

Best Practices for Next-Generation IP Address Management

*By Tim Rooney
Director, Product Management
BT Diamond IP*

Best Practices for Next-Generation IP Address Management

By Tim Rooney, Director, Product Management

Introduction

IP address management (IPAM) can be defined broadly as encompassing three major interrelated functions:

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

Each of these functions is critical to the proper operation of an IP network. Users need at least one IP address to access the network, whether via a wired or wireless LAN interface, VoIP device, video device, etc., and they need to access resources on the network and the Internet to maintain a high level of productivity. Typically, these functions occur without user involvement. In fact, one could argue that the job of an effective IP address manager is to be invisible. In other words, as users attach to various network points, they are automatically configured to communicate and easily access network resources by URL/name.

Effective IP management requires proper allocation of address space so there's adequate address capacity where it's needed when it's needed, accurate configuration of DHCP servers for dynamic address users, including differentiation of employees versus "guests", and accurate configuration of DNS servers so resources can be accessed easily. When these behind-the-scenes tasks are flawlessly executed, network users don't need to contact the help desk with complaints about the network; the IP address manager is invisible. In addition to flawlessly configuring and managing each of these three foundational elements of IP address management, the IP address manager must also "manage" these three areas collectively, and integrate these management functions into the broader IT network management environment.

This white paper provides IT professionals a guide for how to flawlessly execute tasks for IP management, and recommends best practices for simplifying the IP management process. These best practices are derived from the BT Diamond IP leadership team's collective experience in the IP management space obtained through numerous implementations of IP management systems, and interactions with end users and industry analysts. Many members of the team have also been active in the Internet Engineering Task Force (IETF) in evolving IP technology.

IP Address Inventory Management

IP address inventory has several facets in its own right. This function within IP management lays the foundation for the other IP management functions, and impacts other critical IP network functions, not the least of which is routing. Most enterprise organizations will obtain public IP address space from an ISP, though some that have been using the Internet for some time have a legacy relationship with their Regional Internet Registry, e.g., ARIN and RIPE. After a block of public IP address space has been obtained, it can then be allocated to locations across the network. Similarly, private IP address space (RFC 1918) can also be allocated in a similar manner.

Address Planning

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

While the easy answer is to grossly oversize each subnet for each application, in reality this isn't feasible given IP address space constraints. Even for plentiful private address space for large networks requiring a centralized IP management system, licensing and support costs associated with managing many such networks can be prohibitive. Within these address space sizing constraints, administrators must meet the challenge of accurately and optimally allocating address space to each site.

Address Allocation

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

The address space planning process then needs to carefully consider the macro level requirements for address space as well as the rollout of individual address space requirements. For example, a global corporation may wish to subdivide its space among a core backbone of sites covering three continents (Figure 1). It may make sense to subdivide the

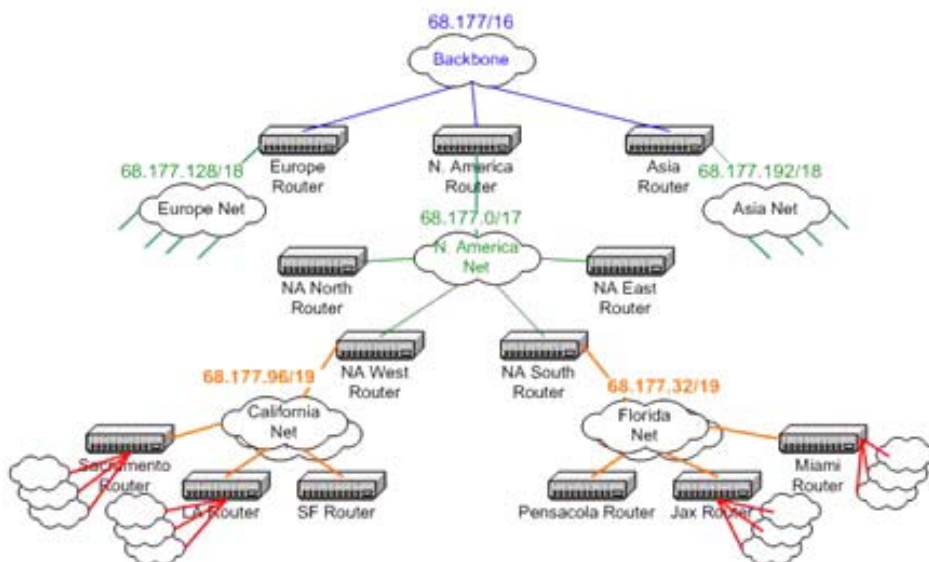


Figure 1: Hierarchical Network Allocation

“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 assigned, networks can be rendered unreachable due to the route summarization example described above, 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.

