



IPv6 Addressing Considerations

IPv6 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. Current models show that the IPv4 address space will be exhausted in the 2010/2011 timeframe.¹

There are several available resources to help build an IPv6 integration plan. The IETF has several RFCs and drafts that lay out integration plans.² There are several books available that give a background on the technology and also layout an integration plan.³

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. Topics covered include:

- [Addressing Introduction](#)
 - [Address Representation](#)
 - [Address Types](#)
- [IPv6 Address Assignment Policies](#)
 - [Address Allocation Model](#)
- [Address Planning](#)
 - [Provider Independent Addressing](#)

1. For information on IPv4 address space exhaustion, see Geoff Huston's IPv4 address exhaustion predictions (<http://www.potaroo.net/tools/ipv4/index.html>), the wikipedia on IPv4 address exhaustion (http://en.wikipedia.org/wiki/IPv4_address_exhaustion), and the Tony Hain article (http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_8-3/ipv4.html).

2. RFCs and drafts for IPv6 integration include RFC 2460 (<http://www.ietf.org/rfc/rfc2460.txt>) and RFC 4291 (<http://tools.ietf.org/html/rfc4291>).

3. IPv6 books include Deploying IPv6 Networks, IPv6 Essentials, Understanding IPv6, Cisco Self Study: Implementing Cisco IPv6 Networks, Migrating to IPv6, and Global IPv6 Strategies. For more information, see [Resources](#).



- ULA Addressing
- Network Level Design Considerations
 - Subnet Planning—Initial Block Request
 - Subnet Planning—Aggregation
 - Subnet Planning—Growth
 - Subnet Planning—Prefix Length
- Building the Addressing Plan
- Assigning Interface Identifiers
- IPv6 Address Plan Case Study

Addressing Introduction

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:

- There is no case sensitivity. Lower case “a” means the same as capital “A”.
- 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 can be omitted, so a field of zeroes can be represented by a single 0.
- Trailing zeroes must be represented.
- Successive fields of zeroes can be shortened down to “::”. This shorthand representation can only occur once in the address.

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

```
2001:0DB8:130F:0000:0000:7000:0000:140B
```

```
2001:DB8:130F:0:0:7000:0:140B      (Leading zeroes)
   ^      ^      ^      ^
```

```
2001:DB8:130F:0:0:7000:0:140B      (Trailing zeroes)
                   ^^^
```

```
2001:DB8:130F::7000:0:140B         (Successive field of zeroes)
                   ^
```

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



The 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.

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. These addresses are typically used when a specific end system needs to communicate with another specific end system (i.e., the conversation is peer to peer). IPv6 unicast addresses also have a scope defined for them—global, unique local, and link local. [Figure 2](#) shows the scope for each defined address.

Figure 2 Address Scopes for IPv6




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 a unique local or globally unique address. The interface could also have both types of addresses.

A link local address is used for communications on a single link and packets with a link local source or destination address are not forwarded by a router off that link. Link local addresses only have meaning 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 is 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.

The recommendation in current RFCs is to use the last 64 bits of an IPv6 address as the interface identifier. There are several methods available to assign the interface identifier—manual, automatic/stateless, and DHCP. These methods are covered in greater detail in [Assigning Interface Identifiers](#).

Figure 3 Link Local Address Representation




Unique local addresses are defined by RFC 4193 (Unique Local IPv6 Unicast Addresses). 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.

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 that 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 Web (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 (e.g., RIPE, ARIN, APNIC). Global addresses are all currently assigned out of the 2000::/3 block.

Figure 5 Global Address Representation



The current globally unique block allocations to the regional registries is shown in [Table 1](#). The full list breakout can be found at <http://www.iana.org/assignments/ipv6-unicast-address-assignments>.

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:

- 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)

Multicast

A multicast address is defined as an identifier for a set of interfaces that typically belong to different nodes. Multicast addresses are normally used to identify groups of interfaces that are interested in receiving similar content (e.g., video). The conversation model in this case is a one-to-many model. Multicast addresses are all assigned 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 only intended for systems on a link and are not to 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.

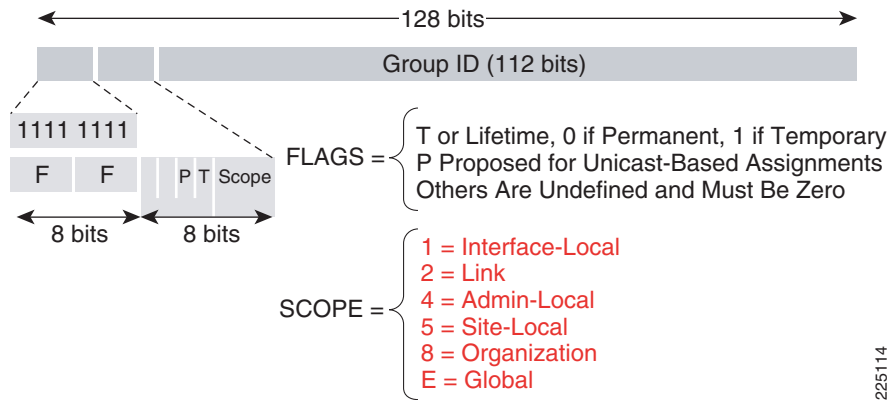
Similar to the unicast address space, there are some reserved or special use multicast addresses. A couple of the more common multicast groups and their intended use are mentioned below. For a more comprehensive list of currently assigned multicast addresses, see:

<http://www.iana.org/assignments/ipv6-multicast-addresses>

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

- FF02::1—Link local, all nodes address
- FF02::2—Link local, all routers address
- FF02:0:0:0:1:FFXX:XXXX—Link local, solicited-node address

Figure 6 Multicast Address Representation



Anycast

The last defined IPv6 address type is anycast (defined for IPv4 in RFC 1546 circa 1993, but rarely used). An anycast address is defined as an identifier assigned to multiple interfaces on different nodes. Anycast communications are similar to multicast communications, but the 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 from unicast when looking at the address (i.e., there are no defined bits that make an anycast address).

