Product Brief

High Availability DHCP, DNS, IPAM (DDI)

by Timothy Rooney

Product management director

BT Diamond IP

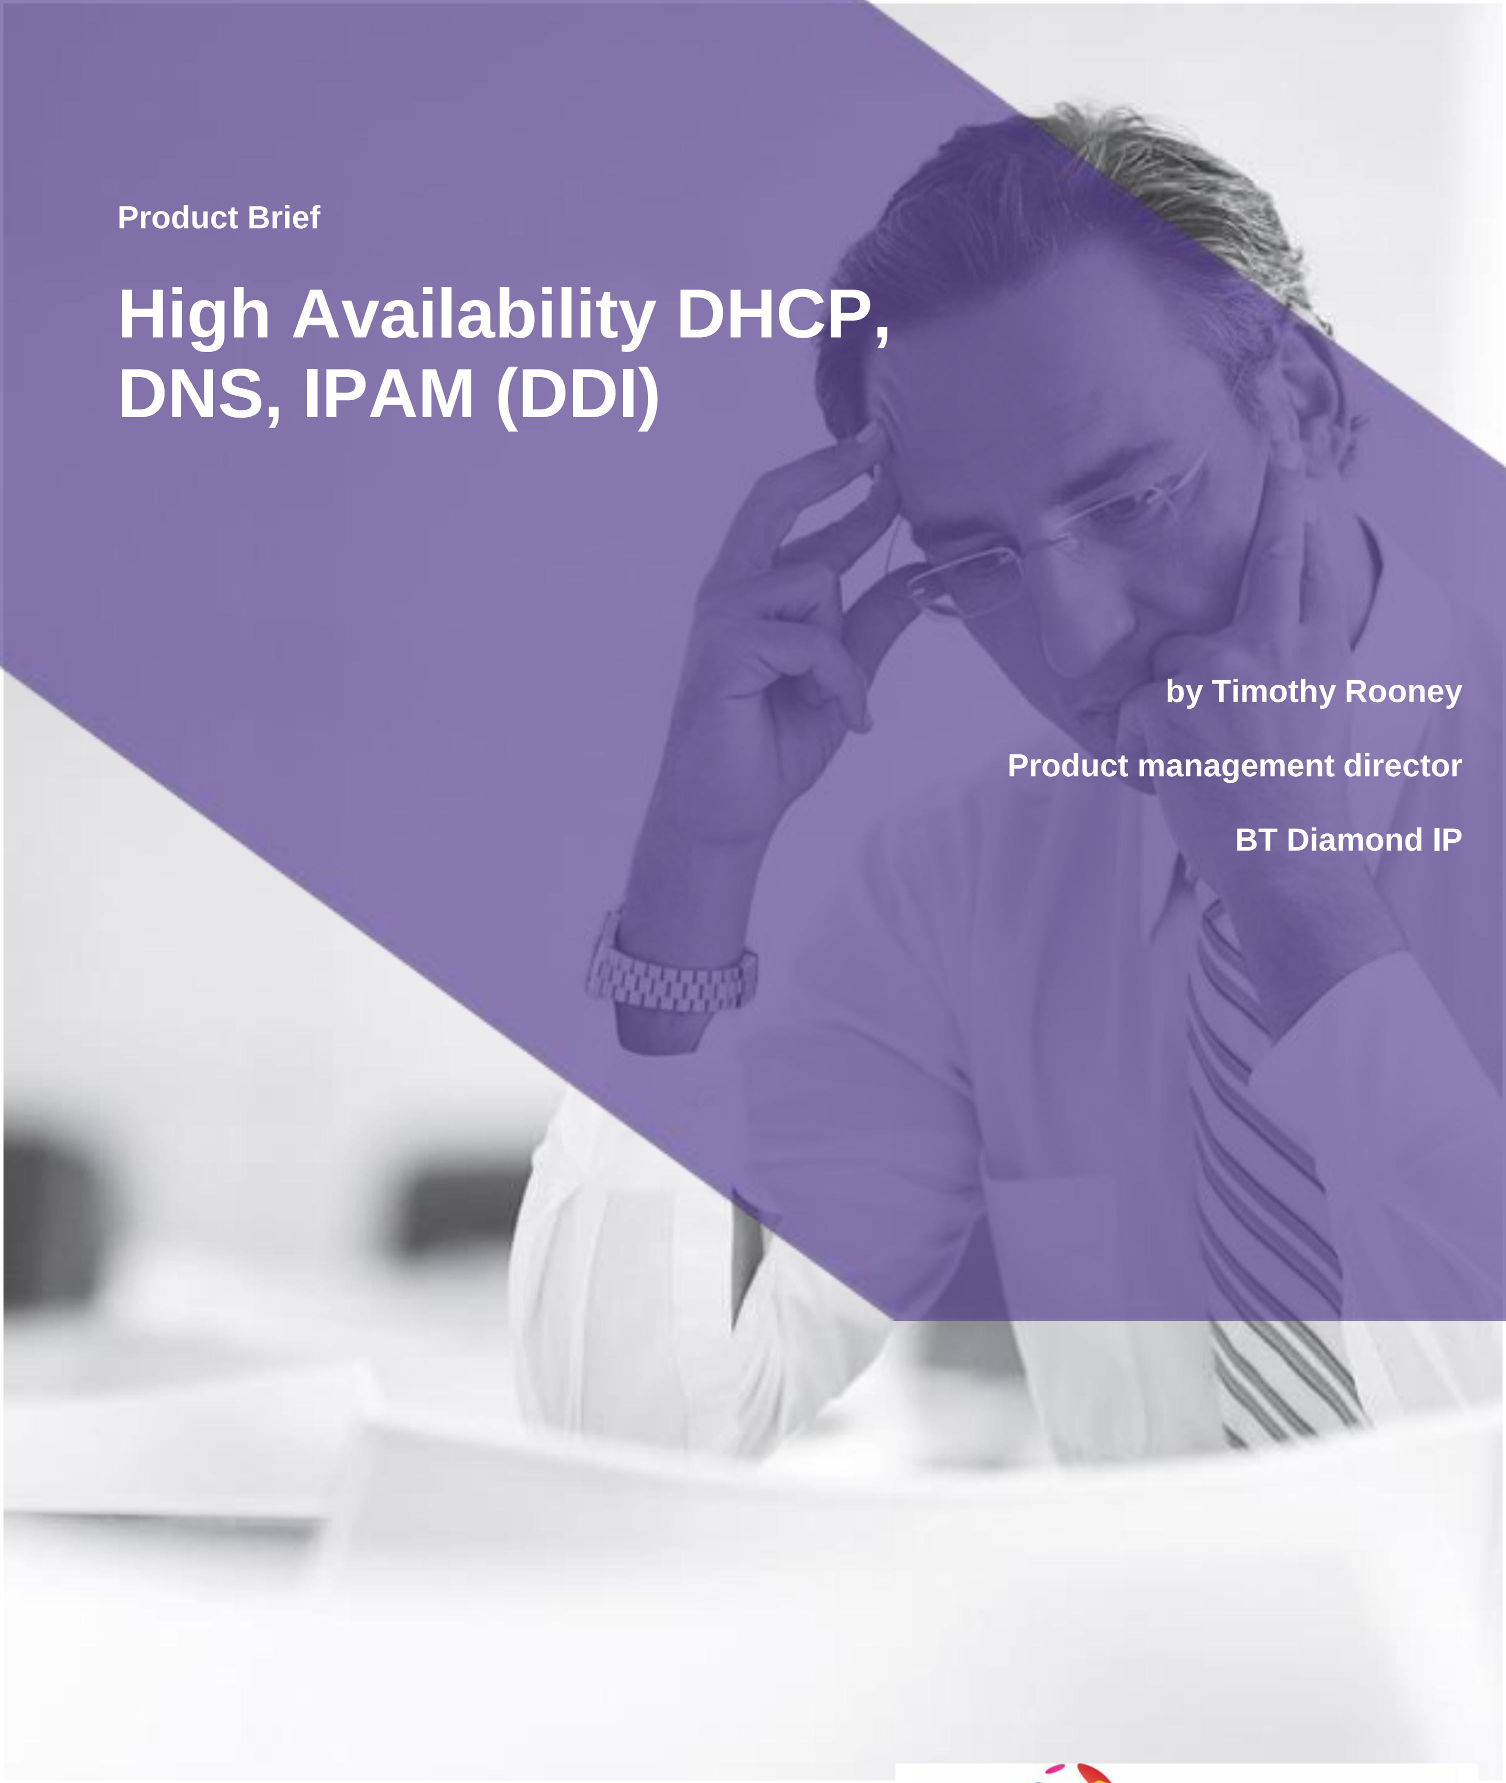

High Availability DHCP, DNS, IPAM (DDI)

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

Introduction

BT Diamond IP's IPControl™ Software provides comprehensive IP address management (IPAM) functionality, including IP address inventory, Domain Name System (DNS) and Dynamic Host Configuration Protocol (DHCP) servers configuration and management. DNS and DHCP services are critical to the operation of any IP network. Without the availability of proper DHCP services, IP hosts or endpoints attempting to obtain a valid IP address via DHCP will be unable to do so. The person using the IP device will be unable to access the network for email, web applications, or other data, voice or multimedia IP-based applications. The net effect will likely be calls to the help desk reporting the "outage" and the resultant impacts on IT or Operations to diagnose and resolve the trouble report. Likewise, without the availability of proper DNS services, users of IP devices attempting to access by "name" will be unable to do so, with similar resultant impacts on the help desk and IT/Operations.

As suggested above, availability of DNS/DHCP services is predicated not only on "availability" in terms of getting a response for an IP address for DHCP and name resolution for DNS, but on "availability of proper DHCP/DNS services" to assure the response given to the client is accurate and appropriate for that client. The availability of "proper" services is greatly facilitated with the use of IPControl inherently; its graphical interface and support for advanced DHCP/DNS features enables the accurate configuration of DNS and DHCP services. This proper configuration is a topic of another paper, namely our *IPAM Best Practices* white paper available on btdiamondip.com. For the purposes of this paper, we will focus on maximizing the availability of IPAM, DHCP and DNS services.

IPControl Architecture

IPControl was designed from the ground up for scalability and high availability. The basic architecture is illustrated in Figure 1. Starting at the bottom of the figure, this architecture promotes the deployment of redundant components for critical DHCP and DNS services, supporting multiple master/slave DNS configurations and multiple DHCP failover configurations. This allows clients on the enterprise or commercial network to access multiple servers as necessary for DHCP and DNS services. The DHCP and DNS servers communicate with the centralized IPControl Executive and database services via lightweight IPControl Agents installed on the servers. These agents provide the secure reliable communications link to the IPControl Executive over an internal or management network.

