

IP Address Management Best Practices

By Timothy Rooney



About Cygna Labs

Cygna Labs is a software developer and one of the top three global DDI vendors. Many Fortune 100 customers rely on Cygna Labs' DDI products and services, in addition to its industry-leading security and compliance solutions to detect and proactively mitigate data security threats, affordably pass compliance audits, and increase the productivity of their IT departments. For more information, visit <https://cygnalabs.com>.

© 2022 Cygna Labs Corp. All Rights Reserved.

Introduction

As an IT manager responsible for keeping your IP network up and running, the discipline of IP address management (IPAM) represents a critical ingredient in your recipe for success. The IPAM discipline entails the design, planning, provisioning, monitoring, and management of IP addresses to assure infrastructure devices and eligible end users can obtain an IP address to access your network. Sounds simple enough, and most of the time, your IPAM discipline successfully yields the desired result with end users and devices able to effortlessly initialize on your network from anywhere including remote offices, within private and public clouds, or even in desolate remote sites with Internet of Things (IoT) devices.

But what appears effortless for end users is made possible only with diligent effort on your part. An IP address must be available for each user. The IP address must be routable to their current location so they can communicate. Thus the IP address must logically roll up in a manner aligned with your networking topology. Certain devices like those with streaming requirements may require special routing treatment and hence be assigned an IP address for which routers can apply such treatment. All in all, your IP space must be allocated according to your topology and application requirements with sufficient capacity for the plethora of end user devices accessing your network.

Once initialized with suitable IP addresses, users need that ability to navigate IP applications by name. The domain name system (DNS) facilitates this navigation with its name-to-IP address resolution function. As IP addresses are assigned, corresponding IP address-to-device name mappings must be updated. Hence DNS updates are closely linked to IP address assignments, and therefore DNS is a core IPAM component.

IT managers need to make sure IP addresses are available and are being assigned, and that DNS is keeping up. Effective IPAM then, can be defined broadly as encompassing three major interrelated functions :

- **IP address inventory** – Obtaining and defining public, private, and local IPv4 and IPv6 address space, and logically allocating that address space to locations, subnets, address pools, and devices to be available for assignment to users accessing the network.
- **IP address assignment** – Once the address space has been properly allocated, individual IP addresses may be assigned to user devices. Since most non-infrastructure devices tend to be mobile or otherwise transient, most devices can obtain IP addresses dynamically for use on a temporary basis while they are “on” the network. This address assignment function entails defining IP address pools containing addresses that can be assigned, tracked, and freed up for reuse by others. These pools and corresponding pool parameters are generally deployed for localized distribution from Dynamic Host Configuration Protocol (DHCP) servers which autonomously supply relevant IP addresses and parameters to requesting devices. Managing DHCP server configurations is aided through the monitoring and allocation of address pool capacity to ensure IP addresses are available for those who need them and are authorized to have them. Assignment of cloud instances and other non-dynamic addresses may be performed via a corresponding Application Programming Interfaces (API) or manually if necessary.
- **IP name services management** – As devices obtain IP addresses statically or dynamically, the mapping of device names to corresponding IP addresses must be tracked and published so other users can navigate to each device by name. This function entails configuring Domain Name System (DNS) servers with this address-to-name and name-to-address information. Managing your domain name space and name services also requires proper design of the namespace, configuration of other relevant DNS records, and many behavioral aspects of DNS as well, particularly relating to securing DNS servers and information.

Each of these three core functions is foundational to the proper operation of an IP network, whether that IP network is a private enterprise network, a private or public cloud, the Internet itself, or all of the above. Users need at least one IP address to access the network, whether via a wired or wireless LAN interface, mobile, cloud instance, IoT device, etc., and they need to access resources on the network and the Internet by name to facilitate usability and scale. As mentioned, these functions occur without user involvement. In fact, one could argue that the job of an effective IP address manager is to be invisible: as users attach to various network points, they are automatically configured to communicate and easily access network resources by name.

¹These functions, essentially “DHCP, DNS and IPAM” are often termed “DDI” which is a term synonymous with IPAM. DHCP and DNS configurations are derived from IPAM so either term may be used interchangeably.

Effective IPAM requires proper allocation of address space across the enterprise including extensions into private and public cloud services, so there is adequate address capacity where it's needed when it's needed. Best practices IPAM also entails accurate configuration of DHCP servers for dynamic address users, including differentiation of employees versus "guests", as well as accurate and timely 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 accessing the network. In addition to flawlessly configuring and managing each of these three foundational elements of IPAM, the IP address manager must also cohesively integrate these three areas collectively and integrate these management functions into the broader IT network management environment.

Another core function of IPAM involves its importance in facilitating overall network security. With nearly every IP transaction beginning with a DNS query, it behooves network security specialists to observe DNS transactions to monitor and detect potentially nefarious intent. IPAM's role and impacts on network security planning are highlighted in this paper and extended in more detail in a companion white paper, *DNS Security Strategies*, and we invite you to read that paper as well for more details on the various DNS vulnerabilities, detection methods and mitigation approaches.

This white paper provides IT professionals a guide for how to effectively execute IPAM tasks and recommends best practices for simplifying the IPAM process. These best practices are derived from the Cygna Labs Diamond IP leadership team's collective experience in the IP management space obtained through numerous implementations of IPAM systems, managing customer IPAM environments, and frequent interactions with end users and industry analysts. Many members of the team have also been active in the Internet Engineering Task Force (IETF) in helping to evolve IP technology. Let's begin by digging deeper into our first core area, IP address inventory management.

IP address inventory management

IP address inventory has several facets in its own right. This IPAM function lays the foundation for the other functions and impacts other critical IP network functions, not the least of which is routing. Most enterprise organizations obtain public IP 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 (RIR), e.g., ARIN, RIPE, or others. After a block of IP address space has been obtained, it can then be allocated to locations across the network. Similarly, private IP address space (RFC 1918) or IPv6 unique local addresses (ULA) can also be allocated in a similar manner. This allocation process is necessary to “carve up” each monolithic block into constituent sub-blocks until IP address capacity has been allocated to meet the IP addressing demands of user devices.

Address planning

When planning to allocate IP address space, whether private or public, administrators must forecast the IP address capacity requirements in each subnet on the network, including in virtual private cloud networks. This is typically based on the number of devices, instances, and/or end users at each site, the number of visitors or mobile users expected at the site, cloud burst capacity needs for elastic address demands, and the number of IP addresses required on average for each end user and device.

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 source address from the VoIP address block segment) versus best-effort data packets (packets with the source address from the normal or data block segment). Such source address-dependent routing (SADR) services necessitate allocation of device (source) addresses accordingly.

A third aspect of address planning relates to your use of the cloud and your anticipated cloud IP addressing needs. For example, if you’re extending your private network to a public cloud provider for overflow capacity, i.e., cloud bursting, you’ll need to allocate sufficient IP address space to accommodate your maximum burst of virtual machines (VMs) or virtualized network functions (VNFs).

While the easy solution is to grossly oversize each subnet for each application, in reality this may not be feasible given IP address space limitations, at least for IPv4. Within such address space sizing constraints, administrators must meet the challenge of accurately and optimally allocating address space to each site. For IPv6 address space, constraints are seemingly unlimited and planning should strive for consistently sized allocations at each layer as we’ll discuss next.