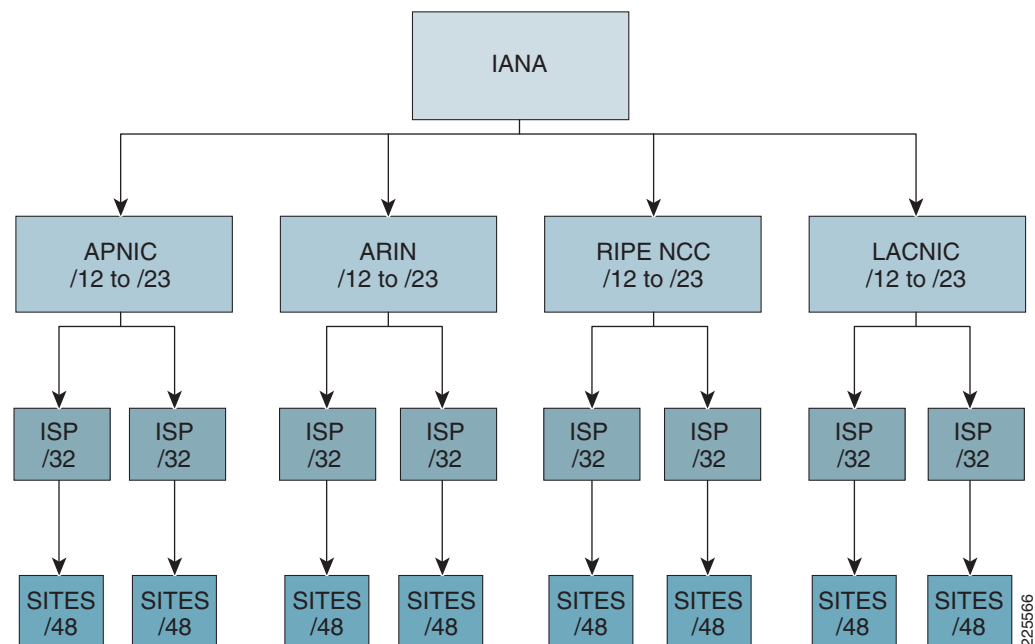
IPv6 Address Assignment Policies

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 their service provider to get IPv6 address space and cannot directly approach the regional registry for IPv6 address space.

Figure 7 graphically represents how this policy is enacted. The prefix lengths that are shown in Figure 7 are recommendations. 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



There is an exception to this policy that some registries have enacted that allows end customers to directly approach registries and request IPv6 address space. This exception is known as provider independent addressing.

The need for provider independent addressing arose because end customers wanted to multihomed to separate service providers. With the proposed allocation model, the customer would be assigned an address block from each service provider. Several approaches have been identified to address multihoming issues.¹ 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 handed out.

Provider independent addressing was adopted by some regional registries as an interim solution to multihoming. In 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.²

1. RFC 4177 Architectural Approaches to Multi-homing for IPv6.

2. ARIN PI policy (http://www.arin.net/policy/proposals/2005_1.html).

Because provider independent addressing has not been adopted by all regional registries, there are some potential issues with provider independent addressing that are discussed in [Address Planning](#).

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 that 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. An IETF draft¹ outlines some issues that also need to be taken into account when building an addressing plan. These considerations are discussed next.

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 their entire network when the provider address space changes.

Provider independent 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 primary issue for provider independent addressing is that it is still a regional concept. This can lead to problems for companies that are multi-national and located in regions that do not support the provider independent addressing model. Another approach for multi-national and large regional companies is discussed in [Subnet Planning—Initial Block Request](#).

There are also potential issues with how service providers handle provider independent address blocks. The current recommendation is to assign /48 prefix blocks 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.

Some recommendations on when to use provider independent address space:

- Your organization is contained within a single region, or multiple regions, that support provider independent space
- Your organization is connecting to multiple different providers
- Agreements are in place with your service providers to accept your IPv6 prefix announcements

ULA Addressing


When building the IPv6 address plan, a question might arise on whether or not to use globally unique addresses or unique local addresses. Remember these alternatives are not mutually exclusive. An IPv6 end point can, and most likely has, multiple IPv6 addresses. Hence both unique local and global addresses can be used. Global addresses must be used if Internet connectivity is desired.

1. draft-ietf-v6ops-addcon-10.txt (<http://www.ietf.org/internet-drafts/draft-ietf-v6ops-addcon-10.txt>).

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 8](#) shows an example of how this scheme might be used.

Figure 8 ULA and Global Address Communications



In [Figure 8](#), 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 8](#). In this example, communications that cross the Internet boundary use the addresses from the 2001:DB8::/32 block shown in [Figure 8](#). It should be noted 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 multicast traffic that is Internet bound and the ability to pass RPF checks. Remember 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.

Some 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 (e.g., printers).
- Do not use when using multicast.