Centralizing IP Inventory

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

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

Managing Address Dynamics

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

Note that these events all impact the IP address space, regardless 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 more servers or devices for performance or other reasons, or by end user behavior in terms of addressing requirements at particular sites. Staying on top of these and other changes, which reflect the organic nature of IP networks, is absolutely necessary for effective IP address space management.

IPv6

Issues with address space management become even more critical as IPv6 begins to make inroads into service provider and enterprise networks. The sheer size and hexadecimal representation of IPv6 addresses invites operator errors, stifling effective IP address management. Because few if any customers will actually deploy IPv6 in a “greenfield” environment and the transition will likely take years to complete, integration of IPv4 and IPv6 address allocation processes is crucial.

IP Address Inventory Management Best Practices

The following are best practices for IP address inventory management.

Best Practice

Inventory address space in a centralized database

IP address space, both public and private, is a precious asset, one that provides the fundamental entity for network communications. Therefore, it must be tracked in a centralized repository to maintain consistency and accuracy. Of course, accuracy requires updates to the database upon address space allocations, “free-ups”, and ideally utilization per allocated block or subnet.

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

Maintaining inventory database accuracy is crucial. If the IP address inventory only tracks top-down allocations of address space entered by administrators manually, how do you know it’s accurate? Comparing the inventory database with network actuals is crucial to identifying discrepancies and tracking IP management processes. For example, if someone circumvented the conventional update process, whereby the inventory database was not “informed” of the change, the identification of this discrepancy would not only highlight an inventory mismatch, but also bring out this network change control issue. 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 discoveries is an effective way to monitor the process and keep inventory accurate.

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

Network allocations should be made in an optimal manner, maximizing utilization of address space, while mapping to the topology model. Allocating address space along a hierarchical structure that models the routing topology facilitates route aggregation to keep routing overhead to a minimum. If exceptions to the aggregation model are necessary, for whatever reason, they can be made knowingly and routes can be proactively updated to maintain reachability. Since routing topology often maps to an organization’s locations, sites, or business unit hierarchy, this hierarchical modeling of address space typically provides the added benefit of tracking address allocations to these entities. Per application address allocation should also be addressed if appropriate to manage address allocations for deployed IP services, e.g., VoIP 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.

Maintain additional information as appropriate per IP device.

Keeping track of what device is occupying each IP address in a subnet is critical to IP management. However, many such devices have other attributes that should be tracked within an IP management solution. Not the least of these attributes is what other IP addresses the device in question occupies. Many devices have multiple IP addresses, whether for virtual networking, IPv4 and IPv6 addressing, or for multi-homed devices. Multi-homed devices have multiple interface cards, each occupying one or more IP addresses. Beyond this critical IP address information, tracking other attributes, including device type, location, switch port, administrative contact, asset information, and associated resource records to name a few, are equally crucial. It is recommended that 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 reasons. Despite the best planning efforts, IP networks seem to have an organic nature, where address needs rise and fall at different times at various locations within the network. 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 to predict potential future address depletion times. This trending analysis provides another decision criterion in the IP address capacity management process.

Additional criteria for even more proactive management of IP address capacity require the use of alerts for notification of pending address depletions before they happen. Alerts should be programmed for address pools (i.e., dynamics only) or networks (i.e., dynamics and statics) approaching full capacity or being underutilized. This proactive measure can assuage potential pending address depletions, which can render end users unable to communicate due to the lack of addresses.

Keep IPv6 in mind when considering IP address investments, even if IPv6 is on the outskirts of your planning horizon.

If your organization is considering adoption of some or all of the best practices outlined in the section to centralize IP inventory, automate allocation processes, etc., consider your long-term plans for IPv6 and if appropriate, require IPv6 support in any tools you invest in to protect your investment for years to come.

Dynamic IP Address Services Management

Following the IP inventory best practices described above can help maintain adequately sized networks and address pools across the network. But sizing address pools to supply IP addresses to end users, critical as it is, is just the beginning of the process of address assignment via DHCP. After all, you don't want just anyone to get any IP address on your network to access network resources! So there's more to configuring DHCP servers than address pool allocations. Additional configuration elements include valid options and policies with associated values for each address pool, valid or invalid devices by MAC address, client class, or user authentication, device software validation, and DHCP failover configuration.

