




Newer Design Guide Available

Cisco Smart Business Architecture has become part of the Cisco Validated Designs program.

For up-to-date guidance on the designs described in this guide, see <http://cvddocs.com/fw/Aug13-195>

For information about the Cisco Validated Design program, go to <http://www.cisco.com/go/cvd>





IPv4 Addressing Guide

● ● ● SMART BUSINESS ARCHITECTURE

February 2013 Series

Preface

Who Should Read This Guide

This Cisco® Smart Business Architecture (SBA) guide is for people who fill a variety of roles:

- Systems engineers who need standard procedures for implementing solutions
- Project managers who create statements of work for Cisco SBA implementations
- Sales partners who sell new technology or who create implementation documentation
- Trainers who need material for classroom instruction or on-the-job training

In general, you can also use Cisco SBA guides to improve consistency among engineers and deployments, as well as to improve scoping and costing of deployment jobs.

Release Series

Cisco strives to update and enhance SBA guides on a regular basis. As we develop a series of SBA guides, we test them together, as a complete system. To ensure the mutual compatibility of designs in Cisco SBA guides, you should use guides that belong to the same series.

The Release Notes for a series provides a summary of additions and changes made in the series.

All Cisco SBA guides include the series name on the cover and at the bottom left of each page. We name the series for the month and year that we release them, as follows:

month year Series

For example, the series of guides that we released in February 2013 is the “February Series”.

You can find the most recent series of SBA guides at the following sites:

Customer access: <http://www.cisco.com/go/sba>

Partner access: <http://www.cisco.com/go/sbachannel>

Comments and Questions

If you would like to comment on a guide or ask questions, please use the [SBA feedback form](#).

Table of Contents

What's In This SBA Guide	1	Managing IP Addresses	12
Cisco SBA Borderless Networks.....	1	IP Addressing in Cisco SBA.....	12
Route to Success	1	Appendix A: Subnet Design Worksheet for Cisco SBA	15
About This Guide	1		
IP Addressing Overview	2		
IP Addressing Basics	3		
Classful IP Addressing.....	3		
Subnetting and Supernetting	4		
VLSM	4		
Summarization.....	5		
IP Multicast	6		
Private IP Addressing	6		
Building an IP Addressing Plan	7		
Maintaining and Growing Network Space	8		

What's In This SBA Guide

Cisco SBA Borderless Networks

Cisco SBA helps you design and quickly deploy a full-service business network. A Cisco SBA deployment is prescriptive, out-of-the-box, scalable, and flexible.

Cisco SBA incorporates LAN, WAN, wireless, security, data center, application optimization, and unified communication technologies—tested together as a complete system. This component-level approach simplifies system integration of multiple technologies, allowing you to select solutions that solve your organization's problems—without worrying about the technical complexity.

Cisco SBA Borderless Networks is a comprehensive network design targeted at organizations with up to 10,000 connected users. The SBA Borderless Network architecture incorporates wired and wireless local area network (LAN) access, wide-area network (WAN) connectivity, WAN application optimization, and Internet edge security infrastructure.

Route to Success

To ensure your success when implementing the designs in this guide, you should first read any guides that this guide depends upon—shown to the left of this guide on the route below. As you read this guide, specific prerequisites are cited where they are applicable.

About This Guide

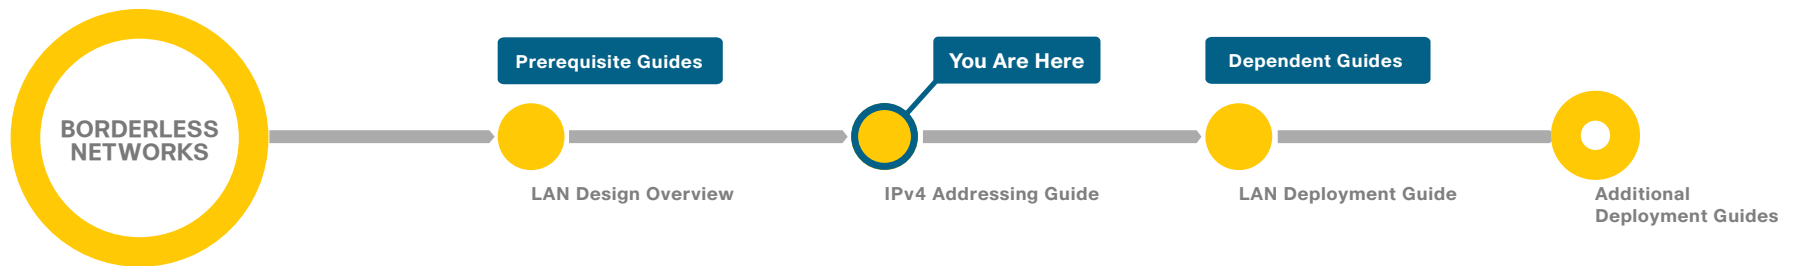
This *design overview* provides the following information:

- An introduction to a Cisco SBA design
- An explanation of the requirements that shaped the design
- A description of the benefits that the design will provide your organization

You can find the most recent series of Cisco SBA guides at the following sites:

Customer access: <http://www.cisco.com/go/sba>

Partner access: <http://www.cisco.com/go/sbachannel>



IP Addressing Overview

An IP address uniquely identifies each device in an IP network. Allocating, recycling, and documenting IP addresses and subnets in a network can become confusing very quickly if you have not laid out an IP addressing plan. A sound plan helps you prepare the network foundation to support additional services such as unified communications, wireless access, and enhanced network security.

IP addressing is a network foundation service, which makes it core to the network design. It provides the base for all other network and user services. Without the foundation, it would not be possible to interact with network and user services, from picking up the phone using the phone service to reading email using the email service.