A security recommendation is to filter ULA addresses at any external boundary to your organization. Unless specifically permitted by a prior agreement (e.g., 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.

Network Level Design Considerations

The vast size of the IPv6 address space gives network engineers a lot of 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 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.

Some things to consider when coming up with what the initial size of the IPv6 address block include:

- 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 it is going to request. The size of the network should take into account the number of subnets which is different 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 that the current assignment policy imposes.

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 (i.e., 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). 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 (e.g., /44) and should be from contiguous address blocks so that aggregation can still occur. This situation would be similar to the above IPv4 scenario.

Organizations could announce the more specific /48 blocks to draw traffic directly to those locations. At the same time they could announce the aggregate to handle redundancy if anything happens to the primary path.

Organizations that span 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.

This strategy works for an organization that is contained within a region (or regions) that supports provider independent space. If an organization is in a region that does not support provider independent space, then it should consider building a case to present to their registry for why it qualifies for a direct IPv6 address block assignment, much like a service provider would obtain. Although each registry has their own requirements, there has been some precedent set in ARIN for companies typically considered enterprise organizations. One example of an organization that has gone through the process and received a direct allocation from ARIN is Bechtel. Bechtel has been directly allocated the 2001:4920::/32 block.¹ The key is to build the request around the requirements that the registry has for the assignment.

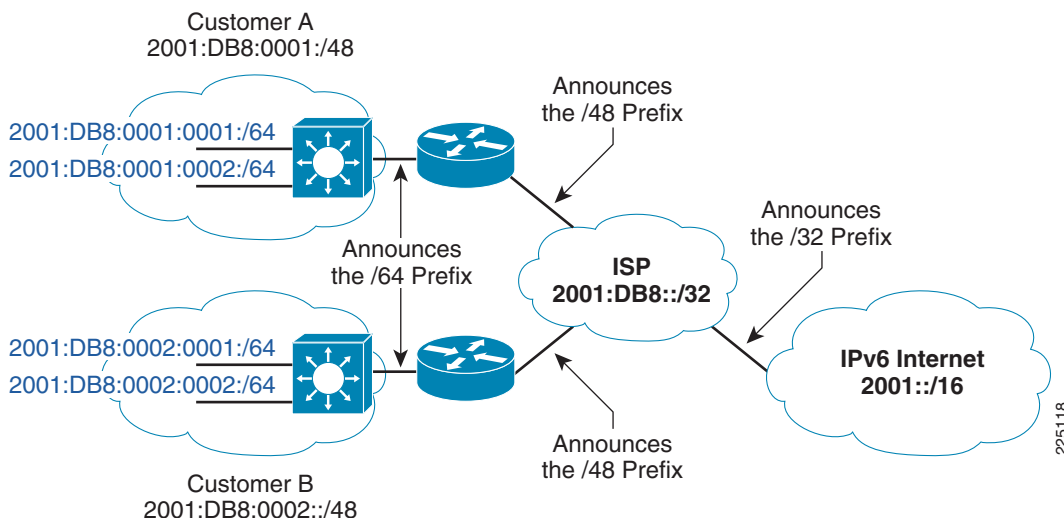
Subnet Planning—Aggregation

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 9 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.

1. For background information on Bechtel's IPv6 deployment, refer to:
http://www.cisco.com/web/strategy/docs/gov/bechtel_cs.pdf.

Figure 9 Hierarchical Addressing



Keep in mind that even though aggregation is key, 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. 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. Room 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. [Table 2](#) shows how this scheme might be implemented. In [Table 2](#), 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.

Table 2 **Address Plan for Growth**

Region (4 bits)	Regional Prefix	Facility (4 bits)	Facility Prefix	Subnets per Facility (8 bits)
1 (0000,0001)	2001:db8:1:0::/52	1 (0000)	2001:db8:1:0::/56	2001:db8:1:0::/64 to 2001:db8:1:ff::/64
	2001:db8:1:1::/52	2 (0100)	2001:db8:1:400::/56	2001:db8:1:400::/64 to 2001:db8:1:4ff::/64
		3 (1000)	2001:db8:1:800::/56	2001:db8:1:800::/64 to 2001:db8:1:8ff::/64
2 (0100)	2001:db8:1:4000::/52	1 (0000)	2001:db8:1:4000::/56	2001:db8:1:4000::/64 to 2001:db8:1:40ff::/64
		1 (0001)	2001:db8:1:4100::/56	2001:db8:1:4100::/64 to 2001:db8:1:41ff::/64
		2 (0100)	2001:db8:1:4400::/56	2001:db8:1:4400::/64 to 2001:db8:1:44ff::/64
		3 (1000)	2001:db8:1:4800::/56	2001:db8:1:4800::/64 to 2001:db8:1:48ff::/64
3 (1000)	2001:db8:1:8000::/52	1 (0000)	2001:db8:1:8000::/56	2001:db8:1:8000::/64 to 2001:db8:1:80ff::/64
		2 (0100)	2001:db8:1:8400::/56	2001:db8:1:8400::/64 to 2001:db8:1:84ff::/64
		3 (1000)	2001:db8:1:8800::/56	2001:db8:1:8800::/64 to 2001:db8:1:88ff::/64
4 (1100)	2001:db8:1:c000::/52	1 (0000)	2001:db8:1:c000::/56	2001:db8:1:c000::/64 to 2001:db8:1:c0ff::/64
		2 (0100)	2001:db8:1:c400::/56	2001:db8:1:c400::/64 to 2001:db8:1:c4ff::/64
		3 (1000)	2001:db8:1:c800::/56	2001:db8:1:c800::/64 to 2001:db8:1:c8ff::/64

Subnet Planning—Prefix Length

There are two areas to consider when looking at prefix lengths—segments that have end stations and infrastructure segments.

