


# IPv6 Deployment

## Case Study



**LINK**  
TECHNOLOGIES INC

# Your Instructor

## ▶ Dennis Burgess

- MikroTik Certified Consultant / Trainer
- Author of “Learn RouterOS” 1<sup>st</sup> and 2<sup>nd</sup> Editions
- MikroTik Dude Consultant
- In the WISP Industry since 2000
- Consulting Since 1997
- Cisco Certified
- Microsoft Certified
- WISP Experience
  - Owned and Operated a WISP Since 2000
  - Deployed Countless MikroTik Based Networks



# Link Technologies, Inc.

- MikroTik System Integrator
  - World Wide Consulting Clients
  - Support MikroTik Networks
  - Cisco, Microsoft, Hotspots, Billing Systems, Network Management, Network Troubleshooting, Network Build outs, and General Network Engineering
  - Engineering/Consulting/Network Design/Support
  - Consulting on Per Hour and Contract basis!
- Manufacturer of the PowerRouter Series of RouterOS Devices – [www.mikrotikrouter.com](http://www.mikrotikrouter.com)
  - Supported/Tested Enterprise & Carrier Grade MikroTik RouterOS x86 Systems made by Link Technologies, Inc..

# What are we going to cover?

---

- What is IPv6
- Common Objections to IPv6
- IPv6 Deployment in 60+ POP Network

# What is IPv6

## IPv4

Deployed 1981

*Address Size:*

32-bit number

*Address Format:*

Dotted Decimal Notation:

192.0.2.76

*Prefix Notation:*

192.0.2.0/24

*Number of Addresses:*

$2^{32} = 4,294,967,296$

## IPv6

Deployed 1999

*Address Size:*

128-bit number

*Address Format:*

Hexadecimal Notation:

2001:0DB8:0234:AB00:0123:4567:8901:ABCD

*Prefix Notation:*

2001:0DB8:0234::/48

*Number of Addresses:*

$2^{128} =$

340,282,366,920,938,463,463,374,

607,431,768,211,456

# What IS IPv6?

- What is the Goal?
  - To Transition from IPv4 to IPv6
- When will that occur?
  - 20+ years!
- What do we do until then?
  - Dual Stack
- Why?
  - Its slim picking's on IPv4 address blocks.
  - Soon there will not be v4 blocks to get from ARIN and alternative methods to get address blocks are expensive . Current rate is \$2500+ per year per /24!

# IPv4 Depletion

- ARIN Region

- As of 4/24 – This is what is left!
- Most assignments are going to be /23s for now
- Going forward there is minimal blocks free
- You can STILL get it, but have to follow their guidelines!

## REMAINING IPV4 INVENTORY

Discrete Block Size (CIDR)	Number of Blocks Available
/11	1
/12	1
/14	1
/16	1
/18	1
/19	1
/21	23
/22	5
/23	185
/24	582

# Issues with IPv6 Deployment

---



# Getting Around Issues

- Can I get IPv6 addressing from ARIN

## ISP Initial Allocation

*/32 minimum allocation  
(/36 upon request)*

[NRPM 6.5.1](#)

- Have a previously justified IPv4 ISP allocation from ARIN or one of its predecessor registries, or
- Qualify for an IPv4 ISP allocation under current policy, or
- Intend to immediately multi-home, or
- Provide a reasonable technical justification, including a plan showing projected assignments for one, two, and five year periods, with a minimum of 50 assignments within five years

# Getting Around Issues

- Can I get IPv6 addressing from ARIN?
  - Sure can
  - Need AS number from ARIN
  - Then can request /32
  - Cost? \$2000 a year for /32 block!
- What about AS number?
  - They want you to add IPv6!
  - You can get your AS on special routing policy just for IPv6

# Getting Around Issues

- Can your Upstream(s) support it?
  - Cogent/HE/Level 3
    - Many higher tiered or multi-national providers support V6 in Dual-Stack configuration already
  - ATT/Charter/others
    - You will need to contact your upstream on this
- What if they don't?
  - HE and some others offer FREE IPv6 peering via 6to4 tunnels!

