IPv6 Addressing

WHITE PAPER

August 2013

Table of Contents

Introduction

The continuous growth of the global Internet requires that its overall architecture evolve to accommodate the new technologies that support the growing numbers of users, applications, appliances, and services. Internet Protocol Version 6 (IPv6) is designed to meet these requirements and enable a global environment where the addressing rules of the network are again transparent to the applications.

Development of IPv6 has been under way since the early 1990s with the initial release RFCs. The primary driver for this development was the recognition that the IPv4 address space is a limited resource and would eventually be used up. On Feb 3rd, 2011 the Internet Address Numbers Authority (IANA), announced that the central pool of IPv4 address was exhausted. Regional Registries are now distributing what they have left in their respective regional pool, which in turn, will soon be exhausted. A popular IPv4 address exhaustion tracking site can be found at- http://www.potaroo.net/tools/ipv4/. It currently shows that Asia Pacific Network Information Center (APNIC) will not have any IPv4 addresses in April 2012, European IP Networks (RIPE) in August 2012, American Registry for Internet Numbers (ARIN) in July 2013, Latin American and Caribbean Internet Addresses Registry (LACNIC) in January 2014, and African Network Information Center (AFRINIC) in October of 2014.

There are several drivers pushing organizations to seriously look at integrating IPv6 into their current environment:

- Business continuity
- Internet evolution
- Government action
- IT consumerization and new applications

From a business continuity perspective, IPv6 is about growth of the network and the ability to provide service wherever the consumers of your service might find themselves. The need for growth feeds directly into the evolution of the Internet.

The Internet is evolving into a platform enabling cloud-based services to a highly mobile user base. Several predictions show that the number of network-attached devices could reach 50 billion by 2020. While this number might seem to be a bit far-fetched, consider how the network is evolving to accommodate smart buildings and offices, the smart grid, virtualization, and mobility.

All these new points of access for the network are leading to a consumerization of IPv6 into the mainstream applications and operating systems. All recent operating systems support IPv6. IPv6 is on by default, and in many cases is the preferred choice. New Internet devices like smartphones run IPv6 by default. Multiple factors are contributing to IPv6 integration:

- Organizations are allowing users to choose their own device (BYOD - Bring Your Own Device). These devices, as already mentioned above, increasingly run IPv6 natively.
- As mobile operators activate IPv6 on mobile devices (smartphones and tablets) in 2012 and beyond, it becomes important that the IT infrastructure is ready to accept these devices in a controlled environment.
- Organizations are adopting Windows 7 and Server 2008, Apple MacOSX, and Linux, which all support IPv6 natively.

There are several available resources to help build an IPv6 integration plan. The IETF has several RFCs and drafts that lay out integration plans. There are several books available that give a background on the technology and also layout an integration plan. The resources section at the end of the paper lists the RFCs and books that can be used to further research IPv6.

Developing addressing plans is touched on in these various documents and books but is not extensively covered. This document describes how to build an IPv6 addressing plan.

Addressing Overview

This section covers some basics related to IPv6 addressing. The addressing overview is meant to be a refresher and in no way a comprehensive primer on IPv6 addressing. For a detailed explanation of the IPv6 addressing architecture, see RFC4291 (http://www.ietf.org/rfc/rfc4291.txt).

One of the most recognizable differences between IPv4 and IPv6 is the size of the address space. IPv4 has 32 bits and allows for approximately 4 billion hosts (4×10^9). IPv6 has 128 bits and allows for approximately 340 undecillion (340×10^{36}) addresses.

Address Representation

The first area to address is how to represent these 128 bits. Due to the size of the numbering space, hexadecimal numbers and colons were chosen to represent IPv6 addresses. An example IPv6 address is:

`2001:0db8:130f:0000:0000:7000:0000:140b`

Note the following:

- The case of the characters does not matter. Lower case "a" means the same as capital "A".
    - RFC 5952 suggests that lower case should be used.
- There are 16 bits in each grouping between the colons.
    - 8 fields * 16 bits/field = 128 bits

There are some accepted ways to shorten the representation of the above address:

- Leading zeroes within each 16-bit group can be omitted, so a field of zeroes can be represented by a single 0.
- Trailing zeroes within each 16-bit group must be represented.
- Successive fields of zeroes can be shortened down to "::". This shorthand representation can occur only once in the address. Note that the :: should be used to shorten the address as much as possible.

To draw an analogy regarding the logic behind keeping trailing zeroes and dropping leading zeroes, consider this: one hundred and sixty dollars can be represented as $0160 (one hundred sixty dollars) and can be shortened to $160 (still a one hundred sixty dollars), but cannot be shortened to $16 (sixteen dollars), unless you want to give up some of your hard earned money.

Taking these rules into account, the address shown above can be shortened to:

`2001:0db8:130f:0000:0000:7000:0000:140b`

After the leading zeroes have been removed the address is shortened to:

`2001:db8:130f:0:0:7000:0:140b`

The trailing zeroes must be represented as shown below:

`2001:db8:130f:0:0:7000:0:140b`

After successive fields of zeros have been shortened the final address is:

`2001:db8:130f::7000:0:140b`

Note that in the last example the zeros after the 7 are significant and cannot be combined with the next field for the double colon shorthand. In any case, only one double colon can be used even if there are multiple groupings of zeros.

The final part to address representation has to do with the prefix notation. A typical IPv6 address uses 64 bits to represent the network and 64 bits to represent the interface identifier or host. Using the above address as an example, the network and host identifier fields are broken out as shown in Figure 1.

Figure 1 - IPv6 address breakdown

Network - 64 Bits	Interface ID - 64 Bits
2001:0db8:130f:0:	0:7000:0000:140b

2164

The classless interdomain routing (CIDR) prefix representation is used to represent the IPv6 address. An example of this notation is:

`2001:db8:130f::870:0:140b/64`

The /64 indicates that the first 64 bits are being used to represent the network and the last 64 bits are being used to represent the interface identifier.

> **i Tech Tip**
>
> RFC 5952 establishes guidelines for the representation of IPv6 addresses.

Address Types

RFC 4291 (IP Version 6 Addressing Architecture) identifies the types of IPv6 addresses that exist:

- Unicast
- Anycast
- Multicast

Unicast

A unicast address is defined as an identifier for a single interface. Typically, you use these addresses when a specific end system needs to communicate with another specific end system (for example, a peer-to-peer conversation). You have a choice of scope when assigning IPv6 unicast addresses: global, unique local, and link local.

Figure 2 shows the scope for each defined address.

Figure 2 - Address scopes for IPv6

Global

Unique-Local

Link-Local

2165

A key difference to note is that an IPv6 interface is expected to have multiple IPv6 addresses associated with it. This model is very different from IPv4, where an interface was typically only assigned a single address. IPv6 interfaces always have a link local address. An IPv6 interface also has one or more of a unique local or globally unique address.

A link local address is used for communications on a single link. Packets with a link local source or destination address are not forwarded by a router off that link. Link local addresses have meaning only on that link. All link local addresses can be identified as starting with the fe80::/10 prefix. As noted previously, all IPv6 interfaces have a link local address assigned to them.

Note in Figure 3 that the last 64 bits are designated as the interface ID. In IPv6 the "host" portion of the IPv6 address is called the *interface identifier*. The interface identifier is a part of all IPv6 addresses whether they are link local, unique local, or globally unique.

There are several methods available to assign the interface identifier–manual, automatic/stateless, and Dynamic Host Configuration Protocol (DHCP). These methods are covered in greater detail in the "Assigning Interface Identifiers" section of this guide.

Figure 3 - Link local address representation

Unique local addresses are reachable outside of a particular link, but they only have meaning inside a limited scope or domain. Unique local addresses are not intended to be routable across the Internet. They should be routable inside a particular site or customer domain. Unique local addresses are analogous to RFC 1918 addresses in IPv4. The main difference between unique local addresses and RFC 1918 space is that the unique local address space is intended to be globally unique.

The intended global uniqueness of the unique local address (ULA) space is a big difference between RFC 1918 and RFC 4193 addresses. Due to the widespread usage of RFC 1918 addresses across organizations, RFC 1918 addresses are not routable on the Internet. The global uniqueness of the ULA space could lead to route leakage and would in theory have less impact because the ULA space should be unique. The issue is that organizations do not have to follow the ULA address block generation algorithm when generating a prefix. Organization X and Y could both choose to use fd00:dead:beef::/48. Ingress and egress filters should be configured to deny ULA prefixes to ensure there are no surprises.

Unique local addresses are recognizable because they are all from the fc00::/7 address block. Figure 4 shows the breakdown of a unique local address. The L bit is set to 1 if the address is locally assigned. RFC 4193 reserves the 0 bit for future usage. This definition of the L bit breaks up the fc00::/7 block into the following two blocks:

- fc00::/8–Reserved for future usage
- fd00::/8–Locally assigned unique local addresses

RFC 4193 specifies a method to assign the 40 bit global ID. A semi-random algorithm is defined in the RFC and offers a very high probability of uniqueness of the global ID. The algorithm for generating unique local addresses has been implemented in several places on the Internet. For instance, see http://www.sixxs.net/tools/grh/ula/.

Figure 4 - Unique local address representation
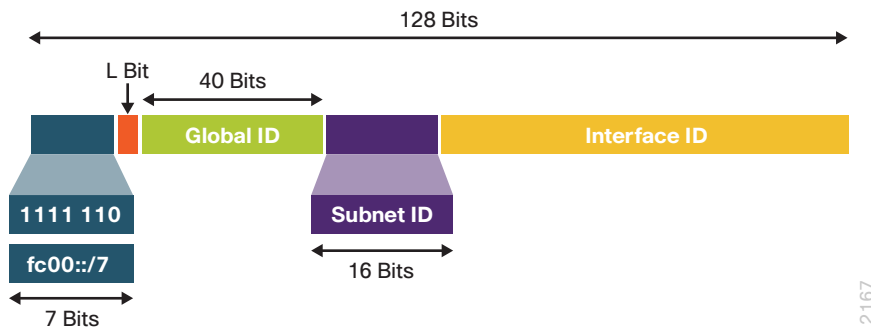
Global addresses are reachable from across the Internet. Global addresses are allocated from the regional registries (for instance, RIPE, ARIN, APNIC). Global addresses are all currently assigned out of the 2000::/3 block.

Figure 5 - Global address representation

Table 1 shows a subset of the current globally unique block allocations to the regional registries. The full list breakout can be found at the Internet Assigned Numbers Authority (IANA) website at http://www.iana.org/assignments/ipv6-unicast-address-assignments/ipv6-unicast-address-assignments.xml

Table 1 - Globally unique block allocations to regional registries

IPv6 address block	Regional registry
2001:/16	Various
2400:0000::/12	APNIC
2600:0000::/12	ARIN
2800:0000::/12	LACNIC
2A00:0000::/12	RIPE NCC
2C00:0000::/12	AfriNIC

There are several reserved or special use blocks of IPv6 address space that have been defined in multiple RFCs.