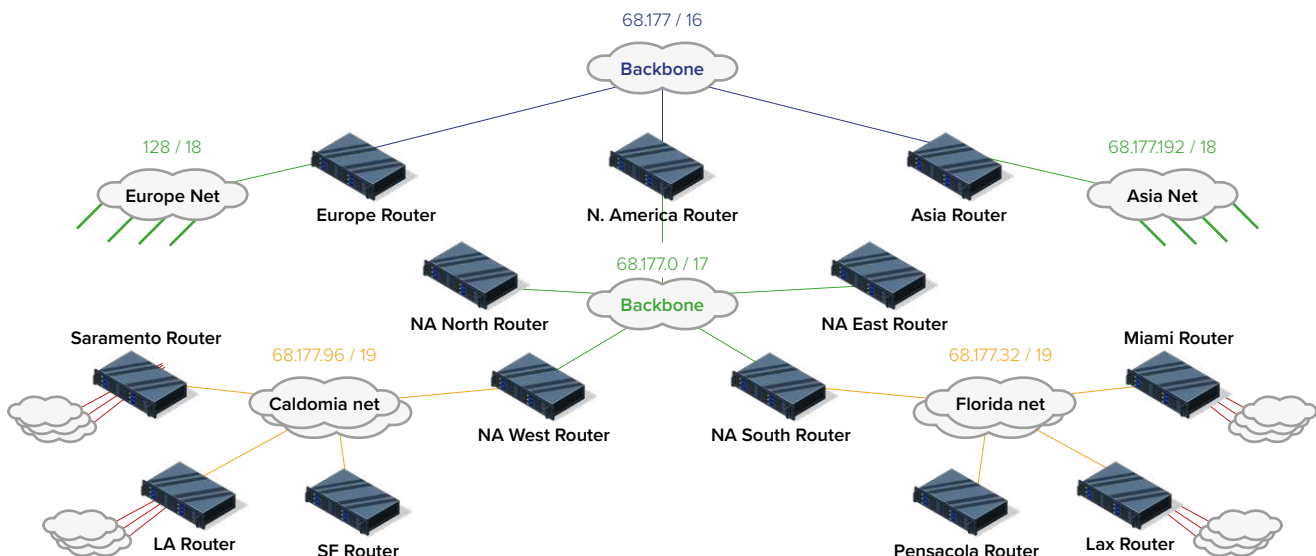


Figure 1: Hierarchical Network Allocation

Address allocation

An additional consideration 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 maximizing address hierarchy in order to facilitate route aggregation for routing protocols such as OSPF (Open Shortest Path First). Route aggregation reduces routing protocol traffic and keeps 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), one of which may not be a continent but a public cloud provider which serves as an extension of enterprise network.

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 mid 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 key to maximizing address utilization in a routing-efficient manner.

If IP network allocation is done improperly, duplicate IP addresses can be unintentionally assigned, networks can be rendered unreachable, 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 via a painful renumbering process before additional address space would be granted by an Internet Registry or ISP.

Your allocation strategy lays the foundation for IP address planning and must be performed effectively to minimize downstream errors and issues and to simplify ongoing network management. A well-designed IP address plan can ease invocation of IP address-based policies such as implementing security controls, routing treatments like SADR, and application policies. This function is so crucial, we wrote an entire white paper focused solely on this topic which we published on the Internet Society website. We invite to read this paper for more details around IP address planning and allocation, particularly for IPv6 (though analogous principles apply to IPv4).

Centralizing IP inventory

Address planning and allocation is best performed using a centralized IP inventory repository. A centralized system provides a single, holistic view of your entire address space deployed over a number of sites on-premises and in the cloud, each potentially with address pools and DNS information deployed on multiple DHCP and DNS servers throughout your 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, and eliminates the requirement of replacing existing DHCP and DNS servers.

You should implement periodic backups and/or replicate to a secondary repository to ensure high availability of this critical IP address information. Another common approach to IP inventory utilizes a decentralized architecture. Such approaches which promote distributed replicated repositories are either not fully replicated (i.e., if each repository member stores only data it needs) or if they are, do provide multiple replicas but can generate tremendous replication traffic on the network in terms of updating all members with all changes. This replication process, with the associated impact on inter-server update performance, can hamper scalability and renders this fully-distributed approach appropriate only for small and single-vendor environments.

Managing address dynamics

After the initial sizing and deployment, even when done perfectly, changes inevitably occur. Virtual private clouds and subnets are allocated and deallocated elastically. 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 and devices such as for IoT are rolled out.

Note that each of these events impact your IP address space, regardless of if they were initiated by business requirements impacting site openings and closures, or by IT in deploying additional IP services such as VoIP and adding cloud VMs 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 necessary for effective IP address space management.

It's important to monitor IP address utilization to track IP capacity and alert IP planners to pending IP address depletion on portions of your network, clouds, or address pools. Proactive monitoring and alerting can help avert an IP addressing crisis. Detecting IP address occupancy beyond address pools also provides feedback related to the integrity of your IP address inventory. Polling your enterprise networks using SNMP, ICMP, DNS, or similar scanning tools should help you identify a snapshot of IP addresses in use. Polling your cloud services platforms via the corresponding cloud API likewise provides IP occupancy data.

By checking the network using the appropriate polling technique then comparing network actuals with your database records, you can identify discrepancies, which may be in the form of "surprise" address assignments, e.g., those assigned locally or of addresses no longer in use, where a device is no longer occupying an address otherwise attributed to it. In the former case, investigation of the potentially rogue device is in order to protect against a security breach and in the latter, the IP address can be reclaimed and added back into a free state for reassignment.

IPv6 deployment

As public IPv4 address space is now effectively depleted, you should plan for deploying IPv6 if you haven't already, at least on your Internet-facing infrastructure. You'll need to make sure your ISP provides IPv6 support, and you should plan to address each Internet-reachable host with both an IPv4 address and an IPv6 address. This dual-stack approach offers the simplest means of implementing IPv6 while offering you valuable experience with the latest Internet Protocol.

Implementing IPv6 will also maximize your Internet presence, affording access to your Internet hosts by either protocol. Eventually you may decide to simplify your IPAM tasks for managing both IPv4 and IPv6 address spaces and decommission your IPv4 network externally and internally. This is the ultimate goal but will likely take some time to transition to this IPv6-only state.

In the meantime, if you're currently in the IPv4-only state and would like more information on where to begin, we invite you to access the following resources (free except for our book).

- [Guidelines for IPv6 address](#) allocation paper on the Internet Society page.
- *IPv6 Addressing and Management Challenges*, a free IPv6 white paper on the cygnalabs.com website.
- Free [videos](#) including several in the IPv6 playlist on my YouTube channel.
- IPv6 Deployment and Management book, published by Wiley/IEEE Press.

IP addressing and security

Your IP address plan plays a key role in facilitating implementation of IP address-based security policies. Access control lists (ACLs) can be defined in one line or several depending on how well your IP address plan corresponds to your "security zones." For example, if you've allocated a single block to a given site, which you've sub-allocated into subnets for users within the site, you could feasibly cordon off that site using the site's block address in a network-wide ACL should you need to impose a quarantine. This is perhaps an extreme example, but you should consider security in your IP address planning tasks.

IP address inventory management best practices

The following are best practices for IP address inventory management.

Inventory address space in a centralized database.

IP address space, public and private, cloud and non-cloud, v4 and v6, is a precious resource, one that provides the fundamental entity for IP network communications. Therefore, it must be tracked judiciously in a centralized though redundant repository to maintain consistency, accuracy, and resiliency. Of course, accuracy requires rigorous updates to the database upon address space allocations and deallocations and IP inventory assurance via network actuals polling across premises and cloud networks.

Rigorously record allocations and periodically reconcile actual IP-related data from the network with the inventory database.

Assuring IP inventory integrity is crucial. During an outage, you may likely need IP address assignment information for local resources if a given site is unreachable directly. Comparing the inventory database with network actuals is crucial to identifying discrepancies and tracking IP management processes. 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 via automated network and cloud discoveries is an effective way to monitor the process and keep inventory accurate.

Perform and track address space allocations in accordance with network topology to model and optimize route aggregation and policy enforcement.

Network allocations should be made in a manner that maximizes utilization of IPv4 address space and is logically consistent for IPv6 space, all while mapping to your network topology model. 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. Per application address allocation should also be addressed if appropriate to manage address allocations for deployed IP services, e.g., voice vs. data.

Implement common allocation policies within address blocks to promote consistent subnet addressing.

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. Provision of a common allocation template promotes consistency in allocation and deployment, and also makes for easier troubleshooting as needed with consistently allocated subnets. Cloud services providers likely reserve addresses as well, and use of corresponding cloud APIs facilitates discovery of these and your IP assignments.

Maintain additional information as appropriate per IP device.