For segments that have 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 some more background and detail to these considerations.

Table 3 **Link Level Prefix Concerns**

64 Bits	< 64 Bits	> 64 Bits
<ul style="list-style-type: none"> • Recommended by RFC 3177 and IAB/IESG • 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 	<ul style="list-style-type: none"> • Enables more hosts per subnet • Considered bad practice • 64 bits offers more space for hosts than current media types and transport can efficiently support 	<ul style="list-style-type: none"> • Address space conservation • Special cases: <ul style="list-style-type: none"> – /126—Valid for p2p – /127—Not valid for p2p (RFC 3627) – /128—Typically used for network infrastructure Loopback addresses • Complicates management • Must avoid overlap with specific addresses: <ul style="list-style-type: none"> – Router Anycast (RFC 3513) – 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 consideration when using prefixes longer than /64 has to do with anycast addresses. Network planners should avoid the use of an all zero interface identifier, which has been defined by RFC 4291 as the subnet router anycast address. The other anycast address to avoid is the reserved IPv6 subnet anycast address defined in RFC 2526. In this case, the last seven bits are reserved for the anycast ID and the other bits of the identifier are set to 1.

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.

A recommended approach for network infrastructure would be to implement /64, /126, and /128 prefixes. A /128 is used for loopback addresses to identify network nodes. A /64 or a /126 is used for point-to-point links such as serial or POS links. RFC 3627 discusses why using a /127 prefix length for point-to-point links is not considered a best practice. A /64 prefix scheme is the simplest scheme to implement. A /126 prefix scheme 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 /126 scheme.

The example plan in [Table 2](#) will be used to demonstrate 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 (1 and 1110) are used to provide infrastructure links. Using this scheme, 8192 (2^{13}) infrastructure subnets can be assigned. This decision presents issues if region 4 experiences growth that requires more address space. In a real world situation, this decision would have to be analyzed against future growth requirements. Infrastructure links allocate bits in a manner similar to that in [Table 2](#). 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. [Table 4](#) shows how this scheme might be implemented.

**Note**

Each facility is actually receiving two infrastructure blocks for use at that facility. This scheme gives each facility 32 infrastructure subnets.

Table 4 /64 infrastructure Prefix Breakdown

Infrastructure prefix— 2001:db8:1:e000::/51				
Region (4 bits)	Regional Prefix	Facility (4 bits)	Facility prefix	Subnets per facility (4 bits)
1 (0000) (0001)	2001:db8:1:e000::/56	1 (0000) (0001)	2001:db8:1:e000::/59	2001:db8:1:e000::/64 to 2001:db8:1:e01f::/64
		2 (0100) (0101)	2001:db8:1:e040::/59	2001:db8:1:e040::/64 to 2001:db8:1:e05f::/64
		3 (1000) (1001)	2001:db8:1:e080::/59	2001:db8:1:e080::/64 to 2001:db8:1:e09f::/64
2 (0100)	2001:db8:1:e400::/56	1 (0000) (0001)	2001:db8:1:e400::/59	2001:db8:1:e400::/64 to 2001:db8:1:e41f::/64
		1 (0010) (0011)	2001:db8:1:e420::/59	2001:db8:1:e420::/64 to 2001:db8:1:e41f::/64
		2 (0100) (0101)	2001:db8:1:e440::/59	2001:db8:1:e440::/64 to 2001:db8:1:e44f::/64
3 (1000)	2001:db8:1:f800::/56	1 (0000) (0001)	2001:db8:1:f800::/59	2001:db8:1:f800::/64 to 2001:db8:1:f81f::/64
		2 (0100) (0101)	2001:db8:1:f840::/59	2001:db8:1:f840::/64 to 2001:db8:1:f85f::/64
		3 (1000) (1001)	2001:db8:1:f880::/59	2001:db8:1:f880::/64 to 2001:db8:1:f89f::/64
4 (1100)	2001:db8:1:fc00::/56	1 (0000) (0001)	2001:db8:1:fc00::/59	2001:db8:1:fc00::/64 to 2001:db8:1:fc1f::/64
		2 (0100) (0101)	2001:db8:1:fc40::/59	2001:db8:1:fc40::/64 to 2001:db8:1:fc5f::/64
		3 (1000) (1001)	2001:db8:1:fc80::/59	2001:db8:1:fc80::/64 to 2001:db8:1:fc9f::/64

An alternative implementation 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:fff::/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. A similar break down is used to identify regions and sites. Four bits are used to identify the region and four bits are used to identify the site. This scheme gives each site $\sim 2^{54}$ infrastructure subnets. [Table 5](#) shows how this scheme might be implemented.

