


# IP Address Management (IPAM) Best Practices

By Timothy Rooney  
Product Management Director  
BT Diamond IP



<a href="#"><u>Introduction</u></a>	<b>3</b>
<a href="#"><u>IPAM Defined</u></a>	<b>3</b>
<a href="#"><u>IP Address Inventory Management</u></a>	<b>4</b>
<a href="#"><u>Address Planning</u></a>	4
<a href="#"><u>Address Allocation</u></a>	4
<a href="#"><u>Centralizing IP Inventory</u></a>	6
<a href="#"><u>Managing Address Dynamics</u></a>	6
<a href="#"><u>IPv6 Deployment</u></a>	7
<a href="#"><u>IP Inventory Assurance</u></a>	7
<a href="#"><u>IP Address Inventory Management Best Practices</u></a>	<b>7</b>
<a href="#"><u>Dynamic IP Address Services Management</u></a>	<b>8</b>
<a href="#"><u>Policy Management</u></a>	9
<a href="#"><u>Discriminatory Address Management</u></a>	9
<a href="#"><u>DHCP Server Redundancy</u></a>	9
<a href="#"><u>DHCP Appliances</u></a>	10
<a href="#"><u>IPv6 Stateless Address Auto-Configuration</u></a>	10
<a href="#"><u>Dynamic IP Address Assignment Management Best Practices</u></a>	<b>11</b>
<a href="#"><u>IP Name Services Management</u></a>	<b>12</b>
<a href="#"><u>DNS Resource Records</u></a>	12
<a href="#"><u>DNS Options</u></a>	13
<a href="#"><u>DNS Security and Availability</u></a>	13
<a href="#"><u>Internationalized Domain Names (IDN)</u></a>	14
<a href="#"><u>DNS Scalability Challenges</u></a>	14
<a href="#"><u>DNS Configuration Verification</u></a>	14
<a href="#"><u>DNS Appliances</u></a>	15
<a href="#"><u>IP Name Services Management Best Practices</u></a>	<b>15</b>
<a href="#"><u>Overall IP Address Management</u></a>	<b>16</b>
<a href="#"><u>Centralized Management</u></a>	16
<a href="#"><u>DHCP/DNS/IPAM (DDI) Automation</u></a>	16
<a href="#"><u>Administrator Access Controls</u></a>	17
<a href="#"><u>High Availability Services</u></a>	17
<a href="#"><u>DHCP/DNS Services Monitoring</u></a>	17
<a href="#"><u>Upgrades and Patch Management</u></a>	17
<a href="#"><u>Adaptation to Your Business</u></a>	17
<a href="#"><u>Integrate IPAM Processes into Broader Enterprise Workflows</u></a>	18
<a href="#"><u>IPAM Reporting</u></a>	18
<a href="#"><u>Overall IPAM Best Practices</u></a>	<b>18</b>
<a href="#"><u>Simplifying Best Practice Implementation with Cisco Prime Network Registrar IPAM</u></a>	<b>19</b>
<a href="#"><u>References</u></a>	<b>19</b>

---

## Introduction

The practice of IP address management (IPAM) has evolved over the last two decades from a set of simple configuration tasks to sophisticated network management-like functions. This evolution testifies to the increasing reliance on IP networks to mobilize today's enterprise applications, not to mention support the very foundation of revenue growth for service providers. This reliance has in turn advanced the requirement for added discipline and rigor in managing IP address space and associated critical IP network services. Meanwhile, underlying IP, DNS, and DHCP\* technologies have likewise evolved rapidly over this time, expanding the scope of IPAM systems to include IPv6 support, DNS security measures, and expanded DHCP support such as advanced client-class support.

This white paper discusses the fundamentals of successful IPAM, which provides a solid foundation for further IT or operation automation in support of advanced IP services management. These fundamentals, or best practices, are derived from numerous implementations of IP management systems over the last 20 years, as well as from interactions with end users and industry analysts, by the BT Diamond IP leadership team. In addition, many members of the team have also been active working within the Internet Engineering Task Force (IETF) in further developing IP technology and have authored books on IPAM and IPv6.

**Note:** Cisco Prime Network Registrar IPAM is an OEM product supplied by BT Diamond IP.

## IPAM Defined

IPAM can be defined broadly as encompassing three major interrelated functions:

- **IP address inventory** - This includes obtaining and defining public and private IPv4 and IPv6 address space, and allocating that address space to locations, subnets, devices, address pools, and users on the network. This function serves as the foundation for the following two functions.
- **Dynamic IP address services management** - This function includes defining the parameters associated with each address pool defined within the IP address space management function. It also includes appropriately configuring Dynamic Host Configuration Protocol (DHCP) servers to supply relevant IP addresses and parameters to requesting users. Lastly, it encompasses effectively managing the capacity of address pools to ensure that dynamic IP addresses are available for those who need them and are permitted to have them.
- **IP name services management** - As devices are assigned IP addresses statically or dynamically, configuring appropriate Domain Name System (DNS) servers with address-to-name and name-to-address resource records so that end users may access hosts and/or applications by name (e.g., by URL) is critical. Managing name space and name services also requires proper design of the name space, configuration of other relevant DNS resource records, and many behavioral aspects of DNS as well.

Each of these functions is critical to the proper operation of an IP network. Each user needs an IP address to access the IP network, whether through a wired or wireless LAN interface, voice over IP (VoIP) device, video device, etc. Users also need to access resources on the network and the Internet in general to maintain a high level of productivity. Typically, these functions occur without user involvement or even awareness. In fact, one could argue that the job of an effective IP address manager is to be invisible.

---

\* Sometimes referred to as DDI (DHCP, DNS, IPAM), we prefer the original term IPAM as encompassing the three constituent DDI components which need to be managed holistically, not as three independently managed functions. DHCP and DNS configurations are based on and derived from the overall IP address plan, and hence, are managed within the scope of **IP address management** (IPAM) as we shall discuss in this paper.

---

In other words, as users attach to various network points, they are automatically configured to communicate and easily access network resources by URL or name.

Effective IP management requires proper allocation of address space so that there is:

- Adequate address capacity where it's needed, when it's needed
- Accurate configuration of DHCP servers for dynamic address users, including differentiation of employees versus "guests" and company issued devices versus user brought
- Accurate configuration of DNS servers so resources can be accessed easily

When these behind-the-scenes tasks are flawlessly executed, network users don't need to contact the help desk with complaints about the network; the IP address manager is invisible. In addition to flawlessly configuring and managing each of these three foundational elements of IP address management, the IP address manager must also manage these three areas collectively, and integrate these management functions into the broader IT network management environment.

## IP Address Inventory Management

The first IPAM function focuses on allocating and tracking IP address space. This IP address inventory function requires several management tasks in its own right. This IPAM function lays the foundation for the other IP management functions and impacts other critical IP network functions, not the least of which is IP routing. Most enterprise organizations will obtain public IPv4 and IPv6 address space from an Internet service provider (ISP), though some that have been using the Internet for some time have a legacy relationship with their Regional Internet Registry, e.g., ARIN and RIPE. a block of public IPv4 or IPv6 address space has been obtained, it can then be allocated to locations across the network. Private IPv4 address space (RFC 1918) can be allocated in a similar manner as can "private," unique, local address (ULA) space for IPv6.