RFC 5156 has a listing of the currently defined special use addresses. Some of the more common blocks and their intended usage include:

- **::1**–Reserved for the loopback interface (RFC 4291)
- **::**–Unspecified address (RFC 4291)
- **2001:db8::/32**–For documentation purposes (RFC 3849])
- **2002::/16**–For 6to4 automatic tunneling (RFC 3964)
- **2001::/32**–For the Teredo tunneling mechanism (RFC 4380)

> **i** **Tech Tip**
>
> Unique local addresses are defined by RFC 4193 (Unique Local IPv6 Unicast Addresses).

Multicast

A multicast address is defined as an identifier for a set of interfaces that typically belong to different nodes. You can use multicast addresses to identify groups of interfaces that are interested in receiving similar content (for example, video). The conversation model in this case is a one-to-many model. You assign multicast addresses out of the ff00::/8 block.

Multicast addresses also have a scope associated with them. The scopes are very similar to the scopes defined for unicast addresses:

- **Link local**–Link local multicast addresses are intended only for systems on a link and should not be forwarded by network equipment off of that link. This behavior is the same as link local unicast addresses.
- **Organization**–Organizational multicast addresses are intended for use within an organization. These addresses are similar to the unicast unique local addresses.
- **Global**–Global multicast addresses are usable across the Internet similar to the unicast globally unique addresses.

There are some additionally defined scopes for IPv6 multicast addresses:

- **Interface local**–Interface local multicast addresses are intended for transmission of multicast within a node.
- **Site local**–Site local multicast addresses are intended for use within a single site.
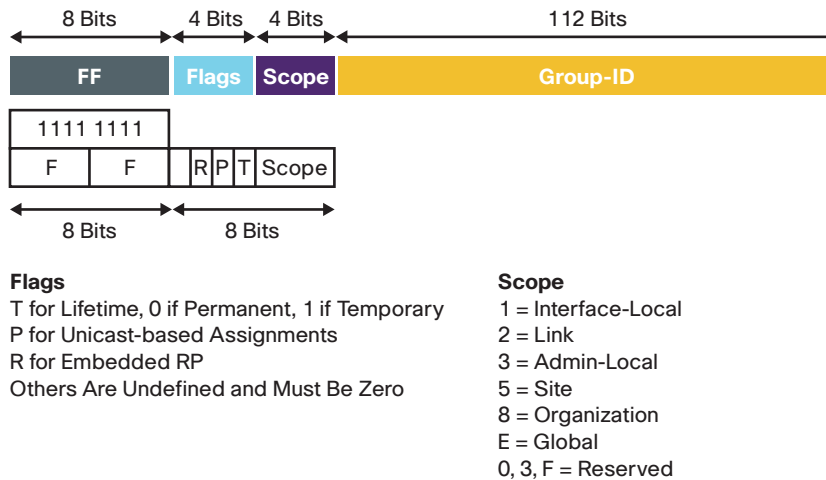
Figure 6 lays out the format of an IPv6 multicast address.

Similar to the unicast address space, there are some reserved or special use multicast addresses. A couple of the more common multicast groups and their intended uses are mentioned below. For a more comprehensive list of currently assigned multicast addresses, see http://www.iana.org/assignments/ipv6-multicast-addresses

Some of the more common, permanently assigned multicast addresses seen on IPv6 systems include:

- **ff02::1**–Link local, all nodes address
- **ff02::2**–Link local, all routers address
- **ff02:0:0:0:0:1:ffXX:XXXX**–Link local, solicited-node address

Figure 6 - Multicast address representation

Flags
T for Lifetime, 0 if Permanent, 1 if Temporary
P for Unicast-based Assignments
R for Embedded RP
Others Are Undefined and Must Be Zero

Scope
1 = Interface-Local
2 = Link
3 = Admin-Local
5 = Site
8 = Organization
E = Global
0, 3, F = Reserved

Anycast

The last defined IPv6 address type is anycast. It was defined for IPv4 in RFC 1546 circa 1993, but it is rarely used except for the provisioning of the Domain Name System (DNS) Root Servers–each of the 13 root server addresses is backed by many machines around the Internet. An anycast address is defined as an identifier assigned to multiple interfaces on different nodes. The anycast communications model is a one-to-the-nearest-of-many. This means that in the anycast communications model, a host communicates to the nearest of many potential nodes. Nearest is a relative term and is typically left to a routing protocol and its associated metrics to decide which anycast address is nearest or best based on the selection criteria. A good example for anycast communications is DNS queries. The host that needs to know what the address is for www.xyz.com does not care which DNS server responds. The host making the query is directed to the topologically closest server. If the DNS server that was responding goes offline, the next nearest server receives the request. Anycast addresses are not distinguishable by address from unicast (for instance, there are no defined bits that make an anycast address).

IPv6 Address Assignment Policies

Address Allocation Model

Currently, IANA allocates address blocks to the regional registries. The registries in turn assign address blocks to service providers. It is the service provider's responsibility to hand out addresses to their respective customers. The current policy varies by region and in the most conservative case dictates that an end user must go through the user's service provider to get IPv6 address space rather than directly approaching the regional registry for IPv6 address space.

Figure 7 graphically represents how this initial policy is enacted. This assignment model is commonly referred to as a provider assigned (PA) or provider dependent (PD) assignment. The prefix lengths that are shown in Figure 7 are recommendations. These initial recommendations followed RFC 3177 (IAB/IESG Recommendations on IPv6 Address Allocations to Sites). This RFC 3177 has since been made obsolete by RFC 6177 (IPv6 Address Assignment to End Sites). The registries and service providers can assign blocks using the processes and procedures that they have established for their regions and customers.

Figure 7 - Provider dependent policy
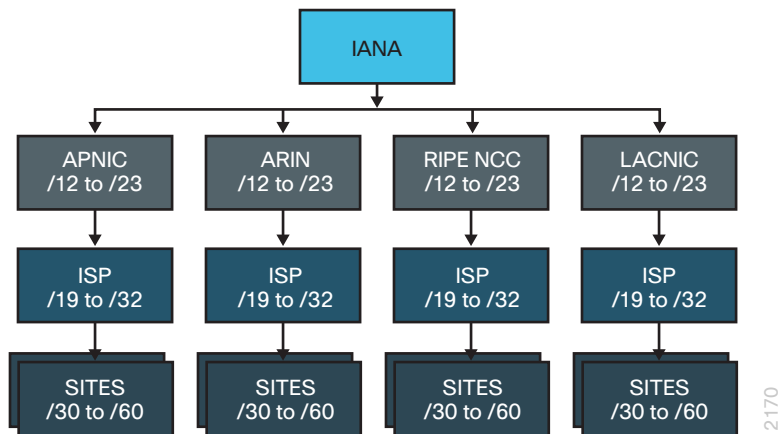
As an example of the policy, IANA has assigned 2600:0000::/12 to ARIN for assignment. This aligns with the top layer of the model. ARIN subsequently has assigned 2600::/29 block to Sprint, 2600:300::/24 to AT&T Mobility, 2600:7000::/24 to Hurricane Electric, etc. These block assignments do not follow the original model defined in RFC 3177. The service providers subsequently assign blocks to their customers based on their customers' needs. The Internet service provider (ISP) has the flexibility to assign a wide range of addresses to its customers. For example, a large enterprise ISP customer might need a /40 assignment while a residential customer would only need a /60 assignment.

There is an exception to this policy enacted by the regional registries that allows end customers to directly approach registries and request IPv6 address space. This exception is known as provider independent (PI) addressing.

The need for provider independent addressing arose to support end customers wishing to multihome to separate service providers, and organizations desiring to control their own resources when requesting address space. With the proposed provider dependent allocation model, the customer would be assigned an address block from each service provider. If an organization switches to a different service, it would have to re-address the entire organization. If that organization was connected to multiple service providers, it would also have to manage multiple different address blocks. Several approaches have been identified to address multihoming issues. Note that multihoming is not new to IPv6; multihoming exists in IPv4. What is new in IPv6 is the policy regarding how IPv6 address blocks are assigned.

Provider independent addressing has been adopted by all regional registries as an interim solution to multihoming. With provider independent addressing, a customer can request that an IPv6 block be directly allocated to their organization. There are requirements that a customer needs to meet to get a block allocated to them. You can find the regional registry policies regarding IPv6 provider independent policy address assignments at the following websites:

- ARIN policy–https://www.arin.net/policy/nrpm.html
- RIPE policy–http://www.ripe.net/ripe/docs/ripe-545
- APNIC policy (called portable assignments)–http://www.apnic.net/policy/ipv6-address-policy
- LACNIC policy–http://lacnic.net/en/politicas/manual.html
- AFRINIC policy–http://www.afrinic.net/en/library/policies/132-afpub-2007-v6-001

There are some potential issues with provider independent addressing. This is discussed next in the "Address Planning" section.

Address Planning

This section covers some guidelines to consider when building an IPv6 addressing plan. There are several RFCs that have been written that discuss IPv6 addresses. Some have been mentioned, such as RFC 4291 and 4193, which define the IPv6 address architecture.

Building an IPv6 addressing plan is a great opportunity to apply all of the lessons learned in building and deploying an IPv4 address plan. RFC 5375 outlines some issues that also need to be taken into account when building an addressing plan. Such considerations as whether or not unique local addresses should be used? Should the organization obtain a provider independent address block or is provider assigned addressing acceptable?

Provider Independent Addressing

The primary attraction to using provider independent space is that an organization is not tied to a specific provider. An organization that is using provider independent space can change providers without having to go through and renumber its entire network when the provider address space changes.

Provider independent (PI) space also allows an organization to connect to multiple service providers with a single IPv6 address block. These multiple connections provide resiliency and redundancy in case a particular service provider network has issues.

The PI model also allows an organization to deploy and manage a single block of address space without worrying about potential source address selection issues. If an organization was multihomed and used the provider assigned (PA) model, it would have to manage two address blocks at a minimum. The end systems would get assigned multiple addresses (for instance, one from each service provider (SP) block), which could lead to higher address management overhead and potential source address selection issues (for example, If SP A is having transit issues, how does an end system know to use SP B addresses?).

The primary question when considering PI addressing is what to do when an organization has presence across regional registries. For example, Company A has sites and a headquarters in North America, and sites and data centers in Europe and Asia. Does Company A get PI space from ARIN, RIPE, and APNIC? Or do they just get PI space from ARIN? This question is not easily answered.

The main focus from an SP perspective has to do with routing table growth. With the expansion of the address space to 128 bits, there is a corresponding potential for the routing tables to grow considerably. This potential for growth in the routing tables, and resulting service provider maintenance, will push SPs to look more closely at how they accept prefixes from their customers and how they advertise prefixes to their peers.

Current SP practices do not have regional PI block filtering in place (for example, an SP in Europe only accepting PI announcements from its customers). This means that an organization that gets a PI assignment from ARIN should be able to split up that block and make announcements in Europe and Asia.