Table 5 /126 Infrastructure Prefix Breakdown

Infrastructure prefix - 2001:db8:1:ffff::/64	Region (4 bits)	Regional Prefix	Facility (4 bits)	Facility prefix	Subnets per facility (54 bits)
1 (0000) (0001)		2001:db8:1:ffff:0::/68	1 (0000)	2001:db8:1:ffff:0::/72	2001:db8:1:ffff::/126 to 2001:db8:1:ffff:00ff:ffff:ffff:fffc::/126
			2 (0100)	2001:db8:1:ffff:0400::/72	2001:db8:1:ffff:0400::/126 to 2001:db8:1:ffff:04ff:ffff:ffff:fffc::/126
			3 (1000)	2001:db8:1:ffff:0800::/72	2001:db8:1:ffff:0800::/126 to 2001:db8:1:ffff:08ff:ffff:ffff:fffc::/126
2 (0100)		2001:db8:1:ffff:4000:/68	1 (0000)	2001:db8:1:ffff:4000::/72	2001:db8:1:ffff:4000::/126 to 2001:db8:1:ffff:40ff:ffff:ffff:fffc::/126
			2 (0100)	2001:db8:1:ffff:4400::/72	2001:db8:1:ffff:4200::/126 to 2001:db8:1:ffff:42ff:ffff:ffff:fffc::/126
			3 (1000)	2001:db8:1:ffff:4800::/72	2001:db8:1:ffff:4400::/126 to 2001:db8:1:ffff:44ff:ffff:ffff:fffc::/126
3 (1000)		2001:db8:1:ffff:8000:/68	1 (0000)	2001:db8:1:ffff:8000::/72	2001:db8:1:ffff:8000::/126 to 2001:db8:1:ffff:80ff:ffff:ffff:fffc::/126
			2 (0100)	2001:db8:1:ffff:8200::/72	2001:db8:1:ffff:8400::/126 to 2001:db8:1:ffff:84ff:ffff:ffff:fffc::/126
			3 (1000)	2001:db8:1:ffff:8400::/72	2001:db8:1:ffff:8800::/126 to 2001:db8:1:ffff:88ff:ffff:ffff:fffc::/126
4 (1100)		2001:db8:1:ffff:c000:/68	1 (0000)	2001:db8:1:ffff:c000::/72	2001:db8:1:ffff:c000::/126 to 2001:db8:1:ffff:c0ff:ffff:ffff:fffc::/126
			2 (0100)	2001:db8:1:ffff:c400::/72	2001:db8:1:ffff:c400::/126 to 2001:db8:1:ffff:c4ff:ffff:ffff:fffc::/126
			3 (1000)	2001:db8:1:ffff:c800::/72	2001:db8:1:ffff:c800::/126 to 2001:db8:1:ffff:c8ff:ffff:ffff:fffc::/126

This example highlights that using the /126 prefix breakdown for infrastructure links provides for greater address conservation by only allowing for 4 addresses per subnet. The example also shows that managing and maintaining this scheme is much more complicated—both in the planning and implementation of the scheme.

Another option exists for customers and their network infrastructure links. This option uses ULAs for network infrastructure. 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 that go 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.

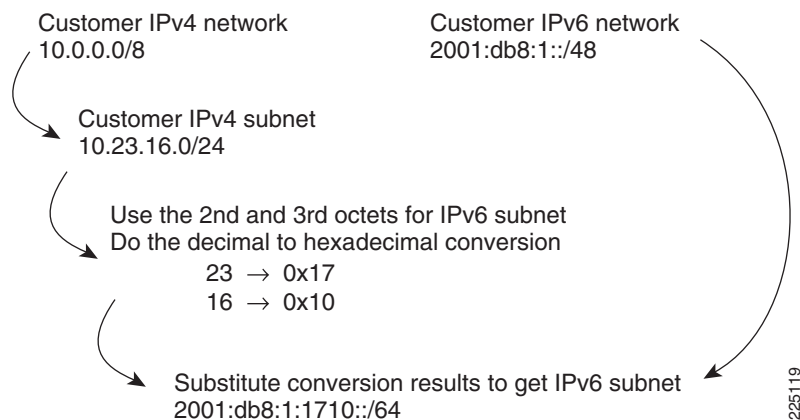
Building the Addressing Plan

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 10](#) 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 10 **Converting IPv4 Subnet to IPv6 Subnet**



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. [Table 2](#) shows how the breakdown might look.

The next method involves assigning prefixes based on organizational boundaries within a customer. In this case, the engineering organization receives a block of addresses, the sales organization a different block, legal another block, and so on. A significant issue with this method is that it does not promote an

efficient aggregation scheme. It is likely that most organizations within a company 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 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.

Some recommendations when building the subnet plan:

- Use only /64 subnets for segments that have end systems/host attached.
- Use only /64, /128, or /126 subnets on network infrastructure.
- 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.
- Leave gaps in the plan for growth.
- Keep the subnet plan simple at first, using /64 prefixes for pilot projects and initial implementations.
- Consider the use of ULAs for network infrastructure.
- Consider /126 prefixes on network infrastructure if there is a compelling need.

Assigning Interface Identifiers

Another consideration when developing the addressing plan is how the interface identifier gets assigned to end stations and network infrastructure. RFC 5157 has some recommendations related to assigning addresses and the implications related to subnet scanning. As mentioned previously, there are several options available when assigning interface identifiers to an end host:

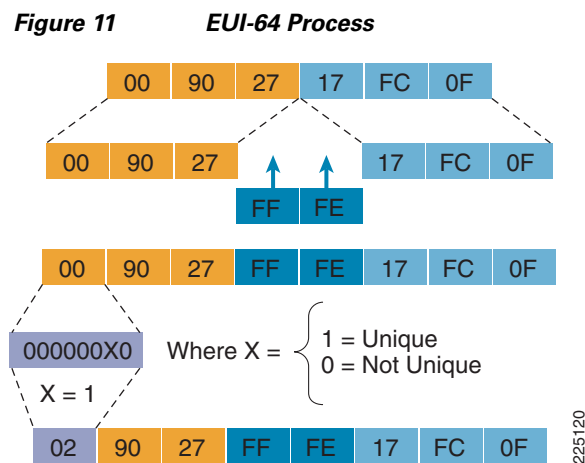
- Manual
- Stateless
- Privacy extensions
- SEND/CGA
- DHCP

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 (e.g., DNS servers, DHCP servers, database servers, Web servers). There are some considerations that need to be accounted for when assigning addresses manually, which are the same ones discussed previously in [Subnet Planning—Initial Block Request](#) related to the “u” and “g” bits, 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 (e.g., DEADBEEF, CAFE, COFFEE, 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. The recommendation when manually assigning addresses is to use a pseudo-random process to generate the interface ID portion of the address.