## Address Planning

When planning to allocate IP address space, whether private or public, IPv4 or IPv6 administrators must forecast the IP address capacity requirements in each end-user accessible subnet on the network. This is typically based on the number of end users located at each site, the number of visitors or mobile users expected at the site, and the number of IP addresses required on average for each end user. Another aspect of address planning is rollout of multiple IP applications requiring address segmentation for routing treatment purposes, such as VoIP. For example, routers may need to be configured to provide priority processing on VoIP packets (packets with a source address from the VoIP address block segment) compared to best-effort data packets (packets with the source address from the data block segment).

While the easy answer to the "How much space to allocate?" question is to grossly oversize each subnet for each application, in reality this isn't always feasible, at least given IPv4 address space constraints. Within these address space sizing constraints, administrators must meet the challenge of accurately and optimally allocating address space to each site. Allocation of seemingly infinite IPv6 address space may not require equally strict rigor, but a disciplined allocation and tracking strategy is nonetheless warranted as address allocations can impact the relative ease of ongoing network management tasks that are based on IP addresses. For example, a requirement to be able to quarantine a site for security purposes should influence the address plan in a way that facilitates implementation of such a requirement, e.g., by configuring an ACL on a router or firewall.

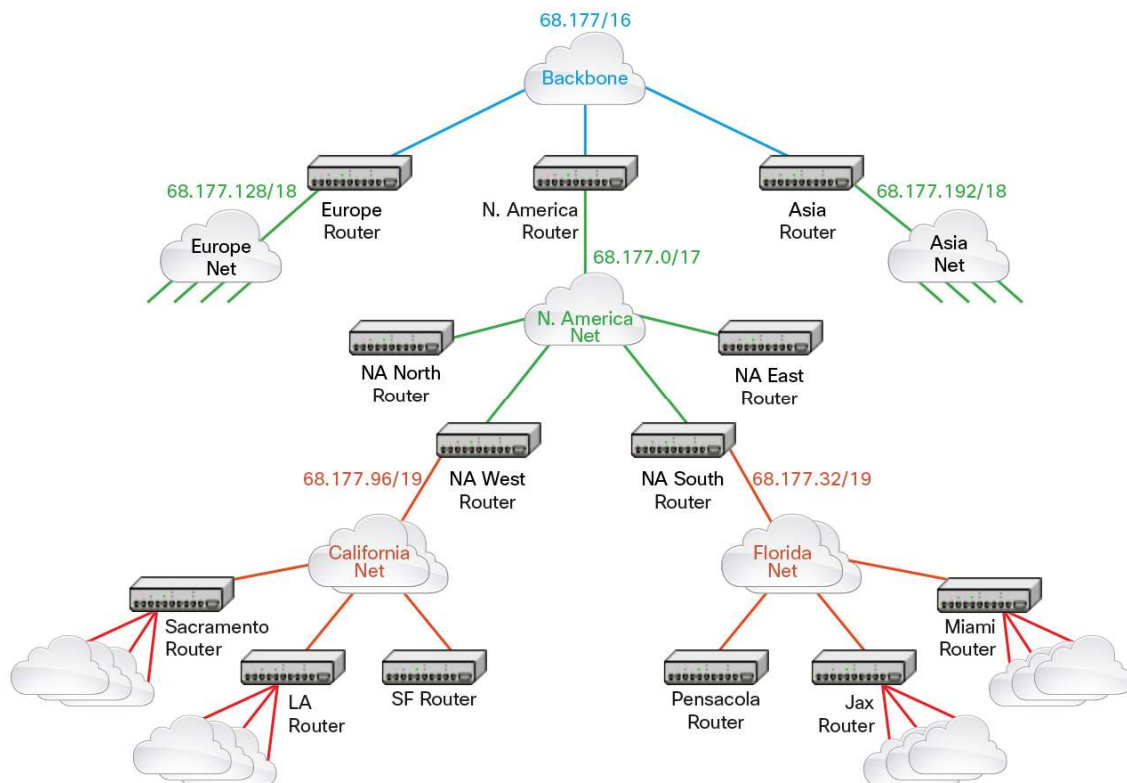
## Address Allocation

An additional constraint is that the allocated address block be appropriate to the routing infrastructure supporting each site. Block allocations at each site must roll up, in terms of taking full advantage of address hierarchy in order to facilitate route aggregation for routing protocols such as Open Shortest Path First (OSPF). Increasing route aggregation helps to reduce routing protocol traffic and keep routing tables manageable. In addition, it helps to reduce the probability of rendering certain networks unreachable. This can occur when an address block from one region is assigned to another region, but the block is included in a higher layer route advertisement, rendering the assigned block unreachable outside the advertising region.

The address space planning process then needs to carefully consider the macro-level requirements for address space as well as the rollup of individual address space requirements. For example, a global corporation may wish to subdivide its space among a core backbone of sites covering three continents (Figure 1). It may make sense to subdivide the root address block into three, in a manner that meets the current and foreseeable capacity needs of each continent.

To size each block properly, planners must define the individual site requirements - perhaps roll these up to regional levels for a middle tier within the routing topology, and then roll up to the tri-continental core routers. Modeling address space in such a hierarchical inheritance-based manner, then allocating space optimally at each hierarchy layer, is critical to fully using address utilization in a routing-efficient manner.

**Figure 1.** Hierarchical Network Allocation





---

If IP network allocation is done improperly:

- Duplicate IP addresses can be assigned
- Networks can be rendered unreachable due to the route summarization example described in the previous paragraph, or
- IP address space itself can be rendered unusable if address allocation is not only performed hierarchically, but in an optimal manner to preserve address space for use elsewhere

Due to the nature of binary arithmetic in subnetting IP networks, errors or suboptimal allocations can occur, resulting in ineffective address capacity utilization. When more address space is needed, such inefficiencies would likely need to be corrected using a tedious renumbering process before additional address space would be granted by an Internet Registry or ISP.

### **Centralizing IP Inventory**

Address planning and allocation is best performed using a centralized IP inventory database. A centralized system provides a single holistic view of the entire address space deployed over a number of sites and with address pools, and DNS information deployed on multiple DHCP and DNS servers throughout the network. Centralized management with distributed deployment also facilitates support of multiple vendor DHCP and DNS environments. For example, many organizations run Microsoft DNS and DHCP for internal clients, while running BIND DNS servers for external queries. A single consistent user interface and view of these multivendor configurations reduces errors, saves time, lowers multi-system training costs, and eliminates the requirement of replacing existing DHCP and DNS servers.