Keeping track of what device is occupying each IP address in a subnet is a critical IPAM function. You should also associate IP addresses to devices, particularly multi-homed or dual stacked devices which are assigned multiple IP address. You may want to track other attributes, including device type, location, switch port,

administrative contact, asset information, and associated resource records to name a few. It is especially useful if these attributes be selectable by device type and ideally by location to maintain relevancy for the IP administrator managing the device.

Monitor address utilization to manage the capacity of the IP address space.

Although initial addressing needs may be impeccably forecast, changes happen in IP networks due to business, IT, or other initiatives. Despite the best planning efforts, IP networks have fluid nature, where address needs rise and fall at different times at various locations within the network and/or cloud. 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 via linear regression models for example to predict future address depletions, which ideally may be alerted up for proactive notification.

Plan for IPv6 if you haven't already

IPv6 penetration is growing and you should plan to support IPv6 addressing, at least externally at first, and ultimately everywhere. While IPv6 address capacity on a standard /64 subnet seem astronomical, we still recommend keeping an eye on IPv6 address occupancy if not for utilization, at least for auditing and monitoring.

Consider security in your IP address planning

Many security policies are enacted by IP addresses, e.g., defining access control lists (ACLs), firewall policies and routing policies. Consider how security policies are generally enacted within your organization; e.g., if by site, such as to quarantine a site, ensure a site-based allocation strategy is "earlier" in the allocation hierarchy. This will enable fewer IP address block specifications to enact a site-based policy. On the other hand, if you need security policy enforcement by application, e.g., to shut down VoIP temporarily for example, define application level allocations at a higher allocation level to ease the enforcement of application based IP policies.

Dynamic IP address services management

Adhering to these IP inventory best practices helps maintain adequately sized subnets and address pools across your network and clouds. But sizing subnets and address pools to supply IP addresses to infrastructure devices and end users, critical as it is, is just the beginning of the process for individual address assignment. For cloud environments, you'll typically assign individual IP addresses to virtualized machines via the cloud platform API. As such, you should link your cloud IPAM functions with your centralized IPAM system to integrate cloud IP assignments within your pan-enterprise IP repository.

For non-cloud enterprise networks, you'll typically assign IP addresses to end users via DHCP . And there's more to configuring DHCP servers than mere address pool allocations. Additional configuration elements include valid DHCP options and DHCP polices with associated values for each address pool, valid or invalid devices by MAC address, device unique identifier (DUID), client class, or user authentication, device software validation, and DHCP failover configuration.

Policy management

Most or all of the DHCP servers in your network will typically require similar DHCP policies, which influence the behavior of the server when processing DHCP transactions. We recommend that you centralize the configuration of these servers to create a single or set number of policies, then deploy the policy(ies) across your servers as appropriate. This practice ensures a consistent and accurate approach to setting these critical policies. Otherwise, you need to enter essentially the same information multiple times into each of your DHCP servers. A similar argument can be made for defining DHCP option sets, which define configuration parameters for DHCP clients, with defined DHCP options and valid values for use on assignment.

Discriminatory address management

Discriminating address assignment refers to the assignment of an IP address to a device relevant to the device's role or permission. For example, a visitor to one of your offices may be afforded an IP address from a subnet whose source IP address is routable only to the Internet and not otherwise within the enterprise. There are several levels of policies or controls most DHCP solutions provide. The first is to simply filter by the MAC address or DUID of the client requesting an address. If the DHCP server has a repository of acceptable (and/or unacceptable) MAC/DUIDs, it can be configured to provide an IP address from a given pool with associated parameters to those clients accordingly. The pool from which the address is assigned should correspond to routing permissions appropriate to the MAC/DUID.

This type of discriminatory IP address and configuration assignment is also possible by filtering the type or 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 option of the DHCP packet. The user class option may also be used. The DHCP server can be configured to recognize the user classes and/or 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/or additional configuration parameters can be assigned to the client via 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 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 outcome entails denying an address assignment; another option is to require the user of the client to login via a secure access web page. This enables easier capture of new MAC addresses for legitimate users of your network. Solutions ranging from simple perl scripts to sophisticated integrated software solutions are available to direct such users to a login/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 via DHCP.

² We'll use the term "DHCP" generically to include both DHCP for IPv4 and DHCPv6.

DHCP resiliency

DHCP redundancy is recommended to support IP address services resiliency across your network. Redundancy can be implemented in the form of split scopes or DHCP failover. Split scopes entails assigning half of the addresses in each address pool to one DHCP server and the other half to a separate DHCP server. If one fails, the other will be able to continue to serve clients. However, as you can surmise, you'd have to provision the full address pool capacity on each server, doubling address requirements to support full redundancy. With DHCP failover, a pair of servers shares common address pools and communicate lease bindings such that if a DHCP server crashes, the failover server can take over and begin processing DHCP transactions for the same set of address pools knowing what has already been assigned.

For IPv6, a standard /64 subnet contains over 1.8×10^{19} IP addresses so you could probably use the split scope approach to provision half of an address pool on a given DHCPv6 server and the other half on another server. These techniques provide resiliency for your end users requesting IPv4 and IPv6 addresses. Remember, if clients cannot get IP addresses, they will be unproductive and will call the help desk!

DHCP appliances

A DHCP appliance implementation can simplify the procurement, deployment, security, monitoring, and upgrading of DHCP services throughout the IP network. Most appliances are self-contained hardware or virtual appliance units featuring a hardened Linux operating system and additional security features. The appliance approach 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 management can be simplified as well, with appliances offering centralized monitoring and patch management features. Many appliances can also be deployed in dual-appliance configurations to provide high availability at the hardware level or via simple virtualized instantiation and configuration via your IPAM solution.

Dynamic IP address assignment best practices

The following are best practices for IP address assignment management.

Link your cloud IP assignment functions with your centralized IPAM system

Maintain a global perspective on all of your IP address space and assignments by linking cloud subnet, IP and DNS assignments with your IPAM solution. Use automation to link virtualized machine instantiation with corresponding IP assignments and periodically poll your cloud platform API to synch its inventory of IP assignments with that of your IPAM system.

Centralize DHCP server configuration to improve configuration accuracy and consistency.

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.

Implement security measures to provide selective DHCP address assignment.

Consider implementing one or more of the following approaches:

- **Device identification via MAC address/DUID** – filter client requests using a list of acceptable and/or unacceptable MAC addresses/DUIDs
- **Device identification via client class** – provide additional configuration information for known client classes configured on your network
- **User identification via authentication** – support user login/password authentication against a database or other authentication scheme
- **Invoke device security scanning or software validation** – scan the requesting device for viruses and/or valid software prior to granting a production IP address

Adopt and use established DHCP option and policy sets across your DHCP servers.

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 to the network.

Track dynamic address assignments and monitor utilization of address pools, including shared subnets, to proactively manage address pool utilization.

As with the address inventory capacity management best practices, this “corollary” best practice recommends monitoring of address assignments and DHCP address pool utilization, including shared pools or shared subnets, to provide both a broad view and drill-down on the capacity impacts from a pool and pool user perspective.

Maintain IP address pool history data to monitor address usage trends and proactively align address space to where it's needed.

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 (Figure 2) helps convey utilization trends at a glance and enables proactive management of address pools including realignment of address pool capacity as necessary to prevent address depletions.

Consider DHCP hardware and/or virtual appliances to streamline deployment, improve security, and simplify upgrades.

A DHCP hardware appliance integrates the hardware, operating system and DHCP service into a simple self-contained platform. Virtual DHCP appliances offer a highly secure appliance solution with the added benefits of flexibility and elasticity. Appliances are secure, purpose-built platforms with restricted operating system permissions, users, ports, and files. Most appliances are deployable in a high-availability configuration, with centralized monitoring and management of administration and updates, which provide lights-out support for distributed appliance deployments.

Configure DHCP resiliency for high availability address assignment services.

IP address assignment is the first basic 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.