By following recommended IP address management standards, you can avoid:

- Overlapping or duplicate addressing
- Unsummarized routes in the network
- Wasted IP address space
- Unnecessary complexity

Notes

IP Addressing Basics

IP version 4 (IPv4) addresses are 32 bits in length, divided into four eight-bit octets. Each octet has a decimal value from 0-255 or a binary value from 00000000-11111111. IPv4 addresses are typically shown in dotted decimal notation. This IP address in binary notation, 11000000.10101000.00000101.00000001, becomes the much more human-readable IP address 192.168.5.1 when converted to dotted decimal. IP addresses consist of a network prefix and a host portion. The bits of the address that make up the network prefix are determined by the subnet mask, and the remaining bits are used for hosts.

Table 1 - Example IP address in binary and dotted decimal notation

Type	Binary notation	Dotted decimal notation
IP address	11000000.10101000.00000101.00000001	192.168.5.1
Subnet mask	11111111.11111111.11111111.00000000	255.255.255.0
Network	11000000.10101000.00000101.00000000	192.168.5.0
Host	00000000.00000000.00000000.00000001	0.0.0.1

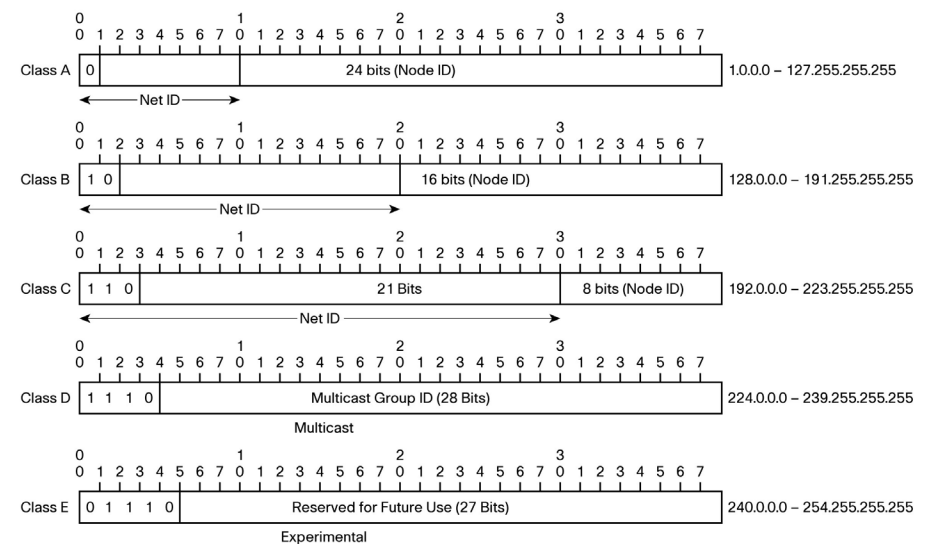
In Table 1, the IP address 192.168.5.1 has a subnet mask of 255.255.255.0. The network prefix portion of the address is 192.168.5 and the host portion of the address is .1.

Classful IP Addressing

IP addresses were initially divided into several different categories, or classes, based on how many bits of the address were used for the network prefix and how many bits were used for hosts. The classful addressing system defined class A, B, C, D, and E address spaces. Classes A, B, and C were used for general networking. Class D was designated for IP Multicast traffic, and class E was set aside for experimental use. As illustrated in Figure 1, the network prefixes for addresses in classes A, B, and C are characterized by the following:

- In class A addresses, the first octet is the network prefix and the remaining three octets are hosts.
- In class B addresses, the first two octets are the network prefix and the remaining two octets are hosts.
- In class C addresses, the first three octets are the network prefix and the remaining octet is the host.

Figure 1 - Classful addresses



In the classful model, the class A address space consists of 127 networks, each with 16,777,214 hosts. A single, flat network with over 16 million hosts is not practical. An organization can *subnet*, or use bits from the host portion of the address as part of the network prefix, to create more networks with fewer hosts. This is desirable in larger, distributed Layer 3routed networks that are used today.

Subnetting and Supernetting

Subnetting allows you to create multiple logical networks that exist within a single class A, B, or C network. Without subnetting you could create only one network from a class A, B, or C network, which would not be very useful.

Each Layer 2 segment on a network must have a unique network prefix, and each device on that segment must have a unique host address. Subnetting allows major networks (class A, B, or C networks) to be divided into smaller subnetworks (subnets). For example, instead of having a single network with 16 million hosts, a class A network could be divided into 65,535 networks, each with 254 host addresses; this is a much more practical network size for most organizations.

To subnet a network, bits from the host portion of the address are used to create additional subnets. In class A, if you take 16 bits from the host portion of the address—specifically the second and third octets—and use them to create subnets, you can make a large number of smaller networks. This gives you a more flexible and practical network design.

The class A network 10.0.0.0 has a mask of 255.0.0.0, commonly called its natural mask. Using the next 16 bits as the subnet gives you a mask of 255.255.255.0, or /24. This saves the last eight bits for host addresses, giving you 10.0.0.0-10.0.0.255 for the addresses for the network 10.0.0.0/24. This is shown in the example below. On the left-hand side the network 10.0.0.0 and subnet mask 255.255.255.0 is shown in decimal notation and on the right their binary counterparts are shown.

```
10.0.0.0      - 00001010.00000000.00000000.00000000
255.255.255.0 - 11111111.11111111.11111111.00000000
-----[ subnet ]-----
```

When planning IP subnetting, sometimes it is easier to visualize the network prefix and host portions of the address by looking at them in binary notation. The subnet mask is also represented in dotted decimal and binary. Any address bits that have corresponding mask bits set to 1 represent the network prefix, and address bits that have corresponding mask bits set to 0 represent the host addresses.



Tech Tip

The slash or prefix/length notation is shorter and easier to read; in this example, you are using the first 24 bits for the subnet mask, so the prefix/length notation is written as "/24". The network prefix/length notation is used throughout the rest of this document unless dotted decimal masks are needed to explain additional concepts.