The other potential issue with how service providers handle provider independent address blocks has to do with prefix lengths that are accepted and further propagated. Each regional registry has its own policy on initial block size assignment. The minimum prefix length that will be assigned is a /48 prefix block for provider independent space. While it might be perfectly acceptable to your service provider to accept that announcement, the downstream service providers that peer to your service provider might not be willing to accept a /48 announcement. In this case, the other service providers are concerned about the size of the IPv6 routing table that their routers might have to carry.

These two issues highlight a critical issue when looking at developing an IPv6 addressing plan. It is the responsibility of the organization that is considering PI space to sit down with its SP to figure out what IPv6 prefix policies the SP has in place. Listed below are questions that should be part of the initial conversations with the SP:

- Do you accept PI advertisements?
- What is the longest prefix length you will accept from customers?
- What is the longest prefix length you will accept from your peering partners?
- What is the longest prefix length your peering partners will accept?

Most providers have settled on the /48 as the longest prefix that they will accept from customers or peering partners. This policy should be verified with the SP. The SP policy can be verified by looking at a route server or looking glass service to see how the SP is handling their current IPv6 prefix announcements. You can find route servers and looking glass servers at http://www.bgp4.as/looking-glasses.

Here are some examples for when it is recommended to use provider independent address space:

- Your organization is connecting to multiple different providers
- Agreements are in place with your service providers to accept your independent IPv6 prefix announcements

Provider Assigned Addressing

Provider assigned addressing is another option for how to acquire IPv6 address space. In the provider assigned model, the service provider that an organization connects to is responsible for providing the address space to that organization. In this model, the organization works with the service provider to determine how large an address space the organization needs. The advantage for the service provider is that they can aggregate several customer blocks into a single announcement. Figure 8 illustrates this concept.

Figure 8 - Aggregation of customer blocks into a single announcement
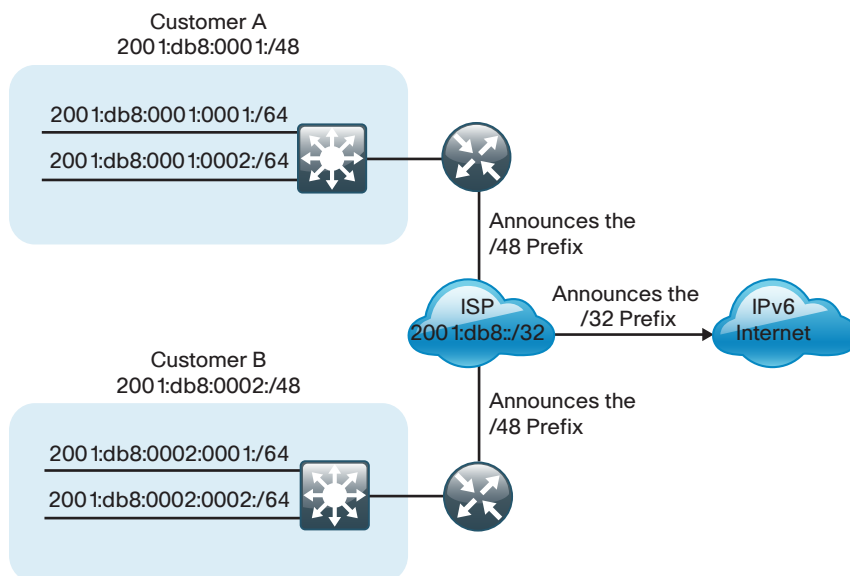
In this example, Customers A and B connect to the SP and each have been assigned a /48 block from the SP /32 block. The SP will carry the more specific /48 blocks inside its autonomous system, but will advertise the /32 to its peering partners.

The advantage of the PA addressing model for the organization is that the SP will handle a majority of the address plan development. In this model the SP is responsible for ensuring that the organization has enough address blocks to properly operate. The details of the levels the SP will drill down into the address plan development will vary from provider to provider.

It is recommended to consider provider assigned address space in these examples:

- Small to mid-size organizations using a single SP
- Small IT shops desiring to outsource

Addressing with Unique Local Addresses

When building the IPv6 address plan, a question might arise on whether or not to use globally unique addresses or unique local addresses (ULA). These alternatives are not mutually exclusive. An IPv6 end point can, and most likely has, multiple IPv6 addresses, so you can use both unique local and global addresses. If you want Internet connectivity, you must use global addresses.

It is worth noting that deploying unique local addresses allows for an addressing scheme to be deployed that is independent of whatever provider assigned or provider independent address space is used. Deploying unique local addresses allows for the internal network to be operational during any global re-addressing event.

One potential application for unique local addresses is to use them for internal communication and to use global addresses when accessing devices outside of the customer domain. In the case where you do have both unique local and global addresses deployed, RFC 3484 (Default Address Selection for Internet Protocol version 6 (IPv6)) should select the appropriate address for communication between the end systems. As with any new design, this application and behavior should be verified in a test environment. Figure 9 shows an example of how this scheme might be used.

*Figure 9 - ULA and global address communications*

In Figure 9, several devices have been given both unique local and globally unique addresses. In the case where internal-only communication occurs, such as to printers or network management systems, then the ULA is used for that session. This communication is indicated by the red line. The communications session is established between the end systems at fd01:1:1:1::3 and fd02:2:2:2::21 respectively. Globally unique addresses are used when the communication has to occur across the organization/site boundary. This session is highlighted by the

blue lines in Figure 9. In this example, communications that cross the Internet boundary use the addresses from the 2001:db8::/32 block shown in Figure 9. Note that certain functions, such as Path MTU Discovery (PMTUD), might not work correctly if unique local addresses are used. Organizations will probably filter packets sourced using unique local addresses. For this reason, globally unique addresses should be considered for use so that features such as PMTUD can work for all communications. As mentioned previously, this behavior should be properly verified for correct operation in a test environment.

A potential drawback for deploying unique local addresses has to do with multicast. Using the current source address selection method defined in RFC 3484, the unique local address is chosen over the global address. This selection can cause issues for Internet bound multicast traffic and disrupt the ability to pass RPF checks. Remember that ULAs are for internal use and not intended for use across the Internet. It is also highly likely that organizations are filtering out packets with a ULA source address on their Internet boundaries.

With the release of RFC 6296, IPv6-to-IPv6 network prefix translation has been standardized. This RFC opens the door for an organization to deploy only ULA addresses on their internal infrastructure. At the edge where they communicate with business partners or the Internet edge, they can deploy IPv6-to-IPv6 Network Prefix Translation (NPTv6) capable translation devices. The figure below demonstrates how this scenario might work.

Figure 10 - Scenario where only ULA addresses are deployed

NPTv6 provides a mechanism to translate the private internal organization prefixes to public globally reachable addresses. The translation mechanism is stateless and provides a 1:1 relationship between the internal addresses and external addresses. The use cases for NPTv6 outlined in the RFC include peering with partner networks, multi homing, and redundancy and load sharing.

Address Planning · August 2013

12

While NPTv6 does do a translation function, it is not the same function as the Network Address Translation (NAT) function that is in widespread use today. NPTv6 is stateless where NAT has to maintain a state table for the translations that have happened. NPTv6 does not do port mapping. NPTv6 does not change the layer 4 port numbers in a packet. NPTv6 will not change the interface identifier. This means that NPTv6 can provide only a direct 1:1 mapping for internal to external addresses. NAT can change the layer 4 port which allows a single global address to be used for multiple internal addresses. The stateless nature of NPTv6 also allows it to support both inbound and outbound connections. The stateful nature of NAT allows some control over the communications establishment process. If state for a translation does exist for NAT, then inbound communications can typically not be established. The following table highlights the differences between NAT and NPTv6.

NAT 44	NPTv6
Stateful	Stateless
1:N address mapping	1:1 address mapping
Layer 4 port manipulation	No Layer 4 port manipulation
Host address manipulation	No host address manipulation
Outbound connections	Inbound and outbound connections

RFC 5520 outlines some other concerns related to source address selection and ULA usage. An internet draft is working its way through the standardization process to update RFC3484 and update the source address selection process. For more information, see draft-ietf-6man-rfc3484bis-01 and draft-ietf-6man-rfc3484-revise-05.

Here are some pointers and recommendations when considering ULAs:

- ULAs are useful during a network wide re-numbering if globally unique addressing has to be changed. They allow for continuous internal communications as everything is being updated.

- Use ULAs for internal network management functions, but allow for proper operation of such features as Path MTU Discovery (PMTUD) by using globally unique addresses for loopback interfaces.

- Use ULAs for access to internal-only resources (for example, printers).

- NPTv6 and NAT do NOT provide similar functionality.

    ◦ NPTv6 provides a stateless 1:1 mapping for an internal and external IPv6 prefix.

    ◦ NPTv6 will NOT provide a 1:N mapping like NAT.

    ◦ Any perceived security benefit that NAT might provide is not present in NPTv6.

A security recommendation is to filter ULA addresses at any external boundary to your organization. Unless specifically permitted by a prior agreement (for example, extranet partner), all traffic that has a ULA source or destination address and is originating from outside your network should not be allowed into the network.

Address Block Recommendation

As a best practice, it is recommended that an organization get provider independent space from the regional registry that the organization is primarily associated with. PI space gives the organization a portable address space allowing connectivity to diverse SPs, without the potential translation issues identified in RFC 6296, and without source selection issues identified in RFC 5220. The PI model should be pursued even if an organization intends to use a single ISP for connectivity.

Network Level Design Considerations

The vast size of the IPv6 address space gives network engineers significant flexibility in designing an address plan. There are two considerations to building the addressing—how to size and assign subnets and how to assign the interface identifiers. This section discusses how to build the IPv6 subnetting scheme.

Subnet Planning—Initial Address Block Request

The initial request for an IPv6 address block deserves some attention when building the addressing plan. This step occurs if an organization is looking to use provider assigned or provider independent space.

Here are some considerations for initial sizing of the IPv6 address block:

- **Overall size of the current network and future growth**—An organization must consider the size of the network when estimating the size of the IPv6 address block to request. The size of the network should take into account the number of subnets—this is a difference from the IPv4 planning based on number of end systems. Is the organization large enough to justify requesting a /32? Would a /44 block work? Can the organization fit everything into a /48?

- **Multihoming strategy**—When formulating the initial request for IPv6 addresses, an organization must consider how it approaches redundancy and failure scenarios when connecting to a single or multiple service providers.

- **Multinational considerations**—Multinational organizations must now consider their approach when requesting IPv6 address blocks due to the strict hierarchy imposed by the current assignment policy.

The following discussion provides some recommendations for organizations to follow when building their initial IPv6 block request.

To size the request of the initial block, the organization should consider how large the current network is (for instance, how many subnets) and anticipated future growth. Another consideration is how to handle failover, traffic engineering, and redundancy. Service providers are continually updating their policies on prefix lengths that they will accept and advertise. Following the current recommendations and policy where an organization is given a /48 for their use, service providers are likely to accept a /48 as the longest prefix length that is advertised to other providers (some providers may accept longer prefixes for users completely contained within their network and advertise an aggregate of the longer prefixes). This policy has some implications for how organizations handle redundancy. With IPv4, an organization can break up their assigned /16 address block into /17 address blocks. They can then advertise these longer address blocks to enforce some routing policy and traffic engineering with their service providers. Subsequently, the organization can send the /16 to handle redundancy if anything happened to the peer announcing the more specific routes.