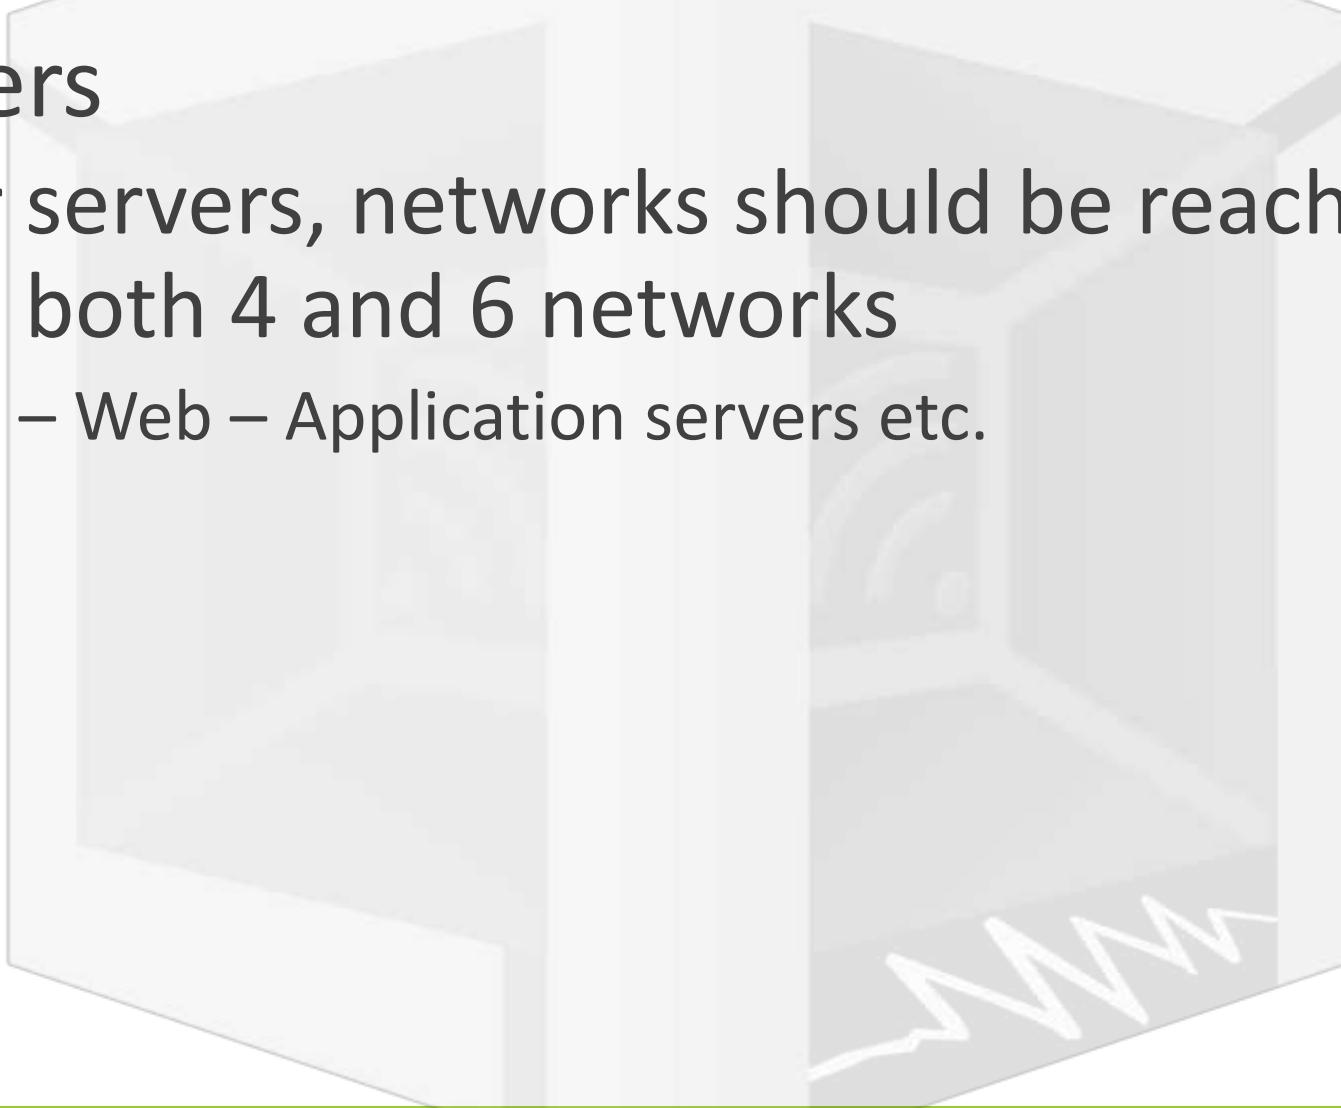
# How do you Implement IPv6

- Get a Assignment from ARIN
  - /32 or /48 would be assigned
  - Allows your lots of growth
- Dual-Stack Networks
  - Ipv6 is NOT backwards compatible with IPv4
  - Both will run at the same time for years to come!