Implementing IPAM database replication or periodic backups, or running a secondary centralized database, are steps that can be taken to help ensure high availability of this critical IP address information. Another approach to IP inventory features a decentralized architecture. Decentralizing IP inventory, typically on each DNS or DHCP server, provides multiple copies of the database. However, doing so can generate tremendous replication traffic on the network in terms of updating all servers with changes. This process, with the associated impact on inter-server update performance, hampers scalability and renders this solely-distributed approach appropriate only for small and single-vendor environments.

### **Managing Address Dynamics**

After the initial IP address space sizing and deployment, even when done perfectly, changes will inevitably occur. New corporate sites are opened and others are consolidated. Perhaps more mobile users require IP addresses on a subnet than initially expected. Several servers are moved to a different subnet without prior notification. New services such as VoIP are rolled out.

Note that these events all impact the IP address space, regardless of whether they were initiated:

- By business requirements impacting site openings and closures
- By IT in deploying additional IP services such as VoIP and adding more servers or devices for performance or other reasons, or
- By end-user behavior in terms of addressing requirements at particular sites

Staying on top of these and other changes, which reflect the organic nature of IP networks, is absolutely necessary for effective IP address space management.

## IPv6 Deployment

Disciplined address space management becomes even more critical when deploying IPv6. With IPv4 address space exhaustion looming, new and growing organizations and end users in the growing mobile segment will soon have access only to IPv6 address space. Consequently, even if you have substantial IPv4 address space, you must support IPv6 at least externally (Internet-facing) in order to communicate with (e.g., sell to) IPv6-only organizations.

Managing IPv6 space requires understanding of the IPv6 addressing structure. When implementing IPv6 within an extant IPv4 network, co-existence technologies should be considered. The sheer size and hexadecimal representation of IPv6 addresses invites operator errors, stifling effective IP address management. Because few, if any, customers will actually deploy IPv6 in a “greenfield” environment and the transition will likely take years to complete, integration of IPv4 and IPv6 address allocation processes is crucial.

## IP Inventory Assurance

The IP address inventory serves as the foundation on which IP planning decisions are made. For instance, the assignment of a free subnet according to inventory is usually made within a new branch office. A new IP address pool is defined on an existing subnet in order to add DHCP pool capacity based on available addresses according to the inventory. Unique resource record information is defined and configured in DNS based on supplementing the inventory. An accurate IP inventory is critical to facilitate these planning and implementation decisions. The goal of IP inventory assurance is just that: to assure the accuracy of IP inventory through periodic discovery, exception reporting, and selective database updates. These discovery litmus tests confirm the integrity of the IP inventory for effective IP planning.

## IP Address Inventory Management Best Practices

Table 1 highlights best practices for IP address inventory management.

**Table 1.** Best Practices for IP Address Inventory Management

Best Practice	
<b>Inventory address space in a centralized database</b>	IPv4 and IPv6 address space, both public and private, is a precious commodity - one that provides the fundamental entity for network communications. Therefore, it must be tracked in a centralized repository to maintain consistency and accuracy. Of course, accuracy requires updates to the database upon address space allocations, “free-ups” and, ideally, ongoing inventory assurance techniques.
<b>Rigorously record IP subnet allocations and IP address assignments</b>	Instituting an IP address change control process, or incorporating IP addressing changes into an existing change control process, can reduce the risk of duplicate address assignments. Allocating, assigning, or freeing up of IP addresses affects the IP network, so these functions should be performed judiciously.
<b>Perform and track address space allocations in accordance with routing topology to model and optimize route aggregation and in consideration of management and security policies</b>	<p>Network allocations should be made in an optimal manner, increasing utilization of address space for IPv4, though this is of lesser concern with IPv6. These allocations should map to your network topology model and consider IP-based management and security policies. Allocating address space along a hierarchical structure that models the routing topology facilitates route aggregation to keep routing overhead to a minimum.</p> <p>If exceptions to the aggregation model are necessary, for whatever reason, they can be made knowingly and routes can be proactively updated to maintain reachability. Since routing topology often maps to an organization's locations, sites, or business unit hierarchy, this hierarchical modeling of address space typically provides the added benefit of tracking address allocations to these entities.</p> <p>Per application address allocation should also be addressed if appropriate to manage address allocations for deployed IP services, e.g., VoIP compared to data. Management and security policies should likewise be considered should a need arise to throttle an application or cordon off a network segment based on IP address policy settings.</p>

Best Practice	
<b>Implement common allocation policies for, and within, address blocks to promote consistent subnet addressing</b>	<p>Many organizations allocate or reserve specific portions of each subnet for ranges of static device addresses and dynamic address ranges. For example, you may reserve addresses .1 and .2 for router addresses on a subnet (or the first and second addresses in general), .3 and .4 for time servers, .15 through .80 for a DHCP pool, etc.</p> <p>Provision of a common allocation template promotes consistency in allocation and deployment, and also makes for easier troubleshooting as needed with consistently allocated subnets. In fact, application of a site template further promotes consistency of multi-block allocations. For example if every branch office or retail store requires a VoIP subnet, data subnet, wireless subnet, and management subnet, a site template can be defined once, then applied for each planned branch or retail office.</p>
<b>Maintain additional information as appropriate per IP device</b>	<p>Keeping track of what device is occupying each IP address in a subnet is critical to IP management. However, many such devices have other attributes that should be tracked within an IP management solution. Not the least of these attributes is what other IP addresses the device in question occupies. Many devices have multiple IP addresses, whether for virtual networking, IPv4 and IPv6 addressing, or multi-homing. Multi-homed devices have multiple interface cards, each occupying one or more IP addresses.</p> <p>Beyond this critical IP address information, tracking other attributes, such as device type, location, switch port, administrative contact, asset information, associated resource records, and others, is equally crucial. We recommend that these attributes be identified by device type and, ideally, by location to maintain relevancy for the IP administrator managing the device.</p>
<b>Monitor address utilization to manage the capacity of the IP address space</b>	<p>Although initial addressing needs may be impeccably forecasted, changes happen in IP networks due to business, IT, or other reasons. Despite the best planning efforts, IP networks seem to have an organic nature, where address needs rise and fall at different times at various locations within the network.</p> <p>Address utilization statistics across subnets and DHCP pools should be collected to provide snapshot and historical tracking of address use. This information can also be trended using linear regression models to predict potential future address depletion times. This trending analysis provides another decision criterion in the IP address capacity management process.</p> <p>Proactive IP address capacity management requires the use of alerts for notification of pending address depletions before they happen. Alerts should be programmed for address pools (i.e., dynamics only) or networks (i.e., dynamics and statics) approaching full capacity or being underutilized. This proactive measure can offset unnecessary, end-user communication problems caused by address depletion.</p>
<b>Plan for IPv6 deployment</b>	<p>With the exhaustion of IPv4 address space already hitting the Asia and Europe regions, and looming for the rest of the world in the next couple of years, the time is ripe for planning for and deploying IPv6. With only IPv6 available to new and growing organizations, particularly in Asia and Europe, the Internet population of IPv6 users will shoulder the mass of continuing IP address growth anticipated on the Internet over the coming years. To communicate with this growing population of IPv6-only users, you must deploy IPv6 (dual stack) at least on your Internet-facing servers for email or web access. The future of the Internet is IPv6; begin planning for it now.</p>
<b>Implement IP inventory assurance techniques</b>	<p>Maintaining accuracy of your IP address inventory database is crucial. If the IP address inventory only tracks top-down allocations of address space entered by administrators manually, it is impossible to verify its accuracy. Comparing the inventory database with network actuals is imperative for identifying discrepancies and tracking IP management processes. For example, if someone circumvented the conventional update process without updating the inventory database, the identification of this discrepancy would not only highlight an inventory mismatch, but also bring out this network change control issue.</p> <p>Whether you employ a top-down or bottom-up approach to allocating subnets to router interfaces, address pools to DHCP servers or individual IP addresses to devices, updating the inventory must be a key step in the process. Periodically reconciling the network actuals with the database plan using automated discoveries is an effective way to monitor the process and keep inventory accurate.</p>