A /48 prefix is the longest prefix length that a service provider is likely to announce to other providers. If an organization needs to do some traffic engineering and has redundancy and failover concerns, then the initial block request should be larger than a /48 (for example, /44) and should be from contiguous address blocks so that aggregation can still occur. This situation is similar to the IPv4 scenario discussed earlier. An organization that receives a /40 block could announce more specific /48 blocks to draw traffic directly to those locations. At the same time they could announce the /40 aggregate to handle redundancy if anything happens to the primary path.

At the time of this writing, there are discussions about enforcing the filtering of the PA prefixes of length /48, so keep this in mind when considering addressing policy. For more information, see:
http://mailman.nanog.org/pipermail/nanog/2012-March/046419.html

Organizations spanning across multiple registries should consider obtaining addresses from each registry where they have presence. Using this strategy, an organization can accommodate the different policies that each registry might have. This approach also allows for some flexibility in the way an organization approaches their redundancy and traffic engineering.

The above considerations can be applied to building a subnet plan for both provider assigned and provider independent space. It does not matter whether or not an organization is using provider assigned or provider independent address space.

Provider independent address space is another consideration when building the initial IPv6 address plan. Organizations need to consider whether or not provider independent address space meets their needs. Can organizational redundancy and traffic engineering requirements be sufficiently handled with the use of provider assigned addresses? If not, then provider independent address space provides a potential solution.

As stated earlier, the recommendation is that organizations pursue provider independent address blocks from the regional registry that the organization is primarily associated with. Each registry is going to have a different justification process for what size address block is initially assigned.

For example, ARIN policy is related to the number of sites that an organization has. The list below shows the breakdown listed in ARIN's policy:

- More than 1 but less than or equal to 12 sites justified, receives a /44 assignment
- More than 12 but less than or equal to 192 sites justified, receives a /40 assignment
- More than 192 but less than or equal to 3,072 sites justified, receives a /36 assignment
- More than 3,072 but less than or equal to 49,152 sites justified, receives a /32 assignment.

The other regional registries start at assigning a /48 block and have processes to justify a larger block.

Subnet Planning–Aggregation

After deciding how to pursue the initial IPv6 address block, there are some other factors to consider when building the address plan. The current size of the network was a primary consideration when building the initial block request and it is also a major factor when looking at the overall subnet plan. Current RFCs suggest that a /48 prefix be handed down to organizations. A /48 prefix gives an organization 2^{16} (65536) /64 prefixes to use. This example highlights a potential for a corresponding increase in the size of the routing table that a network device uses to forward packets. A primary driver when building an IPv6 addressing plan is to take into account aggregation of IPv6 prefixes, which allows the network to scale and grow.

Figure 11 shows a simple application of the aggregation principle. In this case a service provider has acquired the 2001:db8::/32 address block from their regional registry. The service provider then assigns address blocks to their customers. In this example, Customer A gets 2001:db8:1::/48 and Customer B get 2001:db8:2::/48. With this scheme, each customer can assign subnets to their internal network in any scheme they choose. However, they aggregate all their internal subnets to one /48 announcement to the service provider. The service provider, in turn, aggregates all the customer address blocks that they have assigned to a single /32 announcement to their peer providers.

Figure 11 – Hierarchical addressing
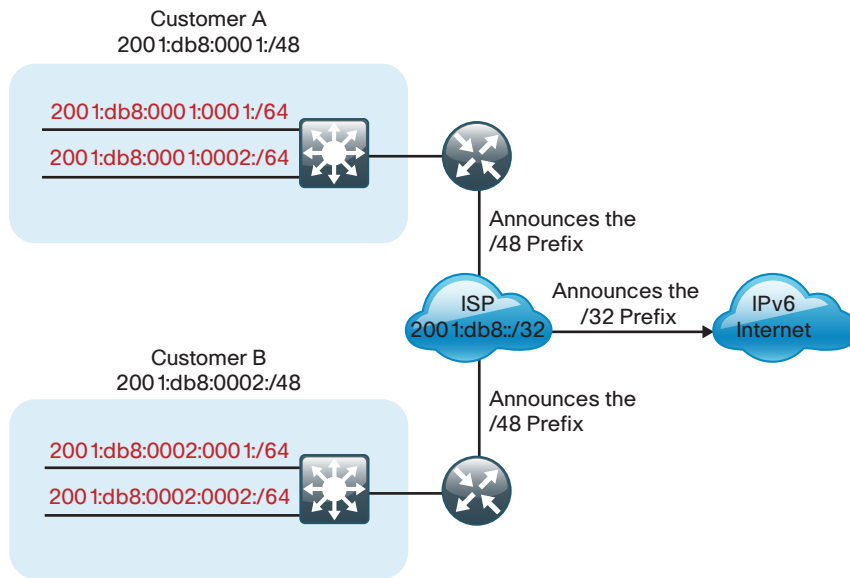
Keep in mind that even though aggregation is key, it is still reasonable to expect that a strict hierarchy is not needed. There might be some reasons that strict hierarchy cannot be followed. Address conservation is still important. Address conservation has a different meaning than it did in IPv4, but it is still something that must be considered when formulating an overall addressing plan.

Subnet Planning—Growth

Growth is another area that must be considered when allocating subnets in the network. Subnets can be assigned and manipulated based on bit boundaries within the organization prefix. Room also needs to be left in the subnet plan to accommodate future growth and the addition of more subnets to the network, which can be accommodated by leaving adjacent blocks of address space reserved.

To illustrate the process, assume that a company received the 2001:db8:1::/48 prefix to build their IPv6 network. The company divided their network up into four regions. The /48 address block they received allows them to use 16 bits to build their subnet plan. These numbers are based on the assumption that a /64 prefix will be used across the entire organization. The first four bits can be used to identify the region, which allows for 16 potential regions. Consecutive blocks can be assigned for regions that might need more subnet space. Gaps can also be left to accommodate potential growth within each region. Within each region, the next four bits can be used to identify facilities or sites within an organization, which allows for up to 16 facilities per region. The last eight bits are applied to each facility, which allows for 256 subnets per facility.
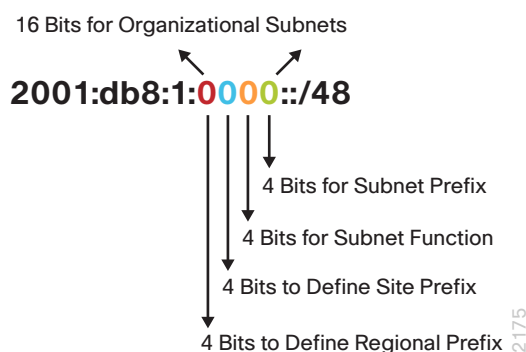
Figure 12 – Prefix breakdown

Table 2 shows the regional prefix breakdown. Note that over half the address space is reserved.

Table 2 – Regional prefix breakdown

Region	Regional prefixes
Reserved	2001:db8:1::/52
Region 1	2001:db8:1:1000::/52
Region 1	2001:db8:1:2000::/52
Reserved for Region 1 growth	2001:db8:1:3000::/52
Region 2	2001:db8:1:4000::/52
Reserved for Region 2 growth	2001:db8:1:5000::/52
Region 3	2001:db8:1:6000::/52
Reserved for Region 3 growth	2001:db8:1:7000::/52
Region 4	2001:db8:1:8000::/52
Reserved for Region 4 growth	2001:db8:1:9000::/52
Reserved	2001:db8:1:a000::/52
Reserved	2001:db8:1:b000::/52
Reserved	2001:db8:1:c000::/52
Reserved	2001:db8:1:d000::/52
Reserved	2001:db8:1:e000::/52
Reserved	2001:db8:1:f000::/52

In the table above, Region 1 has been identified as a larger region and has been assigned two consecutive blocks for use within that region. This assignment to Region 1 allows the region to have 64 facilities with each facility having 256 subnets. The other regions are smaller and do not initially need as large of a block. However, gaps are left in the address plan to accommodate growth. The same can be done for assigning subnets to a facility. Larger facilities can initially be assigned consecutive blocks to accommodate the size of the facility. For example, facility 1 in Region 2 is a larger facility and is assigned consecutive blocks.

> **i** Tech Tip
>
> RFC 3531 presents a plan for assigning subnets based on bit boundaries within the organization's IPv6 prefix and how those boundaries can be manipulated or changed as the network grows and more subnets are needed.

Subnet Planning–Prefix Length

There are two networks to consider when assigning prefix lengths –network segments with end stations and network infrastructure segments.

For segments with end stations connected to them, the addressing RFCs for IPv6 suggest that a /64 prefix length be used. With 2^{64} available addresses per segment, it is highly unlikely that you will see prefix lengths shorter than /64 for segments that host end systems. A /64 segment prefix is also required if stateless autoconfiguration is going to be used to assign the interface ID to the end stations. Secure Neighbor Discovery and privacy extensions also require a /64 prefix.

There are many options available when assigning prefixes for network infrastructure. Network planners could opt to be consistent across the network and deploy /64 prefixes for both network infrastructure and host access segments. Network planners could also opt for a plan that uses prefix lengths longer than /64. With all of these options available, there are no hard and fast rules available for assigning prefixes to network infrastructure. At this stage in the address plan, network planners should keep in mind the principles mentioned above–simplicity, aggregation, and growth. Table 3 summarizes some guidelines to consider when assigning prefixes to a link. The rest of the section adds more background and detail to these considerations.

Table 3 - Link level prefix concerns

Prefix	Concerns
64 bits	Consistency makes management easy
	Must for SLAAC, SEND, and other automatic address assignment methods
	Subnet not aligned with the number of end systems–perceived "waste" of address space
< 64 bits	Enables more hosts per subnet
	Considered bad practice
	64 bits offers more space for hosts than current media types and transport can efficiently support
> 64 bits	Address space conservation
	Special cases:
	/126–Valid for p2p
	/127–Valid for p2p (RFC 6164)
	/128–Typically used for loopback interfaces
	Loopback addresses
	Infrastructure management
	Complicates management
	Must avoid overlap with specific addresses:
	Embedded RP (RFC 3956)
	ISATAP addresses

There are several potential issues when considering the use of prefixes longer than /64. A first area of concern has to do with bit positions 71 and 72 ("u" and "g" bits respectively) in the IPv6 address. These bits have an identified meaning and their value should be correctly set. Bit 71 identifies whether or not the address is globally unique or locally assigned and bit 72 identifies whether the address is unicast or multicast. These bit positions are related to their functions in the MAC address and to the EUI-64 address expansion process. Most IPv6 implementations do not currently account for these bit settings.