Policy Management

As many or all DHCP servers will require similar DHCP policies, 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. This practice ensures a consistent and accurate approach to setting these critical policies. Otherwise, you must be concerned with entering basically the same information multiple times into each of your servers. A similar argument can be made for defining DHCP option sets, with defined DHCP options and valid values for use on assignment.

Discriminatory Address Management

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

This type of IP address and configuration assignment is also possible by filtering on the client class of the client requesting an IP address. Certain clients, such as VoIP phones, provide additional information about themselves when requesting an IP address in the vendor class field of the DHCP packet. The user class field may also be used. The DHCP server can be configured to recognize the user classes and/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 option is to completely deny an address; 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. (Those users sometimes pop in new interface cards!) 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.

DHCP Failover

DHCP failover is recommended to provide address services redundancy across your IP network. If a DHCP server crashes, a failover server can take over and begin processing DHCP transactions. This provides a higher availability service for your end user clients requesting IP addresses. If clients cannot get IP addresses, they will be unproductive and will call the help desk! DHCP vendors implement high-availability DHCP in various ways. For example, the Internet Systems Consortium (ISC) DHCP server supports an interserver failover protocol as described above; however, the popular Microsoft DHCP server does not. Instead, Microsoft recommends splitting each scope or pool across two DHCP servers, with one server supporting 80 percent of the addresses in the pool, and the other server supporting the remaining non-overlapping 20 percent.

DHCP Appliances

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

Dynamic IP Address Assignment Management Best Practices

The following are best practices for IP address assignment management.

Best Practice

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 address assignment.

Implement one or more of the following approaches:

- Device identification via MAC address – filter client requests against a list of acceptable and/or unacceptable MAC addresses
 - 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 into the network.

Configure DHCP failover 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 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.
- 2) *Double Scope Approach*: The second mechanism that can be employed when address space is not overly constrained, e.g., 10.0.0.0 space for some users, is to deploy two address pools of the same size, but of different addresses. This “double scope” approach uses two address pools that can serve the same set of clients independently and alleviates the need for inter-server heartbeat communications, while providing sufficient address capacity for the end users requiring addresses.
- 3) *Split Scope Approach*: A third mechanism, particularly when using Microsoft DHCP, is to implement split scopes, where two DHCP servers manage non-overlapping subsets of each address pool.

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 use of DHCP monitoring of address assignments and address pool utilization, including shared pools or shared subnets, to net out the capacity impacts from a pool and pool user perspective.

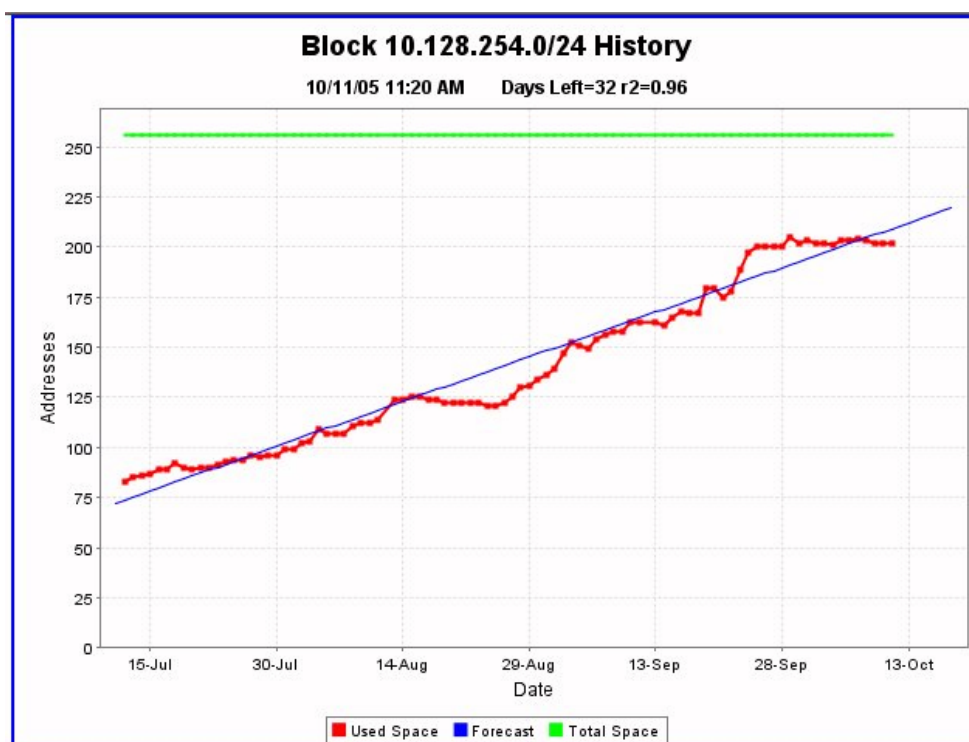
Maintain IP address pool history data to monitor address usage trends and proactively align address space to where it’s needed.

