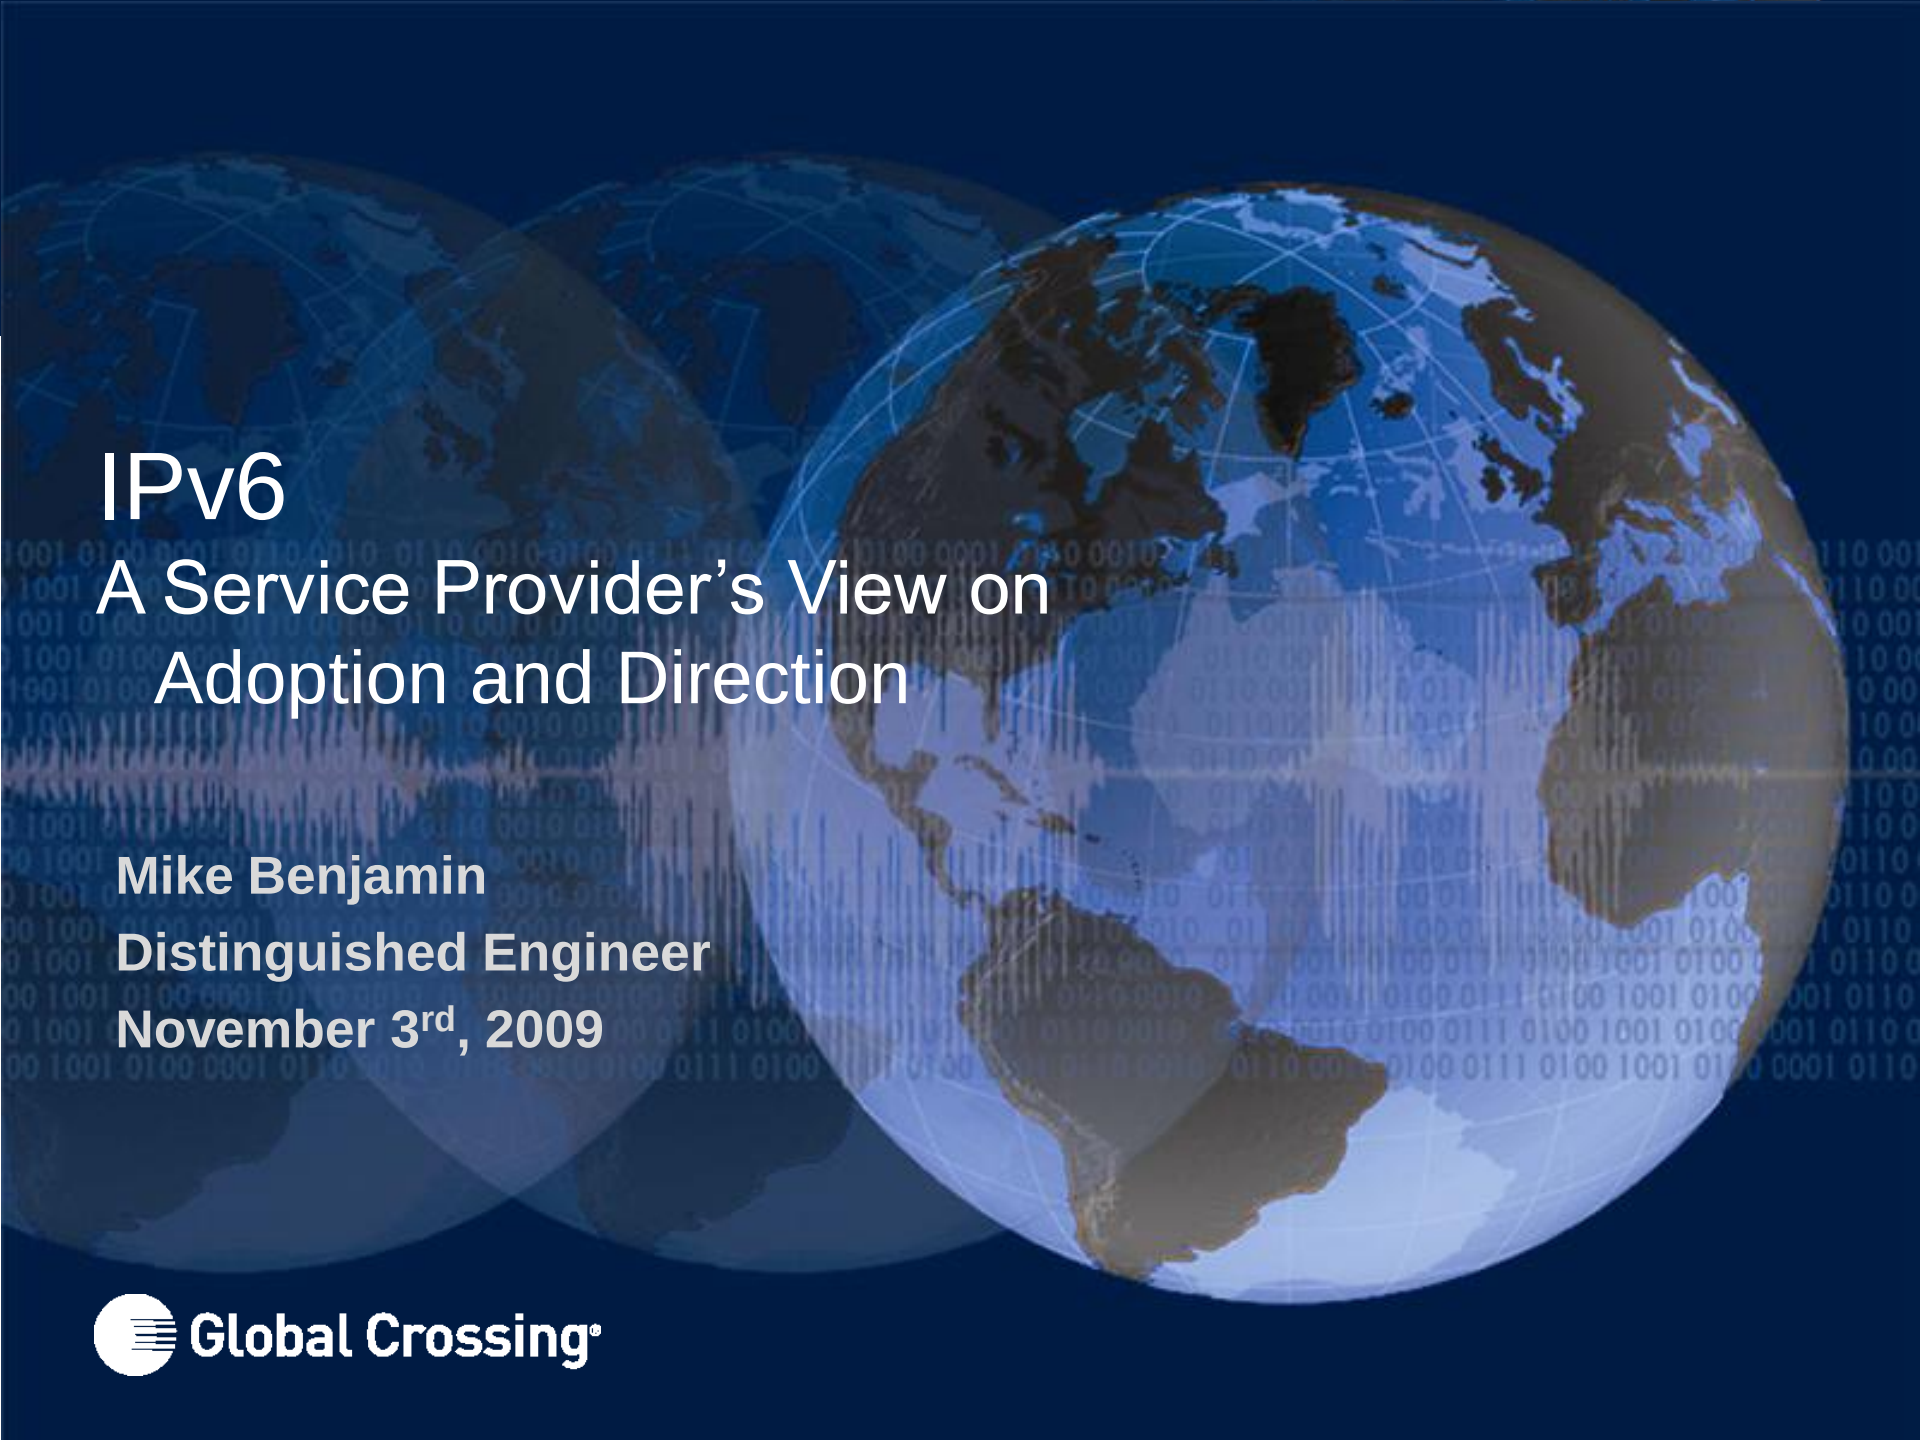

IPv6

A Service Provider's View on Adoption and Direction

Mike Benjamin
Distinguished Engineer
November 3rd, 2009

Global Crossing

A Level Set

IPv6 is not a project with a clear deadline or ROI

The work must be gradual and be ingrained in every project

The need is inevitable

It is a long road ahead

Global Crossing

Why Is It Inevitable?

International Internet growth is exponential

- The U.S. is not the driving force

Projected IANA Unallocated Address Pool Exhaustion: 24-Oct-2011 *

Projected RIR Unallocated Address Pool Exhaustion: 09-Feb-2013 *

Naysayers predict the return of un-routed space

- What if they are wrong?