Another addressing consideration comes into play if multicast is going to be used in the network and rendezvous point (RP) information is going to be embedded in the multicast group per RFC 3956. RFC 3956 requires a prefix length of /64 for the RP. This requirement must be accommodated when developing the overall plan.

A last area of concern has to do with Intra Site Automatic Tunnel Address Protocol (ISATAP) addresses. ISATAP requires a /64 for use and it embeds the IPv4 address in the last 32 bits of the IPv6 address. To complete the host interface identifier, ISATAP uses 0000:5efe. This sequence should be avoided when considering prefix lengths longer than /64.

One recommended approach for network infrastructure is to implement /64, /124, /127, and /128 prefixes. A /128 is used for loopback addresses to identify network nodes. A /64 or a /127 is used for point-to-point links such as serial or Packet over SONET POS links. A /124 or /64 can be used for infrastructure segments that require more than two addresses (for example, segments that will have a firewall and multiple routers). A /64 prefix scheme is the simplest scheme to implement. Alternatively, a scheme that uses more specific prefixes allows for the most address conservation. At this point a choice needs to be made between the simplicity of the /64 scheme and the potential complexity of the specific prefix scheme.

Potentially influencing this decision is an issue identified in RFC 6164. The issue is that on point-to-point links with lots of unused addresses, packets can loop until the hop count expires. In this scenario packets are sent to an address that does not exist on the link. Interfaces such as POS links do not use Neighbor Discovery and will forward the packet. Both routers in this scenario will continue to the forward the packet until the hop count expires. RFC 4443 does resolve the problem by stating that routers must not forward those packets and should generate and ICMPv6 destination unreachable message.

The table below lists some factors to consider when deciding whether to use /64, /126 or /127 prefix lengths.
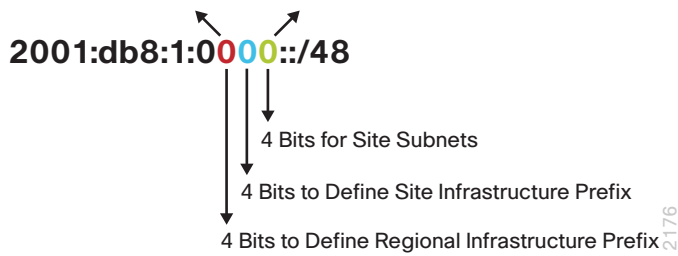
Table 4 – Prefix length factors to consider

/64	/126	/127
Ping pong could occur if a packet is sent to an unspecified address	Theoretically optimal but still could result in a ping pong loop	Old RFC 3627 and 5375 recommends against using /127 due to subnet-router anycast, but newer RFC 6164 recommends using /127
Common use with overall consistency to other LAN blocks – Cisco IOS devices have a fix for ping pong loops	Common use keeping IPv4 type of conservation mentality – Cisco IOS devices have a fix for ping pong loops	Cisco devices disable subnet-router anycast upon configuration of a /127 address
Also, mandated by RFC 4443 to send a Code 3 Destination Unreachable message to the neighbor router	Also, mandated by RFC 4443 to send a Code 3 Destination Unreachable message to the neighbor router	Most vendor equipment do not use or plan to use subnet-router anycast
Use this style, if your operational focus is to keep the same length across the board	Use this style, if your operational focus is to keep the v4 /30 type addressing semantics	Use this style, if your operational focus is to keep the v4 /31 type addressing semantics

The example plan in Table 2 demonstrates how infrastructure addresses might be planned using both /64 and /126 subnets. The /64 case is covered first. In Table 2, the last two regional blocks (2001:db8:1:e000::/52 and 2001:db8:1:f000::/52) are used to provide infrastructure links. Using this scheme, 8192 (2^{13}) infrastructure subnets can be assigned. The further breakdown of the address block is the same as previously identified. Four bits are used to identify the region, four bits are used to identify a site within a region, and four bits are used per site. Note here that the same regional and site numbers that were previously assigned can be re-used when selecting the infrastructure prefixes.

Figure 13 – Infrastructure subnet assignments

12 Bits for Network Infrastructure Subnets

2001:db8:1:0000::/48

4 Bits for Site Subnets

4 Bits to Define Site Infrastructure Prefix

4 Bits to Define Regional Infrastructure Prefix

The table below shows how the 2001:db8:1:e000:::/51 infrastructure prefix could be broken down into regional infrastructure prefixes.

Table 5 – Breakdown of regional infrastructure prefixes

Region (4 bits)	Regional Infrastructure Prefixes
Region 1	2001:db8:1:e100::/56 and 2001:db8:1:f100::/56
Region 2	2001:db8:1:e400::/56 and 2001:db8:1:f400::/56
Region 3	2001:db8:1:e600::/56 and 2001:db8:1:f600::/56
Region 4	2001:db8:1:e800::/56 and 2001:db8:1:f800::/56

Note here that an infrastructure prefix has been assigned to a region from both the 2001:db8:1:e000::/52 and the 2001:db8:1:f000::/52 block. This assignment will allow for 512 infrastructure prefixes per region.

The table below shows a further breakdown for sites that are within region 1.

Table 6 – Breakdown of sites within region 1

Site (4 bits)	Site infrastructure prefixes
Site 1	2001:db8:1:e110::/60 and 2001:db8:1:f110::/60
Site 2	2001:db8:1:e120::/60 and 2001:db8:1:f120::/60
Site 3	2001:db8:1:e130::/60 and 2001:db8:1:f130::/60
Site 4	2001:db8:1:e140::/60 and 2001:db8:1:f140::/60

Taken as a whole, this model allows for site 1 in region 1 to have 32 infrastructure prefixes and 256 subnets for end systems, as shown here:

Region 1 site 1 infrastructure prefixes

    2001:db8:1:e110::/64 thru 2001:db8:1:e11f::/64

    2001:db8:1:f110::/64 thru 2001:db8:1:f11f::/64

Region 1 site 1 end system prefixes

    2001:db8:1:1100::/64 thru 2001:db8:1:11ff::/64

    2001:db8:1:1100::/56 is the site summary prefix

An alternative you can implement is to use /126 subnets for infrastructure links. For this case, a /64 block is used to assign all infrastructure links. Again using the plan developed in Table 2, the 2001:db8:1:ffff::/64 block is used to assign all infrastructure links. Using this block assignment definitively identifies subnets that are being used for network infrastructure and those subnets used for end systems.

You can use a similar breakdown to identify regions and sites. Four bits are used to identify the region and four bits are used to identify the site. Note that the 2001:db8:1:ffff:0::/80 prefix is reserved to avoid any potential conflicts with implementations that use the universal/local bits. The table below shows how the regional infrastructure prefixes break down.

Table 7 – Breakdown of regional infrastructure prefixes

Region (4 bits)	Regional prefix
1	2001:db8:1:ffff:0:1000::/96
2	2001:db8:1:ffff:0:2000::/96
3	2001:db8:1:ffff:0:3000::/96
4	2001:db8:1:ffff:0:4000::/96

This assignment model allows for each region to have 2^{44} infrastructure subnets and each site within a region to have 256 infrastructure subnets. The table below shows a breakdown for sites within region 1.

Table 8 – Breakdown of sites within region 1

Site (4 bits)	Regional prefix
1	2001:db8:1:ffff:0:1100::/112
2	2001:db8:1:ffff:0:1200::/112
3	2001:db8:1:ffff:0:1300::/112
4	2001:db8:1:ffff:0:1400::/112

Taken as a whole, this model allows for site 1 in region 1 to have 2^{40} infrastructure prefixes and 256 subnets for end systems, as shown below.

Region 1 site 1 infrastructure prefixes

      2001:db8:1:ffff:0:1110::/127 thru 2001:db8:1:ffff:0:111f:ffff:fffe::/127

Note that some infrastructure segments might require more than two addresses. In this case contiguous /127 blocks would have to be assigned.

Region 1 site 1 end system prefixes

      2001:db8:1:1100::/64 thru 2001:db8:1:11ff::/64

      2001:db8:1:1100::/56 is the site summary prefix

This example highlights that using the /127 prefix breakdown for infrastructure links provides for greater address conservation by only allowing for two addresses per subnet. The example also shows that managing and maintaining this scheme is a bit more complicated –both in the planning and implementation of the scheme.

ULAs are another option for addressing infrastructure links. This scheme completely separates the network infrastructure prefixes from the end system prefixes by assigning network infrastructure prefixes from a completely different IPv6 address block. This strategy also affords some security for the network infrastructure. ULAs should not be reachable from the Internet, which should screen the network infrastructure from external attacks. In the example above, the ULA network infrastructure prefix could be fd00:2001:db8::/48. Organizations

choosing this route should implement /64 prefixes for ease of management. Consideration should also be given to ensure that PMTUD works for all hosts by using globally unique addresses for loopback interfaces and sourcing responses from that interface. Using this method should prevent any ULA filtering issues that organizations implement.

The last option to be discussed for network infrastructure is to use link local only addresses on network infrastructure but assign a global address to a loopback interface on the network infrastructure. In this model, no global or ULA addresses are assigned to network infrastructure. The routing protocols all use link local addresses for forwarding packets. Protocols such as Simple Network Management Protocol (SNMP) or Terminal Access Controller Access (TACACS) can be sourced from the device loopback interface, which will have a global address assigned. There are some potential drawbacks to consider when analyzing this option. Infrastructure interfaces cannot be pinged anymore because link local addresses only having meaning on that particular link. Traceroute will not report the egress interface as the return ICMP packet would be sourced from the loopback interface. For more about the advantages and disadvantages of this approach, see the IETF Internet-Draft at http://wiki.tools.ietf.org/html/draft-behringer-lla-only-00

Building the Addressing Plan

There are several methods available to develop the IPv6 addressing plan:

- Existing IPv4 based plan is translated into IPv6
- Topologically based
- Organizationally based
- Services based

In the first method, some recognizable and unique part of the existing IPv4 subnet scheme is translated into an IPv6 subnet scheme. For example, a /48 is given to a customer, which gives the customer 16 bits to subnet their internal network. The customer is using the 10.0.0.0/8 network to address their network and has been allocated the 2001:DB8:1::/48 for their IPv6 address block. In this case the customer might choose to use the second and third octets in the IPv4 address to translate into their IPv6 address. For example, the 10.23.16.0/24 subnet would translate to 2001:DB8:1:1710::/64. Figure 14 graphically illustrates this process. This scheme becomes challenging to implement because of the variable length subnet masks that are common in an IPv4 subnet scheme.

Figure 14 – Converting IPv4 subnet to IPv6 subnet

There is an issue to keep in mind when trying to map an existing IPv4 address into an IPv6 subnet. The issue comes up when trying to map an existing /30 subnet into a /127 subnet for point-to-point infrastructure.

