Infoblox®
CONTROL YOUR NETWORK

WHITEPAPER

# IPv6 Addressing Plan Basics

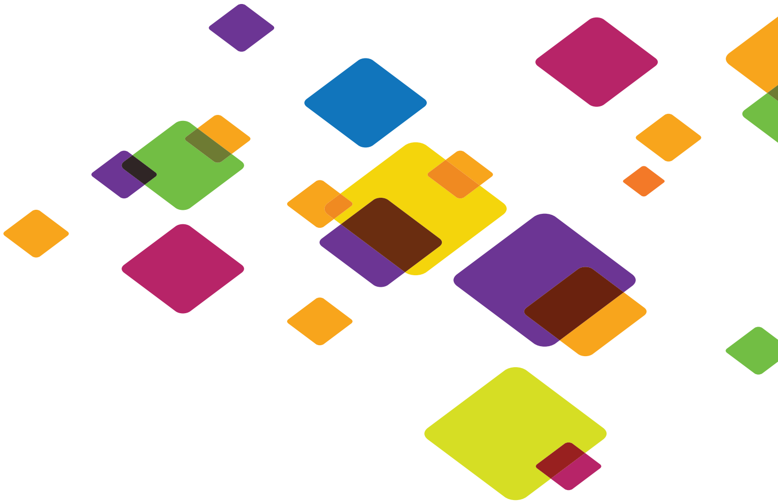

## Your IPv6 addressing plan

Coming up with an IPv6 addressing plan is one of the most important initial tasks for any organization proceeding with IPv6 adoption. And since most IPv6 deployments are green-field, creating an IPv6 addressing plan offers both unique challenges and benefits.

One critical challenge for most organizations is dealing with a new addressing protocol for which little, if any, internal operational or architectural experience may exist. Another challenge includes understanding how to most effectively manage the new complexities and overwhelming abundance of IPv6. This is especially critical given efforts by organizations to leverage emerging automation technologies for their networks (effective management of address resources allows for tighter host control and is perceived as a significant prerequisite for automation).

For most organizations, IPv6 will be deployed in parallel with IPv4 in an existing IPv4 production internetwork. It is perhaps natural to hope that the existing IPv4 addressing scheme might simply be duplicated in some fashion in IPv6. Depending on the organization, the design and administrative ease to deploy and manage IPv6 using such a method might initially prove advantageous. But any temporary advantage gained by such a shortcut will ultimately be dwarfed by the ease and efficiency of operation and design offered by a proper IPv6 addressing plan—one that incorporates the key benefits of the exponentially larger allocations possible with IPv6.

The virtually inexhaustible supply of IPv6 address space allows for an addressing plan no longer constrained by the scarcity of IPv4 addresses. Techniques such as Classless Inter-Domain Routing (CIDR) and Variable Length Subnet Masking (VLSM), previously required in IPv4 to economically match subnet size to host count on a given network segment, become unnecessary and obsolete in IPv6. Instead, a consistent and legible addressing scheme is made possible by the abundance of IPv6 addresses. As we'll see, this abundance also allows for the option of assigning significance to groups of site subnets according to function or location. This can make firewall policies and route aggregation much easier to design and administer. Such techniques, along with standard IPv6 subnet sizes, promise to improve operational efficiency and maximize future network scalability.

## How IPv6 addresses are constructed

Before getting into the details related to an IPv6 addressing plan, it may be useful to briefly review how IPv6 addresses are constructed. Recall that an IPv6 address consists of 128 bits. These bits are bisected to create a boundary between the “network number” and the “host number.” More specifically, the 64 bits of the network number portion of the address are divided into the global routing prefix and the subnet ID. The first three bits of any globally routable address are set to 001. The next 45 bits define the global routing prefix. The 16 bits following that are the subnet ID. The remaining 64 bits of the address are reserved for the interface ID. Also depicted below is which organization is responsible for assigning the relevant bits of the global routing prefix (see figure 1).