# How do you Implement IPv6

- Servers

- Your servers, networks should be reachable with both 4 and 6 networks
  - Mail – Web – Application servers etc.



# How do you Implement IPv6

- ISPs
  - Plan for transition from IPv4 to IPv6
    - Several technologies and methods
    - Dual Stack Most common
    - 6to4 Tunneling for upstream that don't support it
    - Even though this is a "transition" plan, it will be running for many years to come!
  - Your deployment plan will vary based on the software / hardware you use



# IPv6 Deployment


---

# IPv6 Deployment

- Show IPv6 Deployment in 60+ tower MikroTik Network
- Configuration with PPPoE Authentication
  - Using Prefix Delegation
  - OSPFv3
  - BGP with IPv6



# IPv6 Deployment



# Deployment Plan

- Get ARIN Allocation
- Deploy IPv6 across all sites
- Give block to each site for PD (prefix delegation) to each customer if they customer requests PD
- Giving each customer a /64 for Auto configuration in their local network
- Ensure network security and bandwidth control on IPv6 dual-stack just like IPv4
- Add BGP session for transit providers
- Add BGP session for existing peers that support v6
- Larger clients that request larger IPv6 blocks will be handled on a as needed basis.

# Deployment Plan

- Get ARIN Allocation
  - We already had several IPv4 ARIN allocations
  - Work through initial IPv6 ISP Allocation resource with ARIN
  - Received our /32

# BGP IPv6 Configuration

---

# BGP On Network Edge

- Add BGP session for transit providers
  - Ensured IPv6 was deployed on Edge routers
  - Worked with transit providers
  - Received /112 typically with them to peer
- Created BGP Session for IPv6
  - In most cases you will create a v6 peering and v4 peering
  - RouterOS can support v4 and v6 peering in the same session, useful for internal usages, ibgp.

# BGP On Network Edge

- Add BGP session for existing peers that support v6
  - Worked with current v4 peers on peering exchange to add IPv6 peers
  - Peering exchange has a /111
  - Most of the peering exchange request to have two separate sessions, one for IPv4 and one for IPv6

# Test IPv6 BGP Connectivity

- Testing BGP Deployment
  - Created small /64 subnet off network core
  - Configured OSPFv3 from Edge to Core
  - Ensured Edges were advertising /32
  - Used Laptop, to Test for IPv6 connectivity
    - Ended up finding that their DNS server was not upgraded to support quad-a records, so no V6 addresses were being given
  - Further testing to verify that we had IPv6 connectivity to things like:
    - [ipv6.facebook.com](http://ipv6.facebook.com)
    - [ipv6.google.com](http://ipv6.google.com)
    - [www.linktechs.net](http://www.linktechs.net)

# Test IPv6 BGP Connectivity

- Ended up with 2 IPv6 Transit Peering's
- 8 other IPv6 Peering's on Exchange
  - 42,599 IPv6 Routes
  - 22,703 Active IPv6 Routes
- Hurricane Electric on Peering Exchange
  - HE offered FREE IPv6 transit once we had the IPv6 Peer with them.
  - This allowed us to basically have a third Transit provider but only for IPv6
  - Normally providers on peering exchange simply exchange their prefixes, such as us advertising our /32 only. HE did this then said, we will give you every route we have including ones not on our network, thus offering transit on IPv6.



# Test IPv6 BGP Connectivity

- Cogent vs Hurricane Electric
  - Cogent and HE refuse to “peer” with each other
  - So there is are IPv6 routes that are on cogent that you can’t reach from HE and routes from HE you can’t reach from Cogent
    - HE gives 21,477 prefixes – 10,512 are active
    - Cogent Gives 20,378 prefixes – 11,746 are active

# Test IPv6 BGP Connectivity

- Testing went well once we got the DNS servers to had out Quad-A records
- Speed tests were about the same if not better in some cases and ping times were fine.