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 at a glance the general utilization trends and enables proactive management of address pools while facilitating proactive realignment of address pool capacity as necessary to prevent address depletions.

Consider DHCP appliances to streamline deployment, improve security, and simplify upgrades.

A DHCP appliance integrates the hardware, operating system and DHCP service into a simple self-contained platform. Appliances are generally more secure, purpose-built platforms with restricted operating system permissions, users, ports, and files. Most appliances are also deployable in a dual high-availability configuration, and should be capable of being monitored, controlled, and patched from centralized system to streamline ongoing server and services management. The centralized monitoring and patching functions provide lights-out support for remote appliance deployments.

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, hopefully all of which happens seamlessly behind the scenes, most end users will immediately 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 universal resource locator (URL) makes email and web browsing easy and user-friendly. 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 applications could function but would be totally unusable for 99% of your company’s end-user population. Needless to say, DNS services must be configured accurately, and be highly available to users.

DNS Resource Records

It’s up to IP address managers to properly configure the DNS servers in the network with the information needed to resolve host names and URLs into IP addresses. This means that not only statically configured IP devices like routers, web servers, email servers and the like need to have entries in DNS, but also dynamically configured IP devices like printers and even end user devices. In many cases, websites perform a *reverse* DNS lookup for an IP address before continuing a web session to validate that the requesting IP address has some form of legitimacy in DNS. This implies integration between DHCP and DNS, referred to as Dynamic DNS, which is an automated process to update DNS upon address assignment by a DHCP server.

Beyond name-to-address translation and vice versa, DNS provides many other “translation” applications, which we won’t go into here. Each translation type maps to one or more resource record types in DNS. For example, an “A” resource record type is used to translate a text-based host name into an IPv4 address. While all resource record types follow the same basic format in terms of fields within the record, the syntax

is not intuitive nor is it easy to identify errors until problems arise. While DNS does provide a mechanism for a master DNS server to update its slaves via a zone transfer, in some cases, it is desirable to operate in a multi-master mode of operation, whereby each master must be updated individually. This opens the door to potential errors in not only resource record configuration but also in other DNS options and directives, of which there are many.

DNS Options

Configuring these DNS options is critical to properly defining the behavior of the DNS server, in terms of zone transfers, security measures, and other operational parameters. Various directives exist in varying forms in different DNS server versions. For example, logging configuration varies between BIND versions 8 and 9. Other vendors' DNS implementations may have other nuances in configuration. Keeping track of the proper syntax for the particular vendor/version you're running may be tedious, but it's absolutely critical to keeping DNS up and running.

DNS Security and Availability

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

- Configure ACLs – configure which IP addresses or networks can query, notify, update, and transfer to or from each name server. In addition, ACLs on the `ndc/rndc` control channel should be defined along with a security key (see next item) for BIND 9 servers.
- Configure transaction signature keys – sign each update and zone transfer with the use of transaction signature keys (TSIG keys). For deployments to Microsoft Active Directory integrated zones requiring secure updates, sign each update using GSS-TSIG.
- Run the DNS service (named) in `chroot-ed` environment – this provides the name server daemon full file system subtree access at a point below the root; otherwise, root access to the file system is provided by default. Ideally, run the DNS server on a dedicated machine to allow restriction of open TCP/UDP ports.
- “Hide” master DNS servers. If attackers find and infiltrate the master DNS server, slaves will zone transfer from this master, spreading the corruption. Hiding the master can be accomplished by editing the standard NS and A (glue) records pointing to the master DNS server to point to a different (slave) server. The master name (“`mname`”) field of the slave's SOA record should also be edited accordingly.
- Consider running DNS appliances (addressed below) – Appliances are purpose-built 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.

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

- 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
- Consider operating internal root servers
- Use a separate network for queries versus zone transfers and updates.

For additional recommendations, please email diamondip@bt.com.