In the early days of IPv4, organizations requested a class A, B, or C address space block based on their size. Many organizations were too large for the 256 addresses in a class C block, but didn't need the all 65,535 addresses from a class B block. A medium-size organization with a few thousand connected users would request a class B block, and over 90 percent of the address space would go unused. Subnetting in the classful model did not address this problem because you could only create additional, smaller networks out of an address block.

With *supernetting*, or classless interdomain routing (CIDR), you use bits from the network prefix to create a larger block of addresses. For example you can take several class C blocks and make the equivalent of one class B block—that is, 192.168.0.0 can be given a mask of 255.255.0.0 (or /16), effectively killing the old class-based system and turning all IP addresses into a single block that can be divided among organizations much more effectively. The hypothetical organization mentioned above with a few thousand connected users can be given an IP block with a /20 mask, giving them 4096 addresses to work with, which is a much more reasonable number for their size.

VLSM

Until CIDR and variable-length subnet masks (VLSM), when a network was subnetted there was only one subnet mask used for the whole block of addresses. The network 172.16.0.0 subnetted to /24 would consist of 256 /24 subnets. Although this is reasonable for most subnets with users, it is wasteful for several types of links common in networks today (for example, loopbacks and point-to-point WAN links) and smaller networks at remote sites. VLSM allow you to use different masks per subnet and use address space efficiently.

Use VLSM to:

- Create smaller subnets of fewer than 255 hosts for small WAN sites
- Create larger subnets of more than 255 hosts for large LANs
- Create very small subnets for WAN links
- Configure loopback addresses

VLSM Example

Given the 192.168.5.0/24 network and requirements below, you can develop a subnetting plan with the use of VLSM:

- Network A must support 80 hosts for users and devices at a remote site.
- Network B must support six hosts for a point-to-point WAN link supporting Hot Standby Router Protocol (HSRP).
- Network C must support two hosts for a T1 circuit to a remote site.
- Network D must support a single address for a router loopback.

The first step is to determine which mask allows the required number of hosts:

- Network A requires a /25 (255.255.255.128) mask to support 126 hosts
- Network B requires a /29 (255.255.255.248) mask to support 6 hosts
- Network C requires a /30 (255.255.255.252) mask to support 2 hosts
- Network D requires a /32 (255.255.255.255) mask to support 1 address

The easiest way to assign the subnets is to assign the largest first. For example, you can assign the subnets in this manner:

- Network A—192.168.5.0/25 address range .0 to .127
- Network B—192.168.5.128/29 address range .128 to .135
- Network C—192.168.5.136/30 address range .136 to .139
- Network D—192.168.5.140/32 address of .140




Reader Tip

For specific information on IP addressing and VLSM, see “IP Addressing and Subnetting for New Users” at: http://www.cisco.com/en/US/tech/tk365/technologies_tech_note-09186a00800a67f5.shtml.

Summarization

Network summarization allows for the creation of a single summary route made up of a block of contiguous subnets. In a network with no summarization, a route for each subnet exists in the routing process on every router in the network. This can cause several problems: First, networks with large numbers of subnets will have a large number of routes, which can use a lot of memory and CPU and degrade performance on smaller routers. Second, every time a subnet is added or removed from a network, the routers in the network have to recompute the routing table. While the network reconverges to a stable state, IP traffic may not be routed properly. In a stable network, routes are not often added or removed. This is not an issue, but an unreliable WAN link or malfunctioning router can cause a route to “flap,” causing frequent reconvergence of the routing table and network instability.

Figure 2 - Route summarization

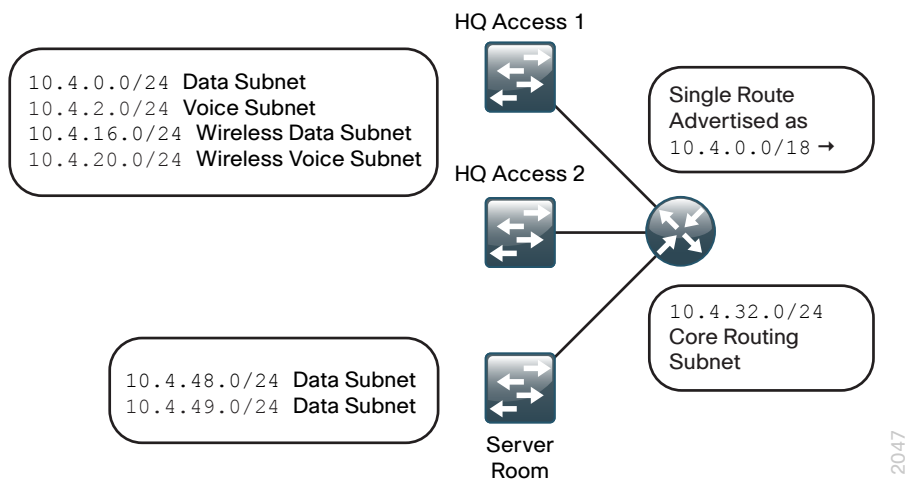


Summarization has a number of benefits. It reduces the size of the routing table, allowing the router to handle more routes. With fewer routes, the routing table reconverges faster, lessening the chance of suboptimal routing. Most importantly, a summary route will only be removed from the routing table if all of the routes it contains are removed—this means that any instability in the network behind the summary will not propagate to the rest of the network.

In smaller networks, summarization is often not necessary because of the small number of routes (although it can minimize the impact of unreliable WAN links, which is helpful in a network of any size). However, as an organization grows, the ability to summarize network routes allows the network to scale.

An example of summarization from the network headquarters out to the remote-site locations is shown in Figure 3. Normal IP routing advertisement would have sent out seven routes in the routing table. With summarization, all seven routes are summarized back to the headquarters as a single route.