1. *Failover Scheme:* The first mechanism is the traditional failover scheme where a common IP 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”, keeping in synch with the primary’s DHCP lease bindings and heartbeat messages. Should the primary server fail, the failover server can kick in and begin handling DHCP address requests.
 1. *Double Scope Approach:* The second mechanism that can be employed if address space utilization is not overly constrained, e.g., for IPv6 or 10.0.0.0 space for modest networks, 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 and alleviates the need for inter-server heartbeat communications, while providing sufficient address capacity for the end users requiring addresses.
 1. *Split Scope Approach:* A third mechanism is to implement split scopes, where two DHCP servers manage non-overlapping subsets of each address pool from the same subnet. For example, half of the addresses on a given subnet could be defined as a pool on one DHCP server, which the other half would be defined as a pool on the other DHCP server. Your routers (“helpers”) would be configured to send DHCP packets from the subnet to both DHCP servers as in the failover scenario. This approach works well for IPv6 subnets or under-utilized subnets but is much simpler to configure and manage than DHCP failover.
-

Block 10.128.254.0/24 History

Days Left = 32 $r^2 = 0.96$

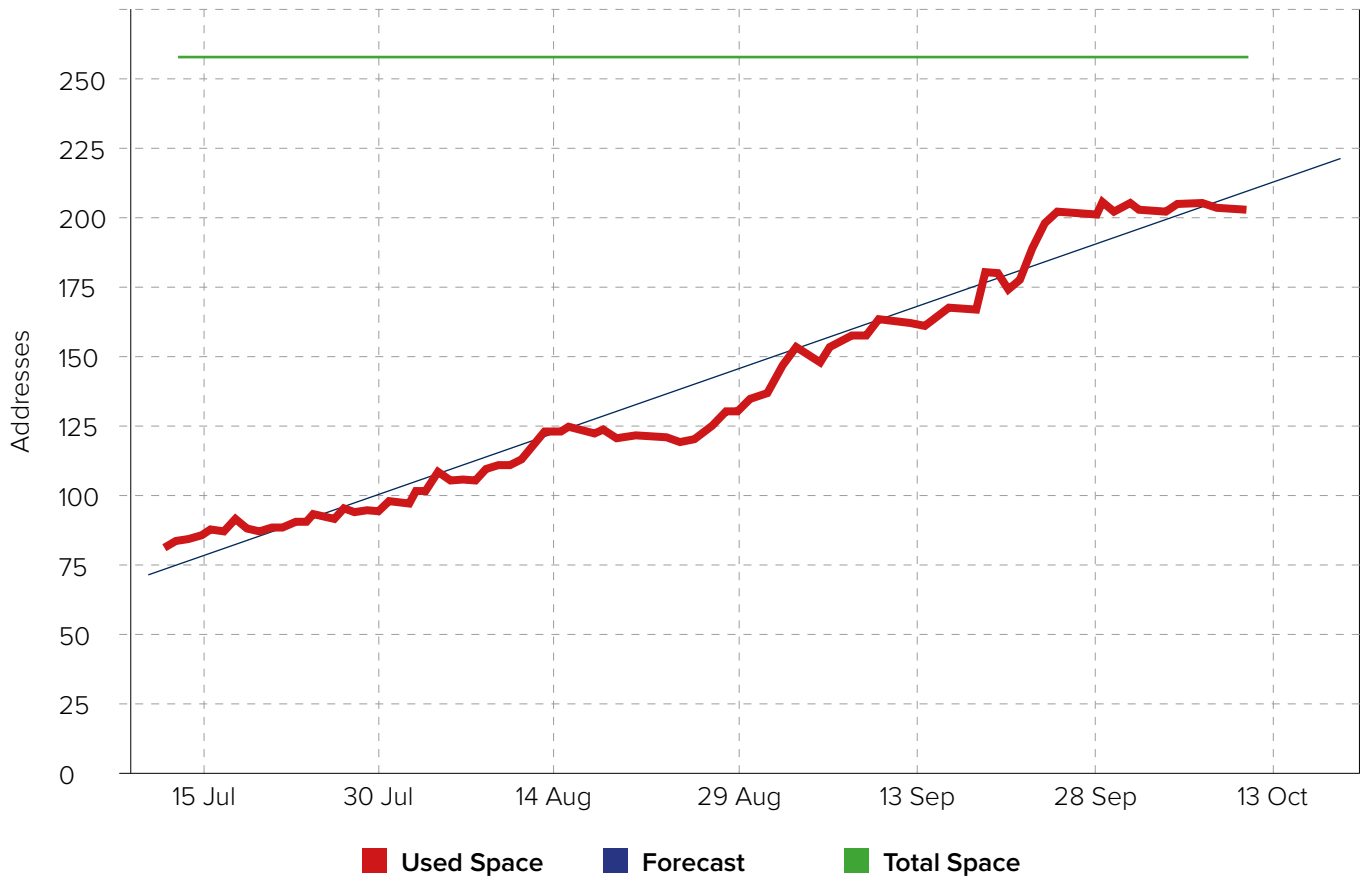


Figure 2: Graphical Address Pool Capacity History and Trending

IP name services management

After a user on your network obtains an IP address and related IP configuration via DHCP or API, hopefully all of which happens seamlessly behind the scenes, most end users will typically access their email and/or the web or intranet. The ability to send email to someone's address at a destination host and browse the web via uniform resource locator (URL) makes email, web browsing, and Internet applications easy and user-friendly. But your computer communicates with the email server and web server via 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 most of these applications could function but would be totally unusable for 99% of your company's end-user population. It follows that DNS services must be configured accurately and be highly available to your users. DNS configuration consists of defining numerous parameters governing the DNS server behavior, zones, security and the like, as well as the translation information, e.g., name to IP address mappings, defined in the form of resource records.

DNS resource records

Configure your DNS servers with resource records needed to resolve domain names and URLs into IP addresses not only for deterministically configured IP devices like routers, web servers, email servers, and cloud virtualized machines, but also for dynamically configured IP devices like office printers and end user devices. As addresses are assigned, DNS needs to be updated with the corresponding domain name to IP address mappings. Many websites perform a reverse DNS lookup for an IP address attempting to make a connection before continuing a web session to validate that the requesting IP address has some form of legitimacy in DNS. This implies integration between DHCP and DNS, which leverages Dynamic DNS (DDNS), an automated process to update DNS upon address assignment by a DHCP server. Otherwise, when you perform address assignment manually or using a cloud API, assure the relevant DNS server(s) is likewise updated.

Beyond name-to-address translation and vice versa, DNS offers many other "translation" applications, which we won't go into here. Each translation type maps to one or more resource record types in DNS. For example, the familiar "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 secondaries via a zone transfer, in some cases, it is desirable to operate in a multi-master mode of operation, though this must be performed carefully given multiple update sources to each "master," and all updates must be synchronized among masters.

Role-based deployment

Deploy DNS servers with specific roles without role-overload. Specific roles include authoritative servers which answer queries directly or recursive or caching servers which query other servers on behalf of clients or resolvers. The authoritative role can be further divided to delineate internal namespace, i.e., for DNS data available to internal enterprise users and external namespace, i.e., your Internet-reachable DNS resolution data. Deploying servers in this manner can improve security and facilitate scalability.

When managing several "sets" of DNS servers for each role, 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 secondaries may be deployed to support external (Internet facing) name resolution; an internal set consisting of a master and five secondaries may be deployed to support internal queries, and three servers for caching. 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 challenge related primarily to scalability entails supporting a common set of resource records across multiple domains. For example, the "www A" record for diamondip.com may be the same as in diamondip.net, cygnalabs.biz and cygnalabsdiamondip.com. Use of a template domain to define and manage these resource records while supporting multiple alias domains that leverage this information can, again, 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 hardware or virtual 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 management can be simplified as well with appliances offering centralized monitoring and patch management features. Many appliances can also be deployed to provide high availability at the hardware or virtual level as well.

IP name services best practices

The following are key best practices for IP name services management.

Centralize the DNS server configuration to improve configuration accuracy and consistency.

Utilizing a single interface and database to configure a number of DNS servers on-premises, in clouds or even as public cloud DNS services provides the ability to enter configuration parameters once, and deploy the appropriate master or secondary configuration to multiple DNS servers, then aggregate dynamic updates to keep the centralized inventory up-to-date. This provides a centralized view into your overall namespace and DNS configuration across your LAN, WAN, SDWAN, cloud and hosted services for DNS servers, domains, zones, and views.

Run multiple DNS servers on different subnets for each zone to maximize availability of critical DNS services to end users.