The IPControl Executive provides a number of services that provide various system functions for scheduling and executing tasks within the system such as configuring DNS/DHCP servers, collecting information from these servers, handling system alerts, logging, and more. The user interface for performing these functions and others within IPControl is a web browser via provided web server, command line interface (CLI) or web services (XML/HTTPS) API interface.

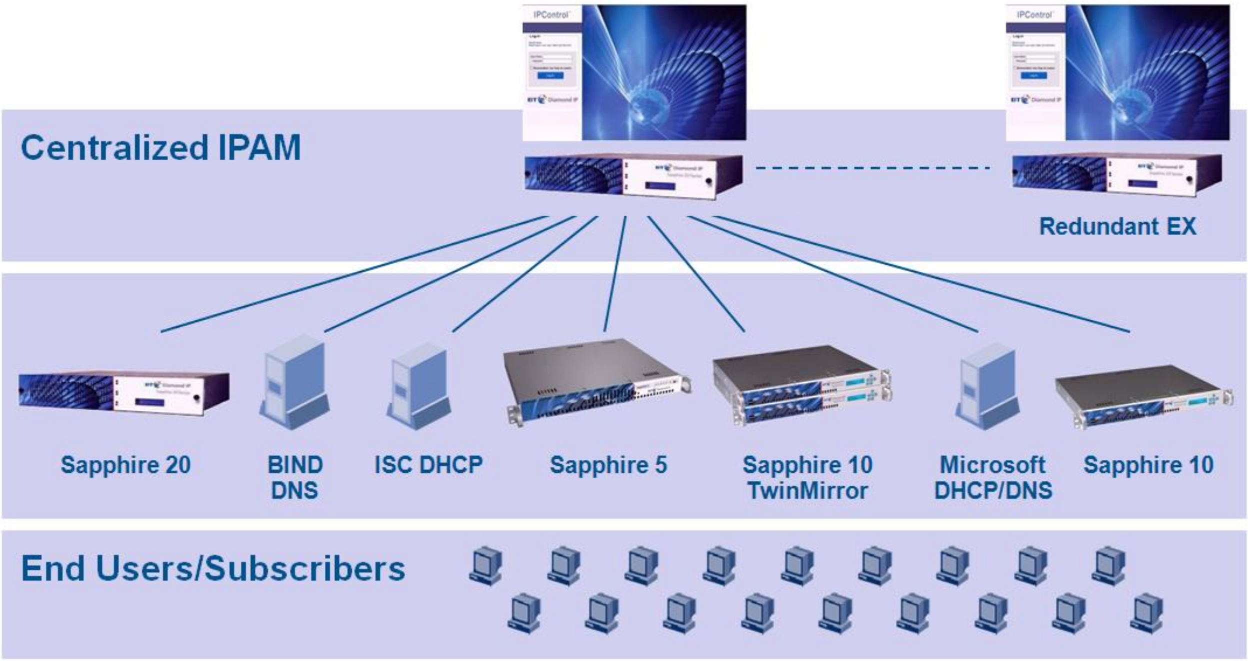

Figure 1: Basic IPControl Architecture Managing Distributed DHCP/DNS Services

IPControl High Availability Components

DHCP and DNS servers can be deployed in a redundant configuration and operate autonomously to provide high availability name resolution and address assignment services to clients. IPControl agents, which provide the communications interface to the centralized IPControl Executive service, can also be deployed in redundant, load balancing configurations. Redundancy can also be implemented within the database and web services. Further details on each component are provided in this section.

Network Services

The network services, or DHCP and DNS servers, provide the critical name resolution and address assignment functions of DNS and DHCP respectively. End users or clients interact directly with these services; therefore high availability at this level is of utmost importance. IPControl enables deployment of these services in a variety of redundant configurations.

DNS

IPControl supports deployment of resource record information to multiple DNS servers, either single or multiple BIND masters and Microsoft Windows Active Directory (AD) multi-master environments. DNS clients (specifically resolvers) must be configured with an IP address, preferably multiple IP addresses, pointing to multiple DNS servers respectively for use for resolving DNS queries. When an application such as a web browser or email client requires name resolution, it invokes the resolver to perform the resolution. The resolver issues an appropriate DNS query to one of these configured DNS server IP addresses. Should the server be unavailable after one or more attempts, the resolver would then issue the query to another DNS server assuming more than one IP address is configured. This resolution attempt to multiple DNS servers occurs without user intervention or even knowledge in most cases. Therefore, the deployment of multiple authoritative DNS servers and associated resolver configuration pointing to them promotes high availability DNS services.

Within DNS, a number of DNS servers can be deployed with authoritative information for a particular zone; for BIND, a single master server is configured and slave servers obtain their zone configuration from the master via zone transfers. The master and all slaves are authoritative for the respective DNS zone data. In a Microsoft Windows AD integrated DNS environment, all DNS servers are masters and authoritative, so updating a particular master results in replication of that zone data to other zone masters via LDAP. IPControl supports either model for configuring BIND based and/or AD based DNS servers to provide high availability DNS services for your clients. These configurations are depicted in Figure 2.