Figure 3 - IP summarization at headquarters



Summarization can be used anywhere the addressing is contiguous, but is most commonly used where many networks connect, like WAN aggregation points, or at the edge of a LAN where it connects to the WAN. It is often difficult to add summarization to existing networks if it was not planned for during the initial network design.

IP Multicast

IP Multicast is a bandwidth conservation technology that reduces traffic and server loads by allowing a single stream of data (typically multimedia) on the network to be received by many users.

Applications that take advantage of multicast technologies include:

- Video conferencing
- Corporate communications
- Music on hold
- Distance learning
- Distribution of software, stock quotes, and news

Internet Assigned Numbers Authority (IANA) has reserved the range of 239.0.0.0/8 as administratively scoped addresses for use in private multicast domains. These addresses are similar to the reserved IP ranges defined in RFC 1918 in that they will not be assigned by IANA to any other group or protocol and are for private network use.

Reader Tip

For more information on IP Multicast, please visit <http://www.cisco.com/go/multicast>.

Private IP Addressing

IANA has reserved three blocks of addresses for use in private intranets.

10.0.0.0	-	10.255.255.255 (10/8 prefix)
172.16.0.0	-	172.31.255.255 (172.16/12 prefix)
192.168.0.0	-	192.168.255.255 (192.168/16 prefix)

These network addresses are expected to be filtered by ISPs per the RFC so that the addresses are not routed on the public Internet. RFC 1918 addresses, as they are commonly called, are reserved for organizations that want to build an internal TCP/IP network but either do not have, or do not want to use, public IP space. The supply of IPv4 addresses is dwindling at such a brisk rate that the only option for most organizations is to use RFC 1918 addresses, and many organizations with publicly routable addresses are renumbering with private addresses and using the public addresses for external-facing services, or even selling them. For access to the Internet from privately addressed networks, organizations can use Network Address Translation (NAT) or Port Address Translation (PAT). These technologies allow for a small pool of public IP addresses to be used by a large number of privately addressed hosts. The downside of private addressing is that because any organization can use addresses from this space, it is possible

for more than one organization to use the same addresses. Normally this is not an issue because these addresses are not routed outside the organization; however, when organizations merge with or are bought by other organizations, they might have difficulty connecting their networks together. One group will be forced to renumber parts of their network or use NAT inside the network, which limits their free access to network resources inside the newly formed organization.

Building an IP Addressing Plan

It is important to approach an IP addressing plan with the entire network in mind. Proper planning of IP address space across each of the layers is critical to ensuring that their interaction is seamless and integrated. Well-planned and documented IP address space can save many hours of troubleshooting time.

Define Addressing Standards

Using consistent standards across the different locations simplifies network maintenance and troubleshooting. Create standards for IP address assignments within each subnet range, and use them throughout the network. You may consider applying the following standards:

- Create addresses for VLANs to match the third octet of the IP subnet plus 100. For example, x.x.71.x is assigned VLAN 171. This results in a self-documenting design.
- Assign the first available addresses within a subnet to routers and HSRP virtual addresses. Routers may be assigned the .2 and .3 addresses, and the HSRP address assigned the .1 address.
- Reserve a small block of addresses for devices that may need static IP addresses (for example, printers).
- Use the remaining addresses for a Dynamic Host Configuration Protocol (DHCP) pool. For example, if the subnet is a /24 with 254 available address assignments for hosts, the DHCP pool could be .10 through .254.

Using a consistent, meaningful naming convention for Domain Name System (DNS) for network and endpoints makes management easier. Include the device type, site name, and interface or network information in the DNS name; this allows you to easily access and identify devices in the network.

Identify DHCP ranges and add them to DNS, including the location of the users. This range may be a portion of the IP address or a physical location. An example might be dhcp-bldg-c2-10 to dhcp-bldg-c2-254, which identifies IP addresses in building C, second floor.

IP Address Range Selection

Determine which address space to use by evaluating user and server requirements. For most organizations, private addressing is the only available option; this discussion assumes you are using RFC 1918 addressing.

The 192.168.0.0 range is used by many network equipment vendors for consumer and home devices. This address range has the lowest number of available addresses, a single /16, which can be quickly consumed in a growing organization. In addition, because this range is so commonly used, the likelihood of overlap is high when organizations merge. The 172.16.0.0/12 range is larger, equal to sixteen /16 networks, and most midsize and enterprise organizations could fit their network in this space. The 172.16.0.0/12 range is also less commonly used in consumer devices, so the risk of overlap or duplicate addressing is slightly lower. The 10.0.0.0/8 range is the largest, and all but the largest organizations could address their entire network inside this range easily.

When addressing a new network, there may be an advantage to starting somewhere other than the beginning of an address range. If organizations merge networks, it is more likely that there will be overlap if you use 10.0.0.0, but there is no guarantee that picking addresses randomly prevents overlap. It is more helpful in this scenario if your address space is from an easy-to-summarize block, say 10.192.0.0/12, and well documented.

Allocate IP Space

Carefully define the size of the IP space that will use public addresses, because it is available only in a finite amount. Be sure to take into account the following:

- Private addresses are not constrained.
- For ease of use, you should use /24 mask for user subnets.
- End devices always grow in number, so there is no reason to set a low limit on the number of private addresses, which are readily available.
- WAN connections have much smaller requirements for IP addresses. In general, a point-to-point network connection between two sites has two IP addresses in use. If HSRP is used with redundant routes on each side, the number of addresses increases to six—three for each side of the link.
 - /30 subnet mask allows for two usable IP addresses
 - /29 subnet allows for six usable IP addresses

Reserve a subnet for physical security. Security requirements can be as simple as a subnet to control door access to a building or something more complex like video surveillance for the entire building. Even if physical security is not required at the initial setup, you should still complete this step.

Reserve a subnet for facilities. This subnet addresses physical plant requirements such as remote power control, air conditioning, and facilities monitoring, which can now be monitored with new technology on the IP network.