Deploy DNS servers to eliminate common points of failure and maximize reachability from internal resolvers and to the Internet. Trade off the simplicity of running a single master DNS server for each zone versus the more complex deployment of multi-master DNS. Single master zones ease configuration by requiring updates to one master server; however, take care to minimize exposure to unauthorized updates to this master. Multi-master configurations have less vulnerability but require careful management of the dynamic update process to eliminate cyclic updates.

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.

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 roll back should a corrupted or misconfigured file end up being deployed. Utilizing a DNS configuration or IPAM tool can help reduce entry errors with data validation.

Configure external, internal, and perhaps other "views" of your name space.

A different version of your internal vs. external domains enables you to resolve a given "www" address differently based on whether the querier resides within your network (internal) or on the Internet (external). This simplest approach to implement such "views" is to configure separate versions of your namespace on separate name servers (e.g., an external set of DNS servers and a separate internal set of DNS servers). Alternatively, albeit less securely, you may deploy multiple domain versions on a single set of DNS servers utilizing the "views" feature of BIND 9. The views feature provides economy of name servers but does require more complex configuration and corresponding issue troubleshooting and resolution.

Role-based DNS server deployment

Deploy DNS servers with one and only one role: internal authoritative, external authoritative or recursive/caching. Implementing this best practice generally necessitates a DNS configuration tool to facilitate grouping of sets of DNS server and any alias domains. However, use of such a tool - or better yet, an integrated IPAM solution - can provide the benefits of managing not only multiple DNS server sets and alias domains, but multiple sets of DHCP servers and a diversified IP address space.

Consider DNS appliance deployment to streamline deployment, improve security, and simplify upgrades.

A DNS hardware appliance integrates the hardware, operating system and DNS service into a simple self-contained platform. Virtual DNS appliances offer a highly secure appliance solution with the added benefits of flexibility and elasticity. Appliances offer enhanced security, purpose-built with restricted operating-system permissions, users, ports, and files. Most appliances are deployable in a high-availability configuration, with centralized monitoring and management of administration and updates, which provide lights-out support for distributed appliance deployments.

Network security

Every server or element on your network represents a potential target for attackers seeking to infiltrate your defenses for the purposes of disrupting your network, bringing down your Internet presence, stealing personal or organizational information or for other nefarious motives. IPAM, DHCP and DNS servers are no exception.

IPAM server security

Like all servers, IPAM servers, including DHCP and DNS servers, should be secured with respect to access controls, rate-limiting to throttle denial of service attacks, and operating system parameters such as services, users, files access, and so on. Hardware and virtual appliances support most if not all of these controls. On top of this foundation, secure the respective IPAM services to minimize exposure to various vulnerabilities.

DHCP service security

Unless you work for a service provider, DHCP services should generally be available for devices within your network. Nevertheless, many organizations welcome guests within the confines of their locations and wireless networks, so exposure to “the outside” is certainly possible if not likely. Besides, a reasonable number of attacks originate by “insiders” or others within a network, so DHCP is not immune from threats. Unfortunately, the DHCP protocol was not designed with security in mind. DHCP authentication provides a token based identity management, but was codified long after major implementations had been deployed and is rarely used. Thus, most solutions recommend layer-2 level controls as highlighted below for corresponding vulnerabilities.

- Thwart pool starvation attacks by limiting the number of MAC addresses that may transmit per switch port. This would limit the number of IP addresses an attacker could lease; such an attack issues several DHCP requests for IP addresses while spoofing the MAC address, thereby appearing as a unique client each time. If successful, such an attack would consume all of the available addresses in the pool.
- Identify from which switch ports a valid DHCP server packet may emanate to clients in an attempt to prevent a rogue DHCP server from dispensing IP addresses and associated IP information. An attacker could supply an IP address and its own DNS server to hijack each client to its website for example to impel the user to supply sensitive information.
- Configure your DHCP servers to assign IP addresses only to “known” clients, such as those with a known MAC address. This can limit starvation attacks and can be used to prevent outsider access, though a separate pool could be configured to supply an IP address to an outsider which enables only Internet access for example.

DNS service security

In terms of security measures for DNS, the following are recommended approaches:

- **Configure DNS 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).
- **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.
- Sign your authoritative DNS information using DNSSEC and configure DNSSEC validation for your recursive servers. DNSSEC provides authentication and data integrity validation for DNS resource records.
- Monitor DNS transactions via query logging or `dnstap` to baseline DNS activity and to identify anomalies and potential DNS tunnelling.
- Consider encrypting DNS transactions using DNS over TLS (DoT) or DNS over HTTPS (DoH). You may want to deploy your own DoT/DoH proxy servers in lieu of leveraging cloud services which would reduce your ability to monitor and raise privacy concerns.
- **Run the DNS service (named) in jailed 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.
- **Consider running DNS appliances** – Appliances are purpose-built hardware or virtual platforms for running DNS (and/or 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.

Identify and stop malware

An attacker may attempt to install malware on devices within your network to enlist them as subject to the control of the attacker. Attackers may gain such access via phishing attacks, social engineering, or other methods. Once within the network, the attacker's malware typically attempts to communicate to a "command and control" (C2) center, from which the attacker can instigate attacks, update malware code, or steal information. Many times, this communications process involves querying DNS to identify the current IP address of the C2 center. Use of DNS enables the attacker to change IP addresses quickly to avoid detection and takedown. But it also provides an opportunity for you to detect and stop such malware by implementing a DNS firewall.

A DNS firewall detects query attempts to known or suspicious malware domains and can block or redirect queriers to a remediation portal. As malware attempts to locate its command and control center for instructions or software updates using DNS, the DNS firewall can stifle the attempt and prevent the communication. A DNS firewall adds a defensive layer on top of your traditional network firewalls to identify and stop malware attacks.

Role-based DNS deployment

We covered role-based deployment in the name services best practices section, but this approach bears repeating from the security perspective. Should a DNS server be infiltrated, its ability for use in further attacks should be limited based on its specific role. In addition, network filters and access control lists can be constrained to a small number of IP addresses for outbound DNS queries and responses for recursive servers and for inbound queries for external authoritative servers. Consider the following DNS server deployment strategies to maximize availability and reduce security vulnerabilities:

- "Hide" master DNS servers. If attackers find and infiltrate the master DNS server, secondaries will zone transfer from this master, spreading the corruption. Hiding the master can be accomplished by editing the standard NS and A/AAAA (glue) records pointing to the master DNS server to point to a different (secondary) server. The master name ("mname") field of all secondaries' SOA records should also be edited accordingly.
- Deploy servers on different networks and on different ISP links to minimize denial-of-service impacts
- Deploy external (Internet-facing) name space on external DNS servers separated from internal DNS servers
- For highly secure internal-only networks, consider operating internal root servers
- Use a separate network for queries versus zone transfers and updates.

IP address-based security policies

Many security policies such as ACLs or ingress/egress filters are configured using IP addresses, either individually or in block (e.g., CIDR) format. Consider this fact when allocating IP address space. While you may not have the luxury of a greenfield deployment, you may for IPv6 address space, if you haven't yet defined your IPv6 address plan. A well-thought-out address plan can simplify security and network management down the road as we mentioned in the IPv6 deployment and management section. Nevertheless, maintaining an accurate IP address repository provides a critical cross-reference when defining security policies required by an attack or related event.

IPAM-related network security best practices

The following are key best practices for IP name services management.

IPAM server security	Implement server common controls to minimize exposure to network-based hack-based, denial of service and related server level attacks. Tactics include server hardening, least-privilege access accounts, encrypted remote access, activity auditing, event monitoring, etc.
Secure DHCP	Configure switches for DHCP snooping to detect rogue DHCP servers and if feasible, limit the number of source MAC addresses per port to defend against pool starvation. Configure separate "guest" address pools for external users and track "known" MACs if feasible.
Secure DNS	<p>Most DNS implementations offer a variety of configurable options that allow specification of ACLs, pair-wise server transaction signatures, IP address/port specifications, and certain forms of rate-limiting. While these options provide the flexibility for configuring these capabilities, the challenge becomes accurately configuring each server with its corresponding ACLs, keys, and IP addresses/port numbers. For a large number of servers, this can be cumbersome and error-prone to configure manually. And don't forget that Microsoft provides a different means of signing updates with GSS-TSIG.</p> <p>Monitor DNS for anomalies and other exploits such as tunneling to prevent data exfiltration and service denial. Implement DNSSEC to digitally sign your DNS data.</p>
Identify and stop malware	Implement DNS firewall capability on your recursive servers to block malware attempts to exfiltrate data and contact C2 centers and to identify malware-infected devices on your network.
Role-based deployment	DNS in particular should be deployed with configurations specific to respective role in order to minimize infiltration bleed and ease trouble identification and resolution.