DNS Scalability Challenges

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

Another scalability challenge relates to supporting a common set of resource records across multiple domains. For example, the “www A” record for ins.com may be the same as in ins.com, ins.biz and intnetsvc.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 units built on hardened Linux operating systems, and provide additional security features. This can simplify the procurement process by eliminating the need to coordinate the operating-system patch level with the DNS service version compatibility for initial deployment. Ongoing patch management can be simplified as well with appliances offering centralized patch management features. Many appliances can also be deployed in dual-appliance configurations to provide high availability at the hardware level.

IP Name Services Management Best Practices

The following are best practices for IP name services management.

Best Practice

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

Utilizing a single interface and database to configure a number of DNS servers provides the ability to enter configuration parameters once, and deploy the appropriate master or slave configuration to multiple DNS servers, then aggregate dynamic updates to keep the centralized inventory up-to-date. This provides a centralized view into the overall DNS configuration across your network for DNS servers, domains, zones, and views.

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 reduce cyclic updates.

<p><input checked="" type="checkbox"/> 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.</p>	<p>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. And for BIND DNS implementations, BIND supplies a pair of verification utilities: named-checkconf and named-checkzone.</p>
<p><input checked="" type="checkbox"/> Configure external, internal, and perhaps other "views" of your name space.</p>	<p>This can be accomplished either by configuring 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 9. This provides an external version of externally exposed domains to keep resolvable hostnames to a manageable number, large or small. Meanwhile, a different version of domains can be provided to those querying DNS from internal networks. This simple dual view example can be extrapolated to multiple views, allowing granular configuration of which host names get resolved with what if any, IP address(es).</p>
<p><input checked="" type="checkbox"/> 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.</p>	<p>BIND offers a variety of configurable options that allow specification of ACLs, pair-wise server transaction signatures, and IP address/port specifications. While these options provide the flexibility for configuring these capabilities, the key is to accurately configure each server with its corresponding ACLs, keys, and IP addresses/port numbers. For a large number of servers, this can be cumbersome and error-prone to configure manually. Microsoft provides a different means of signing updates with GSS-TSIG.</p>
<p><input checked="" type="checkbox"/> For large environments, consolidate management of multiple server sets and alias domains</p>	<p>Implementing this best practice generally requires a DNS configuration tool to support grouping of sets of DNS server and alias domains. However, use of such a tool - or better yet, an integrated IP address management solution - can provide the benefits of managing not only multiple DNS server sets and alias domains, but multiple sets of DHCP servers and a diversified IP address space.</p>
<p><input checked="" type="checkbox"/> Consider DNS appliance deployment to streamline deployment, improve security, and simplify upgrades.</p>	<p>A DNS appliance integrates the hardware, operating system, and DNS service into a simple self-contained platform. Appliances are generally more secure, being built with restricted operating-system permissions, users, ports, and files. Most appliances are also deployable in a dual high-availability configuration, and should be capable of being monitored, controlled, and patched from centralized system to streamline ongoing server and services management. The centralized monitoring and patching functions provide lights-out support for remote appliance deployments.</p>
<p><input checked="" type="checkbox"/> In high performance environments, configure caching-only DNS servers to handle large volumes of DNS queries.</p>	<p>Caching-only servers are simply name servers not configured as authoritative for any zones. All queries to caching-only servers result in a lookup in cache with escalation to the DNS root servers as necessary. Over time, these servers build up a substantial cache and can respond directly from cache for those records with "still alive" TTLs.</p>

Overall IP Address Management

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

Centralized Management

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 also makes sense to 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, and BIND server deployments.

Administrator Access Controls

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

High Availability Services

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

DHCP/DNS Services Monitoring

Accurate and timely deployment of DHCP and DNS configurations is certainly a critical aspect of effectively managing the DHCP and DNS environment. However, it's equally important to monitor these services to ensure they are properly functioning. If end users aren't able to 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

For environments with a number of distributed DHCP and DNS servers, application of upgrades and patches can be tedious and error-prone. Generally, each of the servers must be inventoried from a hardware, OS version, and DHCP/DNS service version perspective. Compatibility issues among these elements must be considered during the upgrade planning process. At times, physical presence at the site is required by knowledgeable resources to successfully deploy the upgrade, adding to the cost and time required to perform the upgrade. However, deployment of DHCP/DNS appliances with centralized patch management can remove many if not all of these headaches. Selective upgrades of OS, kernel, and DHCP and/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 methods of managing IP networks vary just as widely. Employing a system that enables entry of custom attributes for topology elements, subnets, devices, DNS domains, and even resource records 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 Workflows

In addition to adapting to business processes from a data-element, attribute perspective, integrating IPAM-related functions into broader workflows can provide further automation and cost-saving benefits. For example, the allocation of an address block to a site would likely require the associated updating of relevant DHCP and DNS server configurations, but would also require addition 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 workflow shortens the overall implementation interval and reduces miscommunication errors as well as duplicate entry errors. With increasing proliferation of IP services, enterprise IP managers typically need to allocate application specific subnets or VLANs, accurately assign IP addresses and associated configuration parameters 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.