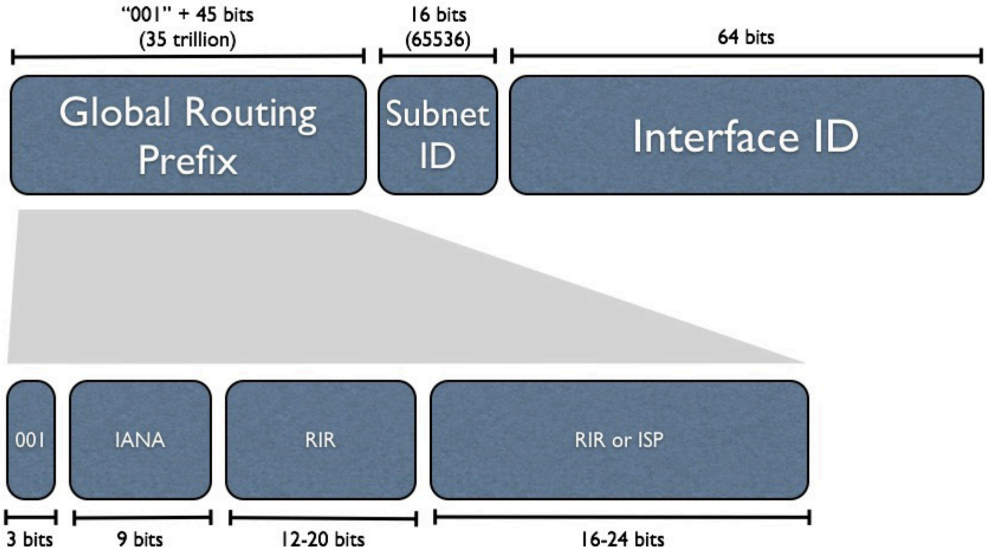

Figure 1

For most organizations, the majority of planning required for their IPv6 addressing scheme will focus on the 16 bits of the subnet ID.

## How many IPv6 addresses will you need?

This question might seem a little peculiar given that the basic interface subnet for IPv6, a /64 (or 1.8x1019 addresses), contains approximately 4 billion times the number of addresses available in all of IPv4 space. But as mentioned earlier, it is precisely this address abundance that provides new opportunities for greater consistency and scalability in IPv6 addressing plans.

First, you'll need to determine the appropriate size allocation to request for a primary allocation. Organizations with single sites typically receive a /48 (or 1.2x1024 addresses). Organizations with multiple sites receive a larger allocation based on the number of overall sites that need addressing. Table 1 shows the number of /48-sized sites supported by the number of bits in the global routing prefix for various allocation sizes.

| Number of /48-sized Sites | Number of bits in the Global Routing Prefix |
|---|---|
| 65,536 | 32 |
| 4,096 | 36 |
| 256 | 40 |
| 16 | 44 |
| 1 | 48 |

Table 1

It's unlikely of course that any single organization would have 65,536 sites in need of addressing! Rather, larger allocations are reserved by the RIRs for assignment to service providers, who will in turn allocate smaller prefixes to organizations based on their addressing requirements.

You may notice that only global routing prefix sizes in multiples of four bits are listed above. This practice, demonstrated in more detail later in the document, is used to preserve prefix legibility and concision (because prefix lengths not in multiples of four are not as immediately legible and may force address expansion to clarify the resulting groups of available subnets).

## Assigning Subnets and Addresses within a Site

The standard single site allocation is a /48 (the maximum number of bits allowed for a prefix that ISPs will accept and re-advertise). The subnet ID portion of the prefix allows for subnet assignments within a given site.

Any subnets derived from this primary site assignment should not be any smaller (i.e., more bits) than the standard interface assignment in IPv6 of a /64.

The more granular subnet assignments derived from the /48 assigned to a site have two primary characteristics. First, as mentioned above, they are typically larger than the standard /64 interface assignment and are appropriately sized to define groups of these interface subnets. Second, these groups will be sized and assigned based on the network function or location they provide addressing for.