## Dynamic IP Address Services Management

Adhering to the IP inventory best practices described in Table 1 can help maintain adequately sized networks and address pools across the network. But sizing address pools to supply IP addresses to end users is just the beginning of the process of address assignment using DHCP. You don't want just anyone to get any IP address on your network to access network resources.

Even for legitimate employee access, bringing their (your) own devices to work with the bring-your-own-device (BYOD) trend, you need to consider your addressing policy for such user devices. There is more to configuring DHCP servers than address pool allocations. Additional configuration elements include valid options and policies with associated values for each address pool, valid or invalid devices by MAC address, client class or user authentication, device software validation, and DHCP failover configuration.



---

## Policy Management

Because many or all of your DHCP servers will require the same or similar DHCP policies, we recommend that you centralize the configuration of these servers to create a single or set number of policies, and then deploy the policies across your servers. This practice ensures a consistent and accurate approach to setting these critical policies and can also save substantial effort in defining and distributing such policies. Otherwise, you must be concerned with entering basically the same information multiple times into each of your servers. A similar approach should be taken in defining a set of DHCP options, with defined DHCP options and valid values for use on assignment.

## Discriminatory Address Management

In terms of discriminating address assignment, there are several levels of policies or controls most DHCP solutions provide. The first is to simply filter by the MAC address of the client requesting an address. If the DHCP server has a list of acceptable (and unacceptable) MAC addresses, it can be configured to provide a certain IP address and associated parameters to those clients with acceptable MAC addresses, and either no IP address or a limited access IP address to those without acceptable MAC addresses. By **limited access IP address**, we mean that the network routing infrastructure is pre-configured to route IP packets with such addresses to only certain networks, such as to the Internet only, or even nowhere.

This type of IP address and configuration assignment is also possible by filtering on the client class of the client requesting an IP address. Certain clients, such as VoIP phones, provide additional information about themselves when requesting an IP address in the vendor class field of the DHCP packet. The user class field may also be used. The DHCP server can be configured to recognize the user classes and vendor classes of devices on your network to provide additional information to the DHCP server when assigning the IP address and configuration parameters. Addresses can be assigned from a certain pool and additional configuration parameters can be assigned to the client through standard or vendor-specific DHCP options.

A third level of discriminating IP address assignment is possible by authenticating the user of the machine requesting an IP address. This function can be used in conjunction with a MAC address and client class discrimination described above. For example, if a client with an unacceptable MAC address attempts to obtain an IP address, one option is to completely deny an address; another option is to require the user of the client to login through a secure access web page. This enables easier capture of new MAC addresses for legitimate users of your network, e.g., for users bringing their own devices. Solutions ranging from simple perl scripts to sophisticated integrated software solutions are available to direct such users to a login or password requesting webpage. A simple lookup against a database of legitimate users then allows access or denial of the client to a production IP address.

Beyond these device identification measures based on MAC addresses, client classes and user authentication, DHCP can also provide additional validation on the machine requesting the IP address. The DHCP process can be used to invoke an external security scanning system to scan the requesting client for viruses or to validate use of acceptable virus protection software. This device scanning step can be used alone or in conjunction with the device identification measures to provide a robust access security solution using DHCP.

## DHCP Server Redundancy

DHCP failover is one approach to providing IP address services redundancy across your IP network. If a DHCP server crashes, a failover DHCP server can take over and begin processing DHCP transactions. This provides a higher availability service for your end-user clients requesting IP addresses. If clients cannot get IP addresses, they will be unproductive and will call the help desk.

---

DHCP vendors implement high-availability DHCP in various ways. For example, the Internet Systems Consortium (ISC) DHCP server supports an inter-server failover protocol to communicate lease binding updates and to check heartbeat status; however, the popular Microsoft DHCP server does not. Instead, Microsoft recommends splitting each scope or pool across two DHCP servers, with one server supporting 80 percent of the addresses in the pool, and the other server supporting the remaining non-overlapping 20 percent.

Variations on this split scope approach from a 50-50 split to an even 100-100 split, doubling address space requirements, and are likewise valid configurations for increasing DHCP server availability. For DHCPv6 configurations, using 100-100 split scopes is the simplest approach given the vast address space within even a single/64 IPv6 subnet. A failover protocol standard is in development for DHCPv6 but as of this writing, has not yet been published.

### **DHCP Appliances**

A DHCP appliance implementation should be considered to simplify the procurement, deployment, security, monitoring, and upgrading of DHCP services throughout the IP network. Most appliances are self-contained units built on hardened Linux operating systems and many provide additional security features. This can simplify the procurement process by eliminating the need to coordinate the operating system patch level with the DHCP service version compatibility for initial deployment. Ongoing patch management can be simplified as well, with appliances offering centralized patch management features. Many appliances can also be deployed in dual-appliance configurations to provide high availability at the hardware level.

### **IPv6 Stateless Address Auto-Configuration**

An innovative feature of IPv6 facilitates plug and play through stateless address auto-configuration (SLAAC). As an IPv6-enabled device initializes, it listens for (or solicits) messages from the serving router, called router advertisements (RAs). These RAs indicate the router's address (default gateway), the IPv6 network or subnet address (also called the prefix), and the availability of DHCP for IPv6 (DHCPv6) services. DHCPv6 works similarly to DHCP for IPv4 in automating IPv6 address assignment. If DHCPv6 is not available for address assignment, the device may auto-configure its IPv6 address by appending its self-derived interface identifier to the advertised network prefix from the RA.

Administrators can configure their routers to send RAs on IPv6-configured interfaces to control the means for IPv6 address assignment. As you can imagine, enabling auto-configuration obviates the need to configure a DHCPv6 server\* but allows any IPv6 device to auto-configure its address on your network. Requiring use of DHCPv6 instead necessitates more work but affords more control over who may obtain an IPv6 address and what IPv6 address is assigned. While enterprise networks may seek to limit the availability of auto-configuration, public networks or special purpose networks for public safety, sensor, or military applications may take full advantage of this capability.