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, which is growing in importance. Reports on which administrators performed certain tasks, or who “owned” an IP address at a given time are critical for periodic audits, for troubleshooting or investigations, and even for regulatory requirements such as Sarbanes-Oxley, HIPAA, and CALEA.

Overall IP Address Management Best Practices

The following are best practices for overall IP address management.

Best Practice

<input checked="" type="checkbox"/> 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.
<input checked="" type="checkbox"/> Enable delegation of IPAM responsibility as desired while controlling access to relevant information	Access control is an important consideration when multiple users have access to the IPAM system. While a core set of users will likely require full access to all system functions and features, it is likely that other administrators would receive a limited set of functionality and scope control based on their respective responsibilities.
<input checked="" type="checkbox"/> 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 to provide continuity during disaster recovery. Consider appliances for intrasite hardware level redundancy for critical servers. If using a centralized IPAM system as recommended, ensure it is not in the critical path to proper DHCP and DNS processing. Consider a high availability deployment of the IPAM system, especially if it is in the critical path.
<input checked="" type="checkbox"/> Monitor IP services to proactively manage services availability	Keep track of the status of DHCP and DNS services operating throughout the network via periodic polling or event notification. Enable drill-down into event logs and remotely control services to facilitate trouble diagnosis and resolution.
<input checked="" type="checkbox"/> Streamline IP services upgrades and patches	With appliance-based deployments, one vendor is responsible for not only the DHCP and DNS services version, but also for the appliance operating system and kernel. Staging patches and upgrades on a centralized system with the ability to deploy to remote servers vastly simplifies the coordination, timing, and resource requirements for this otherwise costly and cumbersome process.
<input checked="" type="checkbox"/> Adapt IP management functions to your business processes	Every organizations’ 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 the systems you use to define your addressing topology, device types, and naming policies, as well as attributes on topology nodes, blocks, subnets, devices, and domains. This enables you to manage your IP address space according to your business processes.
<input checked="" type="checkbox"/> Integrate IPAM processes into broader enterprise workflows	In addition to adapting the 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, tracking asset information on devices, and creating trouble tickets.

Enable reporting for addressing status and audit information

Communications across organizational levels can be simplified with intuitive, highly graphical reporting. Filtering information to particular “hot spots” within the network can highlight and convey information that potentially requires escalation. Audit reports are also required to track user accountability and comply with regulatory requirements.

Simplifying Best Practice Implementation with IPControl™

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

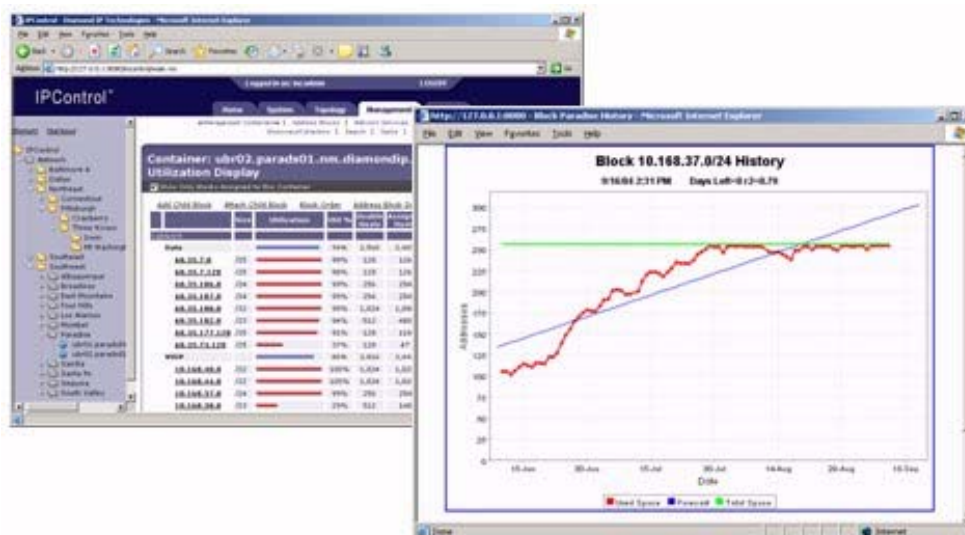
BT Diamond IP’s innovative IPControl™ software and appliance products provide a comprehensive centralized IP management solution for managing IP address space and capacity, as well as DNS and DHCP server configurations. IPControl provides support for sophisticated DNS/DHCP services, including DNS views, logging configuration, transaction signature support, DHCP client classes, and much more. IPControl is available for deployment in a software-only package for installation on customer hardware, an appliance platform for both the centralized management system and DHCP/DNS servers, or a mix of software and appliance deployments.