Stop looking for the killer app and realize exhaustion is the killer need

• Source: potaroo.net

The Commercial Problem

“IPv6 solves a problem that hasn't happened yet.” *

John Curran, president and CEO of ARIN

Our industry has been seeking the killer application for years

Company after company have rejected the business case for IPv6

Meanwhile technical problems have begun the mainstream adoption

- Protocol translators, NATs, and other technology mask it for now

Risk avoidance has to be considered as a business driver

Global Crossing

• Source: InfoWorld

US Federal Government – An Example

They kind of got it right, and we can learn from it

- Initial mandate to deploy IPv6 caused confusion
- Later clarifications ensured readiness only
- Is the next step coming?

Deployment of the new equipment and software is a good first step

- Policies, procedures, RFPs, testing documentation must be included

Look beyond the network for risk of not enabling IPv6 after it is ready

- Consider tools, appliances, third party software, and your customers
  - Does your security solution work with IPv6?

Some Quick Background

IPv6 is a 128bit address displayed as;

- 2001:450:7E8A:B829::/64
- 2001:450:7E8A:B829:3:201:BED:1034

The header is completely incompatible with IPv4

IPv6's header format is extensible, and includes IPSec natively

- This directly impacts forwarding and firewalling

It was officially released as RFC1883 in 1995

A Hybrid World

Global Crossing offers dual stack interfaces to customers

- IPv4 and IPv6 on the same logical or physical interface

Tunneled IPv6 is also an option

- IP in IP or GRE allow IPv6 islands or hosts to tunnel over IPv4 networks

Implementing either solution still focuses on IPv4 as a component

- Is there a way to gracefully migrate networks to IPv6?

NAT Timeline

NAT-PT (RFC2766) was developed as an IPv6 transition method in 2000

RFC4966 outlines the reasons to deprecate NAT-PT in 2007

NAT64 problem statement is released in 2007

NAT64 is currently in draft as draft-ietf-behave-v6v4-xlate-stateful-02

- Last release was October 10th, 2009

DNS64, a required component, is also in draft

Where Did NAT-PT Fail?

An attempt to translate both IPv4 to IPv6 AND IPv6 to IPv4 led to;

- A need to intercept DNS requests
- Dynamic creation of resource pools internally and externally

The need to use a DNS ALG broke;

- Any network topology not following the same routing as the DNS request
- Any static host entries or other protocol methods that do not make use of DNS
- Scale!