---

\* Other networking parameters traditionally provided as DHCP options (e.g., time server addresses) may still necessitate DHCPv6 server configuration.


## Dynamic IP Address Assignment Management Best Practices

Table 2 outlines best practices for IP address assignment management.

**Table 2.** Best Practices for IP Address Assignment Management

Best Practice	
<b>Centralize DHCP server configuration to improve configuration accuracy and consistency</b>	Utilizing a single interface and database to configure a number of DHCP servers provides the ability to enter configuration parameters once, and deploy the master configuration to multiple DHCP servers. This promotes consistency of configuration and simpler address pool allocation and reallocation as necessary for ongoing address pool capacity management, while still allowing for per-server configuration. If the IP network features a variety of DHCP vendors' servers, a centralized configuration tool that supports multiple vendors is recommended.
<b>Implement basic security measures to provide selective address assignment</b>	Implement one or more of the following approaches: <ul style="list-style-type: none"> <li>• Device identification using a MAC address - filter client requests against a list of acceptable and unacceptable MAC addresses</li> <li>• Device identification using a client class - provide additional configuration information for known client classes configured on your network</li> <li>• User identification using authentication - support user login and password authentication against a database or other authentication scheme</li> <li>• Invoke device security scanning or software validation - scan the requesting device for viruses and valid software prior to granting a production IP address</li> </ul>
<b>Adopt and use established DHCP options and policy sets across your DHCP servers</b>	This allows implementation of a consistent set of policies across a variety of DHCP servers, each with its own address pools. This approach allows mobile clients to obtain a consistent IP configuration, no matter where they connect into the network.
<b>Configure DHCP redundancy for high-availability address assignment services</b>	IP address assignment is the first step to communicating on an IP network. Make sure this service is available to your clients in a high-availability configuration. This can be accomplished in at least three ways: <p><b>1) Failover Scheme</b> - The first mechanism is the traditional failover scheme where a common address pool is shared among two DHCP servers. One DHCP server is the primary server and processes DHCP address requests; the other server is a failover server, or hot standby server, keeping in synch with the primary's DHCP lease bindings and heartbeat messages. Should the primary server fail, the failover server can assume control and begin handling DHCP address requests.</p> <p><b>2) Double Scope Approach</b> - The second mechanism that can be employed when address space is not overly constrained, e.g., 10.0.0.0 space for some users, is to deploy two address pools of the same size, but of different addresses. This double-scope approach uses two address pools that can serve the same set of clients independently. It alleviates the need for inter-server heartbeat communications, while providing sufficient address capacity for the end users requiring addresses.</p> <p>With its plentiful address space per subnet, DHCPv6 can be configured for double scopes, with two non-overlapping scopes configured on two different DHCPv6 servers.</p> <p><b>3) Split Scope Approach</b> - A third mechanism, particularly when using Microsoft DHCP, is to implement split scopes, where two DHCP servers manage non-overlapping subsets of each address pool.</p>
<b>Track dynamic address assignments and monitor utilization of address pools, including shared subnets, to proactively manage address pool utilization</b>	As with the address inventory capacity management best practices, this corollary best practice recommends use of DHCP monitoring of address assignments and address pool utilization, including shared pools or shared subnets, to net out the capacity impacts from a pool and pool user perspective.
<b>Maintain IP address pool history data to monitor address usage trends and proactively align address space to where it is needed</b>	While alerting and thresholds provide an effective notification of an impending address depletion based on recent actual utilization data, having the ability to track utilization "snapshots" over time is an effective way to identify address utilization trends. Accessing address pool historical data in a graphical form (see Figure 2) helps convey at a glance the general utilization trends and enables proactive management of address pools to prevent address depletions.
<b>Consider DHCP appliances to simplify deployment, improve security, and simplify upgrades</b>	A DHCP appliance integrates the hardware, operating system, and DHCP service into a simple self-contained platform. Appliances are generally more secure purpose-built platforms with restricted operating system permissions, users, ports, and files. Most appliances are also deployable in a dual high-availability configuration, and should be capable of being monitored, controlled, and patched from a centralized system to simplify ongoing server and services management. The centralized monitoring and patching functions provide lights-out support for remote appliance deployments.
<b>Define IPv6 auto-configuration and DHCPv6 policies</b>	Determine if enabling SLAAC is desirable on your network and configure routers and DHCPv6 servers accordingly to support full IPv6 address assignment, additional parameters only, and prefix delegation.

**Figure 2.** Graphical Address Pool Capacity History and Trending



## IP Name Services Management

After users on your network obtain an IP address and related IP configuration through DHCP, they may then proceed to access their email and the web or intranet. The ability to send email to someone's address at a destination host and browse the web using uniform resource locator (URL) makes email and web browsing easy and user-friendly. Your computer communicates with the email server and web server through IP packets using IP addresses, not names or URLs. Fortunately DNS was invented to allow users to type text-based addresses while providing a mechanism to translate these text-based addresses into IP addresses that computers can communicate. It's not a stretch to say that without DNS, these applications could function but would be totally unusable for 99 percent of your company's end-user population. Needless to say, DNS services must be configured accurately and be highly available to users.

## DNS Resource Records

It's up to IP address managers to properly configure the DNS servers in the network with the information needed to resolve host names and URLs into IP addresses. This means that not only statically-configured IP devices like routers, web servers, email servers, and the like need to have entries in DNS, but also dynamically configured IP devices like printers and even end-user devices. In many cases, websites perform a **reverse** DNS lookup for an IP address before continuing a secure web session to validate that the requesting IP address has some form of

---

legitimacy in DNS. This implies integration between DHCP and DNS, referred to as Dynamic DNS, which is an automated process to update DNS upon address assignment by a DHCP server.

Beyond name-to-address translation and vice versa, DNS provides many other translation applications, which won't be covered in this paper. Each translation type maps to one or more resource record types in DNS. For example, an "A" resource record type is used to translate a text-based host name into an IPv4 address. While all resource record types follow the same basic format in terms of fields within the record, the syntax is not intuitive, nor is it easy to identify errors until problems arise. While DNS does provide a mechanism for a master DNS server to update its slaves through a zone transfer, in some cases it is desirable to operate in a multi-master mode of operation, whereby each master must be updated individually. This opens the door to potential errors in not only resource record configuration, but also in other DNS options and directives, of which there are many.

### **DNS Options**

Configuring these DNS options is critical to properly defining the behavior of the DNS server in terms of zone transfers, security measures, and other operational parameters. Various directives exist in varying forms in different DNS server versions. For example, server configuration parameter definitions vary among different versions of BIND. Other vendors' DNS implementations may have other nuances in configuration. Keeping track of the proper syntax for the particular vendor and version you're running may be tedious, but it's absolutely critical to keeping DNS up and running.