Stateless auto configuration 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. There are three common processes that the end node can use to automatically configure the interface ID:

- EUI-64 process
- Privacy extensions¹
- Secure Neighbor Discovery/Cryptographically Generated Address (SEND/CGA)²

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 11 demonstrates the EUI-64 process.



Stateless address auto-configuration (SLAAC) is another option for interface identifier assignment. Using SLAAC, an end station can automatically assign an address to itself and discover a default gateway. 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. Another potential drawback to SLAAC is the lack of AAA features for tracking who is connecting to the network. This deficiency could lead to some potential security issues where unauthorized users could connect to the network. SLAAC is useful in mobile environments and network segments where “dumb” devices (e.g., sensors) connect.

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, privacy extensions were developed to automatically generate 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.

1. See RFC 4941 Privacy Extensions for Stateless Address Autoconfiguration in IPv6.

2. See RFC 3971 SEcure Neighbor Discovery (SEND) and RFC 3972 Cryptographically Generated Addresses (CGA).

While privacy extensions do address the concerns outlined in RFC 3041, 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 (e.g., 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 12 Privacy Extension Example

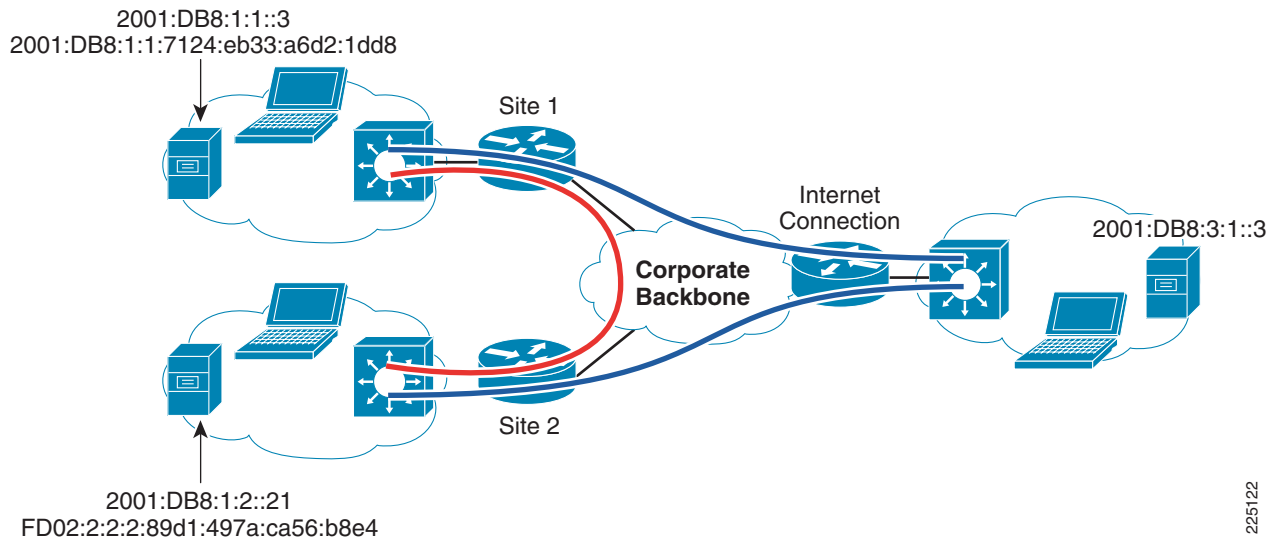


Figure 13 shows a screen capture from a Windows Vista machine. Note that the interface has three globally unique addresses assigned to it. The first address is manually assigned and the second address is an example of privacy extensions in use. As the valid lifetime for this address expires, a new interface ID is generated. The third address that is assigned is the longer lived public address. Note that Microsoft's default behavior for Vista/2008 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. Note also 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 13 Windows Vista IPv6 Addresses

```

Administrator: C:\Windows\system32\cmd.exe - netsh
operable program or batch file.
C:\Users\Administrator>netsh
netsh>interface
netsh interface>ipv6
netsh interface ipv6>show addresses

Interface 1: Loopback Pseudo-Interface 1
Addr Type  DAD State  Valid Life  Pref. Life  Address
-----
Other      Preferred  infinite   infinite   ::1

Interface 10: Local Area Connection* 6
Addr Type  DAD State  Valid Life  Pref. Life  Address
-----
Other      Preferred  infinite   infinite   fe80::5efe:169.254.35.16%10

Interface 9: Local Area Connection* 7
Addr Type  DAD State  Valid Life  Pref. Life  Address
-----
Other      Deprecated infinite   infinite   fe80::100:7f:fffe%9

Interface 8: IPV6 TEST INTERFACE
Addr Type  DAD State  Valid Life  Pref. Life  Address
-----
Manual     Preferred  infinite   infinite   2001:11::101
Temporary Preferred  1d3h28m1s  1d3h28m1s  2001:11::7124:eb33:a6d2:1dd8
Public     Preferred  29d23h57m1s 6d23h57m1s 2001:11::c834:c49b:29e8:2310
Other      Preferred  infinite   infinite   fe80::c834:c49b:29e8:2310%8

Interface 12: Local Area Connection* 10
Addr Type  DAD State  Valid Life  Pref. Life  Address
-----
Other      Preferred  infinite   infinite   fe80::5efe:172.18.82.61%12

Interface 11: MANAGEMENT INTERFACE
Addr Type  DAD State  Valid Life  Pref. Life  Address
-----
Other      Preferred  infinite   infinite   fe80::a5be:945:9d3b:21a1%11
netsh interface ipv6>

