Jeff Doyle and Associates Inc

IPv6 Address Design

A Few Practical Principles

North American IPv6 Summit 2013
19 April, 2013

How Big is the IPv6 Address Space?

- IPv4 developed 1973 – 1977
  - 2^{32} = 4.3 billion addresses
  - More than anyone could possibly use!
- IPv6 developed mid-1990s
  - 2^{128} = 3.4 x 10^{38} addresses
  - More than anyone could possibly use?

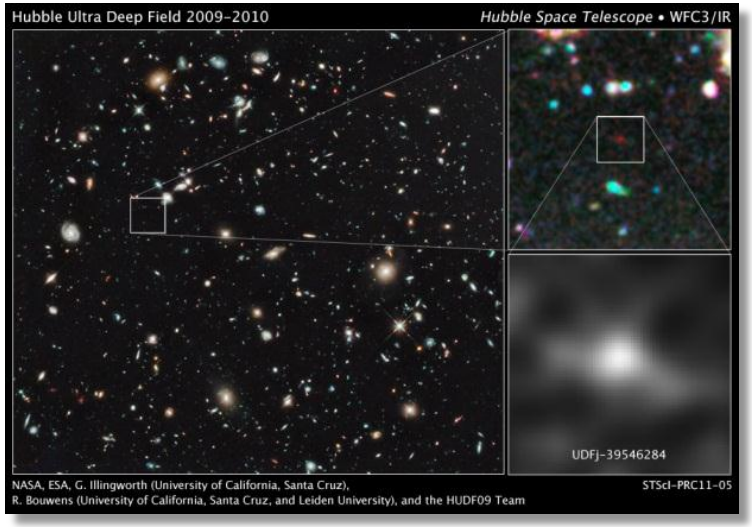

Galaxy UDFj-39546284

Jeff Doyle and Associates Inc

Copyright © 2013 Jeff Doyle and Associates, Inc.

How Big is the IPv6 Address Space?

- IPv4 developed 1973 – 1977
  - 2^{32} = 4.3 billion addresses
  - More than anyone could possibly use!
- IPv6 developed mid-1990s
  - 2^{128} = 3.4 x 10^{38} addresses

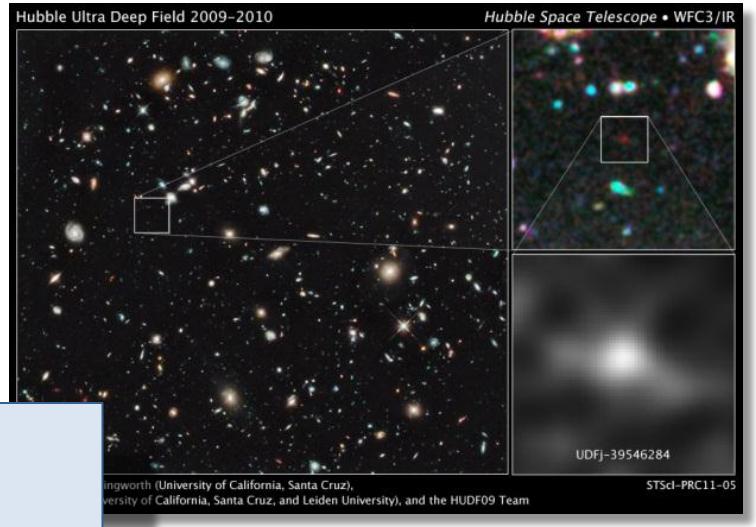

Galaxy UDFj-39546284

Some Perspective:

1 picometer = 10^{-12} (one trillionth) meter

2^{32} picometers = 4.29 millimeters
- length of a small ant

2^{128} picometers = 3.4 x 10^{23} kilometers
- 34 billion light years
- Furthest visible object in universe: 13.2B LYs

Jeff Doyle and Associates Inc

Copyright © 2013 Jeff Doyle and Associates, Inc.

Abandon IPv4 Thinking!

- Foremost IPv4 address design consideration: Address Conservation
- Balancing act between:
  - Number of subnets
  - Number of hosts on *each* subnet
- Result: VLSM
  - Complex
  - Hard to manage
- Legacy "class" categories still sometimes used in IPv4
  - Outdated and misleading
- No such thing as subnet masks in IPv6
  - CIDR-style prefix length notation always used

2001:db8:1234:abcd:5401:3c:15:85/48

Jeff Doyle and Associates Inc

Copyright © 2013 Jeff Doyle and Associates, Inc.