### **DNS Security and Availability**

Recent highly publicized DNS cache poisoning vulnerabilities have accelerated momentum for deployment of DNS security extensions (DNSSEC). DNSSEC is the only definitive mitigation for cache poisoning style attacks, but it does not address other DNS vulnerabilities so other security measures must be implemented as well. Following are some recommended approaches:

- Configure ACLs - configure which IP addresses or networks can query, notify, update, and transfer to or from each name server. In addition, ACLs on the `ndc/rndc` control channel should be defined along with a security key (see next item) for BIND DNS servers.
- Configure transaction signature keys - sign each update and zone transfer with the use of transaction signature keys (TSIG keys). For deployments to Microsoft Active Directory integrated zones requiring secure updates, sign each update using GSS-TSIG.
- Run the DNS service (`named`) in `chroot-ed` environment - this provides the name server daemon full file system subtree access at a point below the root; otherwise, root access to the file system is provided by default. Ideally, run the DNS server on a dedicated machine to allow restriction of open TCP/UDP ports.
- "Hide" master DNS servers - if attackers find and infiltrate the master DNS server, slaves will zone transfer from this master, spreading the corruption. Hiding the master can be accomplished by editing the standard NS and A (glue) records pointing to the master DNS server to point to a different (slave) server. The master name ("`mname`") field of the slave's SOA record should also be edited accordingly.
- Consider DNSSEC implementation, particularly for external (Internet-facing) zone information; determine if parent zone(s) are signed and coordinate with parent zone administrators for key rollover and delegation signer policies.
- Deploy DNS appliances (addressed below) - appliances are purpose-built platforms for running DNS (and DHCP) and associated management services exclusively. Appliances generally offer hardened Linux-

---

based operating systems, restricted services, users and ports, jailed environments, and more, depending on the appliance vendor.

DNS is architected with high availability in mind, with the ability to configure a master DNS server and a set of slave DNS servers that receive resource record updates from the master(s) through zone transfers. Consider the following DNS server deployment strategies to increase availability and reduce security vulnerabilities:

- Deploy servers on different networks in general and external servers on different ISP links if possible to reduce denial-of-service vulnerability
- Deploy external name space on external DNS servers separated from internal DNS servers
- Segregate DNS functions, deploying authoritative servers separately from caching or recursive servers
- Use a separate network for queries instead of zone transfers and updates

### **Internationalized Domain Names (IDN)**

According to DNS specifications, DNS data is stored and communicated in the form of ASCII characters.

Unfortunately, ASCII encoding is insufficient for non-Roman alphabets, which are generally represented using unicode character standards. The International Domain Names for Applications (IDNA) specifications define a mechanism for translating unicode native characters into ASCII format. When a user types in a web URL in native characters for example, the browser first translates the web address into ASCII in accordance with the IDNA specifications, then looks up this ASCII web address in DNS. It is up to the DNS administrator to assure the proper ASCII-mapped DNS names are provisioned in DNS.

With the Internet Corporation for Assigned Names and Numbers (ICANN) recently approving the process to allow international top-level domains (likely beginning deployment in 2012, although international country code top-level domains already exist), organizations will be afforded the opportunity to register international subdomains under new international top-level domains. From a business standpoint, this can make your website and email addresses more easily accessible to end users in particular target countries with associated, native-language top-level domains.

### **DNS Scalability Challenges**

When managing a number of sets of DNS servers, grouping these sets and managing them as individual entities can simplify DNS management. For example, an external set of DNS servers consisting of a master and three slaves may be deployed to support Internet-facing name resolution; an internal set consisting of two masters and five slaves may be deployed to support internal queries, and so on. Managing these sets in terms of supported domains and option settings can simplify overall configuration and reduce entry errors of similar data on multiple servers.

Another scalability challenge relates to supporting a common set of resource records across multiple domains. For example, the “www A record” for example.com may be the same as in example.net, example.org, and example.edu. Use of a template domain to define and manage these resource records while supporting multiple alias domains that use this information can help reduce duplicate entry errors.

### **DNS Configuration Verification**

Certain DNS server products, including BIND, can happily accept erroneously formatted configuration information, yet fail to load and initialize the service or zone file. Deployment of one incorrectly formatted entry could result in the DNS server failing to run and resolve queries. Obviously, this could be a major issue. Having the ability to



validate the configuration information prior to deployment is recommended to reduce the likelihood of the server failing to load the configuration.

## DNS Appliances

A DNS appliance implementation should be considered to simplify the procurement, deployment, security, monitoring, and upgrading of DNS services throughout the IP network. Most appliances are self-contained units built on hardened Linux operating systems and provide additional security features. This can simplify the procurement process by eliminating the need to coordinate the operating-system patch level with the DNS service version compatibility for initial deployment. Ongoing patch management can be simplified as well with appliances offering centralized patch management features. Many appliances can also be deployed in dual-appliance configurations to provide high availability at the hardware level.

## IP Name Services Management Best Practices

Table 3 summarizes best practices for IP name services management.

**Table 3.** Best Practices for IP Name Services Management

Best Practice	
<b>Centralize the DNS server configuration to improve configuration accuracy and consistency</b>	Utilizing a single interface and database to configure a number of DNS servers provides the ability to enter configuration parameters once and deploy the appropriate master or slave configuration to multiple DNS servers, then aggregate dynamic updates to keep the centralized inventory up to date. This provides a centralized view into the overall DNS configuration across your network for DNS servers, domains, zones, and views.
<b>Run multiple DNS servers on different subnets for each zone to increase availability of critical DNS services to end users</b>	Deploy DNS servers to eliminate common points of failure and increase reachability from internal resolvers and to the Internet. Trade off the simplicity of running a single master DNS server for each zone instead of the more complex deployment of multi-master DNS. Single master zones ease configuration by requiring updates to one master server; however, take care to reduce exposure to unauthorized updates to this master. Multi-master configurations have less vulnerability but require careful management of the dynamic update process to reduce cyclic updates.
<b>Periodically validate DNS configuration files to check for syntax errors, lame delegations, and other errors that can reduce the accuracy and effectiveness of the DNS infrastructure</b>	This configuration verification should be done prior to reloading a zone or entire server configuration, as well as on a periodic basis for audit and validation purposes. A backup copy of at least the most recent working version of each server's configuration files should be maintained to allow rollback should a corrupted or misconfigured file be deployed. Utilizing a DNS configuration or IPAM tool can help reduce entry errors with data validation. And for BIND DNS implementations, BIND supplies a pair of verification utilities: named-checkconf and named-checkzone.
<b>Configure external, internal, and other views of your name space</b>	This can be accomplished by configuring different versions of given zones on separate name servers (e.g., an external set of DNS servers and a separate internal set of DNS servers). Alternatively you can configure multiple zone versions on a single set of DNS servers utilizing the "views" feature of BIND 9. This provides an external version of externally exposed domains to keep resolvable hostnames to a manageable number, large or small. Meanwhile, a different version of such domains can be provided to those querying DNS from internal networks. This simple dual-view example can be extrapolated to multiple views, allowing granular configuration of which host names get resolved with what, if any, IP address(es).
<b>Tighten security by: 1) Configuring access control lists (ACLs), transaction signatures for dynamic updates, zone transfers, and control messages and 2) Specifying particular TCP/UDP ports for queries, updates, and zone transfers</b>	BIND offers a variety of configurable options that allow specification of ACLs, pair-wise server transaction signatures, and IP address or port specifications. While these options provide the flexibility for configuring these capabilities, the key is to accurately configure each server with its corresponding ACLs, keys, and IP addresses or port numbers. For a large number of servers, this can be cumbersome and error-prone to configure manually. Microsoft provides a different means of signing updates with GSS-TSIG.
<b>For large environments, consolidate management of multiple server sets and alias domains</b>	Implementing this best practice generally requires a DNS configuration tool to support grouping of sets of DNS server and alias domains. However, use of such a tool - or better yet, an integrated IP address management solution - can provide the benefits of managing multiple DNS server sets and alias domains, as well as multiple sets of DHCP servers and a diversified IP address space.
<b>Deploy DNSSEC to secure resolution of your namespace</b>	Deploy DNSSEC, particularly for external zones, considering your zone's place in the overall DNS chain of trust. Configure internal caching servers which query Internet DNS servers for DNSSEC validation as appropriate.
<b>Consider IDNA implementation to enable native language DNS resolution</b>	Configuring IDNA-compatible ASCII DNS domains and hostnames requires mapping of desired international, character-based domain names in accordance with IDNA standards.