```

Note that 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 you should keep this 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.

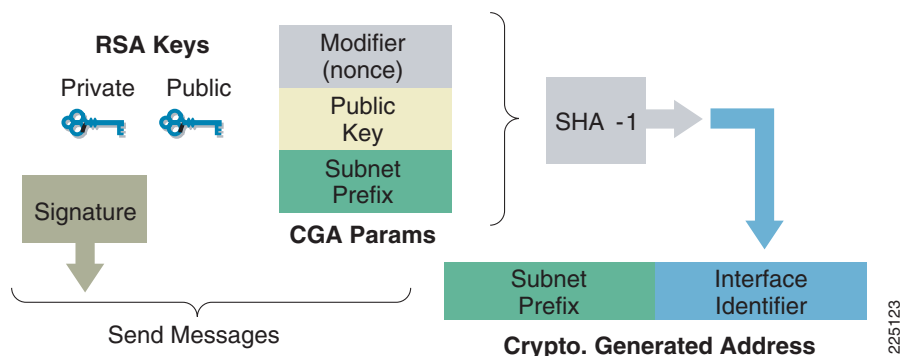
Another method to dynamically generate an interface ID is SEND/CGA, which was developed to address security concerns with the neighbor discovery process as outlined in RFC 3756 (IPv6 Neighbor Discovery (ND) Trust Models and Threats). The key idea behind the SEND/CGA process is to use public/private keys and certificates to ensure 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 14 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 14 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 14 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. The current 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 to include 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 IPv6¹. 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 (e.g., 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.

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.

1. RFC 3315 Dynamic Host Configuration Protocol for IPv6 (DHCPv6) and RFC 3736 Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6.

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 15](#) shows the connectivity between these regional headquarter sites.

Figure 15 **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.

For part of their first line of security, the XYZ corporation decided to use ULAs for network infrastructure. ULAs are not used to address end systems. XYZ corporation uses the `FD00:2001:db8::/48` prefix for infrastructure addressing. They use globally unique addresses for loopback addresses on network infrastructure to ensure that features like PMTUD work. XYZ corporation intends to peer with multiple 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 /32 block of address space from ARIN. The company met the requirements as established by ARIN and was allocated the `2001:db8::/32` 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.

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 for both the ARIN assigned and network infrastructure ULA prefix. The high-level break down for the prefix block it received from ARIN and the ULA prefix block for network infrastructure is shown in [Table 6](#).

Table 6 XYZ Corporation Prefix Assignments

Reserved—2001:db8:0000::/35	fd00:2001:db8:0::/51 (assigned to North America)
North America—2001:db8:2000::/35	fd00:2001:db8:2000::/51
South America—2001:db8:4000::/35	fd00:2001:db8:4000::/51
Europe—2001:db8:6000::/35	fd00:2001:db8:6000::/51
Reserved—2001:db8:8000::/35	fd00:2001:db8:8000::/51 (assigned to Europe)
Africa—2001:db8:a000::/35	fd00:2001:db8:a000::/51
Asia—2001:db8:c000::/35	fd00:2001:db8:c000::/51
Reserved—2001:db8:e000::/35	fd00:2001:db8:e000::/51

**Note**

For the infrastructure prefixes both Europe and North America are assigned consecutive blocks. This assignment recognizes that Europe and North America are larger and need a larger infrastructure block.

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 /40 prefix. Using a /40 per regional headquarters allows for a /48 to be assigned to up to 256 remote locations that connect to the regional headquarters. The break down in [Table 7](#) shows how the North American /35 prefix is further broken down into /40 prefixes.

Table 7 North American Major Facility Prefix Breakdown

Regional Bits (3 bits)	Major Facility Bits (5 bits)	Prefix	Facility
001	00000	2001:db8:2000::/40	San Francisco
	00001	2001:db8:2100::/40	
	00010	2001:db8:2200::/40	San Francisco data center
	00011	2001:db8:2300::/40	
	00100	2001:db8:2400::/40	Atlanta
	00101	2001:db8:2500::/40	
	00110	2001:db8:2600::/40	Atlanta data center
	00111	2001:db8:2700::/40	
	01000	2001:db8:2800::/40	Montreal
	01001	2001:db8:2900::/40	
	01010	2001:db8:2a00::/40	Mexico City

This scheme still has the 2001:db8:2b00::/40 through the 2001:db8:3f00::/40 prefixes unassigned and available for future use.

A /48 block is assigned to each regional headquarters and the remote locations that attach to a regional headquarters. A /48 gives each location approximately 65000 subnets.