IPv6 Global Unicast Address Structure

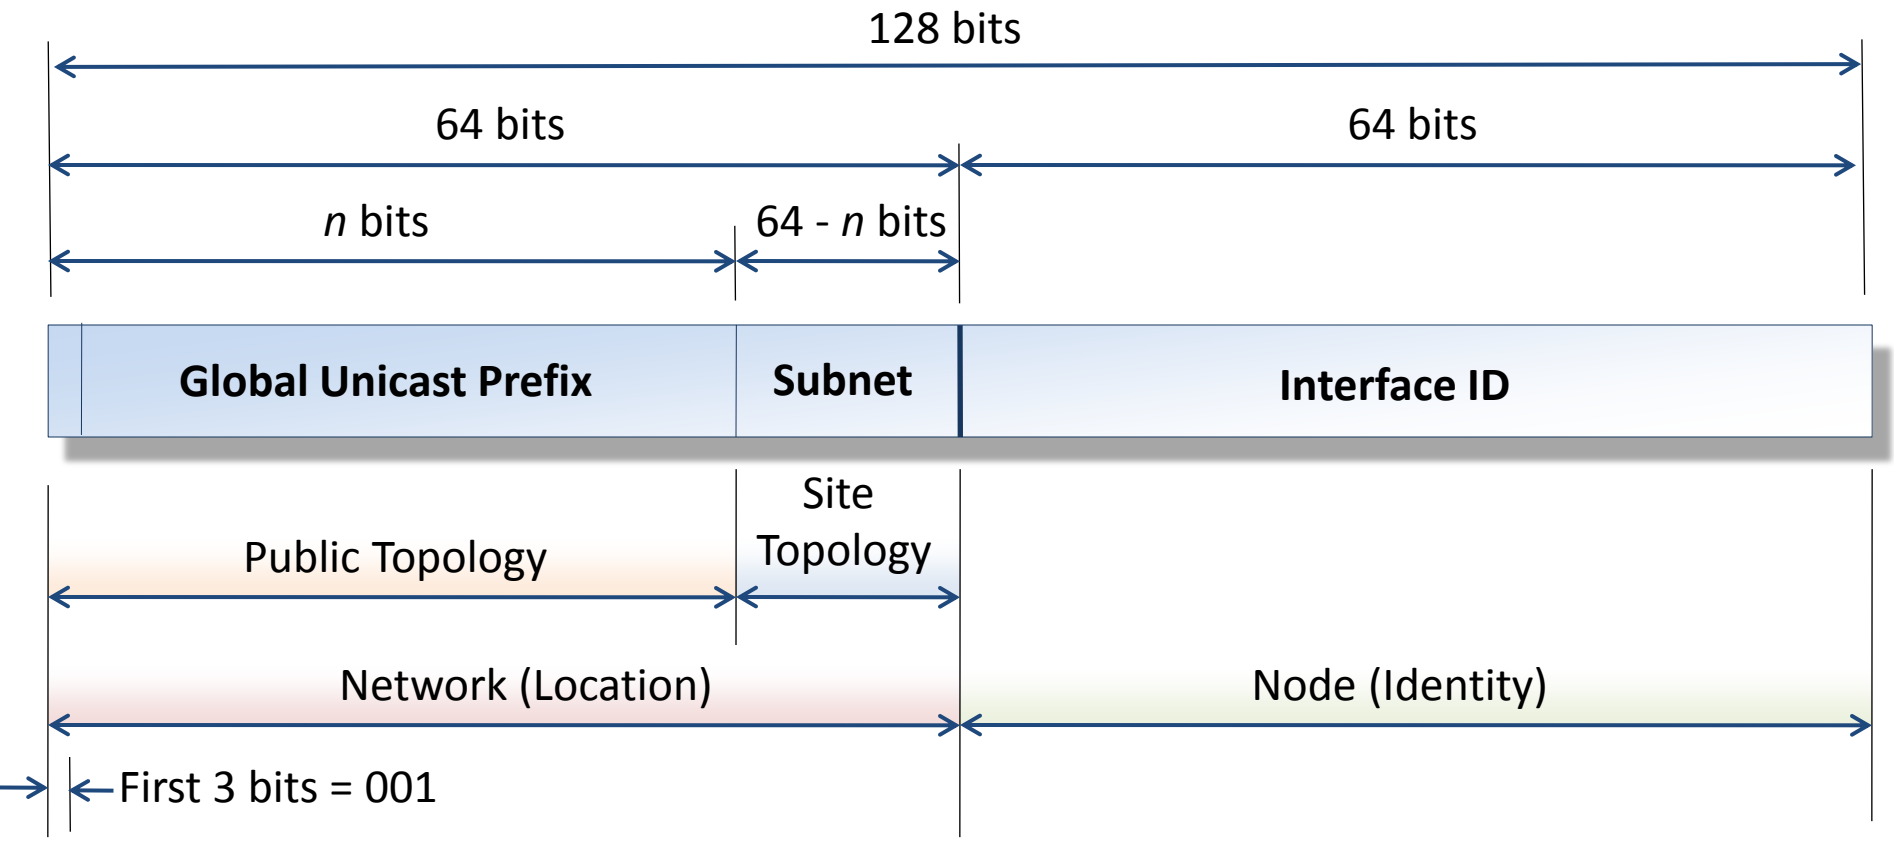

Jeff Doyle and Associates Inc

Copyright © 2013 Jeff Doyle and Associates, Inc.

