ipv6.unitedlayer.com.  
2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1.0.0.0      IN    PTR    Vlan902.br02-200p-sfo.ipv6.unitedlayer.com.  
  
; 2607:f3a0:0:0002/64 Link  
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.2.0.0.0      IN    PTR    Vlan903.br01-paix-pao.ipv6.unitedlayer.com.  
2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.2.0.0.0      IN    PTR    Vlan903.br01-eqx-sjc.ipv6.unitedlayer.com.  
  
; 2607:f3a0:0:0003/64 Link  
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.3.0.0.0      IN    PTR    Vlan904.br01-200p-sfo.ipv6.unitedlayer.com.  
2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.3.0.0.0      IN    PTR    Vlan904.br01-eqx-sjc.ipv6.unitedlayer.com.  
  
; 2607:f3a0:0:0004/64 Link  
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.4.0.0.0      IN    PTR    Vlan804.br02-200p-sfo.ipv6.unitedlayer.com.  
2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.4.0.0.0      IN    PTR    Vlan804.br01-200p-sfo.ipv6.unitedlayer.com.  
  
; 2607:f3a0:0:0005/64 Link  
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.5.0.0.0      IN    PTR    Vlan518.br01-200p-sfo.ipv6.unitedlayer.com.  
2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.5.0.0.0      IN    PTR    Vlan518.cr01-200p-sfo.ipv6.unitedlayer.com.  
  
; 2607:f3a0:0:0006/64 Link  
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.6.0.0.0      IN    PTR    Vlan529.br01-200p-sfo.ipv6.unitedlayer.com.  
2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.6.0.0.0      IN    PTR    Vlan529.cr02-200p-sfo.ipv6.unitedlayer.com.  
  
; 2607:f3a0:0:0007/64 Link  
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.7.0.0.0      IN    PTR    Vlan807.cr01-200p-sfo.ipv6.unitedlayer.com.  
2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.7.0.0.0      IN    PTR    Vlan807.cr02-200p-sfo.ipv6.unitedlayer.com.
```

WHAT COMES AFTER 9?

```
; 2687:f3a8:0:0009/64 Link
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.9.0.0.0    IN     PTR     Vlan519.br02-200p-sfo.ipv6.unitedlayer.com.
2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.9.0.0.0    IN     PTR     Vlan519.cr01-200p-sfo.ipv6.unitedlayer.com.

; 2687:f3a8:0:000a/64 Link
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.a.0.0.0    IN     PTR     Gig2-13.br01-200p-sfo.ipv6.unitedlayer.com.
2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.a.0.0.0    IN     PTR     Gig3-0.cr04-200p-sfo.ipv6.unitedlayer.com.

; 2687:f3a8:0:000b/64 Link
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.b.0.0.0    IN     PTR     Pos2-2.cr04-200p-sfo.ipv6.unitedlayer.com.
2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.b.0.0.0    IN     PTR     S8-3-0.br01-eqx-chi.ipv6.unitedlayer.com.

; 2687:f3a8:0:000c/64 Link
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.c.0.0.0    IN     PTR     Vlan180.br02-200p-sfo.ipv6.unitedlayer.com.
2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.c.0.0.0    IN     PTR     g8-1-0.br01-eqx-ash.ipv6.unitedlayer.com.

; 2687:f3a8:0:000d/64 Link
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.d.0.0.0    IN     PTR     S8-2-0.br01-eqx-ash.ipv6.unitedlayer.com.
2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.d.0.0.0    IN     PTR     S8-1-3.br01-1w-lax.ipv6.unitedlayer.com.

; 2687:f3a8:0:000e/64 Link
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.e.0.0.0    IN     PTR     Vlan800.br01-200p-sfo.ipv6.unitedlayer.com.
2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.e.0.0.0    IN     PTR     Vlan800.br01-1w-lax.ipv6.unitedlayer.com.

; 2687:f3a8:0:000f/64 Link
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.f.0.0.0    IN     PTR     Vlan802.br02-200p-sfo.ipv6.unitedlayer.com.
2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.f.0.0.0    IN     PTR     Vlan802.br01-1w-lax.ipv6.unitedlayer.com.

; 2687:f3a8:0:0010/64 Link
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1.0.0.0    IN     PTR     Vlan851.br01-530w6-lax.ipv6.unitedlayer.com.
2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1.0.0.0    IN     PTR     Vlan851.br01-1w-lax.ipv6.unitedlayer.com.