Allocate public addresses for all production networks in the demilitarized zone (DMZ), which is the network or networks situated between an ISP edge router and corporate firewalls. An alternative is to use NAT.

Allocate a subnet for remote access, which is generally set up as a virtual private network (VPN).

Allocate a subnet for network management to provide access to network devices such as Ethernet switches, firewalls, routers, etc. This subnet allows for easy management with a separate logical network. Cisco SBA uses VLAN 1 for management of network devices.

Create a loopback address to make it easier to manage a single address for a router that has multiple interfaces.

Loopbacks:

- Are always up.
- Are reachable even if a single interface goes down when the router has multiple interfaces.
- Can provide a single source address for voice applications, network management, routing, etc.
- Give continuity to a voice gateway in a router. If the voice gateway is configured for the WAN interface and the WAN interface goes down, the voice gateway also goes down. Loopbacks prevent these problems because they stay up as long as the router stays up and is reachable over an interface.
- Need to be advertised by the IP routing protocol.

Loopback interfaces can be assigned an address of /24 or up to a single /32 (be sure to use summarization if you are using /32 loopbacks). Configure loopback interfaces as the source IP address for traps, Secure Shell (SSH) Protocol, and Simple Network Management Protocol (SNMP).

Documentation

Document the entire IP address space in a worksheet showing site-allocation, usage, and available subnets for each subnet size within the block, along with summary addresses for each particular block of addresses.

An example of a simple IP addressing worksheet is available in *Appendix A*, in this guide.

Maintaining and Growing Network Space

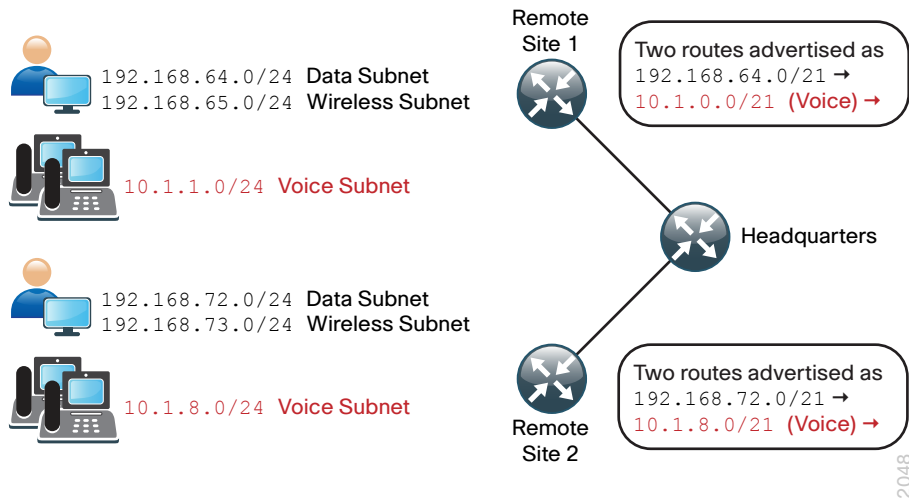
After the IP addressing plan is in place, you are ready to resolve situations as they come up. The following procedures explain how to handle some of the special situations you may encounter.

Overlay Networks

When adding devices to a network for a new service, sometimes it is helpful to overlay a new private IP address range on an existing IP addressing scheme. This can help solve scalability issues with an addressing plan that was not designed to accommodate enough subnets and hosts for each site to support a new service. A common scenario is adding voice to an existing data network. In a network that is currently numbered with 192.168.0.0/16, a voice subnet could be added to each site from the 10.0.0.0/8 or 172.16.0.0/16 blocks. This is simpler than renumbering the sites to include a voice subnet, and has the added benefit that voice traffic is easily identifiable. A simple mask covering all 172.16.0.0/16 or 10.0.0.0/8 addresses could be used to classify voice traffic across all sites for a QoS or security policy.

In Figure 4, the two existing remote sites have wired and wireless data networks, and you need to add voice. There are no addresses available in the 192.168.0.0 subnet space at the sites. To keep from renumbering the WAN, a voice subnet is overlaid in a 10.x.x.x address range at each site, highlighted in red in the illustration.

Figure 4 - Voice overlay subnets



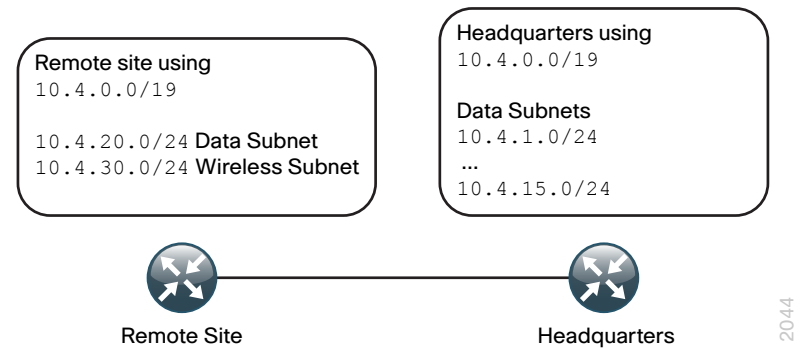
In this example, you could use subnets from 10.1.0.0/16 for all of the remote sites. This makes it easy to apply policy to the voice network.

Resolve Overlapping Address Ranges

When address ranges overlap in a network, it is commonly due to one of two reasons.