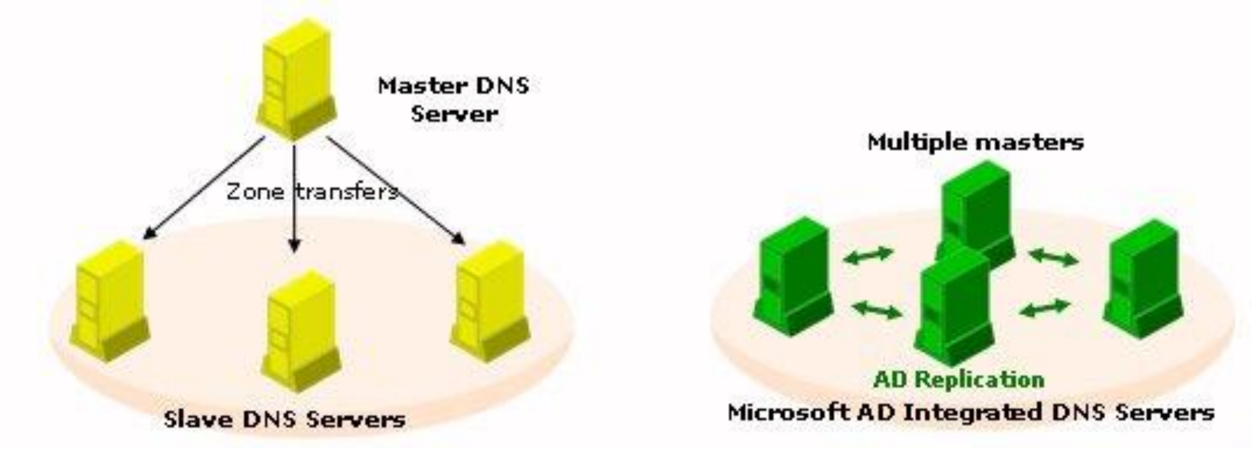

Figure 2: IPControl supports both BIND based and Windows AD High Availability Configurations

The resolver itself can be configured with DNS server addresses directly on the client (IP configuration) or more conveniently and portably via a DHCP server using the "Domain Name Servers" option (DHCP option 6). Using DHCP, when a client accesses a network and utilizes DHCP to obtain an IP address, the client can be configured with its IP address as well as the IP addresses of DNS servers to use corresponding to its point of access or other criteria. IPControl also eases the configuration of these DHCP options to save time and reduce errors in its configuration.

DHCP

From a DHCP perspective, IPControl supports deployment of one-to-one or many-to-one DHCP failover configurations. IPControl greatly simplifies the configuration of failover servers. All that's required is specification of the failover partner server for a primary server along with a few configurable parameters regarding load balancing and failover initiation. IPControl takes care of the rest in building the configuration files for the primary(ies) and failover server and deploys them to the respective servers at your command. In the one-to-one model, shown on the left side of Figure 4, all of the scopes defined on the primary DHCP server are also configured on the failover server. With both DHCP server addresses configured in associated routers[1], both DHCP servers will receive all DHCP packets (e.g., Discover and

[1] The router (or generically, relay agent) is configured with the IP address(es) of the DHCP server to which to

Request packets) per Figure 3. The primary would process the request on a normal basis and communicate lease bindings to the failover as they are created. The primary and failover exchange heartbeat messages to verify that the DHCP server is active. In the event that the failover server detects that the primary server is inactive based on parmeters you can define, the failover server would then begin processing DHCP packets, assuming the role of primary. Given that the failover server had received the lease bindings all along, it has an accurate view of current leases.

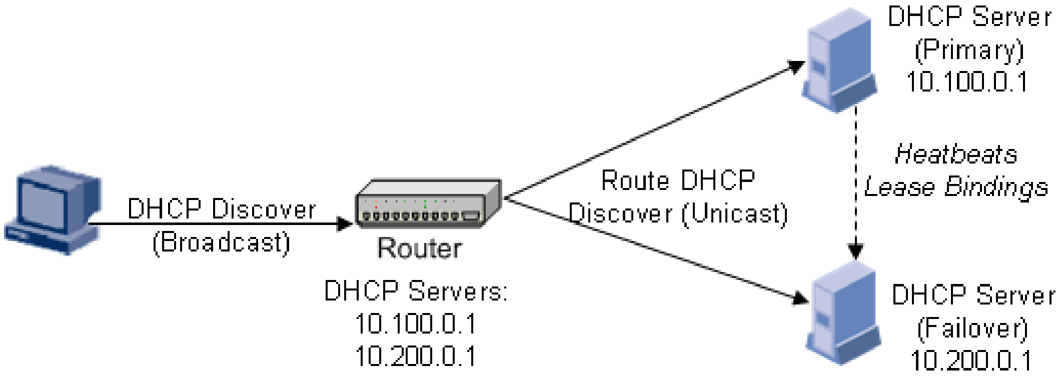

Figure 3: DHCP Helper Addresses for Primary and Failover

When the primary server reactivates, it can retrieve the latest active lease information from the failover server, then once again begin performing functions of the primary DHCP server. Since any intervening router is transmitting any DHCP packets to both DHCP servers due to its DHCP Relay configuration discussed above, there is no change required on the DHCP client to function within a failover configuration.