; 2687:f3a8:0:0011/64 Link
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1.1.0.0    IN     PTR     Vlan883.br01-200p-sfo.ipv6.unitedlayer.com.
2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1.1.0.0    IN     PTR     Gig0-1-0.br01-530w6-lax.ipv6.unitedlayer.com.
```


WHERE ARE WE?

- ✓ Obtain an IPv6 allocation
- ✓ My first IPv6 packets
- ✓ Dual stacking the backbone
- ✓ OSPFv3 implementation
 - IPv6 iBGP implementation
 - IPv6 eBGP implementation
 - IPv6 peering
 - Extending dual stack from the backbone
 - IPv6 for your customers

WHERE ARE WE CONT.

- IPv6 configured on all exchange interfaces
- IPv6 configured on all loopback interfaces
- IPv6 configured on all connected interfaces between backbone routers.
- OSPFv3 configured on Loopbacks (/128s)
- OSPFv3 configured on connected (/64s)

- What's next?
 - + I want to access the rest of the IPv6 world
 - + Next we configure iBGP

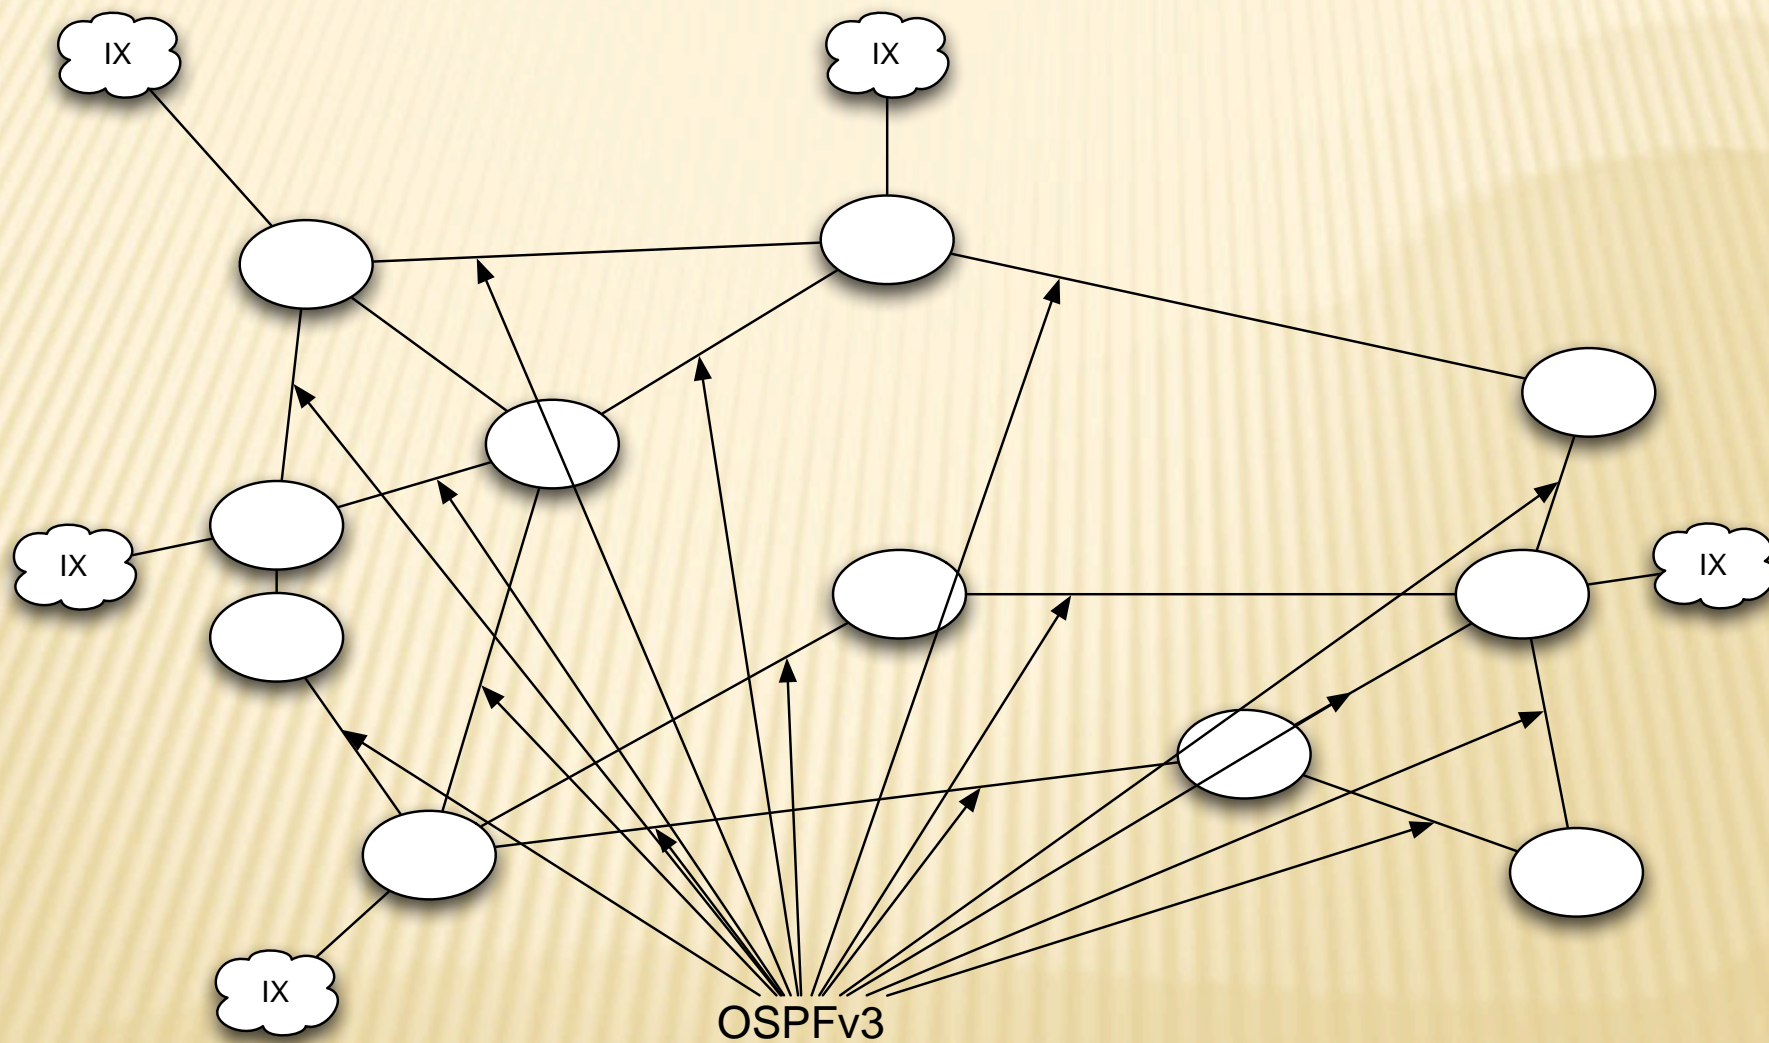
IPV6 PEER PEER-GROUP

- Peer group for an IPv6 peer:
 - + neighbor PEERS-v6 peer-group
 - + neighbor PEERS-v6 soft-reconfiguration inbound
 - + neighbor PEERS-v6 prefix-list Sanity-v6 in
 - + neighbor PEERS-v6 prefix-list Sanity-v6 out
 - + neighbor PEERS-v6 route-map PEER-IN-v6 in
 - + neighbor PEERS-v6 route-map PEER-OUT-v6 out

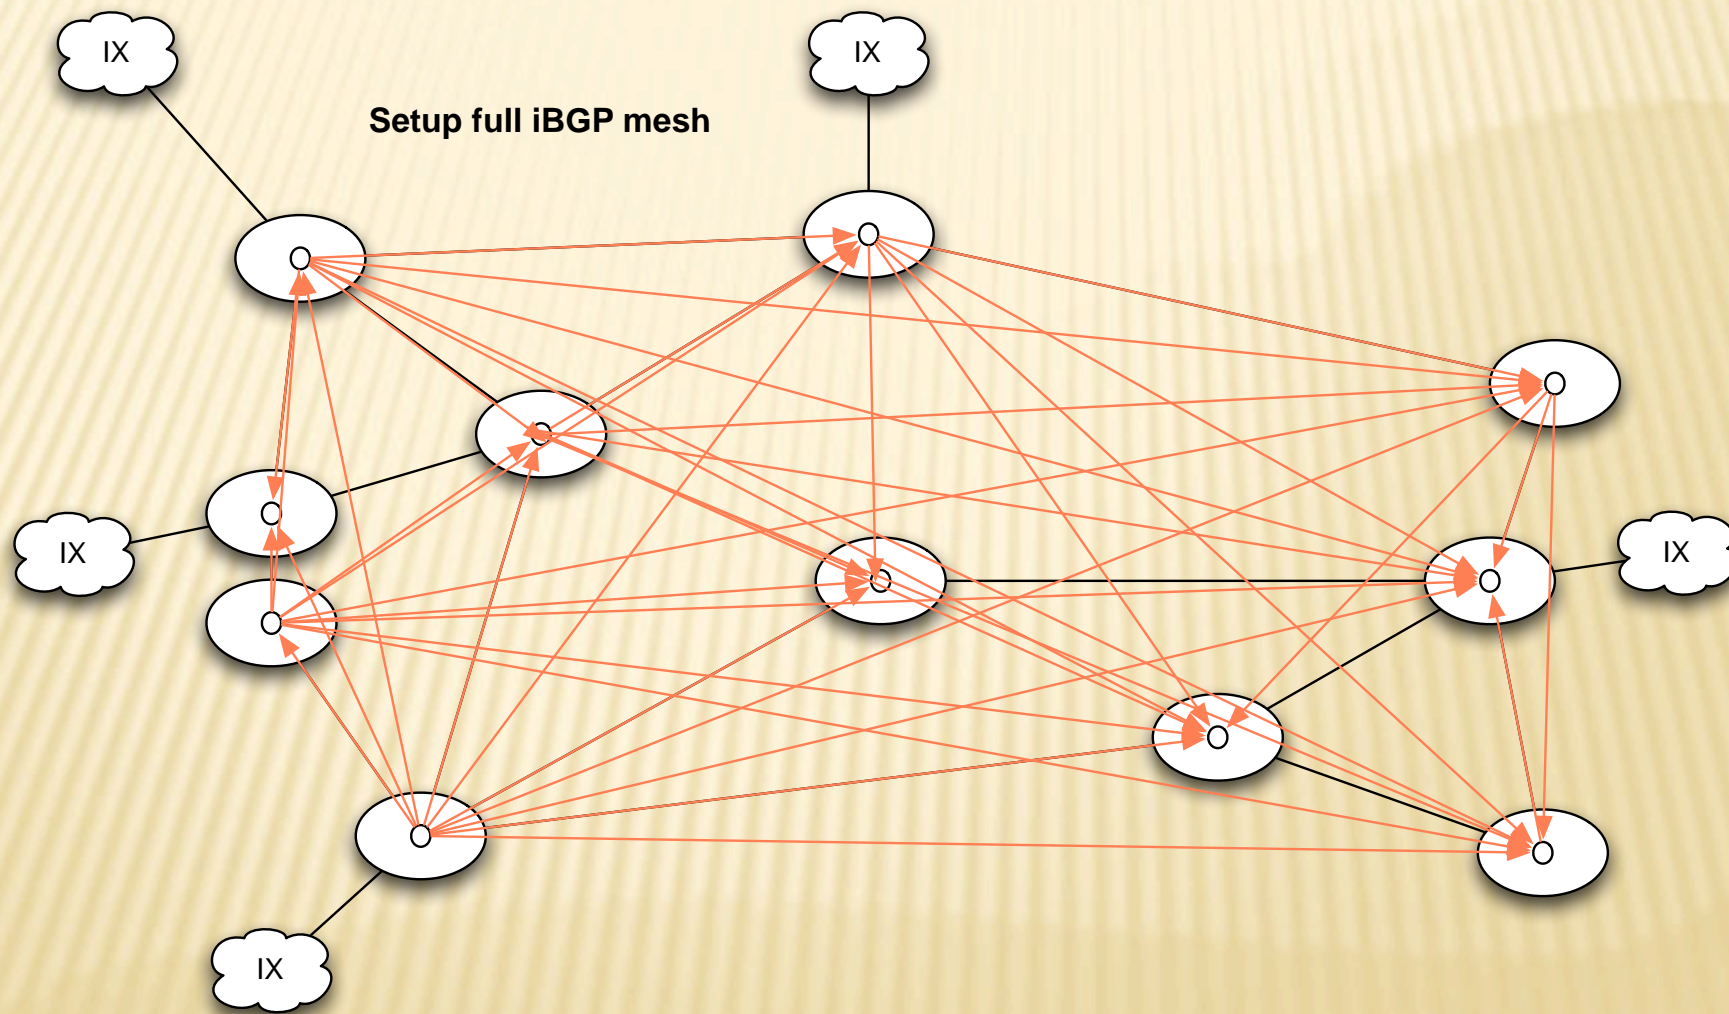
IPV6 IBGP PEER-GROUP

- Peer group for an IPv6 core router:
 - + neighbor CORE-v6 peer-group
 - + neighbor CORE-v6 remote-as XXXX < your ASN
 - + neighbor CORE-v6 soft-reconfiguration inbound
 - + neighbor CORE-v6 update-source Loopback0
 - + neighbor CORE-v6 send-community
 - + neighbor CORE-v6 next-hop-self

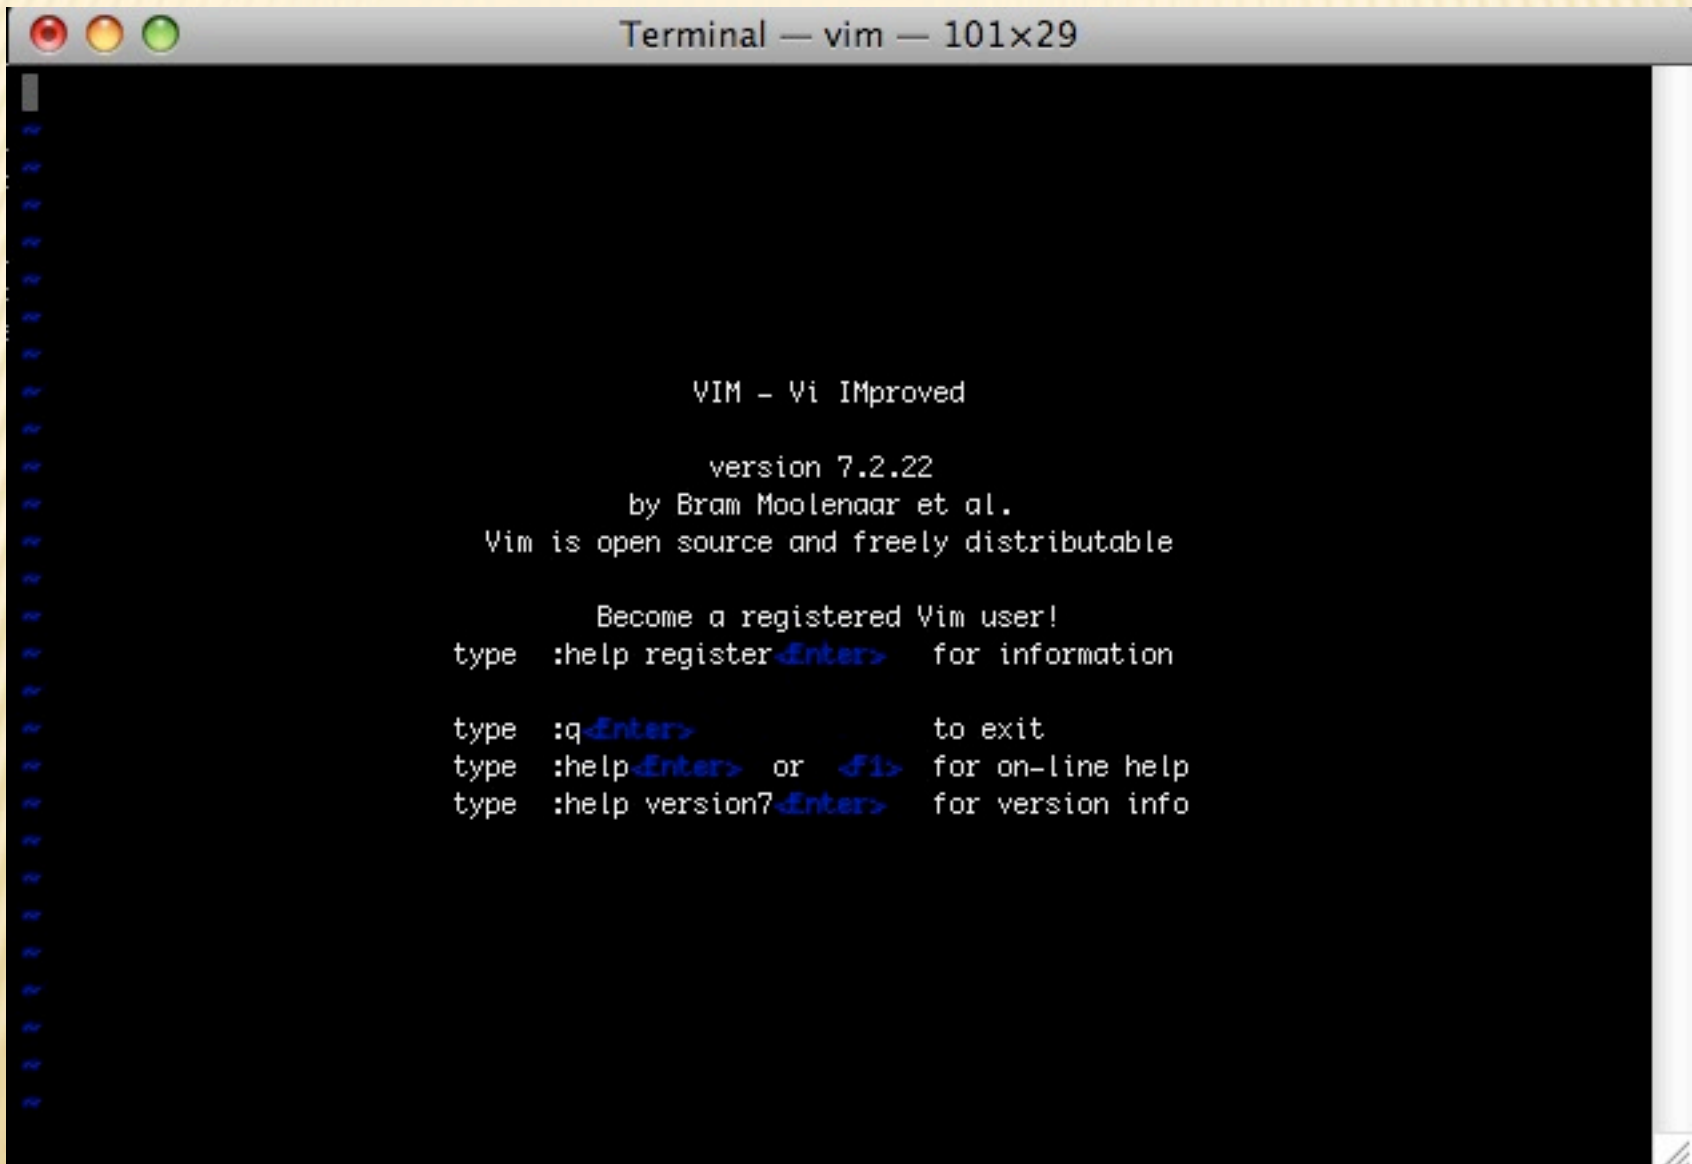
IBGP OVER OSPFV3



IBGP SOURCES OFF LOOPBACKS



OPEN AN EDITOR FOR COMMON CONFIGS

A terminal window titled "Terminal — vim — 101x29" with three window control buttons (red, yellow, green) in the top-left corner. The terminal displays the Vim startup screen with the following text:

```
VIM - Vi IMproved  
  
version 7.2.22  
by Bram Moolenaar et al.  
Vim is open source and freely distributable  
  
Become a registered Vim user!  
type  :help register<Enter>  for information  
  
type  :q<Enter>                to exit  
type  :help<Enter> or <F1>    for on-line help  
type  :help version7<Enter>  for version info
```

IBGP CONFIGURATION

- Remember iBGP is going to handle connected interfaces (except loopbacks)
- We use a route-map to do this:
 - + route-map redist-connected-v6 deny 10
 - × match interface Loopback0
 - + route-map redist-connected-v6 permit 20
 - × match ipv6 address matchall
 - × set community 6:1

BASIC IPV6 BGP CONFIG

- + router bgp XXXX <- your ASN
- + address-family ipv6
- + network 2607:ffff::/32 <- Your block
- + neighbor CORE-v6 peer-group
- + neighbor CORE-v6 remote-as XXXX < your ASN
- + neighbor CORE-v6 soft-reconfiguration inbound
- + neighbor CORE-v6 update-source Loopback0
- + neighbor CORE-v6 send-community
- + neighbor CORE-v6 next-hop-self
- + redistribute connected route-map redistrib-connected-v6
- + no synchronization

MAKE A LIST OF ROUTER LOOPBACKS

- 2607:ffff:66
- 2607:ffff:67
- 2607:ffff:92
- 2607:ffff:95
- 2607:ffff:247
- 2607:ffff:251
- 2607:ffff:252
- 2607:ffff:253
- 2607:ffff:254

BUILD NEIGHBOR CONFIG

- Convert to internal neighbor statements
 - + neighbor 2607:ffff:66 peer-group CORE-v6
 - + neighbor 2607:ffff:67 peer-group CORE-v6
 - + neighbor 2607:ffff:92 peer-group CORE-v6
 - + neighbor 2607:ffff:95 peer-group CORE-v6
 - + neighbor 2607:ffff:251 peer-group CORE-v6
 - + neighbor 2607:ffff:252 peer-group CORE-v6
 - + neighbor 2607:ffff:253 peer-group CORE-v6
 - + neighbor 2607:ffff:254 peer-group CORE-v6

CONFIG FILE TO PUSH

```
config t
!
route-map redist-connected-v6 deny 10
  match interface Loopback0
!
route-map redist-connected-v6 permit 20
  match ipv6 address matchall
  set community 6:1
!
router bgp 23342
  address-family ipv6
  network 2607:F3A0::/32
!
  neighbor ul-inet-core-v6 peer-group
  neighbor ul-inet-core-v6 remote-as 23342
  neighbor ul-inet-core-v6 update-source Loopback0
  neighbor ul-inet-core-v6 send-community
  neighbor ul-inet-core-v6 next-hop-set
  neighbor ul-inet-core-v6 soft-reconfiguration inbound
!
  neighbor 2607:F3A0::66 peer-group ul-inet-core-v6
  neighbor 2607:F3A0::67 peer-group ul-inet-core-v6
  neighbor 2607:F3A0::92 peer-group ul-inet-core-v6
  neighbor 2607:F3A0::95 peer-group ul-inet-core-v6
  neighbor 2607:F3A0::251 peer-group ul-inet-core-v6
  neighbor 2607:F3A0::251 peer-group ul-inet-core-v6
  neighbor 2607:F3A0::252 peer-group ul-inet-core-v6
  neighbor 2607:F3A0::253 peer-group ul-inet-core-v6
  neighbor 2607:F3A0::254 peer-group ul-inet-core-v6
!
  redistribute connected route-map redist-connected-v6
  no synchronization
!
```


PUSHING THE CONFIG FILE

- At this point we can either push the config file as is or wait until we have the peering peer-group defined as well.
- For the sake of simplicity, let's push this now
- Push using ssh, telnet, rancid, etc.

- NOTE: Remove the neighbor statement to yourself for each of the routers.

IBGP SESSIONS COME UP

- At this point you will only see the connected exchange interfaces in the table.