Table 2 shows the number of subnet groups per /48 with the resulting number of /64 subnets.

| Number of bits in the Subnet ID | Number of subnet groups per /48 | Number of /64 subnets |
|---|---|---|
| 48 | 1 | 65,536 |
| 52 | 16 | 4,096 |
| 56 | 256 | 256 |
| 60 | 4,096 | 16 |

Table 2

These values should help guide how many groups of subnets, as well as /64s per group, are required for a given site. In general, the more /64s required for a particular function or location within the site, the fewer the number of larger subnets available for other locations or functions at that same site. Fortunately, most organizations are unlikely to use more than 4,096 /64 subnets for a single site function or location (a limit which would still allow for 15 additional /52s—each with 4,096 /64 subnets).

The determination of how many groups of subnets (and /64s per group) are needed should include current production requirements as well as expected growth. The 4-bit boundary preferred for prefix legibility leaves room in between for unexpected growth (though at the potential future cost of concision and legibility as we'll explore later).

Figure 2 demonstrates /48 site allocation subnet groups when the nibble boundary is adhered to.

IPv6 /48 Site Allocation Subnet Groups*

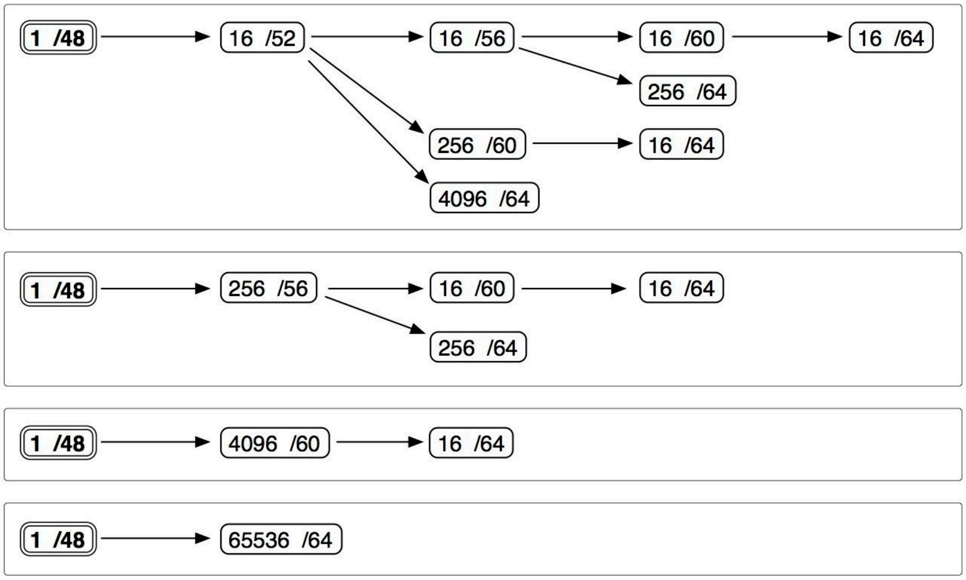

*(*adhering to 4 bit boundary)*

Figure 2

As mentioned, all network interfaces should be assigned a /64 by default. The only exceptions to this rule are for point-to-point links and loopback addresses.

In recent years, there has been much debate in the Internet engineering community and among network operators as to what size subnet should be used for point-to-point links. Initial IPv6 deployments often used /64s for consistency's sake but security issues (i.e., neighbor cache exhaustion and forwarding loops) emerged that made this method controversial. As a result, both /126s and /127s were proposed and deployed as preferable alternatives. More recently, RFC 6164 recommends the use of /127s on point-to-point links.

Figure 3 shows each of the standard IPv6 subnet assignments.