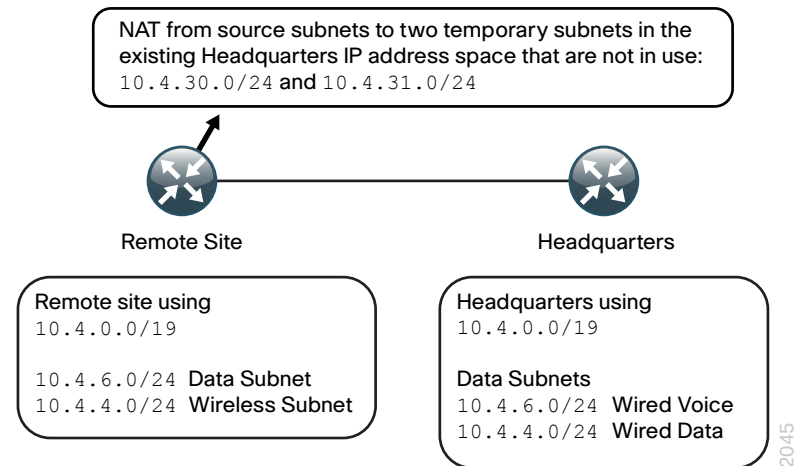
In the first scenario, the ranges are the same but the subnets do not overlap, as seen in Figure 5. This is really a result of over-summarization in the routing table. Removing the summarized routes from the remote-site router allows all hosts to communicate properly. This is not optimal because it means that the address space at the sites is not contiguous, but at least hosts at each site can communicate. Optionally, the addresses at the remote site could be renumbered into a new range that doesn't overlap with the 10.4.0.0/19 space at the headquarters site, and a new summary route could be added to the remote site router.

Figure 5 - Overlapping IP address space, non conflicting



In the second scenario, the same subnets have been allocated to two sites, as shown in Figure 6. This is a much more serious issue, because hosts on these networks will not be able to communicate properly on the network. The duplicate networks need to be readdressed with addresses that are not being used anywhere else in the network. If you need a temporary quick fix, you can use NAT to readdress the duplicate networks at the remote site to an unused address block on the network so they can communicate with the rest of the network. To make the NAT fix work, you must remove the duplicate networks from the routing table so they are not advertised into the network.

Figure 6 - Overlapping IP address space, conflicting



Increase the Number of IP Addresses

As an organization grows, so does its need for IP addresses. Ideally, growth would be built into the organization's addressing plan, but because you cannot build in infinite growth or see into the future, it often becomes necessary to add address space to a network.

A simple way to add space to an existing subnet is to move bits from the network portion of the subnet to the host portion. For example, by changing the subnet mask from /24 to /23, the number of supported hosts in the subnet is doubled. This approach works if the adjacent subnet is available. You need to reconfigure the router interface for the new mask and update the DHCP scope for the network to reflect the new address space. After the change, it may be necessary to have clients release and renew their addresses or reboot so that they have the correct information from the DHCP server.

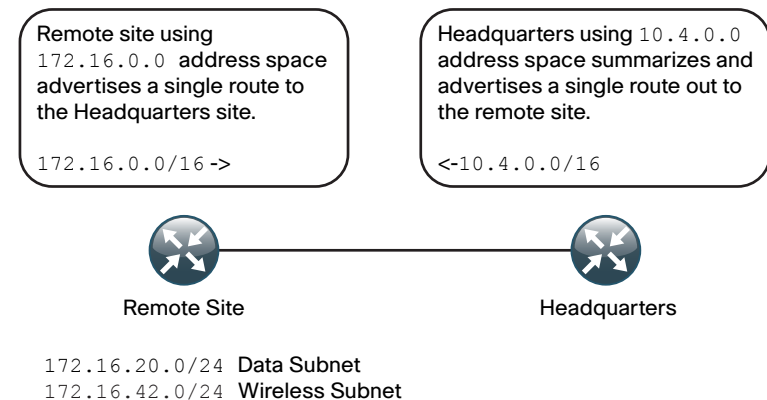
If there are no adjacent address blocks, there are two options:

- You can provision a new block of addresses that is large enough to handle the entire site, and the site can be renumbered. This has the advantage that addresses from that site are easily recognized and the whole site can be advertised to the rest of the network by using a single route summary.
- You can add a new block of addresses to the site. This is typically easier, because most of the subnets at the site will probably not have to be readdressed. This causes more routes to be advertised from the site out to the rest of the network, which typically is not a problem if the number of sites is small. If this is done on a large scale, adding hundreds of entries to the routing table, it can cause memory issues on older routers or stability issues on routers that do not have enough CPU to process a large number of routes.

Merge with Another Organization

When two organizations combine networks, there are several scenarios that might occur. The organization being connected to your network could be using addressing in a completely different block than yours, as shown in Figure 7. This is the simplest case to deal with. The networks can be connected and summarized routes advertised between the networks with very little risk that any issues will arise.

Figure 7 - Merging with a different IP address space



2043

As the new organization sites are migrated into the existing network design, they can be renumbered into the existing IP addressing plan or the plan can be expanded to accommodate more sites as needed.

The second scenario is when the new organization has addresses that are in the same block as your organization but there are no overlapping networks. For example, both groups could have used addresses from the block 10.0.0.0/8, making it difficult to send a single summary route between the networks without running into the over-summarization issue discussed in the previous section. The risk in this scenario is that you would potentially need to send most or all of the new organization's routes into your network unsummarized, which could cause issues on older routers with less memory or slow CPUs. You can minimize the impact by having good summarization in the existing network. After the networks are connected, the overlapping networks can be renumbered so that you can implement proper summarization.

The last and most difficult scenario to deal with is when there are duplicate addresses in use on both networks. When this occurs, it is often impossible to renumber the segments before you need to connect the networks, but you will need a temporary solution because the networks will not function properly unless something is done. You can use NAT at the peering point between the two networks to hide the overlapping addresses as a temporary fix until the networks can be readdressed.