```
br01-paix-pao#sh bgp ipv6 u s | in 23342
BGP router identifier 209.237.224.247, local AS number 23342
2607:F3A0::66 4 23342 2185152 2767004 4441680 0 0 1w2d 27
2607:F3A0::67 4 23342 2396394 2435359 4441680 0 0 25w5d 1527
2607:F3A0::92 4 23342 464884 2432269 4441680 0 0 1w2d 7
2607:F3A0::95 4 23342 464903 2432195 4441680 0 0 1w2d 6
2607:F3A0::251 4 23342 779495 2433652 4441680 0 0 1w2d 11
2607:F3A0::252 4 23342 464920 2432188 4441680 0 0 1w2d 11
2607:F3A0::253 4 23342 2330334 2405038 4441680 0 0 1w2d 1651
2607:F3A0::254 4 23342 464960 2433266 4441680 0 0 1w2d 12
br01-paix-pao#
```

WHERE ARE WE?

- ✓ Obtain an IPv6 allocation
- ✓ My first IPv6 packets
- ✓ Dual stacking the backbone
- ✓ OSPFv3 implementation
- ✓ IPv6 iBGP implementation
 - IPv6 eBGP implementation
 - IPv6 peering
 - Extending dual stack from the backbone
 - IPv6 for your customers

WHERE ARE WE CONT.

- IPv6 configured on all exchange interfaces
- IPv6 configured on all loopback interfaces
- IPv6 configured on all connected interfaces between backbone routers.
- OSPFv3 configured on Loopbacks (/128s)
- OSPFv3 configured on connected (/64s)
- All inter-AS routers are exchanging IPv6 BGP routes
- OSPFv3 is managing iBGP routing based on next-hop (Loopback0s)

CONFIGURING PEERS

- We've done all this work, and still can't reach the outside world!
 - + We need a peering peer-group and a peer!
 - + neighbor PEERS-v6 peer-group
 - + neighbor PEERS-v6 soft-reconfiguration inbound
 - + neighbor PEERS-v6 prefix-list Sanity-v6 in
 - + neighbor PEERS-v6 prefix-list Sanity-v6 out
 - + neighbor PEERS-v6 route-map PEER-IN-v6 in
 - + neighbor PEERS-v6 route-map PEER-OUT-v6 out

CONFIGURING PEERS BASIC SANITY

- Basic sanity prefix-list
 - + ipv6 prefix-list Sanity-v6
 - × seq 5 permit ::/0 ge 16 le 48
 - × seq 10 deny ::/0 le 128
- Don't redistribute peering points connecteds
 - + ipv6 prefix-list PEERINGPOINTS
 - × seq 5 permit 2001:504:0:1::/64
 - × seq 10 permit 2001:504:D:1::/64
 - × seq 15 permit 2001:504:13:1::/64
 - × seq 20 permit 2001:504:0:3::/64
 - × seq 25 permit 2001:504:0:2::/64
 - × seq 30 etc....

CONFIGURE PEERS CONT.

- Create a list of your ASNs IPv6 prefix(es)
 - + ipv6 prefix-list MINE seq 5 permit 2607:ffff::/32
- Create a route-map to apply outbound
 - + route-map PEER-OUT-v6 deny 5
 - × match ipv6 address prefix-list PEERINGPOINTS
 - + route-map PEER-OUT-v6 permit 10
 - × match community ALL-CUSTOMERS
 - + route-map PEER-OUT-v6 permit 20
 - × match ipv6 address prefix-list MINE

CONFIGURE PEERS CONT.

- Create a route-map to apply inbound
 - + route-map PEER-IN-v6 permit 10
 - × match ip address prefix-list Sanity-v6
 - × set local-preference 400
 - × set community 8038:117 Use the same community for peers

TURN UP A PEER!

- Send e-mail to peering@he.net
 - + HE Peering,
 - + I have completed the dual-stack of my backbone and am ready to turn up IPv6 peering. I would greatly appreciate turning up sessions with you at our common locations. Also, I would appreciate the full IPv6 BGP table.
 - + Cheers,
 - + Aaron
- Remember to attach your peering info file

TURN UP A PEER! CONT.

```
Terminal — ssh — 124x84
From: Aaron Hughes <aaron@unitedlayer.com>
To: peering@he.net
Cc: peering@unitedlayer.com
Bcc:
Subject: 6909 HE / 23342 UnitedLayer IPv6 peering
Reply-To:

HE Peering,

I am have completed the dual-stack of my backbone and am ready to turn up IPv6 peering. I would greatly appreciate turning up sessions with you at all of our common locations. Also, I would appreciate a full IPv6 table.

UL Information:
http://as23342.peeringdb.com/
AS: 23342
AS-SET: ALTDB AS-UNITEDLAYER
Equinix Ashburn IPv4 : 206.223.115.154
Equinix Ashburn IPv6 : 2001:504:0:2:0:2:3342:1
Equinix San Jose IPv4 : 206.223.116.45
Equinix San Jose IPv6 : 2001:504:0:1:0:2:3342:1
Equinix Los Angeles IPv4 : 206.223.123.100
Equinix Los Angeles IPv6 : 2001:504:0:3:0:2:3342:1
PAIX Palo Alto IPv4 : 198.32.176.7
PAIX Palo Alto IPv6 : 2001:504:0:1:1007
Any2 Los Angeles IPv4 : 206.223.143.17
Any2 Los Angeles IPv6 : 2001:504:13:130
SFPIX San Francisco : 206.197.187.3
Max-Prefixes-v4: 500
Max-Prefixes-v6: 20
Peering Contact: peering@unitedlayer.com
NOC Contact: support@unitedlayer.com
NOC Phone: +1-415-349-2100
No ad5 requirement.

—

Aaron Hughes
Facility Security Officer
+1-415-349-2120
aaron@unitedlayer.com
http://www.unitedlayer.com/
—
—
—
```

TURN UP A PEER! CONT.

- Receive reply with 'sessions have been configured'

```

From: Bob Fisher <peerings@hawaii.net>
Subject: [REDACTED] Re: IPv6 Transit request for additional locations - AS 2342 (Bridget Layer)
To: walt@hawaii.net
Cc: walt@hawaii.net
Reply-To: peerings@hawaii.net
Date: Tue, 22 Jul 2009 15:29:04 -0700
Subject: RE: Tickets
Message-ID: <226760244.29129@hawaii.net>

Hi Aaron, I have setup sessions to the following addresses:

Equinix SJC: 2081.594.802.812:3342:1
Equinix ADR: 2081.594.802.812:3342:1

If you would like to have any more configured please let us know. Our
information is below, if you could drop us a line when you bring them
live that would be great. Thanks!

Hurricane Electric Peering Information :

Contact Address:
Hurricane Electric
708 Mission Ct
Fremont, CA 94539

NOC:
Phone: 510-598-4100
Fax: 510-598-4155
Email: net@hawaii.net / peerings@hawaii.net

RA Information:
ASN: AS6339
AS-NAME: AS-HURRICANE

Exchange Point Connections:

```

ISP	Status	Speed	IPv4	IPv6
EQUINIX-ADR	UP	10GigE	208.223.115.37	2081:594:802::6929::1
EQUINIX-CHI	UP	10GigE	208.223.119.37	2081:594:804::6929::1
EQUINIX-DAL	UP	10GigE	208.223.119.37	2081:594:805::6929::1
EQUINIX-LAX	UP	10GigE	208.223.123.37	2081:594:803::6929::1
EQUINIX-SJC	UP	10GigE	208.223.116.37	2081:594:801::6929::1
LINK	UP	10GigE	195.80.224.21	2081:718:4:8::1b1b::1
LONAP	UP	GigE	193.203.5.129	2081:718:17::1a5b::1
AMS-IX	UP	10GigE	195.89.145.150	2081:718:1::c500:6929::1
NL-IX	UP	GigE	193.239.156.14	2081:718:13::1a500:6929::1
PAIX Paris Airtel	UP	10GigE	196.32.176.29	2081:594:801:128
PAIX New York	UP	10GigE	196.32.118.57	2081:594:81:139
NYIX	UP	10GigE	196.32.168.61	2081:594:13::c500:6929::1
LAIIX	UP	GigE	196.32.146.50	2081:594:800::c500:6929::1
NYCX	UP	GigE	196.32.229.22	
EDGEPE	UP	100BT		2081:458:25:2::500
SIX	UP	10GigE	196.32.130.40	2081:478:130::40
PARAP	UP	10GigE	62.35.254.111	2081:868:806::6929::1
DE-CIX	UP	10GigE	90.81.192.172	2081:718:100b:88::1
WITA	UP	10GigE	196.32.124.176	2081:478:124::176
AMS2-LAX	UP	10GigE	208.223.143.122	2081:594:13:808:808:14

TURN UP A PEER! CONT.

- Now we configure our side of the session:
 - + router bgp XXXX
 - + address-family ipv6
 - + neighbor 2001:504:D::10 remote-as 6939
 - + neighbor 2001:504:D::10 peer-group PEERS-v6
 - + neighbor 2001:504:D::10 description HE

- Router log should show:
 - + %BGP-5-ADJCHANGE: neighbor 2001:504:D::10 Up

TEST THE SESSION

- Make sure things look good:
 - + br01-paix-pao#sh bgp ipv6 u s | in 2001:504:D::10
 - + 2001:504:D::10 4 6939 446117 168688 4449635 0 0 2m 1793
- Session is up and we see 1793 prefixes!
 - + br01-paix-pao#sh bgp ipv6 u ne 2001:504:D::10 ad
 - + BGP table version is 4449635, local router ID is 209.237.224.247 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, r RIB-failure, S Stale
 - + Origin codes: i - IGP, e - EGP, ? - incomplete
 - + Network Next Hop Metric LocPrf Weight Path
 - + *> 2607:F3A0::/32 :: 0 32768 i
 - + Total number of prefixes 1
- I'm advertising my /32!

TEST REACHABILITY

■ Now we can finally reach the outside world!

■ traceroute ipv6 arin.net. Translating "arin.net."...domain server (209.237.230.11) [OK]

Type escape sequence to abort. Tracing the route to arin.net (2001:500:4:13::80)

```

1 paix.ipv6.he.net (2001:504:D::10) 4 msec 0 msec 0 msec
2 10gigabitethernet2-4.core1.ash1.he.net (2001:470:0:35::2) [AS 6939] 76 msec 76 msec 72 msec
3 equi6ix-ash.arin.net (2001:504:0:2:0:1:745:1) 76 msec 76 msec 76 msec
4 2001:500:4:10::12 [AS 10745] 76 msec 76 msec 76 msec
5 2001:500:4:11::2 [AS 10745] 80 msec 88 msec 80 msec
6 * * *
```

traceroute ipv6 ripe.net. Translating "ripe.net."...domain server (209.237.230.11) [OK]

Type escape sequence to abort. Tracing the route to ripe.net (2001:610:240:11::C100:1319)