Leverage IPAM for IP address-based security policies

Ideally define your IP address plan with security in mind. For existing networks, maintain an accurate IPAM repository to enable rapid implementation of defensive measures upon security incident detection.

IPAM governance

Bringing together and governing IPAM functions within a centralized platform affords several advantages for IP managers. Clearly, the interrelationship among IP inventory, DHCP, and DNS is very close. Automating functions among these three key areas and minimizing duplicate entry of related information can reduce errors and save time. Extending automation beyond this, however, provides additional benefits in terms of automating 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.

Holistic management

We've already discussed the benefits of centralizing IP inventory to enforce change control, enable delegation, and support accurate inventory tracking. Given the closely related functions of DHCP and DNS configuration, it makes sense to also centralize and integrate DHCP and DNS configurations, leveraging the IP inventory information. This enables 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/or with mixed Microsoft, ISC/Kea/BIND, and cloud services deployments.

With networks increasingly incorporating cloud technologies, centralizing your IPAM becomes more critical to providing a holistic pan-enterprise perspective on all of your IP address space. So, whether you enter IP, DHCP and DNS information using a graphical user interface (GUI) or API, the consolidation of these functions can streamline and improve your IPAM processes. Along the way, you'll also need to consider imposing controls on "need to know" access, deploying redundant systems, and fully managing your IPAM systems.

Administrator access controls

For most organizations, responsibility for various aspects of the many requisite 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, cloud services and/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, whether accessed via the GUI or API or either, while "super user" functions can be reserved for the core IPAM team.

High availability services

Clearly, IP address assignment, DHCP, and DNS services are critical to all IP networks. We've previously recommended DHCP failover and deployment of multiple DNS servers to provide high availability. Deployment of redundant IPAM systems is equally critical, especially if you experience a high rate of IP address related changes within your network. Maintaining reliable access to IPAM data, cloud APIs and auditing data enables IP services continuity in the event of an unplanned outage.

Deployment of hardware or virtual appliances can provide an added layer of high availability. Virtual appliances can be instantiated at a moment's notice to add capacity in the face of a failed element. Deploying DHCP, DNS and IPAM appliances in your cloud streamlines the instantiation of these core IP network functions from weeks to minutes. Quickly turn up additional capacity and deploy configurations from your centralized IPAM system.

Most hardware appliances, including those from Diamond IP, are available in “back-to-back” mirrored connections for co-located hardware redundancy. This dual configuration can be deployed in addition to DHCP failover and multiple DNS masters/secondaries 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's equally important to monitor these services to ensure they are properly functioning. If end users can't 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 “kicks in,” it's important to identify and rectify the failed service quickly to minimize vulnerability of service outage should the backup service subsequently fail.

Upgrades and patch management

Keeping up with security and feature patches is important for minimizing vulnerabilities and maximizing feature utility. 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/IPAM appliances with centralized patch management can remove many if not all these headaches. Selective upgrades of OS, kernel, and DHCP and/or DNS version from centralized system can streamline the patch management and rollback process.

Adaptation to your business

Many software tools tend to be rather rigid in terms of IP subnet and device attributes and topology. However, every IP network is different. And techniques for 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 enables 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 automation 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. Integration with cloud automation tools maximizes efficiency for server and cloud deployments. Automating basic IPAM functions such as the allocation of an address block to a site would also likely require the associated updating of relevant DHCP and DNS server configurations, as well as provisioning of the subnet to the corresponding router interface. If the IPAM system supports the passage of subnet allocation information via 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 downstream workflow shortens the overall implementation interval and reduces miscommunication errors as well as duplicate entry errors.

With increasing proliferation of IP services in the enterprise and cloud, enterprise IP managers typically need to allocate application specific subnets or VLANs, accurately assign IP addresses and/or associated configuration parameters via 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. Integrating cloud and non-cloud IPAM integrates these diverse environments while providing a pan-enterprise perspective to manage your IP address space.

IPAM reporting

Communicating the state of the IP network from an addressing perspective 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.

IPAM governance best practices

The following are best practices for IPAM governance.

Centralize management of IP inventory and DHCP and DNS services

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 streamlines IPAM processes.

Enable delegation of IPAM responsibility as desired while controlling access to relevant information

Access control is an important consideration when multiple users have GUI and/or API 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.

Deploy highly available IP services

It goes without saying that DHCP and DNS services are critical to any IP network. Deploy these IP services in site-diverse configurations, cloud and non-cloud, to provide continuity during disaster recovery. Consider hardware appliances for intra-site hardware level redundancy and virtual appliances for rapid deployment of critical IPAM services. Don't overlook IPAM system redundancy with replication and cloud elasticity capabilities.

Monitor IP services to proactively manage services availability

Keep track of the status of IPAM, DHCP and DNS services operating throughout the network via periodic polling, logging, or event notification. Enable drill-down into event logs and remotely control services to facilitate trouble diagnosis and resolution.

Streamline IP services upgrades and patches

With appliance-based deployments, one vendor is responsible for not only the IPAM, 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 coordination, timing, and resource requirements automatically or manually for this otherwise costly and cumbersome yet critical process.

Adapt IPAM functions to your business processes

Every organization's IP network management is unique, despite their common need to effectively manage address space and DHCP and DNS server configurations. To the extent possible, adapt your IPAM systems to align with your addressing topology, device types, and naming policies, as well as attributes on topology nodes, blocks, subnets, devices, and domains. This enables you to manage your IP address space according to your business processes and linguistics.

Integrate IPAM processes into broader enterprise automation workflows

In addition to adapting your IPAM system constructs and attributes to your business processes, consider further automating IPAM-related functions into broader IT workflows, such as deploying a new site, externalizing IP address requests, integrating cloud IP and DNS assignment, tracking asset information on devices, and creating trouble tickets.

Enable reporting for addressing status and audit information

Simplify cross-organizational communications 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.

IPAM governance best practices

The close relationship between IP address space management and its direct impact on DHCP and DNS server configuration across a diverse networking environment warrants utilization of a centralized IPAM system that supports this ever-evolving networking climate. IPAM systems that support current DHCP/DNS server and service technologies simplifies implementation of these best practices and reduces IP management resource requirements and potential configuration errors.

Cygna Labs Diamond IP's innovative IPControl™ software and appliance products provide a comprehensive centralized IPAM solution portfolio for managing IP address space and capacity, as well as DNS and DHCP server configurations. IPControl supports sophisticated DNS/DHCP configurations, including DNS views, response policy zone (DNS firewall) configuration, DNSSEC validation, DHCP client classes, and much more.

IPControl is available for deployment in a software-only package for installation on customer hardware, and as a virtual or hardware appliance platform for both the centralized management system and DHCP/DNS servers, or a mix of software and appliance deployments. Supported virtual platforms include VMware, KVM, Xen, and Hyper-V private cloud platforms as well as Amazon Web Services (AWS), Azure, Google Cloud, and Oracle Cloud public platforms. The Diamond IP Cloud Automation Appliance (CAA) links IPControl with your automation tools to integrate IPAM services with your public and private cloud platforms.

Streamline IP inventory functions

- IPControl provides a centralized comprehensive IP address inventory database, from which IP space can be consistently assigned and managed. Allocate IPv4, IPv6, and multiple subnets with a single mouse click with no spreadsheets. DNS and DHCP servers can be configured via deployment from IPControl. IP discovery capabilities enable periodic collection of actual network information for comparison and reconciliation with the centralized IP address inventory to assure accuracy and identify potential change control violations.
- IPControl automates subnet allocations, simplifying the process to a mouse click, thereby eliminating binary arithmetic errors. The IPControl block-type feature enables definition and allocation of address subspaces to manage multiple IP

address segments for application and administrative purposes.