Reader Tip

For more information on IP addressing, please see the following resources:

Routing and Switching Best Practices: How Cisco IT Deploys IP Addressing in the Enterprise, http://www.cisco.com/web/about/ciscoitatwork/downloads/ciscoitatwork/pdf/Cisco_IT_IP_Addressing_Best_Practices.pdf

Configuration Management: Best Practices White Paper
http://www.cisco.com/en/US/tech/tk869/tk769/technologies_white_paper09186a008014f924.shtml

Notes

Managing IP Addresses

With proper planning, the IP network will be more organized, easier to set up, and easier to troubleshoot.

IP Addressing in Cisco SBA

Cisco SBA uses the 10.0.0.0/8 and 192.168.0.0/16 address ranges. Your IP space requirements determine the range or ranges of addresses you implement. These same examples can apply to public address range, but most midsize organizations will likely deploy private addresses internally and use public addresses from an ISP when connecting to a public network such as the Internet.

Cisco SBA uses RFC1918 IP addresses 10.4.0.0/15 to cover the following main sections of the network:

- Headquarters
- WAN
- Remote sites
- Server room
- Regional site
- Security
- Voice

The Internet DMZ uses addresses from the 192.168.0.0/16 block.

The chosen address space is allocated in contiguous IP address blocks to each of these areas to allow for easy IP summarization:

- Headquarters is assigned addresses in the 10.4.0.0/18 block, allowing for 64/24 subnets.
- The regional site is assigned a separate block.
- Remote sites are assigned a block of 8 /24 subnets.
- For the network foundation, the following guidelines apply:
 - Data subnets are required for wired and wireless users.
 - Wireless access may require additional subnets for guest access.
 - Subnets for end-user access are commonly specified as /24 or /23, and allow for 254 hosts or 510 hosts, respectively, with a single address reserved for the default gateway.
 - Voice subnets are separate from data networks to simplify configuration of QoS and to avoid having IP phones contribute to the lack of space on data subnets that are already heavily populated.
 - One subnet for wired data and one subnet for wireless allows for flexibility. Subnets can be /24 or /23, and allow for 254 hosts or 510 hosts, respectively, with a single address reserved for the default gateway.

Table 2 presents an example of IP address assignment and summarization.