It may be desirable to utilize a single failover server for multiple primary servers to conserve hardware. This configuration is supported with IPControl as well. For each of the primary servers, a common failover server can be associated with them simply by identifying it. This configuration is illustrated on the right side of Figure 4.

"relay" or unicast the DHCP broadcasts that it receives. While Cisco refers to this configuration parameter as the DHCP Helper address, some vendors refer to this as the DHCP Relay address.

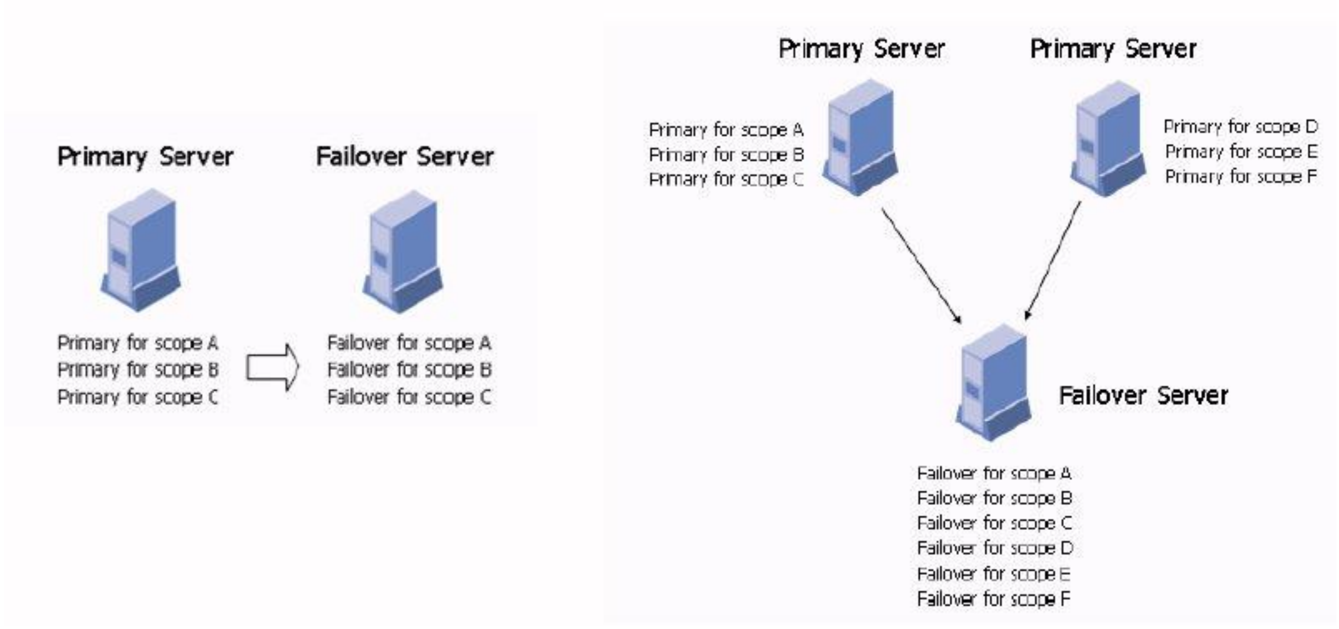

Figure 4: One-to-one and many-to-one DHCP Failover configurations

IPControl Sapphire Hardware Redundancy

On top of these multi-server deployments for DHCP and DNS, which can be deployed across multiple sites, the IPControl Sapphire DHCP/DNS appliance provides an additional level of redundancy. IPControl Sapphire appliances can be deployed in a TwinMirror™ configuration (Figure 5), which features two co-located appliances interconnected via a high-speed interface. In this configuration, the appliances mirrors OS and server-level updates to maintain synchronization. The TwinMirror servers appear to DNS and DHCP clients as one server, with one IP address. However, both appliances are independently addressable from IPControl for monitoring and configuration purposes. In this active-standby configuration, when a failure is detected, the standby unit assumes the active role, providing seamless DHCP and DNS services to clients. Note that in the figures in this paper, the TwinMirror configuration can replace any BT Diamond IP DHCP/DNS server icon.

Figure 5: IPControl Sapphire TwinMirror™ configuration provides an added layer of HA

IPControl Executive

The IPControl Executive supports the scheduling and execution of tasks and database updates. Configuration changes, data collection, and dynamic updates for IPControl-initiated devices (e.g., creation of static devices) are under control of the Executive. While these tasks are important to the ongoing management of the IP network and associated services, the Executive is not critical to the ongoing operation and performance of DNS and DHCP servers continuing to operate in the network. In a sense, the DHCP/DNS servers are the "network elements" providing primary name and address services to end users, while the Executive is the "management tool" for configuring and managing these network elements.

When communications are lost from a DHCP or DNS server to the Executive because the Executive or links are down, the critical DHCP and DNS functions are still performed autonomously by the DHCP and DNS servers serving clients in the network. End users or clients would not be affected by an outage of the Executive. In addition, dynamic DNS updates and any task in progress results ready for reporting back to the Executive will be queued at the Agent for transmission upon availability of the Executive.