- Site allocation templates enable you to define multi-block allocations, such as for a new branch office for instance. Each of your branch offices may require a subnet for data, one for wireless, one for voice, one for IPv6, one for management, etc. With one click, you can allocate all of these subnets of the corresponding templated sizes, types, individual assignments, DHCP pools and DNS resource records.
- IPControl enables you to model your network topology via its innovative, patented container feature. Containers allow you to define a hierarchy and track address space allocations in accordance with routing topology to model route aggregation, security policy hierarchy or other user-definable structure.
- IPControl address allocation templates allow you to reserve and assign subnet addresses for routers, servers, and other elements common to your subnets, enumerating each pre-allocation as static, reserved, dynamic DHCP, automatic DHCP, manual DHCP, or delegated prefix address ranges. Automated creation of associated DNS resource records also saves time and effort.
- IPControl provides unparalleled user definability, including user-defined device types. And you may assign unique attributes to each device type if desired via Information Templates and naming policies for DNS updates. In addition, container policies can be set to define allowable device types and block types per container, and per container Information Templates to allow per device type/per container and per block type/per container attributes.
- IPControl collects data from routers, IP devices, and DHCP servers to gather actual IP address utilization information across the network. User defined alerts warn you of impending address pool exhaustions.
- IPControl natively supports both IPv4 and IPv6 together to enable development of IPv6 address planning, coexistence, and migration strategies, while effectively managing your deployed IPv4 network.

Automate accurate address assignment

- The Sapphire Cloud Automation Appliance (CAA) provides a focal point for IPAM API calls to manage subnets, as well as IP address and DNS assignments in your public and private cloud platforms.
- IPControl enables you to centralize your DHCP server configurations to improve accuracy and consistency for both IPv4 and IPv6.
- Diamond IP solutions provide multiple secure DHCP mechanisms, including client filtering by MAC address, client class, user authentication, and/or device verification callouts.
- IPControl provides user-definable option sets and policy sets, which can be applied across multiple DHCP servers to promote configuration consistency.
- IPControl enables simple configuration of DHCP failover for high availability address assignment services, whether using shared scopes or double scopes. The IPControl subnet display enables simple definition of non-overlapping dynamic address scopes and association with different DHCP servers to reduce errors in configuring split-scopes.
- IPControl automates tracking of dynamic address assignments and monitors utilization of address pools, including shared subnets, to proactively manage address pool utilization.
- IPControl maintains address pool history data along with linear regression trending to present pool utilization in an easy to understand graphical format. At a glance, you can determine address pool usage trends, communicate this among multiple organizational levels, and take proactive action.
- IPControl Sapphire appliances support DHCP (and/or DNS) services on a secure, easily deployed and upgraded hardware or virtual appliance platform. Hardware deployment in a redundant TwinMirror™ configuration provides hardware redundancy and high availability, while virtual deployments on a wide variety of platforms facilitates rapidity and elasticity. Centralized monitoring of hardware and virtual appliances enables services status monitoring and control, as well as patch management.

Streamline DNS configuration while enabling advanced features

- The Sapphire Cloud Automation Appliance (CAA) provides a focal point for IPAM API calls to manage DNS assignments, as well as subnets and IP addresses in your public and private cloud platforms. The CAA also enables deployment of DNS zones and resource records to public cloud services such as Akamai.
- IPControl enables you to centralize your DNS server configuration to improve accuracy and consistency.

● CYGNA LABS DIAMOND IP WHITEPAPER

- IPControl supports nearly any configuration of multiple master/secondary DNS configurations from a few servers to several hundred servers. DNS servers can be deployed throughout your network to facilitate scalable deployments, promoting maximum flexibility and unconstrained DNS server network design while achieving centralized management.
- IPControl provides features to validate your DNS configurations prior to deploying them to production DNS servers in your network. This enables you to gain an extra level of assurance of the validity of your DNS configuration files. Should an erroneous configuration be deployed from IPControl or direct configuration file edits, IPControl enables the added assurance of configuration rollback if needed.
- IPControl was the first IPAM product to support configuration of BIND DNS views from a GUI interface. IPControl enables you to create separate views on separate name servers (e.g., an external set of DNS servers and a separate internal set of DNS servers) or on a single set of DNS servers utilizing the “views” feature of BIND. In fact, virtually any parameter or directive that can be set within a BIND configuration file can be configured through the much simpler IPControl graphical interface.
- IPControl enables you to tighten security by configuring ACLs, transaction signatures for dynamic updates, zone transfers and control messages, and specifying particular TCP/UDP ports for queries, updates and zone transfers. IPControl also eases configuration of DNSSEC and DNS firewalls. IPControl also supports the ability to obtain and use GSS-TSIG keys when signing updates to Microsoft DNS servers.
- IPControl provides a number of unique and innovative features to simplify management of large DNS environments, including support of DNS catalog zones and template/alias domains.
- IPControl Sapphire appliances support DNS (and/or DHCP) services on a secure, easily deployed and upgraded hardware or virtual appliance platform. Hardware deployment in a redundant TwinMirror™ configuration provides hardware redundancy and high availability, while virtual deployments on a wide variety of platforms facilitates rapidity and elasticity. Centralized monitoring of hardware and virtual appliances enables services status monitoring and control, as well as patch management.
- Configuring caching only recursive DNS servers is a snap with IPControl. In general, DNS server templates allows administrators to define DNS servers for virtually any application, whether for caching-only, internal root or authoritative name servers and more.

Secure your IPAM, secure your network

- IPControl supports configuration of complex DNS options natively within its web GUI. Define ACLs, DNSSEC validation, response rate limiting, response policies, and more.
- Diamond IP’s DNS firewall service provides a continually updated response policy zone feed to secure your DNS queries and end user devices.
- Sapphire appliances are purpose-built secure hardware and virtual appliances and support network interface card ACLs and rate limiting features as an added network security layer.
- Sapphire Sx model appliances automate DNSSEC signing for fully automated DNSSEC signing and key management, including support for third party hardware security modules (HSMs).
- Leverage IPControl’s centralized IPAM repository for defining an applying IP address-based security and routing policies.
- The Sapphire A30 Auditor appliance serves as a highly scalable repository of IPAM, DHCP and DNS audit and transaction data. Its rich graphical interface provides visibility and reporting reporting while supporting drill down to the individual DHCP or DNS transaction level for cybersecurity threat investigation (CTI) and forensics.

Bring it all together with IPAM governance

- IPControl leverages a replicatable relational database to provide a high-performance, scalable and centralized IP address inventory database. The inventory consists of IP address, status, device, related DHCP and/or DNS information, and even user-defined attributes for comprehensive IPAM inventory.
- IPControl’s granular, administrator roles enables scoping of administrator or group access to the system by function, topology element, domain, and much more.
- IPControl simplifies definition of multiple DNS servers (masters/secondaries), multiple DHCP servers for one-to-one or many-to-one failover, and split-scope DHCP. Sapphire appliances may be deployed in a multiple hardware deployments to support colocated hardware redundancy or as virtual appliances for rapid deployment and elasticity.

- The appliance dashboard feature provides a summary display of deployed Sapphire appliances, including status of DHCP, DNS, and IPControl services, hardware metrics as well as services reporting, e.g., DNS queries/second. Drill-down to detailed events and the ability to manage the services' states and patch levels provide further control from the centralized IPControl system.
- Diamond IP simplifies the otherwise tedious process of upgrading remote DHCP and DNS servers. Sapphire appliances can be fully patched, upgraded, and rolled back from the centralized IPControl system. Our Sapphire Infrastructure Management (SIM) managed service automates this process with the help of Cygna Labs Diamond IP's managed IPAM services center.
- IPControl provides unsurpassed user definability in enabling user definition of the container hierarchy, block types, device types, naming policies, attributes, and much more. User definability lets you manage your IP space in the manner you desire.
- The extensive APIs and CLIs in IPControl provide the ability to integrate IPAM functions with external systems, with the Sapphire Cloud Automation Appliance (CAA) facilitating such functions out of the box. In addition, the innovative Callout Manger Service can trigger actions or updates to downstream systems based on IPControl-initiated actions.
- IPControl's highly graphical utilization reports facilitate at-a-glance comprehension of IP address status. Extensive audit reports in IPControl and in the Sapphire A30 Auditor appliance enable tracking down of administrator actions and IP address "ownership" status over time for auditing and regulatory compliance.