# Site Router Configuration

---

# Site Router Configuration

- RouterOS has the IPv6 package installed but in many cases it was not enabled.
  - Plan network reboot basically
  - Used Free “Push Script” utility from Link Technologies, Inc. to push configuration out to all POPs - Can download from <http://www.linktechs.net/dl/>
  - Enabled IPv6 package
  - Created scheduled event that would reboot routers at close to 2 AM

# Site Router Configuration

- RouterOS has the IPv6 package installed but in many cases it was not enabled.
  - We also used this time to tell the routers to get the latest version of RouterOS from our FTP server at the NOC
  - So RouterOS fetched the latest version based on its processor
  - Then waited for 2AM

# Site Router Configuration

- RouterOS has the IPv6 package installed but in many cases it was not enabled.
  - 2AM came and went, and many devices rebooted and upgraded without issues
  - IPv6 was turned on
  - Another push script removed the scheduled reboot

# Site Router Configuration

- Now we had all sites running IPv6
- Decided on /56 routes for each site
  - Gave us enough /64s per site for all sites
  - Plus gave us a aligned prefix!

# Site Router Configuration

## Nibble Aligned Prefixes

<u>2001:db8::/32</u>	<u>2001:db8::/32</u>
2001:db8::/36	2001:db8::/40
2001:db8:1000::/36	2001:db8:100::/40
2001:db8:2000::/36	2001:db8:200::/40
2001:db8:3000::/36	2001:db8:300::/40
<snip>	<snip>
2001:db8:c000::/36	2001:db8:9200::/40
2001:db8:d000::/36	2001:db8:9300::/40
2001:db8:e000::/36	2001:db8:9400::/40
2001:db8:f000::/36	2001:db8:9500::/40

## Prefixes Not Nibble Aligned

<u>2001:db8::/32</u>	<u>2001:db8::/32</u>
2001:db8::/37	2001:db8::/41
2001:db8:800::/37	2001:db8:80::/41
2001:db8:1000::/37	2001:db8:100::/41
2001:db8:1800::/37	2001:db8:180::/41
<snip>	<snip>
2001:db8:6800::/37	2001:db8:8000::/41
2001:db8:7000::/37	2001:db8:8080::/41
2001:db8:7800::/37	2001:db8:8100::/41
2001:db8:8000::/37	2001:db8:8180::/41



# Site Router Configuration

- Aligned Prefixes
  - Allowed us to have simple subnetting assignments to each site
- Made IPv6 Deployment simple
  - Think of it as if you had /24s for each site, in our case, we just adjusted by 100.
  - With our /32, we can assign 16.777 MILLION /56s!

Instance	Area	Dst. Address
▶ default		
▶ default	backbone	100::/56
▶ default	backbone	200::/56
▶ default	backbone	300::/56
▶ default	backbone	500::/56
▶ default	backbone	600::/56
▶ default	backbone	700::/56
▶ default	backbone	800::/56
▶ default	backbone	900::/56
▶ default	backbone	1000::/56
▶ default	backbone	1300::/56
▶ default	backbone	1400::/56
▶ default	backbone	1600::/56
▶ default	backbone	1700::/56
▶ default	backbone	1800::/56
▶ default	backbone	1900::/56
▶ default	backbone	2000::/56

BLOCKED


# Site Router Configuration



- Created OSPFv3 Configuration for all sites
  - Used Loopback from IPv4
  - Configured Default for all interfaces as passive – for security
  - Then enabled only interfaces that were backhuls to have Point to Point Network types
  - No networks in OSPFv3

	Area	Interf...	Cost	Priority	Network Type	Instance	Neig...	State
P	backb...	all	10		1 default	default	0	down
	backb...	ether...	10		1 point to point	default	1	point to point
	backb...	ether...	10		1 point to point	default	1	point to point
DP	backb...	PPP...	10		1 broadcast	default	0	passive
DP	backb...	<ppp...	10		1 point to point	default	0	passive
DP	backb...	<ppp...	10		1 point to point	default	0	passive
DP	backb...	<ppp...	10		1 point to point	default	0	passive
DP	backb...	<ppp...	10		1 point to point	default	0	passive

# Site Router Configuration



- Created OSPFv3 Configuration for all sites
  - Created new AREA ID for each site
  - This would be for PPPoE
    - Allows us to filter out smaller blocks in OSPFv3
      - In our case we ended up with less than 70 IPv6 routes in OSPFv3!
      - We did not need to see each /64 for each customer
  - Added Area Range to advertise /56 into backbone area.