Best Practice	
<b>Deploy DNS appliances to simplify deployment, improve security, and simplify upgrades</b>	A DNS appliance integrates the hardware, operating system and DNS service into a simple self-contained platform. Appliances are generally more secure, being built with restricted operating-system permissions, users, ports, and files. Most appliances are also deployable in a dual high-availability configuration, and should be capable of being monitored, controlled, and patched from a centralized system to simplify ongoing server and services management. The centralized monitoring and patching functions provide lights-out support for remote appliance deployments.
<b>In high-performance environments, configure caching-only DNS servers to handle large volumes of DNS queries</b>	Caching-only servers are simply name servers not configured as authoritative for any zones. All queries to caching-only servers result in a lookup in cache with escalation to the DNS root servers as necessary. Over time, these servers build up a substantial cache and can respond directly from cache for those records with "still-alive" times to live (TTLs).

## Overall IP Address Management

Bringing together IPAM functions into a centralized management platform provides a number of advantages for IP managers. Clearly, the relationship among IP inventory, DHCP configuration, and DNS configuration is interconnected. Automating functions among these three key areas and minimizing duplicate entry of related information can reduce errors and save time. Going beyond this, however, can provide additional benefits in terms of extending automation to other related IT systems or functions, reporting on IPAM related information, and generally managing IP inventory, DHCP, and DNS as the critical set of services they represent on your IP network.

### Centralized Management

As discussed in the IP address inventory management section, there are benefits of centralizing IP inventory to enforce change control, help enable delegation, and support accurate inventory tracking. Given the closely related functions of DHCP and DNS configuration, it also makes sense to centralize and integrate DHCP and DNS configurations, taking full advantage of the IP inventory information. A centralized management architecture is also consistent with international standards\* including ITU's TMN and OGC's ITIL.

Centralized management helps enable entry of information once, eliminating the painstaking and error-prone process of entering similar information into multiple systems. For example, for those employing spreadsheets as the centralized inventory, the process of allocating a subnet typically requires calculation and assignment in the spreadsheet, entry of any dynamic addresses within the subnet into a DHCP server's configuration file, and entry of associated resource records for static and even dynamic addresses, if desired, in DNS. Clearly, the entry of information in these three systems is closely related and must be accurate to ensure consistent address assignment and name resolution. Use of a centralized IPAM system can eliminate this duplicate entry, reducing errors and saving time and money, especially in environments with multiple servers and with mixed Microsoft, ISC, and BIND server deployments.

### DHCP/DNS/IPAM (DDI) Automation

A key value proposition for investing in an IPAM solution is to automate otherwise manual processes involved with updating spreadsheets for subnet and address assignments. Then you can manually configure DHCP and DNS servers with corresponding address pools, reverse zones, and resource records respectively.

\* The International Telecommunications Union (ITU)'s Telecommunications Management Network (TMN) M.3010 and M.3400 series standards define a hierarchical pyramid architecture supporting the traditional FCAPS functions, while the U.K.'s Office of Government Commerce (OGC) drove the development of the Information Technology Infrastructure Library (ITIL), which has been depicted graphically in concentric circular form.

---

## **Administrator Access Controls**

For most organizations, responsibility for various aspects of the IPAM functions falls upon more than one person or even one group. In most cases, it's desirable to delegate administration of DHCP or DNS services or overall IPAM functions by geography or business unit, which provides distributed control while controlling the scope of access to particular geographies, domains, or even system functions. By implementing administrator controls, certain functions or areas of network topology can be partitioned to specific administrators, while super-user functions can be reserved for the core IP management team.

## **High Availability Services**

Clearly, DHCP and DNS services are critical to any IP network. One recommendation is DHCP redundancy and deployment of multiple DNS servers to provide high availability. Deployment of appliances can provide an added layer of high availability at the hardware level. Most appliances are available in back-to-back mirrored connections for colocated hardware redundancy. This dual configuration can be deployed in addition to DHCP failover and multiple DNS masters and slaves to provide both hardware-level and site-diverse high-availability services. It may also make sense in your environment to deploy a high-availability IPAM system on top of the DHCP/DNS services, though the IPAM system generally should not be in the critical path to serving up DHCP leases and resolving DNS hostnames. If it is in the critical path, then it must be deployed in a high-availability configuration.

## **DHCP/DNS Services Monitoring**

Accurate and timely deployment of DHCP and DNS configurations is certainly a critical aspect of effectively managing the DHCP and DNS environment. However, it is equally important to monitor these services to ensure they are properly functioning. If end users cannot obtain IP addresses or host names due to a server outage, their productivity and satisfaction will diminish, and they will likely call the help desk. Certainly, deployment of high-availability configurations is recommended per the prior section. But when a failure occurs and the backup system takes over, it is important to identify and rectify the failed service quickly to reduce vulnerability of service outage should the backup service subsequently fail.

## **Upgrades and Patch Management**

For environments with a number of distributed DHCP and DNS servers, application of upgrades and patches can be tedious and error-prone. Generally, each of the servers must be inventoried from a hardware, OS version, and DHCP/DNS service version perspective. Compatibility issues among these elements must be considered during the upgrade planning process. At times, physical presence at the site is required by knowledgeable resources to successfully deploy the upgrade, adding to the cost and time required to perform the upgrade. However, deployment of DHCP/DNS appliances with centralized patch management can remove many, if not all, of these headaches. Selective upgrades of OS, kernel, and DHCP and DNS versions from a centralized system can simplify the patch management and rollback process.