Figure 15 - Mapping an existing /30 subnet into a /127 subnet
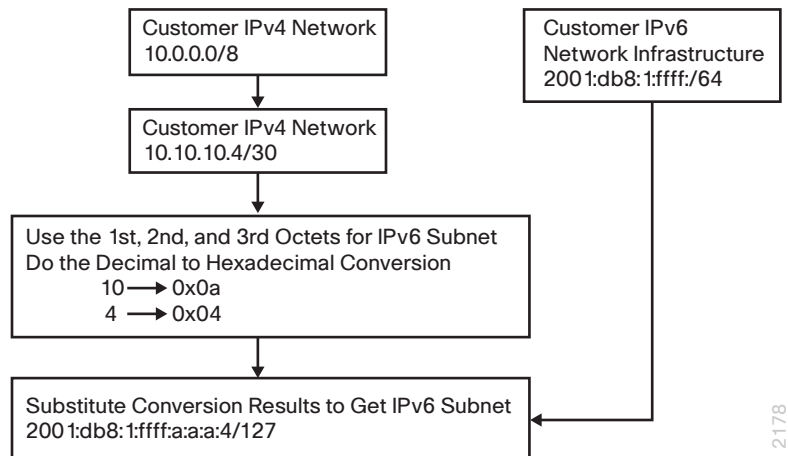
On the IPv4 network, the devices use 10.10.10.5 and 10.10.10.6 to identify the interfaces. Using the above translation method, the IPv6 addresses would be 2001:db8:1:ffff:a:a:a:5 and 2001:db8:1:ffff:a:a:a:6. The issue here is that 2001:db8:1:ffff:a:a:a:5 and 2001:db8:1:ffff:a:a:a:6 are on different subnets. Remember that a /127 only accommodates two addresses similar to a /31 prefix in IPv4. So 2001:db8:1:ffff:a:a:a:4/127 has 2001:db8:1:ffff:a:a:a:4 and 2001:db8:1:ffff:a:a:a:5 as end system addresses in that prefix. When considering using this model for prefix breakdown, keep this issue of overlapping prefixes in mind. Don't let the confusion of your existing IPv4 address plan creep into your IPv6 address plan.

The next method assigns a block of addresses to all locations within the topological constraints of the network. For example, a customer has been allocated the 2001:db8:1::/48 prefix by their provider and they have sites across the country that are topologically broken down into four regions by geography—northwest, northeast, southwest, and southeast. They might choose to use the first four bits of the 16 bits that they have for subnetting to identify the region. With this scheme the network could have sixteen regions and each region could have 4096 (2^{12}) /64 subnets. This scheme could be further pushed down to the facility level where the customer might choose to use the next four bits to identify a facility within a region, which would allow for 16 sites (2^4) per region with each site having a possible 256 (2^8) /64 subnets. This example was previously shown in the "Subnet Planning" section of this guide.

The next method involves assigning prefixes based on organizational boundaries within a company or organization. In this case, the engineering organization receives a block of addresses, the sales organization a different block, the legal department another block, and so on. A potential issue with this method is that it does not promote an efficient aggregation scheme. It is likely that most departments within an organization are located at multiple sites. Because of this organizational dispersion, this scheme is likely to be used in conjunction with a topological breakdown.

The last method is to assign prefixes based on the type of service that is offered, such as devices that provide voice over IP (VoIP) or wireless services. This method has the same aggregation issues as the organizational scheme and is also likely to be used in conjunction with the topological breakdown.

These last two methods could fall under a broader category of templatized addressing where information about the function of the device or subnet is embedded into the address itself. This approach does reduce of the potential for aggregation, but might simplify operations.

From the above example, a customer has received 2001:db8:1::/48 from their provider. They still choose to use the first four subnet bits to identify a region and the next four bits to identify a site. They elect to use the next four bits to identify the subnet function/service. The last four bits would be the different subnets within that function. The new breakout would look something like below.

Figure 16 – Subnet breakdown
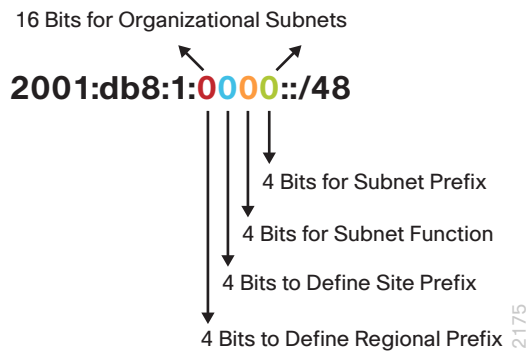
They use the following table to identify the subnet function/service.

Table 9 – Subnet function and service identity

Function/Service	Identifier
Reserved	0
Workstations	1
Reserved	2
Voice	4
Reserved	5
Wireless	6
Reserved	7
Servers	8
Reserved	9
Partners	a
Reserved	b
Reserved	c
Reserved	d
Reserved	e
Reserved	f

Using this model and given the following address—2001:db8:1:1181:227:13ff:fe68:854—the operations team would know that there is a potential issue with a server at site 1 region 1.

Recommendations for Building the Subnet Plan

- Use only /64 subnets for segments that have end systems/host attached.
- Use /128 for loopbacks.
- Use only /64, /126 or /127 subnets for network infrastructure.
    - Have a plan to accommodate network infrastructure segments that require more than two addresses (for example, /124)
    - Keep the subnet plan simple at first, using /64 prefixes for pilot projects and initial implementations. Move gradually into a more specific subnet address plan
- Take advantage of the network topology and the natural aggregation points to summarize prefix information.
    - Consider organization and services-based assignment within the summarization boundaries. Possibly sacrifice some of the strict summarization for operational simplicity.
- Leave gaps in the plan for growth. Consider the use of ULAs for network devices that do not need external connectivity (for example, printers).
- Track the status of link local address usage on network infrastructure.

Assigning Interface Identifiers

Another consideration when developing the addressing plan is how the interface identifier gets assigned to end stations and network infrastructure. As mentioned previously, there are several options available when assigning interface identifiers to an end host:

- Manual
- Stateless
- Privacy extensions
- Secure Neighbor Discovery/cryptographically generated address (SEND/CGA)
- DHCP

Manually configuring addresses on end stations means visiting each network node and configuring an interface identifier for that node. With this consideration in mind, manual address assignment should be reserved for network infrastructure devices and key network servers (for example, DNS servers, DHCP servers, database servers, web servers). There are some considerations that need to be accounted for when assigning addresses manually. The concerns are the same ones discussed previously in the "Subnet Planning" section of this guide—Initial Block Request section, related to the "u" and "g" bits supporting the router subnet anycast address, the IPv6 subnet anycast address, embedded RP addressing, and ISATAP addressing.

For manually assigned interface identifiers, avoiding easily guessed addresses (for example, DEADBEEF, CAFE, C0FFEE, etc.) is a good security practice and helps ensure that hackers are unlikely to find any hosts on a network scan. This recommendation is circumvented a bit for hosts that need to be publicly reachable. For publicly reachable hosts, DNS distributes the address information so that external hosts can communicate. It is still good practice, however, to avoid using easily guessed addresses for these publicly addressable servers.

It is recommended that when you manually assign addresses, you use a pseudo-random process to generate the interface ID portion of the address. Stateless address auto configuration (SLAAC) is a method where the node or device is able to automatically assign an address to itself. In this process, the node listens to specific messages that are sent out by routers on the segment. The node takes the subnet prefix information that the router is advertising and configures an interface ID. In the SLAAC model, there are three common processes that the end node can use to automatically configure the interface ID:

- EUI-64 process
- Privacy extensions
- Secure Neighbor Discovery/cryptographically generated address (SEND/CGA)

The EUI-64 uses the MAC address to build the interface ID. Because the interface ID requires 64 bits and the MAC address is only 48 bits, a method is needed to expand the MAC address. To accomplish this expansion, the MAC address is split in half and FFFE is inserted. The last part of the process is to set the universal/local bit. The universal/local bit is used to identify whether or not the address is universally or locally administered and is the seventh bit in the first octet. Figure 17 demonstrates the EUI-64 process.

Note that RFC 5342 describes the process and explains why FFFE was chosen.

Figure 17 - EUI-64 process

A significant piece of information is not distributed using the SLAAC process—the DNS server. With the expanded address size, DNS is going to be even more critical to overall IPv6 operations. RFC 6106 (IPv6 Router Advertisement Options for DNS Configuration) does specify an option for routers to distribute the DNS server information via router advertisements. This option is not widely supported. Keep in mind also that there are two parts to this feature—support in the router that is advertising the option and support in the host/server that is receiving the option.

Another potential drawback to SLAAC is the lack of a centralized mechanism of address management. Without the additional instrumentation, it is complicated to track back the device based on its IPv6 address. SLAAC is useful in mobile environments and network segments where "dumb" devices (for example, sensors) connect, as well as in the environments where a large number of clients may connect within a short period and a centralized model of DHCP is not acceptable.

Privacy concerns developed because the EUI-64 process is based on mapping the Layer 2 MAC address to the Layer 3 interface ID. The concern stems around the ability to track a device based on the unchanging interface ID. To address the privacy concerns, the IETF defined privacy extensions in RFC 4941to automatically generate pseudo random interface IDs.

Privacy extensions are another way to automatically assign an address to an end host. Using this process, an end host generates a pseudo-random interface identifier that is to be used for a specified time frame. When that time expires, the host generates another address that is used for communications, and so on. While privacy extensions do address the concerns outlined in RFC 4941, they put an administrative burden on the network operations staff. Processes and procedures for troubleshooting, accounting, authorization, access, etc. need to be developed to accommodate the changing end station addresses.

When considering privacy extensions, it is recommended that you use them for originating external communications to end systems outside of the organization's network (for example, the Internet), and use non-privacy assigned addresses for internal communications. Figure 12 shows how these communications might occur. Communication between site 1 and site 2 uses permanently assigned addresses and communication outside of the organization uses temporary addresses.

Figure 18 - Privacy extension example

Figure 19 shows a screen capture from a Windows 7 machine. Note that the interface has three globally unique addresses assigned to it. The first address is manually assigned and the second address is an example of privacy extensions in use. As the valid lifetime for this address expires, a new interface ID is generated. The third address that is assigned is the longer lived public address. Microsoft's default behavior for Windows 7 is to generate random numbers for use as interface identifiers. This choice can be seen below where the public address does not follow the EUI-64 process. Also observe that the lifetime on this address is much longer than the temporary address. The Public address is intended to be a much longer lived address.

Figure 19 - Windows 7 IPv6 addresses

To use these temporary addresses, the default behavior for source address selection has to be overridden. RFC 3484 specifies that public addresses are preferred over temporary addresses. The RFC specifies that applications must provide a mechanism to override this behavior, which should be kept in mind when considering using privacy extensions in this way. Another recommendation for privacy extensions is to keep the time frame for generating new addresses to a "reasonable" period. "Reasonable" is relative to the organization building the address plan and depends on the sensitivity of the organization to privacy. The time period recommended in the RFC is to change the address daily, which should meet the requirements for most organizations.

Not all operating systems adhere to keeping the temporary addresses constant for the entirety of the period. For example, iPhone and iPad devices tend to generate a new temporary address each time they re-associate with the wireless network, thus potentially dramatically shortening the lifetime of an Ipv6 address. This property may affect the design and configuration of various features in the network.