Area	Range	Cost	Advertise
 pppoe	Blocked :2000:0:0:0:0/56	calcula...	yes

Instance	Area Name	Area ID	Type	Default C...	Interfac...	Active I...	Neighb...
 default	backbone	0.0.0.0	default		7	7	2
 default	pppoe	10.100.0.20	stub	1	0	0	0

# Site Router Configuration

- OSPF Neighbors
  - We let the routers use their FE80 – Local Link Address for OSPF communication
  - We could use /127 by allocating a single /64 for a bunch of /127 point to point interfaces
    - End the end this would create more routes in our OSPFv3 table
    - When tracerouting we would see the /56 block IP anyways

Instance	Router ID	Address	Interface	State Changes
 default	10.100.0.19	fe80::d6ca:6dff:fe77:d643	ether3 -...	5
 default	10.100.0.21	fe80::d6ca:6dff:fe3d:b9dc	ether4 -...	5

# Traceroute in IPv6

- Local Link addresses not present in outbound path

Hop	#	Host	Loss	Sent	Last	Avg.	Best	Worst	Std. Dev.	History
1		4300:0:0:0:1	0.0%	22	1.9ms	3.0	1.6	4.6	0.7	
2		4200:0:0:0:1	0.0%	22	3.4ms	7.7	3.4	13.6	3.0	
3		4100:0:0:0:1	0.0%	22	11.4ms	16.0	6.2	33.0	6.9	
4		2000:0:0:0:1	0.0%	22	37.9ms	27.1	14.8	86.8	15.7	
5		1800:0:0:0:1	0.0%	22	56.1ms	40.6	19.8	82.2	15.0	
6		200:0:0:0:1	0.0%	22	86.9ms	84.9	39.9	158.6	31.9	
7		100:0:0:0:1	4.5%	22	146.2ms	86.2	32.5	146.2	32.5	
8		300:0:0:0:1	0.0%	22	102.4ms	84.4	50.1	116.9	18.2	
9		400:0:0:0:1	0.0%	22	126.0ms	81.9	32.4	142.4	27.6	
10		3500:0:0:0:1	0.0%	22	132.0ms	89.4	47.1	259.1	44.9	
11		5200:0:0:0:1	0.0%	22	145.1ms	97.8	53.6	171.5	36.0	
12		5100:0:0:0:1	0.0%	22	154.8ms	97.4	63.1	189.0	33.0	
13		5000:0:0:0:1	0.0%	22	118.6ms	98.9	59.0	159.4	29.3	
14		1	0.0%	22	96.6ms	97.5	58.9	159.5	27.3	
15		2001:504:1a::34:201	0.0%	22	180.2ms	100.8	56.8	186.4	38.2	
16		2001:478:245:1::6	0.0%	22	96.0ms	91.3	50.4	180.4	33.3	
17		2001:4860::1:0:713a	0.0%	21	120.9ms	99.6	50.7	165.5	26.3	
18		2001:4860::8:0:9150	0.0%	21	112.9ms	120.3	63.6	202.8	30.6	
19		2001:4860::1:0:84b4	0.0%	21	128.1ms	121.5	74.2	211.7	33.8	
20		2001:4860::1:0:0:1249	4.8%	21	102.0ms	117.5	68.8	189.9	30.2	
21		2607f8b0:4009:809::200e	0.0%	21	122.5ms	117.4	71.0	178.8	29.3	