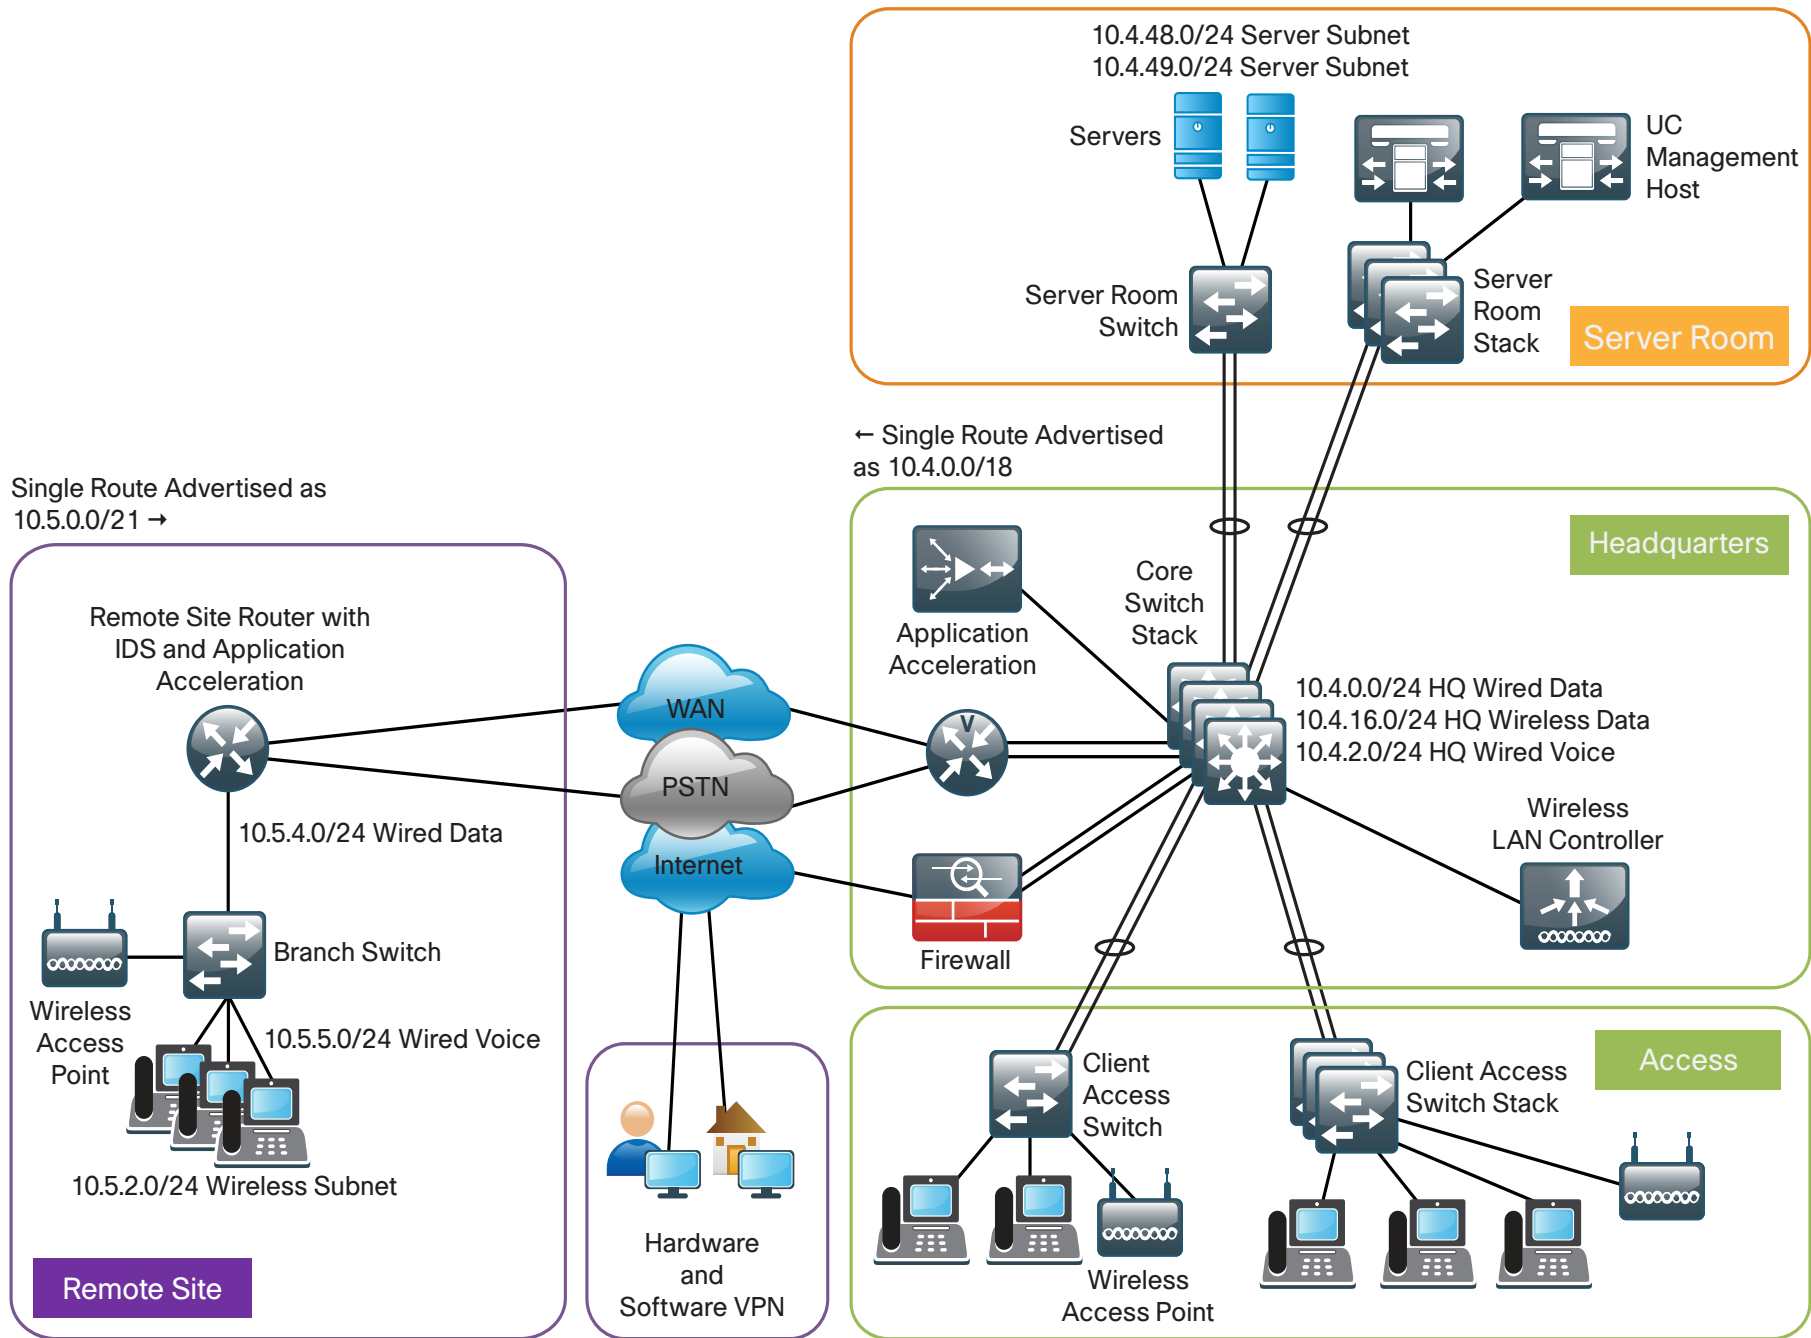
In Figure 8 following the table, you can find a diagram of Cisco SBA with sample IP subnet assignment and summarization. Summarization can be extrapolated to each remote site and across the LAN as required. To add another remote site, simply assign another block of eight /24 subnets. For example, the second remote site would be assigned the address block 10.5.8.0/21.

Table 2 - Cisco SBA assigned IP addresses

Location	VLAN	Subnet	Department	Summary address
Headquarters				10.4.0.0/18
	100	10.4.0.0/24	HQ Wired Data	
	102	10.4.2.0/24	HQ Wired Voice	
	104	10.4.4.0/24	HQ Wired Data	
	106	10.4.6.0/24	HQ Wired Voice	
	115	10.4.15.0/25	HQ Management	
	116	10.4.16.0/22	HQ Wireless Data	
	120	10.4.20.0/22	HQ Wireless Voice	
	127	10.4.27.0/25	Internet Edge	
	132	10.4.32.0/24	HQ WAN	
	148-63	10.4.48.0/24	HQ Server Room	
Remote Site 1				10.5.0.0/21
	N/A	10.5.0.0/21	Infrastructure	
	64	10.5.4.0/24	Wired Data	
	65	10.5.2.0/24	Wireless Data	
	69	10.5.5.0/24	Wired Voice	
	70	10.5.3.0/24	Wireless Voice	
Guest & DMZ				192.168.64.0/19
	1164	192.168.64.0/24	Internet DMZ	
	1176	192.168.76.0/22	Guest Wireless	

Notes

Figure 8 - Cisco SBA design with assigned IP addresses and summarization



Appendix A: Subnet Design Worksheet for Cisco SBA



Location	VLAN	Subnet	Department	Summary address

Feedback

Please use the [feedback form](#) to send comments and suggestions about this guide.



SMART BUSINESS ARCHITECTURE



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

© 2013 Cisco Systems, Inc. All rights reserved.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned 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. (1110R)