Global IPv6 Unicast Prefix Allocations

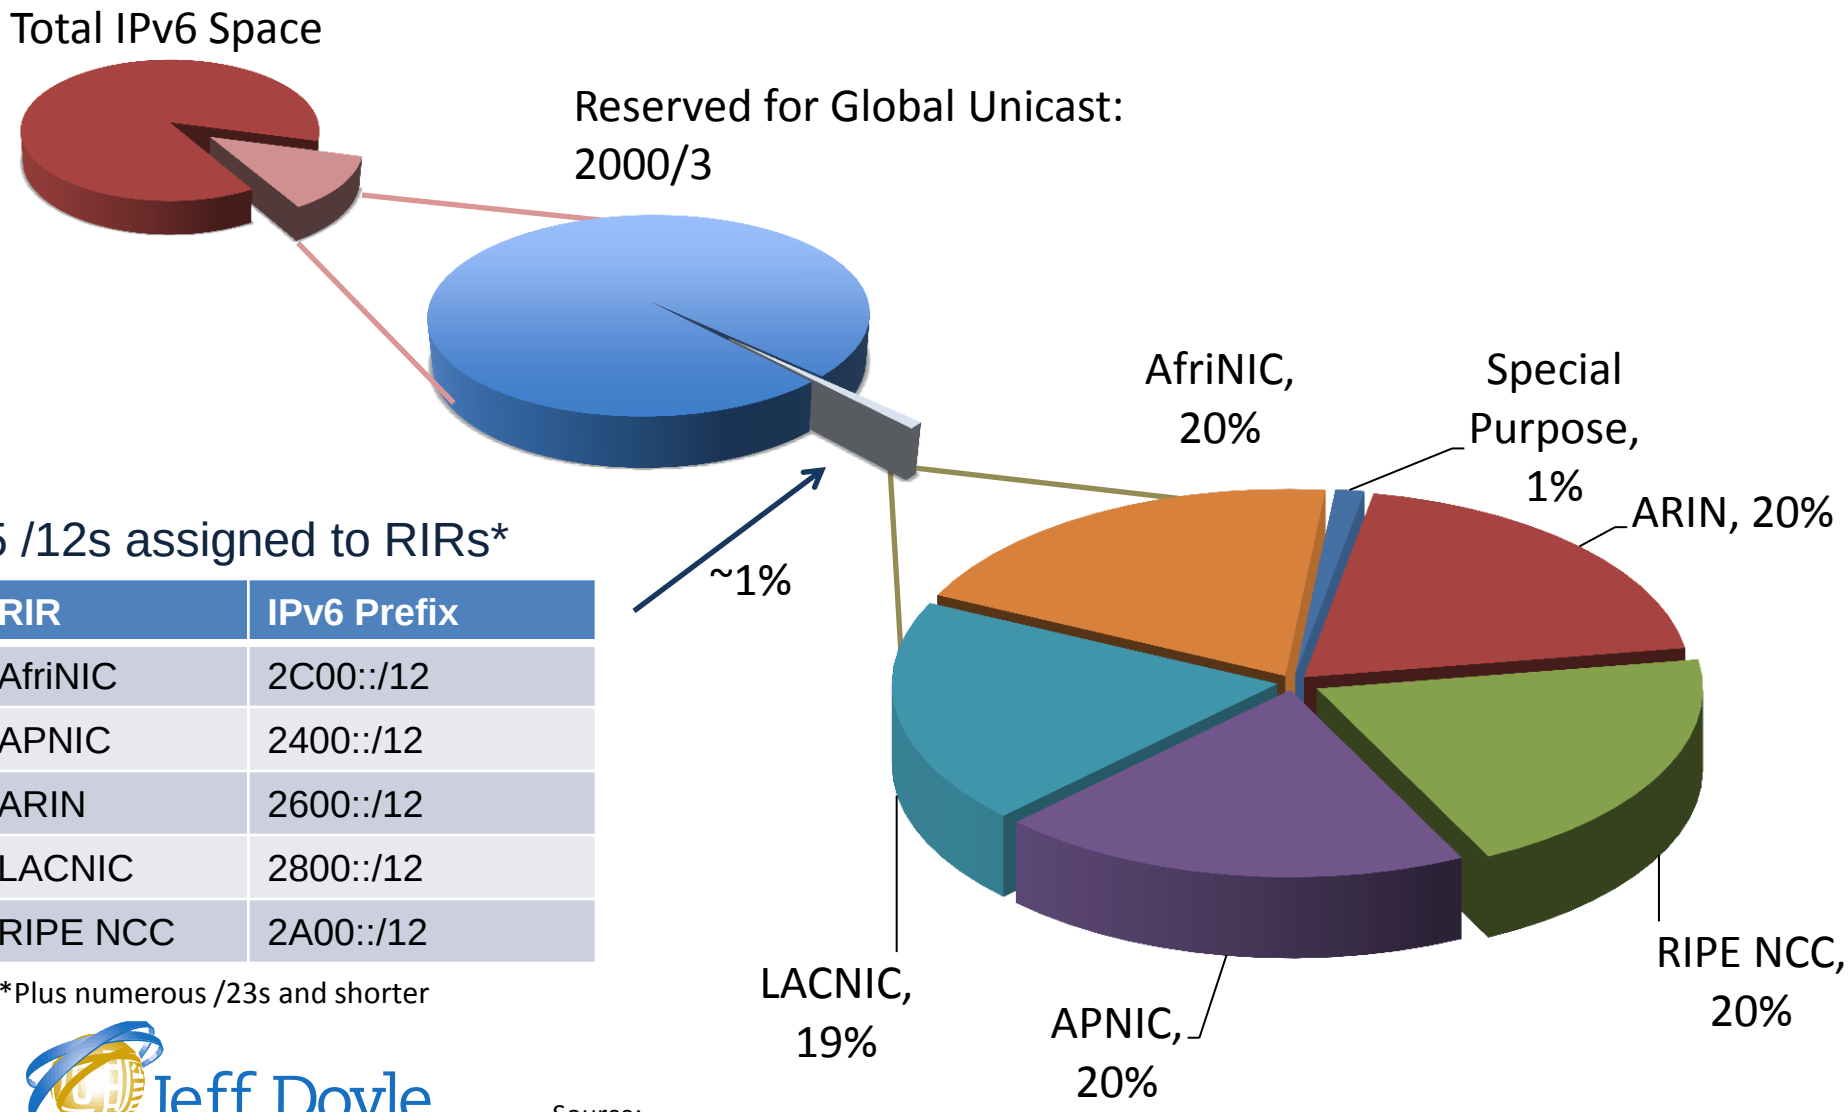

5 /12s assigned to RIRs*

RIR	IPv6 Prefix
AfriNIC	2C00::/12
APNIC	2400::/12
ARIN	2600::/12
LACNIC	2800::/12
RIPE NCC	2A00::/12

*Plus numerous /23s and shorter

Jeff Doyle and Associates Inc

Source:
http://www.iana.org/assignments/ipv6-unicast-address-assignments/ipv6-unicast-address-assignments.xml

Copyright © 2013 Jeff Doyle and Associates, Inc.

In Practical Terms…

- Typical IPv6 prefix assignments:
  - Service provider (LIR): /32 → 2^{32} /64 subnets
  - Large end user: /48 → 65,536 /64 subnets
  - Medium end user: /56 → 256 /64 subnets
  - Small/ Home/ SOHO: /64 or /60 → 1 or 16 /64 subnets
- Addres ation
  - Is th
  - Yes!
- If you have the right p

Is this really practical?