… and of course – the same things every other NAT breaks

- Except rewriting payload addresses is more complex now

NAT64

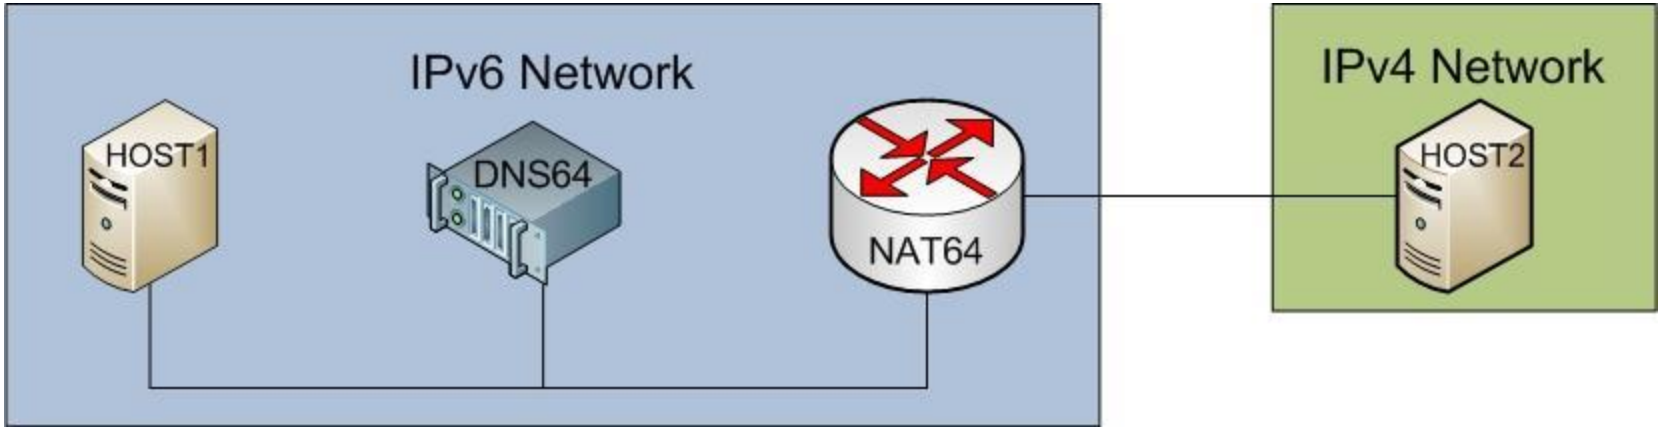

NAT64 only aims to NAT IPv6 to IPv4

The IPv4 network is represented by Pref64::/n on the NAT64

- The block is statically defined on the internal interface

DNS64

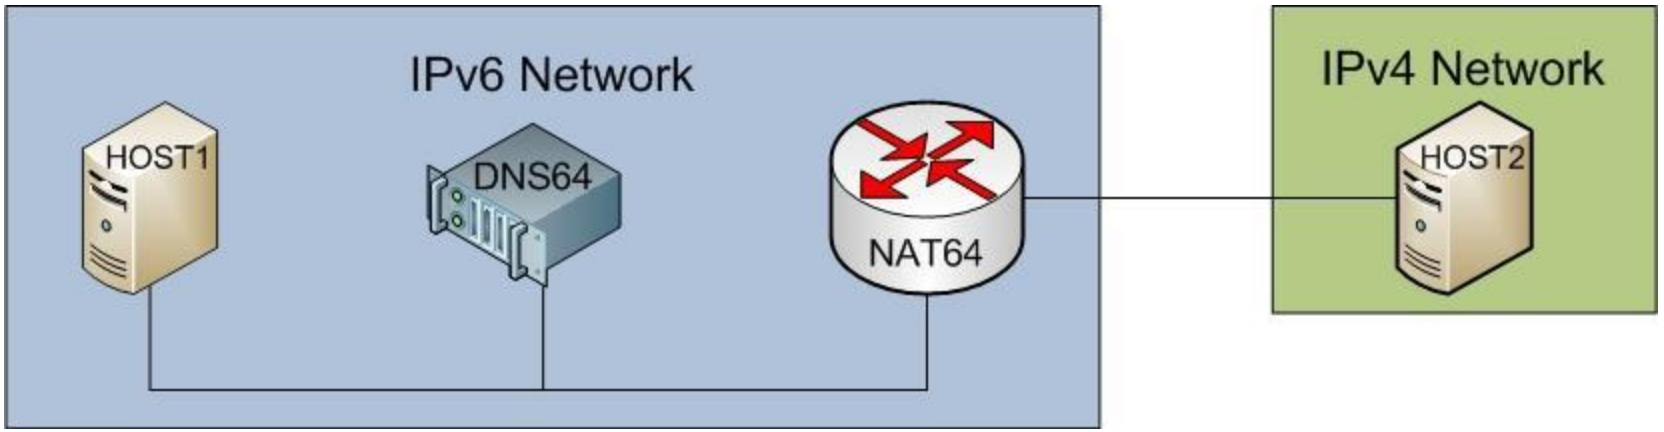

DNS64 responds to all AAAA queries normally and A queries as Pref64::<IPv4 Host>

Multihomed hosts must compensate for DNS server behavior

NAT64 is still NAT!

IPv6 seeks to remove NAT complexities but the transition method includes NAT

- This could cause application incompatibilities depending on adoption rate

Protocols still embed addresses and ALGs must be written to compensate

- Rewriting length information is also now required

Protocols lacking demultiplexing attributes may still fail

IPv4 fragmentation to IPv6 next header conversion can be resource intensive

IPv6 Multihoming – The Devil In The Details

The memory and processor power required to carry the IPv6 table in its entirety with the current model of BGP topology will not scale

RFC1887 from 1995 set the guidelines on restricting table bloat

- Hierarchy is their solution

Many different drafts, RFCs, and mailing list arguments have ensued

There is STILL no clear direction on how to scale IPv6

A War Of Ideas

Verizon (AS 701) holds true to RFC and does not permit non-provider aggregate (PA) routed prefixes

ARIN has been assigning Provider Independent (PI) /48s

Those routing PI space for multihoming are unreachable from AS 701

Who is wrong?

Shim6 – A Solution Or A New Problem?

RFC5533 – released June 2009

Provides a host level mechanism to handle multihoming

- Each host receives a separate IP for each redundant network uplink

This removes all complexity from the network and moves it to the host

- Is that right?
- Is it our only choice?

Can higher layer applications handle the redundancy where required?

- Are updated RFC2462 Router Announcements acceptable for user hosts?

The Route Forward On Multihoming

The industry must agree on a direction NOW

- This direction may not solve the long term scaling problems

Solving the problem should not be hindrance to IPv6 deployment

- Table growth is a concern with IPv4 as well

Minimize PI space

Aggressively filter out smaller blocks at peering points

Continue to evaluate network and host level technology options

Thank You

Global Crossing