Standard IPv6 Address Assignments

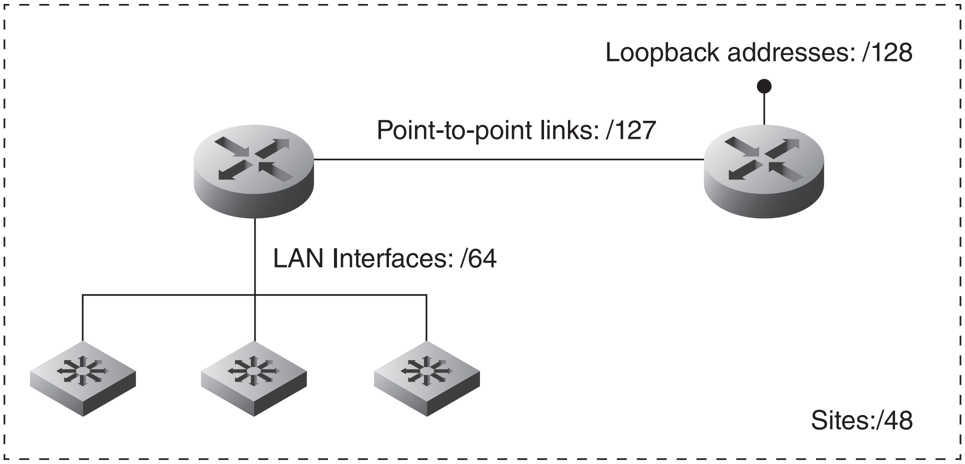

Figure 3

To preserve address plan consistency, every point-to-point link can be allocated a /64 but configured with a single /127 from that allocation. Alternatively, a /64 can be allocated per routing domain from which all of the /127s for that domain will be assigned.

/128s can be assigned from one /64 allocated per routing domain exclusively for loopback addresses.

## Getting IPv6 Addresses

Now that you have some information to determine how many IPv6 addresses you might need, the next step is to obtain an IPv6 address allocation. This is typically done in one of two ways: an organization will either request IPv6 space from their ISP or from one of the Regional Internet Registries (or RIRs). The decision to request an IPv6 allocation from an ISP or a RIR is based on whether an organization needs a Provider Aggregatable (PA) allocation or a Provider Independent (PI) allocation.

IPv6 allocations from ISPs or IP transit providers are typically PA allocations. These allocations are usually good for the duration of the contract for IPv6 connectivity from the ISP or transit provider. As a result, renumbering the network will be required if a new provider is chosen. If your organization's network connectivity and addressing requirements are based on a single site and you're planning on requesting IPv6 addresses from your ISP, now would be an excellent time to inquire as to whether your ISP supports IPv6 and what the requirements are for requesting IPv6 addresses.

Alternatively, IPv6 allocations from the RIRs are typically PI allocations and are considered permanent once allocated. Such allocations are "portable" meaning that the PI allocation will generally be accepted and re-announced by any ISP or IP transit provider. Both PI and PA allocations are globally registered and globally routable. In general, multi-site networks connected to multiple ISPs or transit providers require PI allocations.

PI allocations obtained through a RIR will require justification for the address space by filling out and submitting documentation detailing current and anticipated IPv6 address requirements. If an organization has obtained IPv4 space from a RIR in the past, the process is similar for IPv6. Most organizations should have no difficulty obtaining IPv6 addresses.

## A Sample Plan

Let's say you're creating an address plan for a campus LAN that has groups of network segments supporting particular functions (i.e., voice, wired and wireless data, etc.). In addition to the addressing required for these functions, you'll also need addresses for infrastructure. Finally, additional subnets for future use should also be included:

- Wired data
- VoIP
- Wireless
- Wireless guest
- Finance
- Infrastructure
- Future use

A total of six subnets would be required to support the above functions (along with two set aside for future use). This would require the reservation of three bits in the prefix to provide the necessary subnets (shown below by 'N', while 'X' are unspecified):

**2001:db8:abcd:[NNNXXXXXXXXXXXXX]::/51**

Note that while this provides sufficient subnets, the resulting prefix isn't as legible because the bit boundary doesn't align with the 4 bits used to define the hexadecimal character in the address:

**2001:db8:abcd:0000::/51**
**to**
**2001:db8:abcd:1FFF::/51**

**2001:db8:abcd:2000::/51**
**to**
**2001:db8:abcd:3FFF::/51**

Continuing with our example, the abundance of addresses available in IPv6 allows us to use 4 bits (instead of only 3), which makes the hexadecimal representation of the resulting subnets less ambiguous:

**2001:db8:abcd:0000::/52**
**to**
**2001:db8:abcd:0FFF::/52**

**2001:db8:abcd:1000::/52**
**to**
**2001:db8:abcd:1FFF::/52**

For each subnet group, only one value is possible for the hexadecimal character that corresponds to the 4-bit boundary in the IPv6 prefix (in this case, a /52). This makes the resulting prefix more immediately legible.