IP Address Inventory – streamline IP inventory functions

- IPControl provides a centralized IP address inventory database, from which IP space can be consistently assigned and capacity-managed, and DNS and DHCP servers can be configured. 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 helps automate subnet allocations, simplifying the allocation process to a few mouse clicks and reducing 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.
- IPControl enables you to model your network topology via its innovative, patent-pending container feature (Figure 3). Containers allow you to define a hierarchy and track address space allocations in accordance with routing topology to model route aggregation. Actual configuration information can be collected from network devices to enable reconciliation with the database plan versus network actuals.
- IPControl address allocation templates allow you to reserve subnet addresses for each subnet for routers, servers, and other elements common to your subnets, breaking down each pre-allocation by static, reserved, dynamic DHCP, automatic DHCP and manual DHCP address ranges. Automated creation of associated resource records also saves time and effort.
- IPControl provides unparalleled user definability, including user-defined device types. Each device type can have its own attributes 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 supports both IPv4 and IPv6 to enable development of IPv6 address planning, coexistence, and migration strategies, while effectively managing the deployed IPv4 network.

Figure 3: IPControl Innovative Containers and Graphical Interface



Request a demo of IPControl by visiting bt.ins.com/software/ipcontrol-demo.

IP Address Assignment – automate accurate address assignment

- IPControl enables you to centralize your DHCP server configuration to improve configuration accuracy and consistency.
- IPControl provides 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 x10 appliances support DHCP (and/or DNS) services on a secure, easily deployed and upgraded 1u appliance platform. Deployment in a redundant TwinMirror™ configuration provides hardware redundancy and high availability. Centralized monitoring of x10 appliances enable services status monitoring and control, as well as patch management.

Request a demo of IPControl by visiting bt.ins.com/software/ipcontrol-demo.

Name Services Configuration – simplify accurate DNS configuration while enabling advanced features

- IPControl enables you to centralize your DNS server configuration to improve configuration accuracy and consistency.
- IPControl can support nearly any configuration of multiple master/slave DNS configurations from a few servers to several hundred servers. IPControl agents can be deployed throughout your network to facilitate scalable deployments, promoting maximum flexibility and unconstrained DNS server network design.
- IPControl provides a feature to validate your DNS configuration files prior to deploying them to production DNS servers in your network. This enables you 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 is the first software product to support configuration of 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 9. 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 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 galaxies and template/alias domains.
- IPControl Sapphire x10 appliances support DNS (and/or DHCP) services on a secure, easily deployed and upgraded 1u appliance platform. Deployment in a redundant TwinMirror configuration provides hardware redundancy and high availability. Centralized monitoring of x10 appliances enable services status monitoring and control, as well as patch management.
- Configuring caching only 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.

Request a demo of IPControl by visiting bt.ins.com/software/ipcontrol-demo.

Overall IP Address Management – Bring it all Together

- IPControl leverages a 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/slaves), multiple DHCP servers for one-to-one or many-to-one failover, and split-scope DHCP. In addition, the IPControl Sapphire x10 appliance may be deployed in a TwinMirror configuration to support colocated hardware redundancy. While not critical to processing DHCP and DNS requests, the IPControl Sapphire EX appliance housing the centralized management features can optionally be configured in a primary/backup disaster-recovery configuration.