Jeff Doyle and Associates Inc

Copyright © 2013 Jeff Doyle and Associates, Inc.

What Prefix Size is Right for You?

- How do you define large, medium, and small?
  - Are these arbitrary boundaries?
- Why not assign /48 *per site?*
  - Site = one building
  - 30 buildings per campus = 30 /48s
  - And yes, a home (apartment or house) is a site!
  - It's not about waste, it's about consistency
- /48 for all allocations was original policy (RFC 3177)
  - *"Home network subscribers, connecting though on-demand or always-on connections should receive a /48."*
- Obsoleted by RFC 6177
  - Concerns about waste
  - Intention is that IPv6 should last for 100 years
- ARIN supports a liberal allocation policy

Jeff Doyle and Associates Inc

Copyright © 2013 Jeff Doyle and Associates, Inc.

What Prefix Size is Right for You?

ARIN Number Resource Policy Manual:

2.10. End Site

"The term End Site shall mean a single structure or service delivery address, or, in the case of a multi-tenant structure, a single tenant within said structure (a single customer location)."

6.5.8.2.1.Standard Sites

"An organization may request up to a /48 for each site in its network, and any sites that will be operational within 12 months."

- ARIN supports a liberal allocation policy

Jeff Doyle and Associates Inc

Copyright © 2013 Jeff Doyle and Associates, Inc.

Are You Ready for IPv7?

- All current global unicast IPv6 prefixes start with 001
  - This is 1/8 of the entire IPv6 space
  - 2^{45} = 35 trillion /48 prefixes
- UN projections for world population in 2100:
  - Median figure: 10 billion
  - High end: 16 billion
- 2^{45} / 16 billion = 2199 /48s *per person*
  - And, we still have 85% of the IPv6 space held in reserve
- Opinion: IP will become obsolete before IPv6 is depleted

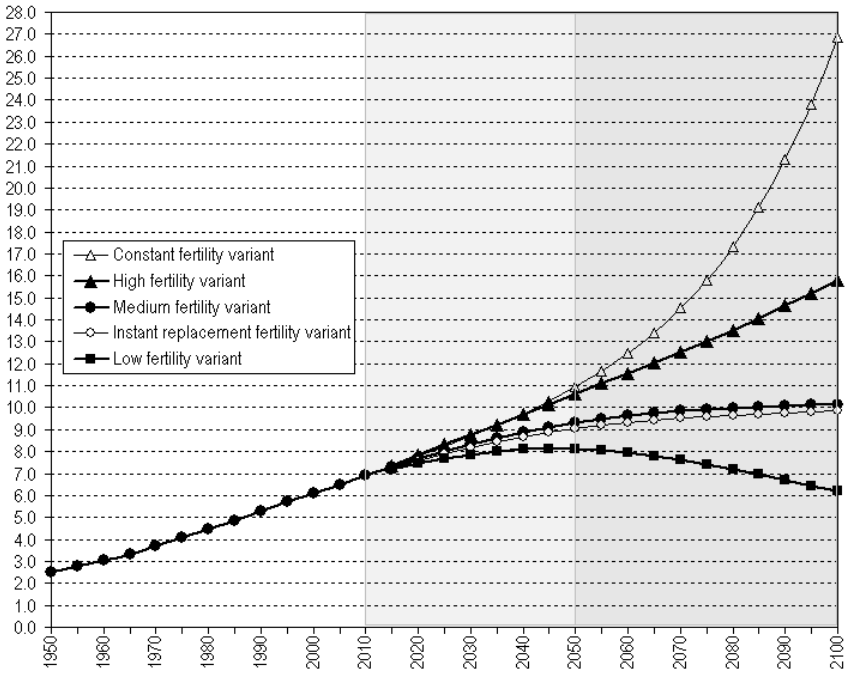

Jeff Doyle and Associates Inc

Copyright © 2013 Jeff Doyle and Associates, Inc.

Subnet Assignments

- RFC 4291 specifies that Interface-ID is 64 bits
  - Several IPv6 functions depend on this
- All subnets should be /64
  - Simplifies address management
  - Random addressing improves security
- Trend is to use stateful address assignments (DHCPv6)

Jeff Doyle and Associates Inc

Copyright © 2013 Jeff Doyle and Associates, Inc.

What About Point-to-Point Links?