Obviously, the use of the more bits gives us more groups of subnets—16 in this case, 6 of which we'll use immediately and 10 for future. But fewer host ID bits also reduces the number of available /64 subnets in each group. In our above example, we went from 8,192 /64s available per /51 to only 4,096 /64s available with a /52.

As mentioned previously, we'll want to reserve a /64 each for both /127 point-to-point subnets as well as any /128 loopback addresses we may need.

Figure 4 below demonstrates the application of these principles depicting part of an IPv6 addressing plan for a ExampleCorp, a fictional enterprise.

**ExampleCorp ARIN Allocation:**
**2001:db8:cc0::/48**

| Location | Assignment |
|---|---|
| New York Office (HQ) | 2001:db8:cc0::/48 |

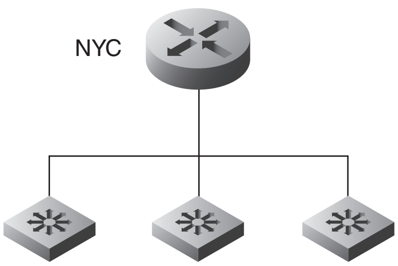

**Infrastructure Addressing**

| Function | Assignment |
|---|---|
| /64 for Loopback addresses | 2001:db8:cc0:f000::/64 |
| NYC router loopback address | 2001:db8:cc0:f000::1/128 |
| NYC SW1 loopback address | 2001:db8:cc0:f000::2/128 |
| NYC SW2 loopback address | 2001:db8:cc0:f000::3/128 |
| NYC SW3 loopback address | 2001:db8:cc0:f000::4/128 |
| /64 for Point-to-point links (future use) | 2001:db8:cc0:f001::/64 |

**New York Office (HQ)**

| Function and/or VLAN | Assignment |
|---|---|
| Wired Data | 2001:db8:cc0:0000::/52 |
| VoIP | 2001:db8:cc0:1000::/52 |
| Wireless | 2001:db8:cc0:2000::/52 |
| Wireless Guest | 2001:db8:cc0:3000::/52 |
| Finance | 2001:db8:cc0:4000::/52 |
| Reserved for future use | 2001:db8:cc0:[5-e]000::/52 |
| Infrastructure | 2001:db8:cc0:f000::/52 |

Figure 4

But what if your organization has more than one site? Since the standard site allocation is a /48 any organizations that have more than one site will receive a larger allocation from the RIR or ISP. The graphic below (figure 5) shows our fictional corporation expanding to three sites. Note that though a /46 would be sufficient to provide three /48s with one left over for future expansion, the RIRs, for now, typically adhere to the 4 bit boundary to obtain the benefits of legibility and efficiency already discussed. Thus, the /44 used for our example below.