Another method to dynamically generate an interface ID is SEND/CGA, which was developed to address security concerns with the neighbor discovery process as outlined in RFC 3756 - IPv6 Neighbor Discovery (ND) Trust Models and Threats. The key idea behind the SEND/CGA process is to use public/private keys and certificates to validate the identities of all equipment associated with the neighbor discovery process. CGA is a lightweight mechanism that provides good protection against Layer 3 address spoofing. SEND is a PKI-based mechanism that provides good protection against router spoofing. The two work in combination to alleviate the threats identified in RFC 3756.

Figure 20 shows the CGA process in action. As noted in RFC 3972:

> *The basic idea is to generate the interface identifier (i.e., the rightmost 64 bits) of the IPv6 address by computing a cryptographic hash of the public key.*

As Figure 20 shows, the CGA process uses the public key and other parameters to generate a cryptographic hash that is used as the interface identifier. The device can then use the private key to sign messages sent from the address. Other devices can then verify the address using the public key.

Figure 20 – Cryptographically generated address process
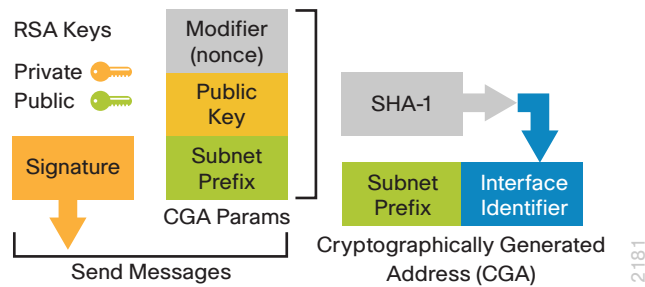
The overall goal of this process is to ensure that neighbors on a segment are who they say they are and provide some security for the neighbor discovery process. SEND/CGA is another feature that has two parts to it for a full solution. This feature requires both routers and hosts/servers support these address types.

The limiting factor for deployment of SEND/CGA is operating system support. Current workstation deployment of SEND/CGA is limited to most Linux systems. Support for SEND/CGA is a work in progress for most end system and network infrastructure vendors, including Cisco with Cisco IOS. The lack of support for SEND and CGA limits the deployment of this addressing method. The recommendation is to use SEND and CGA as more devices integrate support for these features.

DHCP has also been extended to support IPv6. DHCPv6 gives network administrators more control over how interface addressing is handled and includes both improvements based on lessons learned from DHCP and some new features (for example, DHCP prefix delegation). There are two states that DHCPv6 can operate in:

- **Stateful**—In this mode, DHCPv6 operates just like DHCPv4. It hands out and tracks address usage for segments. Network operators can track where addresses are and what addresses are in use.

- **Stateless**—In this mode, DHCPv6 is used primarily to hand out such things as DNS server and domain name information. This information is used to supplement the information that is discovered during the SLAAC process discussed previously. In the stateless model, the DHCPv6 server does not hand out any address information and therefore does not have to maintain any state, such as tracking address leases or end node status.

RFC 5157 has some recommendations related to assigning addresses and the implications related to subnet scanning. Here are some overall recommendations when assigning interface identifiers:

- Use SEND/CGA where and when it is available.

- Manually assign interface identifiers to network infrastructure and key network servers.

- Use DHCPv6 where user accounting and tracking are important concerns.

- When using stateless autoconfig, use stateless DHCPv6 to hand out key information such as DNS servers.

- Do not use easily guessed interface IDs.

IP Address Management (IPAM)

With the expansion of the address space from 32 bits to 128 bits, the number of IPv6 addresses and the number of IPv6 prefixes has greatly increased. Gone are the days when the address plan can be conveniently managed and maintained with a spreadsheet. Any wide scale implementation of IPv6 must include an IPAM tool. This tool can help manage and maintain the IPv6 address space, and can also consolidate the management of DNS and DHCP services.

DNS translates domain names to numerical identifiers for all networking equipment in order to locate and address these devices worldwide. Domain names help enable Internet resource and user information to remain consistent, regardless of their physical locations. With the expanded address space, DNS is going to become an even more critical service. What user, or network operator for that matter, is going to want to remember 2001:db8:ef46:12da:bcde:4987:c:15c0? Without a fast, reliable, and secure DNS service, subscribers' broadband Internet access will be compromised.

DHCP is an automatic configuration protocol used in IP networks. It allows a computer to be configured automatically and maintains a database for keeping track of computers that have been connected to the network, which prevents two computers from accidentally being configured with the same IP address. Without DHCP, every device must be manually assigned a unique address, which today is a virtually impossible task. With highly mobile users, converged third- and fourth-generation (3G/4G) networks, and the growing number of end-user network-capable devices, high-capacity DHCP servers are indispensable.

IPv6 Address Plan Case Study

This section discusses how a fictional company might develop their IPv6 addressing plan. Company XYZ is a multinational corporation headquartered in San Francisco with regional headquarters in Dubai, Johannesburg, Hong Kong, Sydney, Tokyo, Budapest, Paris, Buenos Aires, Mexico City, Atlanta, and Montreal. Figure 21 shows the connectivity between these regional headquarter sites.

Figure 21 - Company XYZ backbone topology

These regional headquarters sites serve up to 100 remote office locations. The company's primary data center is co-located with the company headquarters in San Francisco. The backup data center is located in Atlanta. Internet connectivity is provided via the San Francisco, Atlanta, and Paris sites via three different service providers.

To address the potential issues with multihoming and the complexity of managing three different IPv6 address blocks, XYZ Corporation decided to request a provider independent block from ARIN. The company met the requirements as established by ARIN and based on the number of sites in the network the company was allocated the 2001:db8:1000::/36 block for their use. XYZ Corporation did not approach RIPE or any of the other registries to request address space after securing agreements with their service providers regarding announcements of their IPv6 prefixes. However, they were explicitly advised by their service providers that the longest prefix that they would accept is a /48.

XYZ Corporation decided to assign prefixes based on the regions in which sites are located and defined five regions—North America, South America, Europe, Africa, and Asia. They decided to use the first three bits to identify the region. Note that a block has been reserved between Asia and North America. These are the two markets where the largest growth is expected. The reserved block will allow both regions to expand into this space if needed. There are two other blocks reserved for growth in the other theaters. The high-level break down for the prefix block it received from ARIN is shown in Table 10.

Table 10 - Breakdown for prefix block for corporation XYZ

Region (3 bits)	Prefix
Reserved	2001:db8:1000::/39
North America	2001:db8:1200::/39
Reserved	2001:db8:1400::/39
Asia	2001:db8:1600::/39
Latin America	2001:db8:1800::/39
Africa	2001:db8:1a00::/39
Europe	2001:db8:1c00::/39
Reserved	2001:db8:1e00::/39

The next step is to develop a plan for each region. The North American region is used as an example and the resulting plan can then be applied to other regions.

For regional planning, the next five bits are used to identify the regional headquarters and major facilities, such as data centers or large user locations. This decision allows for 32 of these facilities to be identified. Each of these locations can then assign subnets out of this /44 prefix. The break down in Table 11 shows how the North American /39 prefix is further subdivided into /44 prefixes.

Table 11 – Subdivision of prefixes

Site	Prefix
Reserved	2001:db8:1200::/44
San Francisco	2001:db8:1210::/44
Reserved (San Francisco)	2001:db8:1220::/44
San Francisco data center	2001:db8:1230::/44
Reserved (San Francisco DC)	2001:db8:1240::/44
Atlanta	2001:db8:1250::/44
Reserved (Atlanta)	2001:db8:1260::/44
Atlanta data center	2001:db8:1270::/44
Reserved (Atlanta DC)	2001:db8:1280::/44
Montreal	2001:db8:1290::/44
Reserved (Montreal)	2001:db8:12a0::/44
Mexico City	2001:db8:12b0::/44
Reserved (Mexico City)	2001:db8:12c0::/44
Reserved	2001:db8:12d0::/44

This scheme still has the 2001:db8:12d0::/40 through the 2001:db8:13f0::/44 prefixes unassigned and available for future use. These blocks can be used for very large sites that are expected to consume more than the /52 prefix assigned to sites.

A /52 block is assigned to each regional headquarters site and the remote locations that attach to a regional headquarters. Using a /52 for each site attached to a regional HQ allows for 256 sites to attach to that regional hub. A /52 prefix gives each location 4096 /64 subnets. The table below shows a breakdown for a couple of sites connecting to the San Francisco regional hub.

Table 12 – Breakdown for sites connecting to the San Francisco regional hub

Site	Prefix
Reserved	2001:db8:1210::/52
San Francisco HQ	2001:db8:1210:1000::/52
Reserved	2001:db8:1210:2000::/52
Los Angeles	2001:db8:1210:3000::/52
Phoenix	2001:db8:1210:4000::/52
Las Vegas	2001:db8:1210:5000::/52
Seattle	2001:db8:1210:5000::/52
Denver	2001:db8:1210:6000::/52
Chicago	2001:db8:1210:7000::/52

To give an idea of the number of subnets that a regional site can support, the total range of /52 subnets for sites connecting to the San Francisco regional hub is 2001:db8:1210::/52 through 2001:db8:121f:f000::/52.

The next step is to breakdown the prefix block per facility. The first four bits are used to identify the building or floor where the subnet is located within that site. The next 4 bits identify the function of the subnet. The table below shows the various functions. Note that this table will be used across regions.

Table 13 – Identifying the function and service of the subnet

Function/Service	Identifier
Reserved	0
Workstations	1
DMZ	2
Voice	3
Lab	4
Wireless	5
Servers	6
Guest access	7
Reserved	8
Reserved	9
Reserved	a
Reserved	b
Reserved	c
Reserved	d
Reserved	e
Reserved	f

This scheme allows each site to support up to 16 buildings or floors. Within each building or floor, up to 16 subnets per service can be supported. To ease the overall management of the subnet plan, all facilities will use /64 subnets. The breakdown for a facility is shown in Table 14.

Table 14 – Breakdown for building 1 at the San Francisco regional HQ site

Building (4bits)	Service (4bits)	User/service subnets (4bits)
0001	Reserved	2001:db8:1210:1100::/64 to 2001:db8:1210:110f::/64
–	Workstations	2001:db8:1210:1110::/64 to 2001:db8:1210:111f::/64
–	DMZ	2001:db8:1210:1120::/64 to 2001:db8:1210:112f::/64
–	Voice	2001:db8:1210:1130::/64 to 2001:db8:1210:113f::/64
–	Lab	2001:db8:1210:1140::/64 to 2001:db8:1210:114f::/64
–	Wireless	2001:db8:1210:1150::/64 to 2001:db8:1210:115f::/64
–	Servers	2001:db8:1210:1160::/64 to 2001:db8:1210:116f::/64
–	Guest Access	2001:db8:1210:1170::/64 to 2001:db8:1210:117f::/64

For the infrastructure prefixes, /127 prefixes will be used for point-to-point links, /112 will be used for infrastructure that requires more than two addresses, and /128 prefixes will be assigned to loopback addresses. The 2001:db8:1fff:ffff::/64 prefix, from the reserved 2001:db8:1e00::/39 prefix, will be used to assign infrastructure addresses. 2001:db8:1fff:ffff:0::/80 will be reserved to avoid any conflicts with the universal/local bits. The first 3 bits after the /80 will be used to identify the regional infrastructure prefixes. This choice follows the breakdown for the user prefix assignment model. The breakdown for the regional infrastructure prefixes is shown in Table 15.