Blocked

# Traceroute in IPv6

- Local Link addresses not present in inbound path

```
1      <1 ms      <1 ms      <1 ms      2001:400:5::1
2      <1 ms      <1 ms      <1 ms      2001:400:3::1
3       3 ms       3 ms       3 ms      ::1c::1:1
4      12 ms      10 ms      10 ms      2001:1000::9a36:50b9
5      17 ms      18 ms      17 ms      2001:1000::9a36:2bb2
6      24 ms      24 ms      24 ms      2001:1000::9a36:1f5a
7      25 ms      25 ms      24 ms      2001:1000::9a36:6f2
8      24 ms      24 ms      24 ms      2001:1000::9a18:12e6
9      24 ms      24 ms      23 ms      ::5d::a:2
10     24 ms      24 ms      24 ms      ::1
11     25 ms      25 ms      25 ms      2001:5000::1
12     26 ms      26 ms      26 ms      2001:5100::1
13     30 ms      32 ms      30 ms      2001:5200::1
14     35 ms      37 ms      33 ms      2001:3500::1
15     37 ms      33 ms      36 ms      2001:400::1
16     41 ms      39 ms      46 ms      2001:300::1
17     41 ms      46 ms      42 ms      2001:100::1
18     41 ms      38 ms      40 ms      2001:200::1
19     75 ms      55 ms      52 ms      2001:1800::1
20    106 ms      69 ms      60 ms      2001:2000::1
21    155 ms      90 ms      82 ms      2001:4100::1
22    114 ms     120 ms     103 ms     2001:4200::1
23    113 ms     122 ms     114 ms     2001:4300::1
24    113 ms     134 ms     113 ms     2001:4400::1
```

Blocked

# IPv6 to the Customer

- Network uses PPPoE for authentication
- PPP PD is the simplest way to distribute /64s to customers
  - Create IPv6 Pool
  - ON the PPP Profile, add Pool to DHCPv6 PD!

IPv6 Pool <pv6>

Name:

Prefix:

Prefix Length:

PPP Profile <pppoe>

General | Protocols | Limits | Queue

Name:

Local Address:

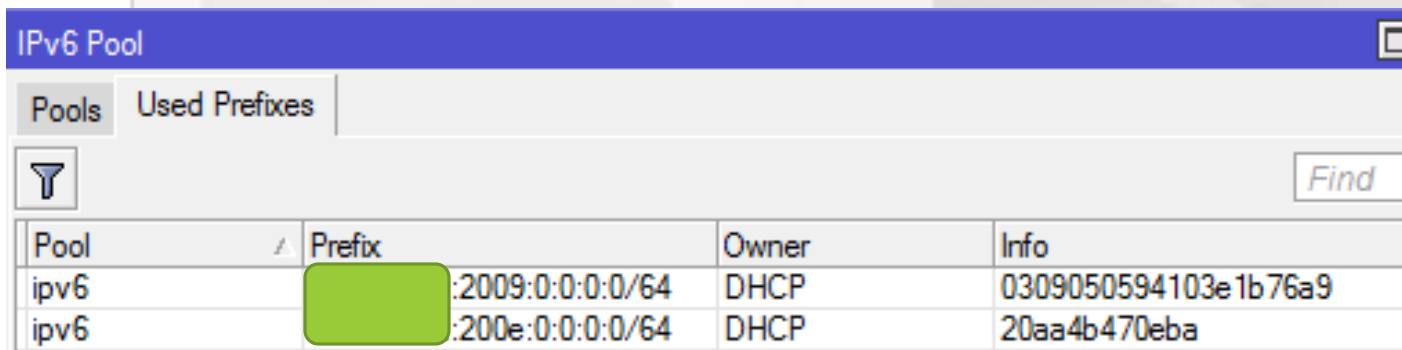
Remote Address:

Remote IPv6 Prefix Pool:

DHCPv6 PD Pool:

# IPv6 to the Customer

- Bumped all PPPoE clients from each POP location
  - If their device (router) supported IPv6 PD, it would request a subnet
  - IPv6 Pool → Used Prefixes
    - Shows IPv6 Prefixes requested via PPP system .
    - /64 prefixes not shown in the OSPF routing table



Pool	Prefix	Owner	Info
ipv6	[redacted]:2009:0:0:0:0/64	DHCP	0309050594103e1b76a9
ipv6	[redacted]:200e:0:0:0:0/64	DHCP	20aa4b470eba