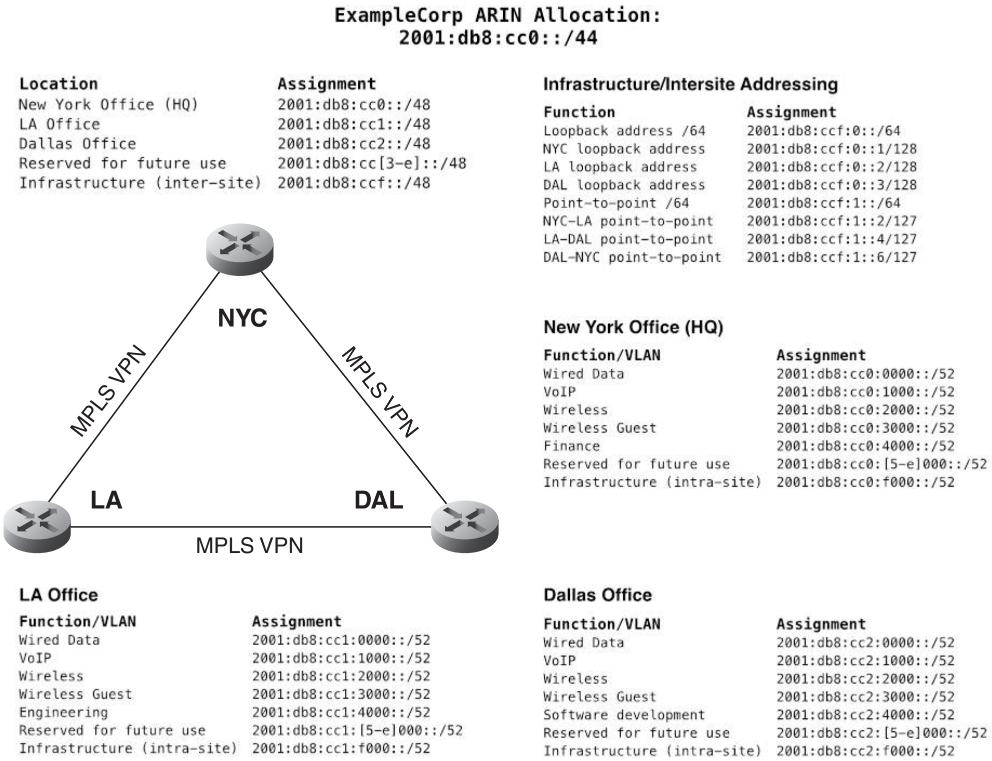

**ExampleCorp ARIN Allocation:**
**2001:db8:cc0::/44**

| Location | Assignment |
|---|---|
| New York Office (HQ) | 2001:db8:cc0::/48 |
| LA Office | 2001:db8:cc1::/48 |
| Dallas Office | 2001:db8:cc2::/48 |
| Reserved for future use | 2001:db8:cc[3-e]::/48 |
| Infrastructure (inter-site) | 2001:db8:ccf::/48 |

**Infrastructure/Intersite Addressing**

| Function | Assignment |
|---|---|
| Loopback address /64 | 2001:db8:ccf:0::/64 |
| NYC loopback address | 2001:db8:ccf:0::1/128 |
| LA loopback address | 2001:db8:ccf:0::2/128 |
| DAL loopback address | 2001:db8:ccf:0::3/128 |
| Point-to-point /64 | 2001:db8:ccf:1::/64 |
| NYC-LA point-to-point | 2001:db8:ccf:1::2/127 |
| LA-DAL point-to-point | 2001:db8:ccf:1::4/127 |
| DAL-NYC point-to-point | 2001:db8:ccf:1::6/127 |

**New York Office (HQ)**

| Function/VLAN | Assignment |
|---|---|
| Wired Data | 2001:db8:cc0:0000::/52 |
| VoIP | 2001:db8:cc0:1000::/52 |
| Wireless | 2001:db8:cc0:2000::/52 |
| Wireless Guest | 2001:db8:cc0:3000::/52 |
| Finance | 2001:db8:cc0:4000::/52 |
| Reserved for future use | 2001:db8:cc0:[5-e]000::/52 |
| Infrastructure (intra-site) | 2001:db8:cc0:f000::/52 |

**LA Office**

| Function/VLAN | Assignment |
|---|---|
| Wired Data | 2001:db8:cc1:0000::/52 |
| VoIP | 2001:db8:cc1:1000::/52 |
| Wireless | 2001:db8:cc1:2000::/52 |
| Wireless Guest | 2001:db8:cc1:3000::/52 |
| Engineering | 2001:db8:cc1:4000::/52 |
| Reserved for future use | 2001:db8:cc1:[5-e]000::/52 |
| Infrastructure (intra-site) | 2001:db8:cc1:f000::/52 |

**Dallas Office**

| Function/VLAN | Assignment |
|---|---|
| Wired Data | 2001:db8:cc2:0000::/52 |
| VoIP | 2001:db8:cc2:1000::/52 |
| Wireless | 2001:db8:cc2:2000::/52 |
| Wireless Guest | 2001:db8:cc2:3000::/52 |
| Software development | 2001:db8:cc2:4000::/52 |
| Reserved for future use | 2001:db8:cc2:[5-e]000::/52 |
| Infrastructure (intra-site) | 2001:db8:cc2:f000::/52 |

Figure 5

Notice that in both examples the function/VLAN subnet assignments are easily identifiable by the prefix given that the assignments conform to a 4-bit boundary. Also, note the assignment of function or location significance to an IPv6 subnet, which is the topic of the next section.