Regardless, the Executive can be operated in an active-standby configuration with two Executives deployed at different locations for disaster recovery. The secondary location's Executive services would need to be configured similarly to the primary. While a single database could be tapped by either Executive, it makes more sense to have a secondary database at the backup location as described in the next section. Similarly, redundant web servers may be deployed locally or across different locations for load balancing and/or redundancy for administrator access. This is discussed further in that section.

IPControl Database

The IPControl database is a MySQL relational database, and customers may also utilize customer-provided Oracle RDBMS. These databases provide various forms of backup functions. MySQL provides the ability to dump the database (which can be "cron'd") to periodically save a snapshot of the database for archiving and external storage. When the primary database fails, the backup copy can be invoked on the backup database to initialize it to the currency of the last primary database dump.

For tighter granularity of primary/backup database information, MySQL also supports one-way master/slave data replication capabilities. All updates must be made on the master database, which tracks all changes in a binary log. Once the primary and slave databases are initialized to equivalent configurations, the binary log on the master can be retrieved by the slave(s) to enable the application of these 'deltas' to the backup database. In addition, backups may be made of the secondary database to minimize performance impacts on the primary database if desired.

Whether sending nightly database dumps or data replication to a second database at a remote location, this functionality can be enabled to allow the backup database to be reasonably synchronized with the primary. Oracle databases provide similar capabilities for data backups and replication as well though configuration of Oracle replication is left entirely to the customer.

IPControl Administrator Interfaces

The IPControl administrator graphical interface is purely a web browser interface. There is no IPControl-specific software or applets required on the client. Administrators can access IPControl via a URL that you can define (in DNS) or directly via its IP address. Multiple instances of the web server provided with IPControl can be run on multiple servers in diverse locations to provide multiple IP address reachability. In addition, a set of web servers can be front-ended with a customer-provided load balancer to provide single entry to multiple web servers accessing an IPControl Executive.

Should there be a need to reach a backup IPControl Executive, administrators would point their browsers to a different URL/IP address associated with the backup Executive. If accessing by URL, the DNS entry for the URL (A record) can be changed to reflect the standby Executive's web server IP address. This record change may also incorporate a relatively short TTL (time to live) parameter so the client will re-query and re-cache the address periodically to capture the change back to the primary server upon restoration. In some cases, it may be desirable to use the IP address explicitly, so it's obvious that administrators are accessing the backup Executive web server by its unique IP address.

Note that multiple web servers can also be employed within the same or different locations to solely provide web server backup as well, if both point to the same IPControl database. The scenarios highlighted above for switching the DNS record for URL-based access or having administrators use the backup web server's explicit IP address would apply. The scenario is not explicitly highlighted in the following section describing an example high availability configuration and process, but it is noted here for completeness.

Example High Availability Configuration and Operation

This section presents an example IPControl deployment configuration to support backup and high availability of IPControl components. The base configuration consists of multiple DHCP servers with failover enabled and multiple DNS servers authoritative for each zone. In addition a primary IPControl Executive, web server and database is installed for normal operation at the primary site and a backup Executive with backup web server and database is deployed at a backup site. This base configuration is illustrated in Figure 6. Note that the dotted-dashed line from the primary site to the backup site in the figure generically represents the selected database backup process to either replicate or provide periodic database dumps to the backup database.