- The appliance dashboard feature provides a summary display of critical IP services running on x10 appliances, including status of DHCP, DNS, and IPControl services. Drill-down to detailed events and the ability to manage the services' states and patch levels provide further control from the centralized IPControl system.
- IPControl simplifies the otherwise tedious process of upgrading remote DHCP and DNS servers. The x10 appliances can be fully patched, upgraded, and rolled back from the centralized IPControl system.
- 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. In addition, the innovative Callout Manger Service can trigger actions or updates to external systems based on IPControl-initiated actions.
- IPControl's highly graphical utilization reports facilitate at-a-glance comprehension of IP address status. Extensive audit reports enable tracking down of administrator actions and IP address "ownership" status over time for auditing and regulatory compliance.

Request a demo of IPControl by visiting bt.ins.com/software/ipcontrol-demo.

Key IPControl Differentiators

Previous generation products were of great assistance for managing IP address space at the time they were introduced, that is, when IP networks were relatively unsophisticated. However, IP networks have evolved to support more hierarchical topologies and multimedia services beyond best-effort data, and in the level of sophistication of DHCP and DNS technologies they require. These previous generation products have not kept pace with the way you need to manage the evolved IP networks of today. IPControl was developed with this evolved IP network in mind. IPControl provides the following unique features, which enable incomparable next-generation IP address management:

- Hierarchical, multi-tiered, centralized IP address inventory
 - IPControl is the only product that enables modeling of multi-tiered, hierarchical IP networks in its centralized inventory. This is enabled by its patent-pending container structure.
 - Block types enable delineation of address space by application, by administrative domain, or other user definable partition.
- Simplified, multivendor DNS/DHCP configuration
 - While many other tools support centralized or distributed configuration of DNS servers, IPControl is the only product today that supports BIND 9 views, TSIG, DDNS, controls, logging, option/server templates, and much more—all within the GUI interface. This simplifies and improves the accuracy of complex DNS configurations.
 - IPControl also provides a unique DHCP policy set 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
 - Only 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 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.
- Most flexible deployment options
 - IPControl’s functionality is available as software-installable on your hardware or as an appliance platform. This applies not only for DHCP and DNS servers you deploy throughout your network, but for the centralized IPControl Executive server as well. So whether you’d prefer an all-software deployment approach, an all-appliance approach, or a mix of the two, IPControl delivers. The IPControl Sapphire x10 appliance provides DHCP and DNS services, while the Sapphire EX appliance supports centralized management functionality.
- 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.
- Integrated IPv4 and IPv6
 - IPControl is the only IP address management product that supports both IPv4 and IPv6 in one integrated product. This enables you to plan out and either experiment with IPv6 in a non-production environment or fully initiate and complete a migration from IPv4 to IPv6 over time. IPv6 is coming, so any IP management software investment you make today should incorporate this “next generation” IP.
- Affordability
 - With this superior feature set, you might expect to pay more for IPControl than you would for previous generation products. But next generation thinking goes beyond product features. IPControl provides exceptional value by providing these key differentiating features for next generation IP address management, at less than half the price of comparable competing solutions.

Request a demo of IPControl by visiting bt.ins.com/software/ipcontrol-demo.

Conclusion

Leveraging expertise gained through years of experience working with customers, industry analysts, and various software implementations, this white paper recommends numerous best practices for effectively managing your IP address space. Building on this experience and these recommendations, BT Diamond IP has developed IPControl software and appliance products to simplify the fulfillment of these recommendations.

BT Diamond IP IPControl products provide the advanced, next-generation IP management solution you need to automate many tedious, error-prone, yet critical IP management functions. IPControl 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 diamondip@bt.com to learn more about how IPControl can automate more of the IP management functions you need at an exceptional ROI.

About BT Diamond IP

BT Diamond IP is a leading provider of software and appliance products that help customers effectively manage complex IP networks. Our next-generation IP management solutions help businesses more efficiently manage IP address space across mid-to-very large sized enterprise and service provider networks. These products include IPControl™ for comprehensive IP address management, Sapphire Appliances for DNS/DHCP services deployment and NetControl™ for full-cycle IP address block management and utilization. Our cable firmware management product, ImageControl™, helps broadband cable operators automate and simplify the process of upgrading and maintaining firmware on DOCSIS devices in the field. Our customers include regional, national and global service providers and enterprises in all major industries. For additional information, please visit bt.diamondip.com or contact BT Diamond IP at 1-800-390-6295 in the U.S. or 1-610-423-4770 worldwide.

IPControl and TwinMirror are trademarks of BT INS, Inc.

Copyright © 2007, BT INS, Inc.

This is an unpublished work protected under the copyright laws.
All trademarks and registered trademarks are properties of their respective holders.
All rights reserved.