```

1 paix.ipv6.he.net (2001:504:D::10) 0 msec 4 msec 0 msec
2 10gigabitethernet4-1.core1.sjc2.he.net (2001:470:0:32::1) [AS 6939] 0 msec 0 msec 0 msec
3 10gigabitethernet1-3.core1.nyc4.he.net (2001:470:0:33::2) [AS 6939] 80 msec 80 msec 96 msec
4 10gigabitethernet1-2.core1.lon1.he.net (2001:470:0:3E::2) [AS 6939] 156 msec 148 msec 148 ms
5 10gigabitethernet1-1.core1.amsl.he.net (2001:470:0:3F::2) [AS 6939] 156 msec 156 msec 156 ms
6 gw.ipv6.amsix.nikrtr.ripe.net (2001:7F8:1::A500:3333:1) [AS 1200] 156 msec 156 msec 160 msec
7 gw.ipv6.transit.nsrp.ripe.net (2001:610:240:101::1) [AS 3333] 156 msec 156 msec 156 msec
8 ripe.net (2001:610:240:11::C100:1319) [AS 3333] 156 msec 156 msec 160 msec
```

SETUP FIRST PEER CONT.

- Setup the others sessions and look at the table

```

br01-paix-pao#sh bgp ipv6 u
BGP table version is 4449688, local router ID is 209.237.224.247
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop           Metric LocPrf Weight Path
 * 2001::/32
 *> |             2001:504:::4C           400           0 12909 24785 12573 ?
 * |             2001:504:::10           1 400           0 6939 i
 * |             2687:F3A8::67           1 400           0 6939 i
 * |             2687:F3A8::253         1 400           0 6939 i
 *> | 2001:200::/32 2001:504:::10           400           0 6939 2500 i
 * |             2687:F3A8::67           0 400           0 6939 2500 i
 * |             2687:F3A8::253         0 400           0 6939 2500 i
 *> | 2001:200:136::/40
 * |             2001:504:::10           400           0 6939 2516 7668 9367 i
 * |             2687:F3A8::67           0 400           0 6939 2516 7668 9367 i
 * |             2687:F3A8::253         0 400           0 6939 2516 7668 9367 i
 * | 2001:200:600::/40
 * |             2687:F3A8::67           0 400           0 6939 2516 7667 i
 *> |             2001:504:::10           400           0 6939 2516 7667 i
 * |             2687:F3A8::253         0 400           0 6939 2516 7667 i
 * | 2001:200:900::/40
 * |             2687:F3A8::67           0 400           0 6939 2516 7668 i
 *> |             2001:504:::10           400           0 6939 2516 7668 i
 * |             2687:F3A8::253         0 400           0 6939 2516 7668 i
 * | 2001:200:A000::/35
 * |             2001:504:::10           400           0 6939 3257 2497 4690 i
 *> |             2687:F3A8::67           0 400           0 19151 2497 4690 i
 *> | 2001:200:C000::/35
 * |             2001:504:::10           400           0 6939 2500 23634 i
 * |             2687:F3A8::67           0 400           0 6939 2500 23634 i
 * |             2687:F3A8::253         0 400           0 6939 2500 23634 i
 *> | 2001:200:E000::/35
 * |             2001:504:::10           400           0 6939 4635 7668 i
 * |             2687:F3A8::67           0 400           0 6939 4635 7668 i
 * |             2687:F3A8::253         0 400           0 6939 4635 7668 i
 * | 2001:200::/32 2687:F3A8::67           0 400           0 6939 23911 9000 30035 7610 i
 *> |             2001:504:::10           400           0 6939 23911 9000 30035 7610 i
 * |             2687:F3A8::253         0 400           0 6939 23911 9000 30035 7610 i
 * | 2001:218::/32 2687:F3A8::67           0 400           0 19151 2914 i
 *> |             2001:504:::10           400           0 6939 2914 i
 * |             2687:F3A8::253         0 400           0 6939 2914 i
 *> | 2001:220::/35 2001:504:::10           400           0 6939 2500 7668 9278 i
 * |             2687:F3A8::253         0 400           0 6939 2500 7668 9278 i
 * |             2687:F3A8::67           0 400           0 6939 2500 7668 9278 i
 *> | 2001:220:2000::/35
 * |             2001:504:::10           400           0 6939 4635 23911 7668 9278 38128 i
 * |             2687:F3A8::253         0 400           0 6939 4635 23911 7668 9278 38128 i
 * |             2687:F3A8::67           0 400           0 6939 4635 23911 7668 9278 38128 i

```


WHERE ARE WE?

- ✓ Obtain an IPv6 allocation
- ✓ My first IPv6 packets
- ✓ Dual stacking the backbone
- ✓ OSPFv3 implementation
- ✓ IPv6 iBGP implementation
- ✓ IPv6 eBGP implementation
- ✓ IPv6 peering
 - Extending dual stack from the backbone
 - IPv6 for your customers

ATTACHING A HOST TO THE V6 NETWORK

- Now that we have a functioning IPv6 network, let's get a host on-line to play with.
 - + Something non-production?
 - + A small segment of the office?
 - + A development environment?
 - + Your desktop or laptop?

- Keep in mind that you are not yet monitoring or supporting your IPv6 network.

NETWORK CONFIG FOR FIRST HOST

- Find the interface on the network the host is connected to via IPv4. In this case we have chosen ns0 as our dev box:

```
br02-sf7-200p-sfo#sh ip route ns0
Routing entry for 209.237.230.32/28
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Redistributing via ospf 23342
  Advertised by ospf 23342 metric-type 1 subnets
  Routing Descriptor Blocks:
  * directly connected, via Vlan705
    Route metric is 0, traffic share count is 1
```


NETWORK CONFIG FOR FIRST HOST

- Look at the existing v4 config

```
interface Vlan705
description [UL:VLAN] Dev nameservers
ip address 209.237.230.44 255.255.255.240
no ip redirects
no ip proxy-arp
!
```


NETWORK CONFIG FOR THE FIRST HOST

- Add the IPv6 config to the interface of the router

```
config t
int vl705
ipv6 enable
ipv6 address 2607:F3A0:0:1002::2/64
```

- Wait a few seconds..

```
eth0    Link encap:Ethernet  HWaddr 00:30:48:42:59:EA
        inet addr:209.237.230.37  Bcast:209.237.230.47  Mask:255.255.255.240
        inet6 addr: 2607:f3a0:0:1002:230:48ff:fe42:59ea/64 Scope:Global
        inet6 addr: fe80::230:48ff:fe42:59ea/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        RX packets:69869405 errors:0 dropped:0 overruns:0 frame:0
        TX packets:63543470 errors:1160015 dropped:0 overruns:0 carrier:1160015
        collisions:1379164 txqueuelen:100
        RX bytes:936819857 (893.4 MiB)  TX bytes:1187390989 (1.1 GiB)
        Base address:0xa000 Memory:ec000000-ec020000
```

- Poof!

NETWORK CONFIG FOR FIRST HOST

■ A little testing...

```
root@ns0:/home/aaronh> ping6 bind.com
PING bind.com(trace.ipv6.bind.com) 56 data bytes
64 bytes from trace.ipv6.bind.com: icmp_seq=1 ttl=61 time=46.7 ms
64 bytes from trace.ipv6.bind.com: icmp_seq=2 ttl=61 time=44.9 ms
64 bytes from trace.ipv6.bind.com: icmp_seq=3 ttl=61 time=45.8 ms
64 bytes from trace.ipv6.bind.com: icmp_seq=4 ttl=61 time=40.0 ms
64 bytes from trace.ipv6.bind.com: icmp_seq=5 ttl=61 time=43.0 ms
64 bytes from trace.ipv6.bind.com: icmp_seq=6 ttl=61 time=46.9 ms
64 bytes from trace.ipv6.bind.com: icmp_seq=7 ttl=61 time=43.6 ms

--- bind.com ping statistics ---
7 packets transmitted, 7 received, 0% packet loss, time 9416ms
rtt min/avg/max/mdev = 40.092/44.482/46.971/2.274 ms
root@ns0:/home/aaronh> ping6 ripe.net
PING ripe.net(aquila.ripe.net) 56 data bytes
64 bytes from aquila.ripe.net: icmp_seq=1 ttl=56 time=151 ms
64 bytes from aquila.ripe.net: icmp_seq=2 ttl=56 time=151 ms
64 bytes from aquila.ripe.net: icmp_seq=3 ttl=56 time=151 ms

--- ripe.net ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2001ms
rtt min/avg/max/mdev = 151.486/151.610/151.700/0.330 ms
root@ns0:/home/aaronh> traceroute6 arin.net
traceroute to arin.net (2001:500:4:13::81) from 2607:f3a0:0:1002:230:48ff:fe42:59ea, 30 hops max, 16 byte packets
 1  Vlan705.br02-200p-sfo.ipv6.unitedlayer.com (2607:f3a0:0:1002::2)  1.64 ms  0.31 ms  0.444 ms
 2  g0-1-0.br01-eqx-ash.ipv6.unitedlayer.com (2607:f3a0:0:c::2)  70.646 ms  70.597 ms  70.554 ms
 3  equi6ix-ash.arin.net (2001:504:0:2:0:1:745:1)  71.954 ms  71.704 ms  71.91 ms
 4  2001:500:4:10::12 (2001:500:4:10::12)  72.392 ms  72.363 ms  72.683 ms
 5  2001:500:4:11::2 (2001:500:4:11::2)  75.807 ms  75.738 ms  75.55 ms
 6  *
```

NETWORK CONFIG FOR FIRST HOST

- From the routers perspective:

```
br02-sf7-200p-sfo#ping 2607:F3A0:0:1002:230:48FF:FE42:59EA

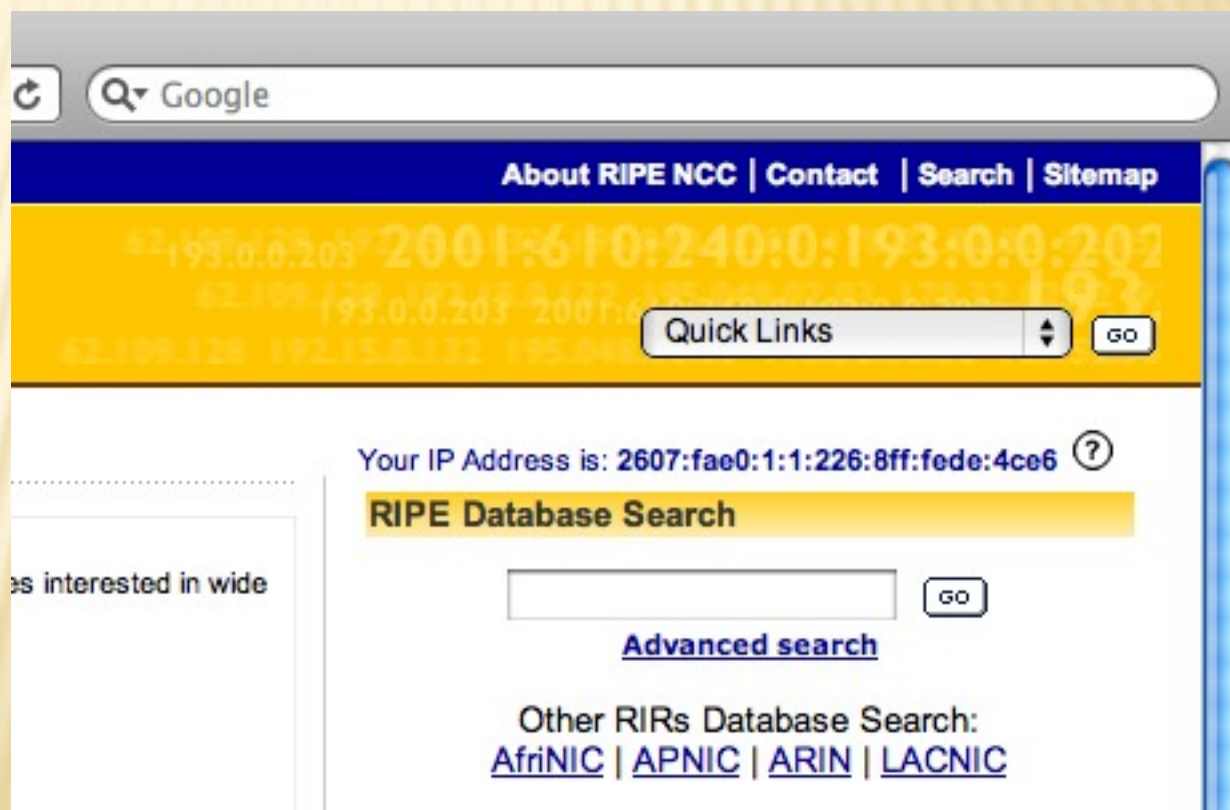
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2607:F3A0:0:1002:230:48FF:FE42:59EA, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/1/8 ms
br02-sf7-200p-sfo#sh ipv6 neighbors | in V1705
2607:F3A0:0:1002:230:48FF:FE42:59EA      0 0030.4842.59ea REACH V1705
FE80::230:48FF:FE42:59EA              1 0030.4842.59ea STALE V1705
br02-sf7-200p-sfo#
```

- This will be the first connected IPv6 prefix:

```
br02-sf7-200p-sfo#sh bgp ipv6 u 2607:F3A0:0:1002::/64
BGP routing table entry for 2607:F3A0:0:1002::/64, version 1236
Paths: (2 available, best #1, table Global-IPv6-Table)
  Advertised to update-groups:
    2
  Local
  :: from 0.0.0.0 (209.237.224.252)
    Origin incomplete, metric 0, localpref 100, weight 32768, valid, sourced, best
    Community: 6:1
  Local, (received & used)
  2607:F3A0::254 (metric 1) from 2607:F3A0::254 (209.237.224.254)
    Origin incomplete, metric 0, localpref 100, valid, internal
    Community: 6:1
br02-sf7-200p-sfo#
```


TESTING THE HOST

- If you have a web-browser on this machine, try <http://ripe.net/> Look at the top right corner.



A QUICK SLAAC NOTE

- It is important to note that all hosts setup for autoconfig (default on all UNIX OS's and any other current OS) will receive an IPv6 address if they are connected to the same subnet. This means if your dev box is on the same subnet as production boxes, they too will autoconfig via SLAAC and receive a public IPv6 address.

TIME TO ADD NAMESERVICE

■ Add DNS

+ Reverse:

× a.e.9.5.2.4.e.f.f.f.8.4.0.3.2.0.2.0.1 IN PTR ns0.ipv6

+ Forward:

× ns0 IN A 209.237.230.37

× ns0 IN AAAA 2607:f3a0:0:1002:230:48ff:fe42:59ea

+ rndc reload and test!

ADDING DNS CONT.

- Our first host on IPv6!

```
br02-sf7-200p-sfo#ping ns0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2607:F3A0:0:1002:230:48FF:FE42:59EA, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/1/8 ms
br02-sf7-200p-sfo#
```


SECURITY NOTE

- This machine is now globally accessible on the IPv6 Internet with no filters in place. It is listening on all ports daemons are running on.
- Everything connected to this VLAN or interface with autoconf enabled has an IPv6 address.
- Use `show ipv6 neighbors` to view configured hosts.
- If you have a security policy implemented for IPv4 you will need to implement the IPv6 version of that at this time.

WHAT'S NEXT? MORE PEERING!

Reset 118 sessions on this page | [Peers only](#) | [Customers only](#) | [Upstream only](#) | [Show Down Sessions](#) | [IPv6 Only](#) | [IPv4](#)

ASN	Company Name	Location	IP Address	AS-SET	Router	Peer IP	Peer AS
42	Packet Clearing House	PAIX PAO	2001:504:d::35	AS-PCH	br01-paix-pao	1001:504:0:2::7385:1	br01-eqx-ash
42	Packet Clearing House	Equinix ASH	2001:504:0:2::42:1	AS-PCH	br01-eqx-ash	1001:504:0:1::7385:1	br01-eqx-sjc
293	Energy Sciences Network	Equinix SIC	2001:504:0:1::293:1	AS-ESNET	br01-eqx-sjc	1001:504:D::1D	br01-paix-pao
293	Energy Sciences Network	PAIX PAO	2001:504:d::a	AS-ESNET	br01-paix-pao	1001:504:d::32	br01-paix-pao
293	Energy Sciences Network	Equinix ASH	2001:504:0:2::293:1	AS-ESNET	br01-eqx-ash	1001:504:13::1e	br01-1w-lax
558	Net2EZ	ANY2 LAX	2001:504:13::48	AS-NET2EZ	br01-1w-lax	1001:504:D::B1	br01-paix-pao
1280	Internet Systems Consortium, Inc.	Equinix SIC	2001:504:0:1::1280:1		br01-eqx-sjc	1001:504:13::23	br01-1w-lax
1280	Internet Systems Consortium, Inc.	PAIX PAO	2001:504:D::12		br01-paix-pao	1001:504:0:2::7784:1	br01-eqx-ash
1280	Internet Systems Consortium, Inc.	PAIX PAO	2001:504:D::1		br01-paix-pao	1001:504:0:2::8001:1	br01-eqx-ash
2516	KDDI	PAIX PAO	2001:504:D::D		br01-paix-pao	1607:F3A0:0:1A::2	AS8038
3043	Internet Operating Services LLC	ANY2 LAX	2001:504:13::26		br01-1w-lax	1001:504:13::46	AS-MICROSOFT
3043	Internet Operating Services LLC	Equinix ASH	2001:504:0:2::3043:1		br01-eqx-ash	1001:504:13::61	br01-1w-lax
3320	Deutsche Telekom	Equinix ASH	2001:504:0:2::3320:1	AS3320-AS-DTAG-V6	br01-eqx-ash	1001:504:13::2A	br01-1w-lax
3320	Deutsche Telekom	PAIX PAO	2001:504:D::39	AS3320-AS-DTAG-V6	br01-paix-pao	1001:504:13::29	br01-1w-lax
3856	Packet Clearing House	Equinix ASH	2001:504:0:2::3856:1	AS-PCH	br01-eqx-ash	1001:504:0:1::8121:1	br01-eqx-sjc
3856	Packet Clearing House	PAIX PAO	2001:504:d::f9	AS-PCH	br01-paix-pao	1001:504:0:2::8121:1	br01-eqx-ash
4323	TWTelecom	Equinix ASH	2001:504:0:2::4323:1		br01-eqx-ash	1001:504:D::57	br01-paix-pao
4323	TWTelecom	Equinix LAX	2001:504:0:3::4323:1		br01-1w-lax	1001:504:0:2::8218:1	AS-NEOT
4323	TWTelecom	Equinix SIC	2001:504:0:1::4323:1		br01-eqx-sjc	1001:504:d::5f	AS-NEOT
4589	Easynet	Equinix ASH	2001:504:0:2::4589:1	AS-EASYNET	br01-eqx-ash	1001:504:d::60	AS-OTEL-SET
4589	Easynet	ANY2 LAX	2001:504:13::1d	AS-EASYNET	br01-1w-lax	1001:504:0:2::8781:1	AS-OTEL-SET
4589	Easynet	Equinix SIC	2001:504:0:1::4589:1	AS-EASYNET	br01-eqx-sjc	1001:504:0:2::8881:1	AS-VT-TRANSIT
4589	Easynet	PAIX PAO	2001:504:d::4a	AS-EASYNET	br01-paix-pao	1001:504:13::39	AS-RETN
4648	Telecom New Zealand	PAIX PAO	2001:504:d::61		br01-paix-pao	1001:504:0:2::9002:1	AS-RETN
4648	Telecom New Zealand	Equinix LAX	2001:504:0:3::4648:1		br01-1w-lax	1001:504:D::AE	br01-paix-pao
4657	StarHub SiX	PAIX PAO	2001:504:d::2c		br01-paix-pao	1001:504:13::33	br01-1w-lax
4739	Intermode Systems	Equinix LAX	2001:504:0:3::4739:1		br01-1w-lax	1001:504:13::40	br01-1w-lax
4739	Intermode Systems	Equinix SIC	2001:504:0:1::4739:1		br01-eqx-sjc	1001:504:13::40	br01-1w-lax
4739	Intermode Systems	PAIX PAO	2001:504:d::4f		br01-paix-pao	1001:504:2:0:1:310:1	AS-YAHOO
4739	Intermode Systems	ANY2 LAX	2001:504:13::13		br01-1w-lax	1001:504:1:0:1:310:1	AS-YAHOO
4826	Vocus Communications	Equinix SIC	2001:504:0:1::4826:1		br01-eqx-sjc	1001:504:0:1:0:1:310:1	AS-YAHOO
4826	Vocus Communications	PAIX PAO	2001:504:d::86		br01-paix-pao	1001:504:2:0:1:745:1	br01-eqx-ash
4826	Vocus Communications	ANY2 LAX	2001:504:13::3B		br01-1w-lax	1001:504:13::1f	br01-1w-lax
6939	Harcane Electric, Inc.	ANY2 LAX	2001:504:13::1A		br01-1w-lax	1001:504:2:0:1:1666:1	br01-eqx-ash
6939	Harcane Electric, Inc.	Equinix SIC	2001:504:0:1::6939:1		br01-eqx-sjc	1001:504:13::65	br01-1w-lax
6939	Harcane Electric, Inc.	PAIX PAO	2001:504:D::10		br01-paix-pao	1001:504:2:0:1:2989:1	AS-HIGHWINDS

MORE PEERING!