Table 15 – Breakdown of regional infrastructure prefixes

Region (3 bits)	Infrastructure prefix
Reserved	2001:db8:1fff:ffff:0::/83
North America	2001:db8:1fff:ffff:0:2000::/83
Reserved	2001:db8:1fff:ffff:0:4000::/83
Asia	2001:db8:1fff:ffff:0:6000::/83
Latin America	2001:db8:1fff:ffff:0:8000::/83
Africa	2001:db8:1fff:ffff:0:a000::/83
Europe	2001:db8:1fff:ffff:0:c000::/83
Reserved	2001:db8:1fff:ffff:0:e000::/83

The following table further breaks down the North American infrastructure prefix.

Table 16 – Further breakdown of North American infrastructure prefix

Function	Prefix
Point-to-point segments	2001:db8:1fff:ffff:0:2000::/96
Multiaccess segments	2001:db8:1fff:ffff:0:2e00::/96
Loopback	2001:db8:1fff:ffff:0:2f00::/96

The scheme can be further broken down to allow for a specific site to be identified. The following table shows how that breakdown might happen for the San Francisco site infrastructure.

Table 17 – Further breakdown of San Francisco site infrastructure

Function	Prefix
Point-to-point segments	2001:db8:1fff:ffff:0:2000:1000::/100
Multiaccess segments	2001:db8:1fff:ffff:0:2e00:1000::/100
Loopback	2001:db8:1fff:ffff:0:2f00:1000::/100

This scheme allows for 2^{17} point-to-point subnets, 2^{12} multiaccess subnets and 2^{18} network infrastructure devices per site.

For the individual address assignments to end systems, XYZ Corporation uses DHCPv6 to assign interface identifiers to user machines. Devices that are not capable of supporting DHCPv6 will use SLAAC. Because SLAAC will be used, the processes and procedures used to identify who is connecting to the network will have to be modified to accommodate these devices. A remediation plan for upgrading these devices to an operating system that supports DHCPv6 will be developed.

Key servers and all network infrastructure use manually-assigned interface identifiers. Privacy extensions are not used on the network. As CGA/SEND implementations become available, they will use SEND/CGA to help improve the overall security in the network. XYZ Corporation will work with both their network and end system vendors to get this feature integrated into products that are used in the XYZ network.

This example was shown using an initial block size of /36. The concepts presented here can also be used with a longer initial assignment such as a /40, /44 or even a /48. There are fewer bits available when the initial prefix assignment is longer than requested. The key is to adapt the plan to fit the assignment.

As an example, if the initial assignment in the above example was for a /40, the same breakdowns could be used. In this case though there would be more usage of the reserved blocks to account for the smaller space. The table below shows how the regional breakdown would look for 2001:db8:1100::/40. Three bits are still used to identify the regions.

Table 18 - Regional breakdown for 2001:db8:1100::/40

Region (3 bits)	Prefix
Reserved	2001:db8:1100::/43
North America	2001:db8:1120::/43
Reserved	2001:db8:1140::/43
Asia	2001:db8:1160::/43
Latin America	2001:db8:1180::/43
Africa	2001:db8:11a0::/43
Europe	2001:db8:11c0::/43
Reserved	2001:db8:11d0::/43

For the major sites within a region, five bits can still be used to identify those sites.

Table 19 – Major sites within a region

Regional Bits (3 bits)	Major Site Bits (5bits)	Site	Prefix
001	00000	Reserved	2001:db8:1120::/48
–	00001	San Francisco	2001:db8:1121::/48
–	00010	Reserved (San Francisco)	2001:db8:1122::/48
–	00011	Reserved (San Francisco)	2001:db8:1123::/48
–	00100	Reserved (San Francisco)	2001:db8:1124::/48
–	00101	Reserved (San Francisco)	2001:db8:1125::/48
–	00110	San Francisco data center	2001:db8:1126::/48
–	00111	Reserved (San Francisco DC)	2001:db8:1127::/48
–	01000	Reserved (San Francisco DC)	2001:db8:1128::/48
–	01001	Atlanta	2001:db8:1129::/48
–	01010	Reserved (Atlanta)	2001:db8:112a::/48
–	01011	Reserved (Atlanta)	2001:db8:112b::/48
–	01100	Reserved (Atlanta)	2001:db8:112c::/48
–	01101	Reserved (Atlanta)	2001:db8:112d::/48
–	01110	Atlanta data center	2001:db8:112e::/48
–	01111	Reserved (Atlanta DC)	2001:db8:112f::/48
–	10000	Reserved (Atlanta DC)	2001:db8:1130::/48
–	10001	Montreal	2001:db8:1131::/48
–	10010	Reserved (Montreal)	2001:db8:1132::/48
–	10011	Reserved (Montreal)	2001:db8:1133::/48
–	10100	Mexico City	2001:db8:1134::/48
–	10101	Reserved (Mexico City)	2001:db8:1135::/48
–	10110	Reserved (Mexico City)	2001:db8:1136::/48
–	10111	Reserved	2001:db8:1137::/48
–	11000	Reserved	2001:db8:1138::/48
–	11001	Reserved	2001:db8:1139::/48
–	11010	Reserved	2001:db8:113a::/48
–	11011	Reserved	2001:db8:113b::/48
–	11100	Reserved	2001:db8:113c::/48
–	11101	Reserved	2001:db8:113d::/48
–	11110	Reserved	2001:db8:113e::/48
–	11111	Reserved	2001:db8:113f::/48

The big difference in this plan is that more blocks are reserved so that the sites that attach to a regional hub can be accommodated. For the sites that attach to a hub, a /52 can still be assigned and the breakout is the same at that point. In this example, the San Francisco regional hub can accommodate up 80 sites with a /52.

Conclusion

With IPv4 address depletion looming on the horizon, integration of IPv6 into enterprise and service provider networks is coming. The regional registries have acknowledged that IPv4 address depletion is a reality and encouraged organizations to start the IPv6 integration process. A key step in that integration process is acquiring address and subsequently building a plan to deploy those addresses. This guide has outlined several approaches to acquiring IPv6 address space and building an addressing plan. The way that an organization approaches acquiring and deploying IPv6 address space depends on the needs of that organization, but planning for that process needs to start now. IPv6 is here—get ready!

Resources

IPv6 Resources

- Cisco.com IPv6 information at http://www.cisco.com/ipv6
- The IPv6 Forum. Cisco is a founding and active member of the IPv6 Forum. The mission is to promote the use of IPv6 protocol. http://www.ipv6forum.com/
- IPv6 Task Force around the World: http://www.ipv6tf.org/
  - North-America IPv6 Task Force: http://www.nav6tf.org/
  - European IPv6 Task Forces: http://www.ipv6tf.org/meet/tf/eutf.php
  - Japan IPv6 Promotion council: http://www.v6pc.jp/en/
  - Asia Pacific IPv6 Task Force: http://www.ap-ipv6tf.org/
- IPv6 books:
  - IPv6 for Enterprise Networks, Shannon McFarland, Muninder Sambi, Sanjay Hooda, Nikhil Sharma, ISBN 1587142279
  - Deploying IPv6 networks, Ciprian Popoviciu, Erik Levy-Abegnoli, and Patrick Grossetete, ISBN 1587052105
  - IPv6 Essentials, Silvia Hagen, ISBN 0596100582
  - Understanding IPv6, Joseph Davies, ISBN 0735624461
  - Cisco Self Study: Implementing Cisco IPv6 Networks, Regis Desmeules, ISBN 1587050862
  - Migrating to IPv6: A Practical Guide to Implementing IPv6 in Mobile and Fixed Networks, Marc Blanchet, ISBN 0471498920
  - Global IPv6 Strategies: From Business Analysis to Operational Planning, Patrick Grossetete, Ciprian Popoviciu, Fred Wettling, ISBN 1587053438

Addressing Resources

- IPv6 Addressing Architecture (RFC 4291): http://www.ietf.org/rfc/rfc4291.txt
- IPv6 Global Unicast Address Format (RFC 3587): http://www.ietf.org/rfc/rfc3587.txt
- Deprecating Site Local address (RFC 3879): http://www.ietf.org/rfc/rfc3879.txt
- Unique Local IPv6 Unicast Addresses (RFC 4193): http://www.ietf.org/rfc/rfc4193.txt
- Special-Use IPv6 Addresses http://www.ietf.org/rfc/rfc5156.txt
- Requirements for Address Selection Mechanisms http://www.ietf.org/rfc/rfc5221.txt
- SEcure Neighbor Discovery (SEND) http://www.ietf.org/rfc/rfc3971.txt
- Cryptographically Generated Addresses (CGA) http://www.ietf.org/rfc/rfc3972.txt
- IPv6 Address Prefix Reserved for Documentation http://www.ietf.org/rfc/rfc3849.txt
- Default Address Selection for Internet Protocol version 6 (IPv6) http://www.ietf.org/rfc/rfc3484.txt
- A Flexible Method for Managing the Assignment of Bits of an IPv6 Address Block http://www.ietf.org/rfc/rfc3531.txt

- Use of /127 Prefix Length Between Routers Considered Harmful http://www.ietf.org/rfc/rfc3627.txt
- Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address http://www.ietf.org/rfc/rfc3956.txt
- IPv6 Implications for Network Scanning http://www.ietf.org/rfc/rfc5157.txt
- Privacy Extensions for Stateless Address Autoconfiguration in IPv6 http://www.ietf.org/rfc/rfc3041.txt
- Reserved IPv6 Subnet Anycast Addresses http://www.ietf.org/rfc/rfc2526.txt
- IPv6 Unicast Address Assignment Considerations: http://tools.ietf.org/html/rfc5375
- IPv6 Top Level Aggregator (TLA) Assignment: http://www.iana.org/assignments/ipv6-tla-assignments
- IPv6 Multicast Address Assignment: http://www.iana.org/assignments/ipv6-multicast-addresses
- AFRINIC IPv6 Policies: http://www.afrinic.net/en/services/ipv6-programme
- ARIN Number Resource Policy Manual: https://www.arin.net/policy/nrpm.html
- ARIN IPv6 Info Center: https://www.arin.net/knowledge/ipv6_info_center.html
- LACNIC IPv6 Registration Services: http://lacnic.net/en/registro/ipv6.html
- RIPE NCC Registration Services: http://www.ripe.net/rs/ipv6/index.html

Feedback

Please use the feedback form to send comments and suggestions about this guide.

B-0000200-1 08/13