The next step is to breakdown the prefix block per facility. The first four bits are used to identify the building on the site. The next four bits are used to identify the floor of the building. This plan leaves the last eight bits for use as user subnets on the floor. To ease the overall management of the subnet plan, all facilities use /64 subnets for user and infrastructure links. /128 subnets are assigned to loopback interfaces on infrastructure devices.

The breakdown for a facility is shown in [Table 8](#).

Table 8 San Francisco Facility Prefix Breakdown

San Francisco Facility Prefix—2001:db8:2000::/48		
Building (4 bits)	Floor (4 bits)	User subnets (8 bits)
0000	0000	2001:db8:2000:0000::/64 to 2001:db8:2000:00ff::/64
	0001	2001:db8:2000:0100::/64 to 2001:db8:2000:01ff::/64
	0010	2001:db8:2000:0200::/64 to 2001:db8:2000:02ff::/64
	0011	2001:db8:2000:0300::/64 to 2001:db8:2000:03ff::/64
	0100	2001:db8:2000:0400::/64 to 2001:db8:2000:04ff::/64
	0101	2001:db8:2000:0500::/64 to 2001:db8:2000:05ff::/64
	0110	2001:db8:2000:0600::/64 to 2001:db8:2000:06ff::/64
	0111	2001:db8:2000:0700::/64 to 2001:db8:2000:07ff::/64
	1000	2001:db8:2000:0800::/64 to 2001:db8:2000:08ff::/64
	1001	2001:db8:2000:0900::/64 to 2001:db8:2000:09ff::/64
	1010	2001:db8:2000:0a00::/64 to 2001:db8:2000:0aff::/64
	1011	2001:db8:2000:0b00::/64 to 2001:db8:2000:0bff::/64
	1100	2001:db8:2000:0c00::/64 to 2001:db8:2000:0cff::/64
	1101	2001:db8:2000:0d00::/64 to 2001:db8:2000:0dff::/64
	1110	2001:db8:2000:0e00::/64 to 2001:db8:2000:0eff::/64
	1111	2001:db8:2000:0f00::/64 to 2001:db8:2000:0fff::/64

For the infrastructure prefixes, the next 8 bits after the regional bits are used to identify facilities within that region. Larger facilities receive consecutive blocks to accommodate more infrastructure prefixes for that site. The breakdown for infrastructure prefixes is shown in [Table 9](#).

Table 9 Infrastructure Prefix Breakdown

Regional bits (3 bits)	Facility bits (8 bits)	Prefix	Prefixes available per facility (5 bits)
000 (North America)	00000000 (San Francisco)	fd00:2001:db8::/59	fd00:2001:db8::/64 to fd00:2001:db8:1f::/64
	00000001 (San Francisco)	fd00:2001:db8:20:/59	fd00:2001:db8:20::/64 to fd00:2001:db8:3f::/64
	00000010 (San Francisco)	fd00:2001:db8:40:/59	fd00:2001:db8:40::/64 to fd00:2001:db8:5f::/64
	00000011 (San Francisco)	fd00:2001:db8:60:/59	fd00:2001:db8:60::/64 to fd00:2001:db8:7f::/64
	00000110 (San Francisco data center)	fd00:2001:db8:c0:/59	fd00:2001:db8:c0::/64 to fd00:2001:db8:df::/64
	00000111 (San Francisco data center)	fd00:2001:db8:e0:/59	fd00:2001:db8:e0::/64 to fd00:2001:db8:ff::/64
	00001010 (Atlanta)	fd00:2001:db8:140:/59	fd00:2001:db8:140::/64 to fd00:2001:db8:15f::/64
	00001011 (Atlanta)	fd00:2001:db8:160:/59	fd00:2001:db8:160::/64 to fd00:2001:db8:17f::/64
	00001110 (Atlanta data center)	fd00:2001:db8:1c0:/59	fd00:2001:db8:120::/64 to fd00:2001:db8:13f::/64
	00001111 (Atlanta data center)	fd00:2001:db8:1e0:/59	fd00:2001:db8:140::/64 to fd00:2001:db8:15f::/64
001 (North America)	00000000 (Remote Site #1)	fd00:2001:db8:2000:/59	fd00:2001:db8:2000::/64 to fd00:2001:db8:201f::/64
	00000010 (Remote Site #2)	fd00:2001:db8:2040:/59	fd00:2001:db8:2040::/64 to fd00:2001:db8:205f::/64

XYZ corporation uses DHCPv6 to assign interface identifiers to user machines. 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 use SEND/CGA to help improve the overall security in the network.

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 paper 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 is going to depend 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/index.html>
 - Korea IPv6 Task Force: <http://www.ipv6.or.kr/eng/index.html>
- IPv6 books:
 - 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 (Draft): <http://www.ietf.org/internet-drafts/draft-ietf-v6ops-addcon-10.txt>
- 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>
- IPv6 Allocation List Regional Registries: <http://www.ripe.net/rs/ipv6/stats/index.html>
- AFRINIC IPv6 Policies: <http://www.afrinic.net/docs/policies/afpol-v6200407-000.htm>
- APNIC IPv6 Resources Guide: http://www.apnic.net/services/ipv6_guide.html
- ARIN IPv6 Registration Services: <http://www.arin.net/registration/ipv6/index.html>
- LACNIC IPv6 Registration Services: <http://lacnic.net/en/registro/ipv6.html>
- RIPE NCC Registration Services: <http://www.ripe.net/rs/ipv6/index.html>

CCDE, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0812R)