## Assigning Function & Location Significance to IPv6 Subnets

As mentioned, one major benefit resulting from the abundance of IPv6 addresses is the ability to assign function and location significance to groups of subnets derived from any larger allocation. While this technique works for any size allocation, we'll continue our focus on the /48 allocation typical for one site.

Groups of subnets with function or location significance can allow for easier route aggregation and more manageable design and enforcement of security policy. Having geographic or functional significance encoded into a group of subnets can also improve operational efficiency.

The first step for this technique is to choose a group of networks derived from our overall assignment for the site (in this case, a /48). These networks become the primary group of subnets and will be assigned function or location significance based on the requirements of the site.

Recalling our campus LAN example from earlier, most sites will have functions defined by groups of segments distributed across the entire site. With logical and well-defined boundaries between groups of subnets per function, security policy will be much easier to plan and administer. Thus, for most sites it's usually preferable to assign a group of primary subnets functional significance.

**2001:db8:abcd:[FFFFXXXXXXXXXXXX]::/52**

In this example, the characters in brackets represent individual bits rather than hexadecimal values with "F" indicating function significance (while the bits denoted by X are reserved for the /64s in each function and location group).

Thus, this example shows a subnet-grouping scheme that would support up to 16 functions:

**2001:db8:abcd:{0-F}XXX::/52**

With 4,096 /64 networks per function:

**2001:db8:abcd:X{0-F}{0-F}{0-F}::/64**

By contrast, if the site contains many sub-sites whose individual assignments will be aggregated at the parent site level (perhaps to allow individual intra-site location control of subnet definition and assignment) the primary group of subnets would be assigned location significance:

**2001:db8:abcd:[LLLLFFFFXXXXXXXX]::/NN**

In this example, the characters in brackets represent individual bits rather than hexadecimal values with “L” and “F” indicating location and function significance respectively. The bits denoted by “X” are reserved for the /64s in each location and function group (and the prefix length “NN” is unspecified for now).

Thus, this example shows a subnet-grouping scheme that would support up to 16 locations:

**2001:db8:abcd:{0-F}XXX::/52**

With up to 16 functions per location:

**2001:db8:abcd:X{0-F}XX::/56**

And 256 /64 networks per function:

**2001:db8:abcd:XX{0-F}{0-F}::/64**

This design choice makes sense where the groups of subnets defined and assigned to multiple locations will be aggregated and announced to an upstream (e.g., an ISP, into the core of a large enterprise network, or a hosting service).

Here’s an example of a fictional hosting service, CloudCo, that has chosen to encode location then functional significance into their IPv6 subnet assignments (figure 6).

**CloudCo ARIN Allocation:**
**2001:db8:dd0::/44**

| Location | Assignment |
|---|---|
| New York POP | 2001:db8:dd1::/48 |

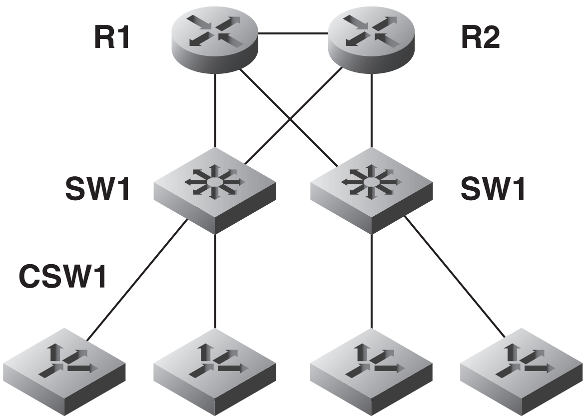

**New York POP**

| Location | Assignment |
|---|---|
| Customer rack 1 | 2001:db8:dd1:100::/56 |
| Customer rack 2 | 2001:db8:dd1:200::/56 |
| Customer rack 3 | 2001:db8:dd1:300::/56 |
| Customer rack 4 | 2001:db8:dd1:400::/56 |
| Customer rack 5 | 2001:db8:dd1:500::/56 |
| Customer rack n | 2001:db8:dd1:[05-ff]::/56 |