- 18 million trillion addresses in a /64 link
  - And I will only *ever* use 2 of them?
  - Are you kidding???
- People have a very hard time accepting this
  - Again: This is not IPv4!
  - What else are you going to do with those addresses?
- It's a matter of comprehending the scale
  - 5000 out of 2^{64} is not really any bigger than 2 out of 2^{64}

Jeff Doyle and Associates Inc

Copyright © 2013 Jeff Doyle and Associates, Inc.

Point-to-Point Subnets (Battling RFCs)

- Reasons for using /64
  - RFC 3627
  - RFC 5375 => /64 usage endorsed and encouraged
  - Design consistency
  - Anycast problems are not significant on PtP links
    - Subnet-Router Anycast
    - MIPv6 Home Agent Anycast
- Reasons for using /127
  - RFC 6164
  - Ping-pong vulnerability
    - This is an issue with older version of ICMPv6 (RFC 2463)
    - Issue is corrected in newer version of ICMPv6 (RFC 4443)
    - Vendors: Upgrade your code!
  - Neighbor cache exhaustion vulnerability

Copyright © 2013 Jeff Doyle and Associates, Inc.

Point-to-Point Subnets (cont.)

- Insist that your vendors use current ICMPv6!
- Don't use /126
  - This is IPv4 thinking
  - "Subnet number" is meaningless in IPv6
  - IPv6 does not use broadcast addresses
- Potential compromise:
  - Assign /64 per PtP subnet
  - Address /127 out of the /64

Jeff Doyle and Associates Inc

Copyright © 2013 Jeff Doyle and Associates, Inc.

What Do I Get in Exchange for Waste?

- Simplicity
  - One-size-fits-all subnets
- Manageability
  - Hex is much easier to interpret at binary level than decimal
- Scalability
  - Room to grow
- Flexibility
  - Room to change

Jeff Doyle and Associates Inc

Copyright © 2013 Jeff Doyle and Associates, Inc.

Designing for Simplicity

- Start by mapping “working” bits
  - Generally the bits between assigned prefix and Interface-ID
- Group by hex digit (nibble)
  - 4 bits per hex digit
- Define “meanings” you need to operate
  - Geographic area? Logical topology? Type designation? User ID?
- Try to keep “meanings” on hex boundaries
  - Defined meanings will then be some multiple of 2^{4n}
  - Ex: 16, 256, 4096, 65536…
- Don’t get carried away with meanings
  - No need for 10 layers of address hierarchy if 4 will do

Jeff Doyle and Associates Inc

Copyright © 2013 Jeff Doyle and Associates, Inc.

Designing for Simplicity (continued)

- Use zero space as much as possible
  - Which address is easier to read?
    - 2001:DB8:2405:83FC:72A6:3452:19ED:4727
    - 2001:DB8:2405:C::27
- Benefit: Operations quickly learns to focus on meaningful bits
  - Ignore public prefix (usually)
  - Ignore Interface-ID (usually)
  - A few hex digits tell operations most of what they need to know

Jeff Doyle and Associates Inc

Copyright © 2013 Jeff Doyle and Associates, Inc.

Designing for Scale

- Leave “zero” space whenever possible
  - Designate as Reserved
- Insert between “meaningful” digits or bits
  - Allows future expansion in two directions

Jeff Doyle and Associates Inc

Copyright © 2013 Jeff Doyle and Associates, Inc.

Designing for the Future

- Trying to anticipate the unanticipated
  - A challenge for any kind of design
- Another reason for well-placed Reserved (zero) space
  - Horizontal Reserved space
  - Vertical Reserved space
- Do not integrate IPv4 into an IPv6 design!
  - Reading IPv4 in hex is (almost) meaningless
  - IPv4 will (eventually) go away

Jeff Doyle and Associates Inc

Copyright © 2013 Jeff Doyle and Associates, Inc.

Other Issues

- DNS design and management is critical
  - DNS issues are well documented
- IP Address Management is critical
  - IPv6 design is not easy to manage via spreadsheets
  - Good luck finding integrated DNS and DHCPv6 management
- Abandon IPv4 thinking!

Jeff Doyle and Associates Inc

Copyright © 2013 Jeff Doyle and Associates, Inc.

Questions?

jdoyle@doyleassociates.net

www.doyleassociates.net

+1-303-428-4680

Jeff Doyle and Associates Inc

Copyright © 2013 Jeff Doyle and Associates, Inc.