```
br01-paix-pao>sh bgp ipv6 u s
BGP router identifier 209.237.224.247, local AS number 23342
BGP table version is 8756433, main routing table version 8756433
2131 network entries using 300471 bytes of memory
9838 path entries using 747688 bytes of memory
201507/1627 BGP path/bestpath attribute entries using 28222100 bytes of memory
19 BGP rrinfo entries using 456 bytes of memory
93144 BGP AS-PATH entries using 2444746 bytes of memory
7200 BGP community entries using 573030 bytes of memory
10 BGP extended community entries using 1370 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 32209941 total bytes of memory
2255 received paths for inbound soft reconfiguration
BGP activity 5188551/4888696 prefixes, 254129137/252934605 paths, scan interval 60 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
2001:504:D::1	4	1200	399077	394839	8756433	0	0	6w6d	10
2001:504:D::A	4	293	575769	542651	8756433	0	0	11w4d	3
2001:504:D::D	4	2516	333532	330018	8756433	0	0	4w1d	21
2001:504:D::10	4	6939	3364618	394918	8756433	0	0	20w0d	1905
2001:504:D::12	4	1200	236227	221549	0	0	0	28w1d	Active
2001:504:D::10	4	7305	332415	366299	8756433	0	0	15w3d	4
2001:504:D::1F	4	15169	820665	776323	8756433	0	0	19w1d	13
2001:504:D::2C	4	4657	130779	129939	8756433	0	0	3w5d	1
2001:504:D::2E	4	38061	168730	130140	8756433	0	0	6w0d	0
2001:504:D::30	4	9924	0	0	0	0	0	never	Active
2001:504:D::32	4	7473	819559	776594	8756433	0	0	26w0d	17
2001:504:D::35	4	42	77171	76621	8756433	0	0	7w4d	15
2001:504:D::38	4	33517	0	0	0	0	0	never	Active
2001:504:D::39	4	3320	332770	330434	8756433	0	0	11w4d	4
2001:504:D::4A	4	4509	370406	365502	8756433	0	0	3w5d	11
2001:504:D::4C	4	12909	384346	336225	8756433	0	0	3d16h	37
2001:504:D::55	4	32354	397626	394915	8756433	0	0	26w0d	1
2001:504:D::57	4	8121	4771007	4658258	8756433	0	0	26w0d	1
2001:504:D::5B	4	64597	363767	1021363	8756433	0	0	11w4d	0
2001:504:D::5F	4	8210	847623	776592	8756433	0	0	26w0d	22
2001:504:D::60	4	8781	200300	278229	8756433	0	0	1w4d	1
2001:504:D::61	4	4648	560906	524339	8756433	0	0	26w0d	2
2001:504:D::70	4	18500	430049	1309564	8756433	0	0	3d23h	2
2001:504:D::86	4	4826	217672	195710	8756433	0	0	3w5d	26
2001:504:D::AE	4	9264	797740	711140	8756433	0	0	26w0d	20
2001:504:D::B1	4	7575	826209	776592	8756433	0	0	26w0d	20
2001:504:D::C1	4	26415	348411	346074	8756433	0	0	6d22h	1
2001:504:D::F9	4	3856	81003	96343	8756433	0	0	4w2d	0
2607:F3A0::66	4	23342	2997638	4784819	8756433	0	0	5w3d	1797
2607:F3A0::67	4	23342	3528416	4347757	8756433	0	0	11w4d	1677
2607:F3A0::92	4	23342	694447	4345015	8756433	0	0	19w3d	6
2607:F3A0::95	4	23342	694335	4344947	8756433	0	0	19w3d	7
2607:F3A0::251	4	23342	1102036	4346399	8756433	0	0	19w3d	13
2607:F3A0::252	4	23342	694373	4344932	8756433	0	0	19w3d	9
2607:F3A0::253	4	23342	3448942	4318983	8756433	0	0	19w3d	1833
2607:F3A0::254	4	23342	694395	4346011	8756433	0	0	19w3d	12

```
20521 BGP AS-PATH entries using 567390 bytes of memory
1176 BGP community entries using 40370 bytes of memory
9 BGP extended community entries using 1346 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 6534297 total bytes of memory
4163 received paths for inbound soft reconfiguration
BGP activity 1124032/1047355 prefixes, 6632409/6524326 paths, scan interval 60 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
2001:504:0:2::42:1	4	42	77190	77171	622503	0	0	7w4d	15
2001:504:0:2::293:1	4	293	240320	232735	622503	0	0	1w4d	3
2001:504:0:2::2641:1	4	26415	100292	100295	622503	0	0	7w3d	3
2001:504:0:2::3043:1	4	3043	93244	93647	622503	0	0	1w4d	1
2001:504:0:2::3320:1	4	3320	113573	113593	0	0	0	200h	Active
2001:504:0:2::3056:1	4	3056	77004	77020	622503	0	0	7w3d	0
2001:504:0:2::4323:1	4	4323	241097	232910	622503	0	0	11w3d	12
2001:504:0:2::4509:1	4	4509	117140	116492	622503	0	0	2w2d	11
2001:504:0:2::6939:1	4	6939	2594280	115950	622503	0	0	3w2d	1904
2001:504:0:2::7305:1	4	7305	104480	115947	622503	0	0	11w3d	4
2001:504:0:2::7704:1	4	7704	241359	231859	622503	0	0	11w3d	1
2001:504:0:2::8001:1	4	8001	231622	231846	622503	0	0	20:16:43	3
2001:504:0:2::8121:1	4	8121	115993	115947	622503	0	0	11w3d	1
2001:504:0:2::8218:1	4	8218	244207	232900	622503	0	0	11w3d	22
2001:504:0:2::8701:1	4	8701	0	0	0	0	0	never	Active
2001:504:0:2::8001:1	4	8001	233144	116472	622503	0	0	11w3d	12
2001:504:0:2::9002:1	4	9002	240596	231851	622503	0	0	9w6d	26
2001:504:0:2:0:1:310:1	4	10310	242563	231700	622503	0	0	7w0d	1
2001:504:0:2:0:1:745:1	4	10745	115266	115301	622503	0	0	1d16h	2
2001:504:0:2:0:1:5666:1	4	11566	115961	115964	622503	0	0	7w4d	1
2001:504:0:2:0:1:2909:1	4	12909	132116	116465	622503	0	0	3d18h	37
2001:504:0:2:0:1:5169:1	4	15169	230925	231800	622503	0	0	11w3d	11
Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
2001:504:0:2:0:1:9000:1	4	9000	0	0	0	0	0	never	Active
2001:504:0:2:0:1:9151:1	4	10151	332187	116409	622503	0	0	5w3d	1811
2001:504:0:2:0:2:940:1	4	20940	0	0	0	0	0	never	Active

MORE PEERING

- It's IPv6 peering so today:
 - + Policies are more flexible
 - + Almost all networks have an open IPv6 peering policy
 - + Others will want to peer with you
 - + E-mail everyone regardless of IPv4 peering policy
 - + Watch the IX lists for new IPv6 peers

- Today, every bit moved over IPv6 is FREE!

WHERE ARE WE?

- ✓ Obtain an IPv6 allocation
- ✓ My first IPv6 packets
- ✓ Dual stacking the backbone
- ✓ OSPFv3 implementation
- ✓ IPv6 iBGP implementation
- ✓ IPv6 eBGP implementation
- ✓ IPv6 peering
- ✓ Extending dual stack from the backbone
 - IPv6 for your customers

TURN UP YOUR FIRST CUSTOMER

- I have a customer picked, what now?
- You guessed it, back to router configs!

V4->V6 ROUTE-MAPS