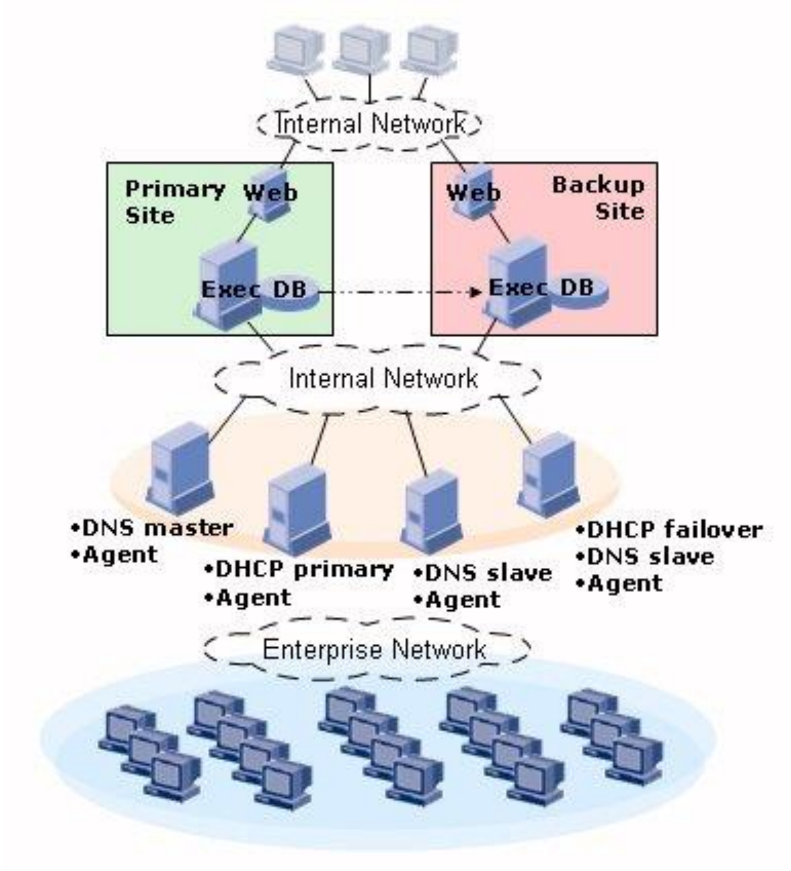

Figure 6: Base High Availability Deployment

In terms of connectivity among the ‘layers’ of the architecture, the end users or clients utilize the normal IT or service provider network to access DNS and DHCP servers, termed “Enterprise Network” in Figure 6. Communications among DNS servers for updates and zone transfers as well as among DHCP servers for failover heartbeats and dynamic DNS updates to DNS servers may also utilize the enterprise network or a separate “internal” or management network. Likewise, communications to/from the DNS/DHCP server agents from/to the Executive may utilize this internal network. These agents would be configured to point to the primary Executive under normal circumstances. The internal or management network would also be recommended for database replication/backups and for administrator access via web browsers, though SSL connections over public networks may also be utilized for browser access.

Scenario 1: DNS and/or DHCP Server Failure

This outage scenario is potentially the most detrimental as end users rely on these servers for configuration and name resolution. Fortunately, as we have discussed previously, IPControl facilitates the deployment of multiple authoritative DNS servers, and client resolvers can be configured to point to multiple DNS servers via manual configuration or via DHCP. Likewise, IPControl simplifies deployment of multiple DHCP servers in various failover configurations to provide seamless backup of DHCP services. In either the DNS or DHCP server failure scenario, no end user or client interaction is required, nor would any outage likely even be noticed. This keeps end users productive and minimizes help desk calls and subsequent

Operations/IT trouble resolution resources. Please refer to Figure 7. IPControl Executive communications pending to or from the failed server agent would be queued for subsequent delivery upon restoration of the server.

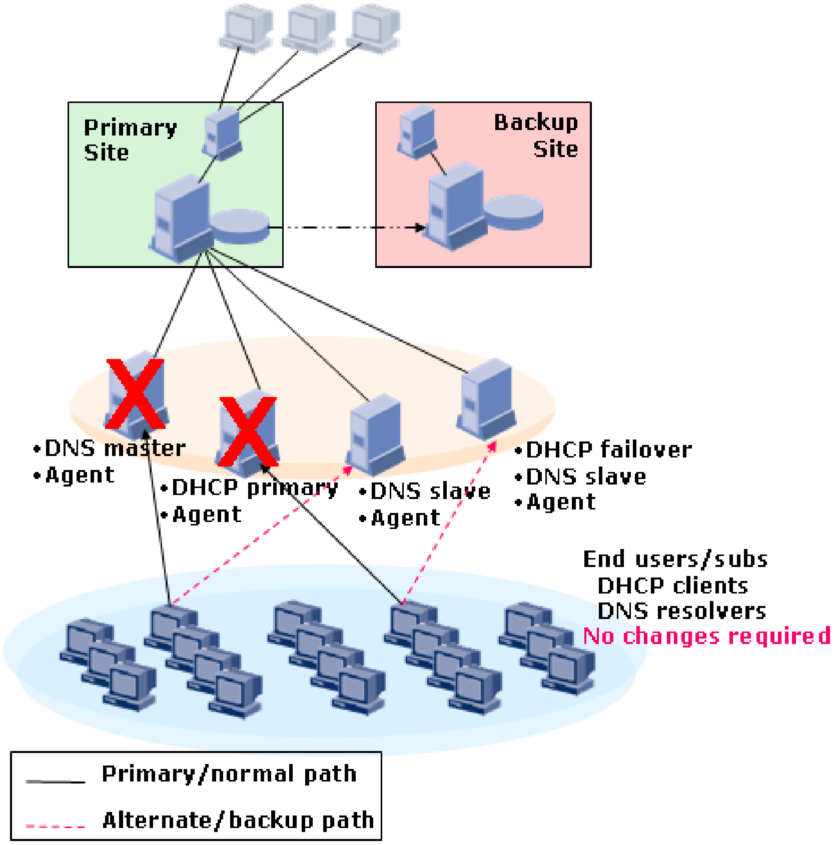

Figure 7: DHCP Failover and DNS Redundancy Maintains Critical Services' Availability

Note that Figure 7 is greatly simplified to illustrate the backup processes. Note that multiple servers and even sets of servers can be deployed in a network based on overall network design objectives to meet customers' specific performance and redundancy vs. cost requirements. In any case, this simple example can be extended for more sophisticated DHCP/DNS server deployments.

Scenario 2: IPControl Executive Failure

In the event of an extensive IPControl Executive failure, the backup Executive could be activated. Note that the critical DHCP and DNS servers will continue operating autonomously serving end users. Hence for intermittent outages, invocation of the backup strategy may not be necessary. However, assuming the disaster recovery plan is activated and assuming replication or periodic backups of the database have been performed, the backup database would need to be activated as the master. Note that MySQL also supports chained master/slaves for replication so a secondary backup server could also be deployed for additional redundancy and to facilitate restoration of the primary site's database when it returns to service.

All IPControl tasks, results and configuration per the latest backup or replication would be accessible using the backup Executive accessing the backup database.   In order to "re-point" the IPControl agents to the new Executive, a script can be run on each server to update the Executive IP address in each agent's configuration. This script can be kicked off via manual command line, a single script for all servers, or via even a higher layer management tool such as HP-Openview. Once this process has completed, tasks can be sent and results collected to/from the agents respectively from the backup Executive.