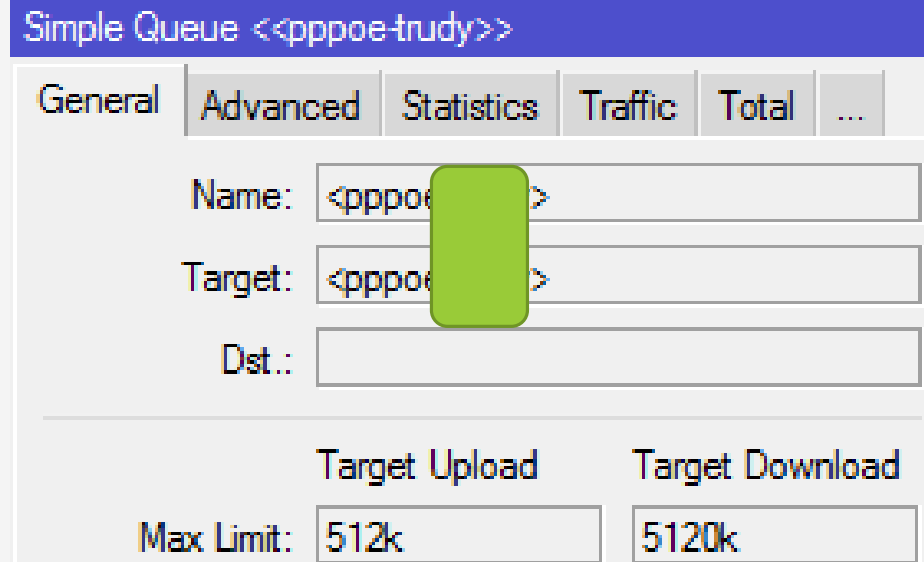


# Things we did.....

- Get ARIN Allocation
- Create BGP peering as needed with transit and peers
  - Test
- Updated routes and enabled IPv6 package (if required)
  - Required reboot on those to enable that
- Configured OSPF on each router
- Configure PPP on each router
  - Bumped all PPPoE users to force them to reconnect and get IPv6 if they requested it

# Bandwidth Control

- Unlike other methods, PPPoE works quite well, as it creates a queue on the interface.
- Nothing is really needed to be done, as IPv6 is limited by the simple queue when created.



The screenshot shows a configuration window titled "Simple Queue <pppoe-trudy>". It has several tabs: "General", "Advanced", "Statistics", "Traffic", "Total", and "...". The "General" tab is selected. Below the tabs, there are three input fields: "Name:" with the value "<pppoe-trudy>", "Target:" with the value "<pppoe-trudy>", and "Dst.:". Below these fields, there are two columns of settings: "Target Upload" and "Target Download". Under "Target Upload", the "Max Limit:" is set to "512k". Under "Target Download", the "Max Limit:" is set to "5120k". A green rectangular highlight is placed over the "Name:" and "Target:" input fields.

# How long?

- **Prep was key.**
- 1-2 weeks to get /32 allocation from ARIN
  - Maybe 4 hours actual work if that
- 1 day to prep Edge and Core devices
  - Required reboot
  - This took at most 2 hours, 2 core and 2 edge devices
- 2 weeks to enable IPv6 BGP Sessions on transit and peers that offered it
  - Time on this as more like 10 hours, mostly just spending time dealing with transits and peers. Peering exchange was simple, but getting the IPv6 enabled with some companies requires going to sales and they are slow!

# How long?

- 1 month of testing off core
  - This was the customer office with IPv6, testing to make sure all sites worked and IPv6 could be reached. A part of this was also troubleshooting their DNS Servers.
  - Actual work was around 4 hours, on and off over the month testing as they brought their DNS servers up to date.
- 2 days to upgrade routers and enable IPv6 on 60+ sites
  - Actual time was around 14 hours, both on ensuring they were enabled and verifying afterward's
- 1 more day to deploy OSPFv3 and IPv6 across the network
  - About 7 hours total

# How long?

- Less than 2 months from; “Hey lets deploy IPv6!” to “its running”
- Consulting time: Less than 40 hours to complete upgrades
  - Also included upgrading all routers in network to common version
  - Adding BGP sessions
  - Testing time and troubleshooting others (not our issues)

# Things I would do Differently?

---

- Umm,, nothing, other than have their DNS guy ensure Quad-A records could be resolved prior to us starting our testing
- Overall went quite well, simple, and worked just as promised!

# QUESTIONS?

I appreciate your business and look forward to working with you on future projects!

# Thank you!

[WWW.LINKTECHS.NET](http://WWW.LINKTECHS.NET)

[SALES@LINKTECHS.NET](mailto:SALES@LINKTECHS.NET)

314-735-0270