## **Adaptation to Your Business**

Many software tools tend to be rather rigid in terms of IP subnet, device attributes, and topology. However, every IP network is different. And methods of managing IP networks vary just as widely. Employing a system that enables entry of custom attributes for topology elements, subnets, devices, DNS domains, and even resource records supports adaptation of the IPAM software to the organization's business processes. These additional attributes should enable definition of a variety of data types, e.g., text boxes, drop-down lists, and URLs, and they should also be searchable to quickly locate elements containing these user-defined attributes.

## Integrate IPAM Processes into Broader Enterprise Workflows

In addition to adapting to business processes from a data-element, attribute perspective, integrating IPAM-related functions into broader workflows can provide further automation and cost-saving benefits. For example, the allocation of an address block to a site would likely require the associated updating of relevant DHCP and DNS server configurations, but would also require the addition of the subnet to the corresponding router interface. If the IPAM system supports the passage of subnet allocation information through an integration point such as a callout, then this information transfer could be automated to update the router directly or a configuration management system. This workflow shortens the overall implementation interval and reduces miscommunication errors as well as duplicate entry errors.

With increasing proliferation of IP services, enterprise IP managers typically need to allocate application-specific subnets or VLANs, accurately assign IP addresses and associated configuration parameters using DHCP, and manage resource records in a common or application-specific set of domains. Integrating these processes into a broader workflow for “deploy VOIP LAN”, “add support for XYZ video device,” etc. can simplify the overall processes for executing these workflows.

## IPAM Reporting

Communicating the status of address assignments in relation to your IP network is an important aspect of managing IP space, just as it is for other network management functions. Reports that convey information graphically can facilitate communication of information across the organization from top to bottom. Tabular reports are also important for managing address allocations and server configurations. These reports should be provided for address allocation and capacity hot spots (e.g., networks or servers nearing address depletion), services status, and audit information. Reports on which administrators performed certain tasks, or who owned an IP address at a given time are critical for periodic audits, for troubleshooting or investigations, and even for regulatory requirements such as Sarbanes-Oxley, HIPAA, and CALEA.

## Overall IPAM Best Practices

Table 4 outlines best practices for overall IP address management.

**Table 4.** Best Practices for Overall IP Address Management

Best Practice	
<b>Centralize management of IP inventory and DHCP and DNS services</b>	Centralizing the management of IP inventory with DHCP and DNS configuration simplifies and automates the closely related functions of IP inventory, DHCP, and DNS. A centralized umbrella function promotes consistency among these key elements and simplifies IPAM processes.
<b>Automate DDI functions</b>	Mouse-click allocation uses templates for subnet and IP address assignments, as well as corresponding DHCP pool assignments and DNS domain and resource record, to reduce manual steps and errors associated with entering the same data in multiple user interfaces.
<b>Help enable delegation of IPAM responsibility as desired while controlling access to relevant information</b>	Access control is an important consideration when multiple users have access to the IPAM system. While a core set of users will likely require full access to all system functions and features, it is likely that other administrators would receive a limited set of functionality and scope control based on their respective responsibilities.
<b>Deploy highly available IP services</b>	It goes without saying that DHCP and DNS services are critical to any IP network. Deploy these IP services in site-diverse configurations to provide continuity during disaster recovery. Consider appliances for intrasite hardware-level redundancy for critical servers. If using a centralized IPAM system as recommended, ensure it is not in the critical path to proper DHCP and DNS processing. Consider a high-availability deployment of the IPAM system for operational continuity.
<b>Monitor IP services to proactively manage services availability</b>	Keep track of the status of DHCP and DNS services operating throughout the network through periodic polling or event notification. Help enable navigation into event logs, and remotely control services to facilitate trouble diagnosis and resolution.

Best Practice	
<b>Align IP services upgrades and patches</b>	With appliance-based deployments, one vendor is responsible for not only the DHCP and DNS services version, but also for the appliance operating system and kernel. Staging patches and upgrades on a centralized system with the ability to deploy to remote servers vastly simplifies the coordination, timing, and resource requirements for this otherwise costly and cumbersome process.
<b>Adapt IP management functions to your business processes</b>	Every organization's IP network management policies are unique in some way, despite their common need to effectively manage address space and DHCP and DNS server configurations. To the extent possible, adapt the systems you use to define your addressing topology, device types, and naming policies, as well as attributes on topology nodes, blocks, subnets, devices, and domains. This helps you to manage your IP address space according to your business processes.
<b>Integrate IPAM processes into broader enterprise workflows</b>	In addition to adapting the IPAM system constructs and attributes to your business processes, consider further automating IPAM-related functions into broader IT workflows. Examples include deploying a new site, externalizing IP address requests, tracking asset information on devices, and creating trouble tickets.
<b>Help enable reporting for addressing status and audit information</b>	Communications across organizational levels can be simplified with intuitive, highly graphical reporting. Filtering information to particular hot spots within the network can highlight and convey information that potentially requires escalation. Audit reports are also required to track user accountability and comply with regulatory requirements.

## Simplifying Best Practice Implementation with Cisco Prime Network Registrar IPAM

Given the tight relationship between IP address space management and its implication on DHCP and DNS server configuration, deployment of a centralized IP management tool that supports modern DHCP/DNS server technologies can simplify and automate implementation of these best practices. It can also reduce IP management resource requirements while reducing configuration errors.

Cisco Prime™ Network Registrar IPAM provides a comprehensive, centralized IP management solution for managing IP address space and capacity, as well as DNS and DHCP server configurations. The solution provides support for sophisticated DNS/DHCP services, including DNSSEC, DNS views, IDNA, logging configuration, transaction signature support, DHCP client classes, and much more.

Cisco® Prime Network Registrar IPAM provides the advanced, next-generation, IP management solution you need to automate many tedious and error-prone, yet critical IPAM functions. With unsurpassed extensibility and user-definability, you can manage your IP address space the way you want to manage it. Send email to [ask-networkregistrar@cisco.com](mailto:ask-networkregistrar@cisco.com) to learn more about Cisco Prime Network Registrar IPAM and how it can automate more of the IPAM functions you need at an exceptional ROI.

## References

Dooley, Michael; Rooney, Timothy, **IPv6 Deployment and Management**, IEEE Press/Wiley, 2013.

Rooney, Timothy, **Introduction to International Domain Names for Applications (IDNA)**, BT Diamond IP White Paper, 2011.

Rooney, Timothy, **Introduction to IP Address Management**, IEEE Press/Wiley, 2010.

Rooney, Timothy, **IP Address Management Principles and Practice**, IEEE Press/Wiley, 2011.

Rooney, Timothy, **IPv4-IPv6 Transition and Co-Existence Strategies**, BT Diamond IP White Paper, updated 2011.

Rooney, Timothy, **Securing Domain Name Resolution with DNSSEC**, BT Diamond IP White Paper, 2010.

Rooney, Timothy, **Service Provider IPv6 Deployment Strategies**, BT Diamond IP White Paper, 2011.



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](http://www.cisco.com/go/offices).

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](http://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)

Printed in USA

C11-728639-00 09/13