Administrators requiring access to the backup Executive would need to login. They can either continuing to use the IPControl URL used internally, after the DNS record for the URL has been updated, or they can use the backup IPControl web server's IP address as described previously.

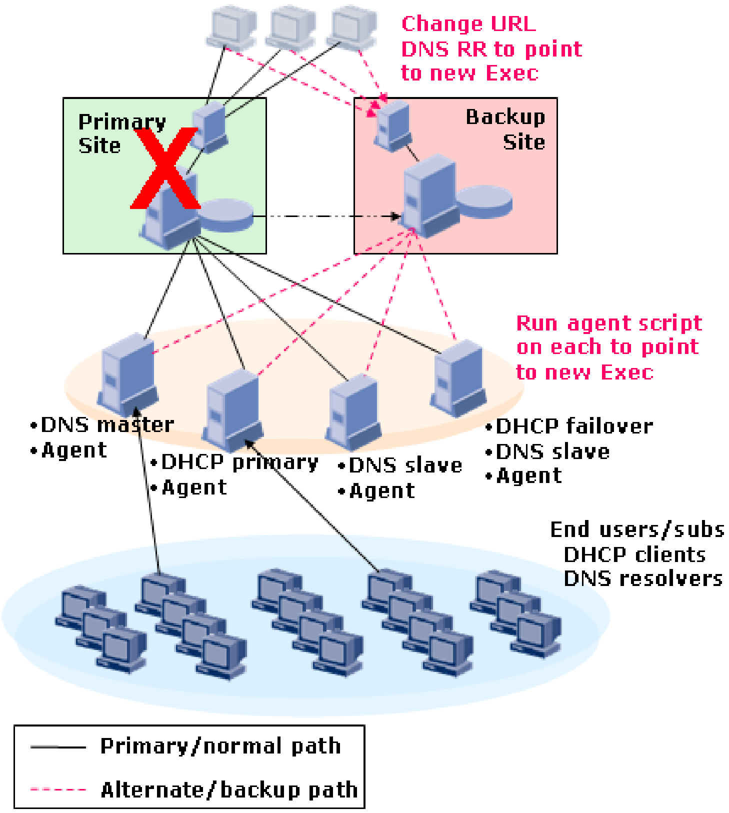

Figure 8: IPControl Executive Failure Scenario

Conclusion

BT Diamond IP IPControl software provides an advanced next generation IP management solution that enables you to provide high availability DNS and DHCP services to your end users. IPControl also supports additional high availability capabilities for full disaster recovery planning and to keep this valuable management tool up and running. IPControl provides many features to automate many tedious, error-prone, yet crucial IP management functions across the entire life cycle of an IP address. InControl provides unsurpassed extensibility and user-definability to enable management of IP address space the way customers want to manage it, all at an affordable price. Figure 8 highlights the redundancy capabilities at each level of the IPControl architecture.

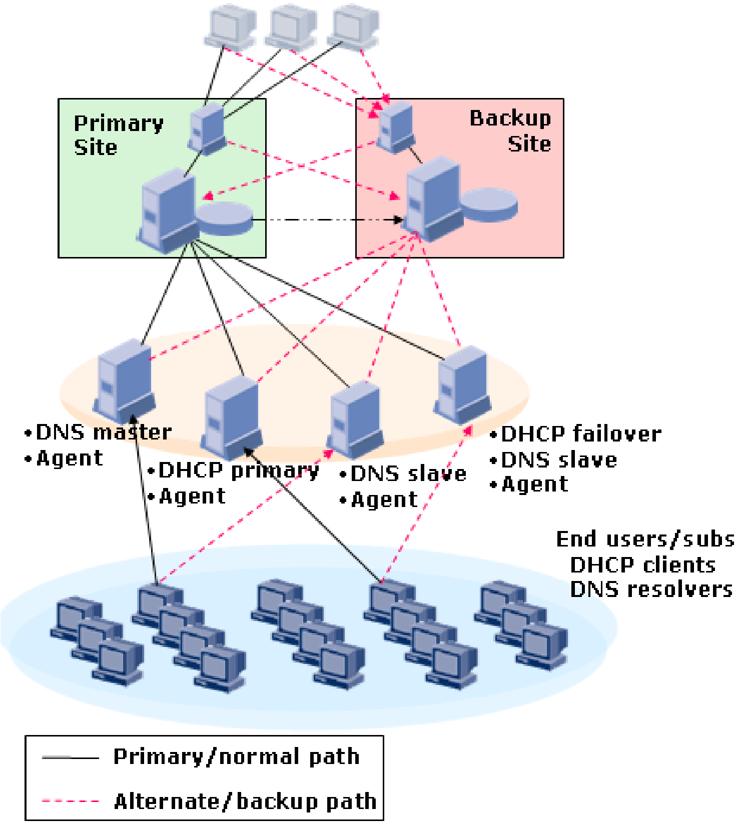

Figure 9: High Availability Summary

Please email us at diamondip@bt.com to learn more about how IPControl products 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 and services 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 and Sapphire Appliances for DNS/DHCP services deployment. 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 more information, please contact us directly at +1-610-321-9000 worldwide, email to btdiamondip-sales@bt.com or consult www.diamondipam.com.

IPControl and ImageControl are trademarks of BT Americas, Inc.