Key Diamond IP differentiators

IT managers are facing the challenge of managing IP networks that are rapidly growing in breadth and sophistication without a corresponding growth in resource budgets. IP networks have evolved to supporting nearly all forms of communications including multimedia services beyond wired and wireless data, requiring more hierarchical topologies and higher levels of sophistication of DHCP and DNS technologies.

Meanwhile, budget constraints beckon IT managers to automate as much as possible while adding support for evolving IP network features. The need for elastic capacity at minimal cost has spurred expansion of traditional enterprise networks into private and public clouds, which in turn has further driven the need for automation.

Cygna Labs Diamond IP developed its IPControl and Sapphire products with this evolved IP network in mind. Our managed IPAM services utilize our IPControl and Sapphire products to offer an automated IPAM option. Diamond IP provides the following unique features, which enable incomparable IPAM and adherence to IPAM best practices.

- Manage IPAM on your own or with our help
 - Diamond IP's IPAM portfolio comprises the industry's broadest offering, enabling you to fully manage IPAM yourself, leverage our services, or any where in between.
 - Diamond IP is the only IPAM vendor offering software, virtual and hardware appliances and managed IPAM services, affording you the most choices for your critical IPAM implementation. Our managed services even include support for multiple vendors other than our own products.
- Holistic pan-enterprise IPAM encompassing cloud and non-cloud networks
 - Manage and track subnets, IP assignments and DNS updates for on-premises and cloud-based network functions or elements.
 - Automate subnet, IP and DNS assignments for private and public cloud platforms with our Sapphire Cloud Automation Appliance (CAA).
 - Deploy secure Sapphire appliances as hardware or as virtual appliances supported on VMware, Oracle VM, Xen, Hyper-V, AWS, Azure and CYGNA LABS Cloud Compute.
- Improve overall network security
 - Plan for and track IP allocations and assignments to facilitate security policy enforcement, particularly for IP address based filters, policies and ACLs.
 - Secure your DHCP and DNS services with a variety of appliance and server security features including rate-limiting, ACLs, DNSSEC, and much more.

- Detect and stop malware proliferation with our DNS firewall service and recursive DNS firewalls.
- Reduce the risk of critical DHCP, DNS and IPAM outages with DHCP failover or split scopes, multiple DNS server deployments and deployment of multiple centralized IPAM systems. The Sapphire Executive Monitor Appliance (EMA) supports monitoring of a pair of IPControl Executive systems and automates failover upon detection of a user-definable outage condition.
- Extensible automation
 - Our REST APIs facilitate automation.
 - Incorporate IPAM tasks into broader IT or operations workflows.
 - Integrate IPAM functions with cloud provisioning with our Sapphire Cloud Automation Appliance (CAA).
 - Our unique callout manager enables downstream calls based on IPAM event triggers for flow-through workflows.
- Our products offer logical, multi-tiered, centralized IPv4/IPv6 address inventory
 - IPControl is the only product that enables modeling of multi-tiered, hierarchical IP networks within its centralized inventory. This is enabled by its patented container structure with corresponding policies governing respective IPAM functions. For example, easily model branch offices, remote sites, data centers, virtual private clouds (VPCs) as containers within IPControl.
 - User-definable block types enable delineation of address space by application, by administrative domain, or other user definable partition. For example, to allocate a new VoIP subnet, simply define select the VoIP block type and size within the desired the container and click Submit.
- Simplified, multi-vendor DNS/DHCP configuration
 - Protect your investment in your current DNS/DHCP infrastructure as Diamond IP supports DNS/DHCP configuration for not only our Sapphire hardware and virtual appliances, but also for stock ISC/BIND, Microsoft and Cisco PNR DNS/DHCP servers. And with our Sapphire CAA, we support DNS cloud services such as Akamai.
 - While many tools support centralized or distributed configuration of DNS servers, IPControl is the only product today that supports sophisticated ISC BIND features including catalog zones, views, TSIG, statistics, controls, logging, options, server templates, and much more—all within the GUI interface. This simplifies and improves the accuracy of complex DNS configurations.
 - IPControl also enables definition of unique DHCP policy sets to define behavioral aspects of DHCP servers under management. In addition, DHCP failover configuration is vastly simplified using either the traditional approach on a per-server or per-subnet basis
- Network data collection and reconciliation
 - IPControl integrates the automated data collection of interface and IP address configuration information from routers and MIB-II devices to reconcile the inventory database's version of network ("planned") and router/switch configuration ("actual") configuration.
 - IPControl integrates the automated data collection of configuration and active lease information from network services to reconcile the inventory database's version of network and server configuration ("planned") versus the server ("actual") configuration.
 - IPControl supports active subnet discovery to reconcile individual IP address inventory with network actuals. Discrepancies can be highlighted and selectively accepted as inventory updates.
 - IPControl is the only product that tracks historical address utilization data via intuitive, easy-to-read graphical utilization and trending reports. User-defined thresholds and alerts enable you to define conditions for alerting you of impending address depletions so you can proactively allocate address capacity before it runs out.
 - The Sapphire A30 Auditor appliance collects DHCP and DNS transaction data from deployed Sapphire DHCP/DNS appliances to provide centralized reporting, auditing, and forensics.

- Unsurpassed user definability
 - IPControl enables you to manage your IP address space the way you want to manage it, not in accordance with a rigid software tool. You can define user defined fields of various data types (text, radio button, text box, drop down list, URL, and more), whether required or not, along with other attributes. Groups of user-defined fields, called Information Templates, can then be associated with containers, subnets, and devices to enable you to track this additional information with the corresponding system element.
 - IPControl also supports user defined device types and the most flexible device naming policies on the market. You can define and concatenate free text, IP address, incrementors, and more to define policies per device type.
 - Policies for containers, which can map to your network topology or geography, can be established in terms of allowable address types and device types, and associated Information Templates to enable you to attach different information to a particular device type in one area differently from another if desired.
 - Thresholds and alerts can be activated to define the conditions required to fire an alert, along with the associated criticality, for container-level, address block level, and DHCP server level alerts.
- Affordability
 - Despite this expansive feature set, you don't need to purchase licenses for each major feature individually, unlike some other appliance solutions. Diamond IP integrates these features into the base product to maximize your value. You'll find that CYGNA LABS offers exceptional value by providing these key differentiating at a lower price than that of comparable competing solutions.

We are the IPAM experts

Each Diamond IP leadership team member boasts over two decades of experience managing IP networks and developing IPAM products and services.

We literally wrote the book on IPAM with five books on IPAM related topics published by Wiley/IEEE Press.

- Introduction to IP Address Management
- IP Address Management Principles and Practice
- IPv6 Deployment Management
- DNS Security Management
- IP Address Management, Second Edition

Conclusion

Leveraging expertise gained through more than two decades of experience working with customers, industry analysts, and various software implementations, this white paper recommends numerous best practices for effectively managing your foundational IPAM services. Building on this experience and these recommendations, Cygna Labs Diamond IP has developed IPControl software and appliance products as well as a managed IPAM service to simplify the fulfillment of these recommendations.

Cygna Labs Diamond IP products and services offer an advanced, comprehensive IPAM solution you need to automate many tedious, error-prone, yet critical IPAM functions. Diamond IP provides unsurpassed extensibility and user-definability to enable you to manage your IP address space the way you want to manage it, all at an affordable price. Please email us at sales@cygnalabs.com to learn more about how IPControl can automate more of the IPAM functions you need at an exceptional return on investment (ROI).

Toll Free: **(844) 442-9462**
International: **+1 (305) 501-2430**
Fax: **+1 (305) 501-2370**

Sales: sales@cygnalabs.com
Support: support@cygnalabs.com
Billing: finance@cygnalabs.com

cygnalabs.com