```

!
route-map Customer-In permit 10
description Prepend_CU1
match ip address prefix-list Sanity
match community prependCU1
set local-preference 500
set as-path prepend last-as 1
set community 23342:418
!
route-map Customer-In permit 20
description Prepend_CU2
match ip address prefix-list Sanity
match community prependCU2
set local-preference 500
set as-path prepend last-as 2
set community 23342:418
!
route-map Customer-In permit 30
description Prepend_CU3
match ip address prefix-list Sanity
match community prependCU3
set local-preference 500
set as-path prepend last-as 3
set community 23342:418
!

```

```

!
route-map Customer-In permit 40
description Prepend_UL1
match ip address prefix-list Sanity
match community prependUL1
set local-preference 500
set as-path prepend 23342
set community 23342:418
!
route-map Customer-In permit 50
description Prepend_UL2
match ip address prefix-list Sanity
match community prependUL2
set local-preference 500
set as-path prepend 23342 23342
set community 23342:418
!
route-map Customer-In permit 60
description Prepend_UL3
match ip address prefix-list Sanity
match community prependUL3
set local-preference 500
set as-path prepend 23342 23342 23342
set community 23342:418
!
route-map Customer-In permit 70
description No prepending
match ip address prefix-list Sanity
set local-preference 500
set community 23342:418
!


```

V4 -> V6 ROUTE-MAPS

```

!
route-map Customer-In-v6 permit 10
description Prepend_CU1
match ip address prefix-list Sanity-v6
match community prependCU1
set local-preference 500
set as-path prepend last-as 1
set community 23342:418
!
route-map Customer-In-v6 permit 20
description Prepend_CU2
match ip address prefix-list Sanity-v6
match community prependCU2
set local-preference 500
set as-path prepend last-as 2
set community 23342:418
!
route-map Customer-In-v6 permit 30
description Prepend_CU3
match ip address prefix-list Sanity-v6
match community prependCU3
set local-preference 500
set as-path prepend last-as 3
set community 23342:418
!

```




```


!
route-map Customer-In-v6 permit 40
description Prepend_UL1
match ip address prefix-list Sanity-v6
match community prependUL1
set local-preference 500
set as-path prepend 23342
set community 23342:418
!
route-map Customer-In-v6 permit 50
description Prepend_UL2
match ip address prefix-list Sanity-v6
match community prependUL2
set local-preference 500
set as-path prepend 23342 23342
set community 23342:418
!
route-map Customer-In-v6 permit 60
description Prepend_UL3
match ip address prefix-list Sanity-v6
match community prependUL3
set local-preference 500
set as-path prepend 23342 23342 23342
set community 23342:418
!
route-map Customer-In-v6 permit 70
description No prepending
match ip address prefix-list Sanity-v6
set local-preference 500
set community 23342:418
!

```


V4 -> V6 ROUTE-MAPS



```
!  
route-map Customer-Out deny 10  
  match community BLACKHOLE no-export  
!  
route-map Customer-Out permit 20  
  match ip address prefix-list UL  
  set community none  
!  
route-map Customer-Out permit 30  
  match community ALL-CUSTOMERS ALL-PEERS ALL-TRANSIT  
!
```



```
!  
route-map Customer-Out-v6 deny 10  
  description ->BGP->Match Blackhole  
  match community BLACKHOLE NO-EXPORT  
!  
route-map Customer-Out-v6 permit 30  
  match community ALL-CUSTOMERS ALL-PEERS ALL-TRANSIT  
!  
route-map Customer-Out-v6 permit 40  
  match ipv6 address prefix-list UL  
!
```


V4 -> V6 PEER-GROUPS

```
neighbor CUSTOMERFULL peer-group
neighbor CUSTOMERFULL version 4
neighbor CUSTOMERFULL activate
neighbor CUSTOMERFULL send-community
neighbor CUSTOMERFULL soft-reconfiguration inbound
neighbor CUSTOMERFULL route-map Customer-In in
neighbor CUSTOMERFULL route-map Customer-Out out
neighbor CUSTOMERFULL maximum-prefix 2000
```

```
neighbor CUSTOMERFULLv6 peer-group
neighbor CUSTOMERFULLv6 activate
neighbor CUSTOMERFULLv6 activate
neighbor CUSTOMERFULLv6 send-community
neighbor CUSTOMERFULLv6 soft-reconfiguration inbound
neighbor CUSTOMERFULLv6 route-map Customer-In-v6 in
neighbor CUSTOMERFULLv6 route-map Customer-Out-v6 out
neighbor CUSTOMERFULLv6 maximum-prefix 100
```

IPV6 CUSTOMER TURN UP

IPv4:

```
conf t
router bgp XXXXX
neighbor x.x.x.x remote-as 8038
neighbor x.x.x.x description Customer-Bind
neighbor x.x.x.x peer-group CUSTOMERFULL
```

IPv6:

```
conf t
router bgp XXXXX
neighbor 2607:F3A0:0:1F remote-as 8038
neighbor 2607:F3A0:0:1F description Customer-Bind
neighbor 2607:F3A0:0:1F peer-group CUSTOMERFULLv6
```



```

br01-1w-lax#sh bgp ipv6 u s | be 1F::2
2607:F3A0:0:1F::2
      4 8038 119098 325951 11423932    0    0 1w5d          2
br01-1w-lax#sh bgp ipv6 u ne 2607:F3A0:0:1F::2 route
BGP table version is 11424013, local router ID is 209.237.224.253
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop             Metric LocPrf Weight Path
*> 2607:F3A0:BEEF::/48
                2607:F3A0:0:1F::2
                                0    500    0 8038 i
*> 2607:FAE0::/32   2607:F3A0:0:1F::2
                                0    500    0 8038 i

Total number of prefixes 2
br01-1w-lax#sh bgp ipv6 u 2607:F3A0:BEEF::/48
BGP routing table entry for 2607:F3A0:BEEF::/48, version 11065134
Paths: (4 available, best #3, table Default)
  Advertised to update-groups:
    1          2          3          4
6939 8038
  2001:504:13::1A (FE80::212:F2FF:FE97:5901) from 2001:504:13::1A (216.218.252.178)
  Origin IGP, localpref 400, valid, external
  Community: 23342:117
6939 8038, (received-only)
  2001:504:13::1A (FE80::212:F2FF:FE97:5901) from 2001:504:13::1A (216.218.252.178)
  Origin IGP, localpref 100, valid, external
8038
  2607:F3A0:0:1F::2 (FE80::203:E3FF:FE03:1F1) from 2607:F3A0:0:1F::2 (207.7.159.26)
  Origin IGP, metric 0, localpref 500, valid, external, best
  Community: 23342:418
8038, (received-only)
  2607:F3A0:0:1F::2 (FE80::203:E3FF:FE03:1F1) from 2607:F3A0:0:1F::2 (207.7.159.26)
  Origin IGP, metric 0, localpref 100, valid, external
br01-1w-lax#

```


IPV6 CUSTOMER TURN UP CONT

What are we advertising?

```
br01-1w-lax#
br01-1w-lax#sh bgp ipv6 u ne 2607:F3A0:0:1F::2 ad
BGP table version is 11424103, local router ID is 209.237.224.253
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 2001::/32        2001:504:13::1A      1    400    0 6939 i
*> 2001:200::/32    2001:504:13::1A      400    0 6939 2500 i
*> 2001:200:A000::/35
                    2001:504:13::1A      400    0 6939 2497 4690 i
*> 2001:200:C000::/35
                    2001:504:13::1A      400    0 6939 2500 23634 i
*> 2001:200:E000::/35
                    2001:504:13::1A      400    0 6939 4635 7660 i
*> 2001:218::/32    2001:504:13::1A      400    0 6939 2914 i
*> 2001:220::/35    2001:504:13::1A      400    0 6939 2500 7660 9270 i
*> 2001:220:806::/48
                    2001:504:13::1A      400    0 6939 4635 23911 24489 9270 9687 i
*> 2001:220:2000::/35
                    2001:504:0:3:0:1:9151:1
                                                0    400    0 19151 9304 4635 17579 7660 9270 38128 i
*> 2001:220:4000::/34
                    2001:504:0:3:0:1:9151:1
                                                0    400    0 19151 9304 4635 17579 7660 9270 38128 i
*> 2001:220:8000::/33
                    2001:504:0:3:0:1:9151:1
                                                0    400    0 19151 9304 4635 17579 7660 9270 38128 i
*> 2001:238:700::/41
                    2001:504:13::1A      400    0 6939 3549 9505 i
*> 2001:240::/32    2001:504:13::1A      400    0 6939 2497 i
*> 2001:250::/32    2001:504:13::1A      400    0 6939 23911 4538 i
*> 2001:250:204::/48
```

OTHER WAYS TO DUAL STACK CUSTOMERS

- Connected interface
 - + Colo customers
 - + Metro E
- Static routes
 - + no BGP
 - + Connected /64 + Static /64 -> /48
 - + Add route-map for redist static

WHERE ARE WE?

- ✓ Obtain an IPv6 allocation
- ✓ My first IPv6 packets
- ✓ Dual stacking the backbone
- ✓ OSPFv3 implementation
- ✓ IPv6 iBGP implementation
- ✓ IPv6 eBGP implementation
- ✓ IPv6 peering
- ✓ Extending dual stack from the backbone
- ✓ IPv6 for your customers

NEXT STEPS

- Address your security policy
- Get your operations staff to use IPv6
- Dual stack your NOC
- Get IPv6 objects in monitoring
- Dual stack your offices
- Start updating tools
- Work with all departments on education
- Get a test customer up and running
- Make a longer term plan from here..

CONCLUSION

- Obtaining IPv6 space is easy
- IPv6 BGP multihoming works
- IPv6 generally well supported on routers
- Transit providers are not needed to implement IPv6
- Lots of IPv6 peers out there to connect with
- Dual stacking the backbone will not impact your edges until you are ready
- Implementation is not hard
- Config mostly intuitive and same as IPv4

QUESTIONS?

- Get IPv6 implemented today!