**Customer rack 1**

| Function | Assignment |
|---|---|
| HTTP | 2001:db8:dd1:110::/60 |
| Streaming Media | 2001:db8:dd1:120::/60 |
| Storage | 2001:db8:dd1:130::/60 |

**Infrastructure Addressing**

| Function | Assignment |
|---|---|
| /56 for local infrastructure | 2001:db8:ddf:100::/56 |
| /64 for Loopback addresses | 2001:db8:ddf:100::/64 |
| NYC R1 loopback address | 2001:db8:ddf:100::1/128 |
| NYC R2 loopback address | 2001:db8:ddf:100::2/128 |
| NYC SW1 loopback address | 2001:db8:ddf:100::3/128 |
| NYC SW2 loopback address | 2001:db8:ddf:100::4/128 |
| Customer rack CSW1 loopback address | 2001:db8:ddf:101::5/128 |
| /64 for Point-to-point links | 2001:db8:ddf:101::/64 |

Figure 6

## Host Address Assignment

Three primary methods for host assignment exist in IPv6. Two of these methods (static addressing and DHCP) should be familiar from IPv4, while one (Stateless Address Autoconfiguration, or SLAAC) is unique to IPv6.

As with IPv4, static addressing is typically utilized for servers, routers, switches, firewalls, and network management interfaces for any appliances—or any instance where address assignments are unlikely to change over time.

Stateless Address Autoconfiguration (SLAAC) is available on router interfaces that support IPv6 and will allow hosts on such a segment to self-assign a unique address (default router information is provided via ICMPv6 Router Advertisements). Because SLAAC does provide any authentication mechanism and allows a host to connect to the network and communicate with other nodes, this addressing method is not recommended where security is required or preferred. Lab environments or totally isolated networks where tight host control isn't a requirement are good candidates for the exclusive use of SLAAC.

Another issue with the use of SLAAC may arise where privacy extensions are enabled on the host. Privacy extensions allow the interface ID portion of a SLAAC-assigned address to be randomized in an effort to increase privacy for traffic originating from the host (otherwise the SLAAC assigned host address will always contain the traceable hardware address of the host's network interface). Privacy extensions are configurable in most host operating systems and may need to be disabled on the host if strict tracking and control of hosts is desired.

By contrast, Stateful DHCPv6 provides dynamic host address assignment but also includes the ability to pass additional options to the client. These options include information such DNS recursive name servers or the default domain name. Stateless DHCPv6 is yet another configuration option. With Stateless DHCPv6, SLAAC is used to provide host address assignment and default router information while DHCPv6 provides a list of DNS recursive name servers or the default domain name.

Finally, RFC6106 proposes including DNS server and search list information in RAs to provide host configuration options for SLAAC currently provided by DHCPv6, but it isn't widely implemented among host operating systems.

## Conclusion

Most organizations will find that over time their IPv6 addressing plan evolves as the best practices recommended in this document brush up against the particular technical and business requirements of the deploying organization. Of course a similar process often takes place with IPv4 addressing plans but IPv6 offers distinct advantages whenever growing and changing an IP addressing plan might become necessary. These advantages include the ability to easily obtain a larger IPv6 allocation from an ISP or RIR, the support for renumbering built into the protocol, and the considerable amount of unused IPv6 addresses when the initial plan adheres to a four bit boundary.

Such advantages are most significantly leveraged where organizations have deployed a DDI (DHCP, DNS, and IP Address Management) solution with robust IPv6 support. DDI allows for cost effective and safe modification and evolution of any IPv6 addressing plan. In any case, this flexibility should help dispel any presumed need for the perfect plan before you begin your own IPv6 adoption initiative.

CORPORATE HEADQUARTERS:
+1.408.986.4000
+1.866.463.6256
(toll-free, U.S. and Canada)
info@infoblox.com
www.infoblox.com

EMEA HEADQUARTERS:
+32 (0) 3.259.04.30
info-emea@infoblox.com

APAC HEADQUARTERS:
+(852) 3793.3428
sales-apac@infoblox.com