Infoblox®
CONTROL YOUR NETWORK

WHITEPAPER

DNS Appliance Architecture: Domain Name System Best Practices

A Practical Look at Deploying DNS Appliances in the Network to Increase Simplicity, Security & Scalability

Cricket Liu, Chief Infrastructure Officer

In the first white paper in this series, *The DNS Appliance Imperative,* I argued that the evolution of the Domain Name System demanded that it move to an appliance platform. DNS is now critical to the operation of nearly every non-trivial networked application. DNS has also become dauntingly complex, both in theory and in its implementations. And unfortunately, hackers increasingly target DNS infrastructure. This confluence of factors isn't unprecedented—it's the same combination that drove IP routing, network storage, and firewalls to appliance implementations, because only appliances could deliver the requisite simplicity of management, reliability, and security.

But how are appliances deployed effectively in the real-world? In this case study, let's take a look at a corporate DNS architecture based on industry best practices. The design will consider the company's requirements for availability, security, and disaster recovery; network topology; administrative staff and structure; other protocols, including DHCP; and the need to support Active Directory.

Over the course of the paper, we'll examine how the features of appliances enhance yet simplify this DNS infrastructure, making it more secure, more reliable, and less costly to manage and maintain.

While reading this paper, consider your own DNS infrastructure. Does it provide the same level of resiliency and security? Could the use of appliances enhance your infrastructure? Are there aspects of this design that you could adapt to your own use?

The Company

We were recently asked to design a new DNS architecture for a multinational company. This company has three classes of sites connected to its corporate network: Its headquarters site, which houses the largest number of users (about 2000); regional offices, which each accommodate up to 200 local users; and branch offices, which usually support roughly a dozen users each. Branch offices are divided into geographical regions, each supported by a regional office. Each branch office is connected to the network via a link to the regional office that supports it. Only the headquarters site has a connection to the Internet.

The headquarters site supports all of the company's critical business applications, used by employees throughout the company. The majority of the site's users have Windows-based desktops, which are part of the company's single Active Directory domain. The Active Directory domain has the same name as the company's parent forward-mapping zone, which we'll refer to as 'company.com'. The company's executive staff, notoriously intolerant of downtime or delays, is also based at headquarters.

The regional offices host administrative users of the headquarters-based business applications. These users run Windows-based desktops, which belong to the company's Active Directory domain.

Branch offices support up to a dozen users each. They use a mix of Windows-based desktop and notebook computers to access headquarters-based business applications. These computers are also members of the company's Active Directory domain.

The company's IT support staff is concentrated at headquarters. Each regional office has one IT staffer, mostly to provide local support for Windows and to act as remote "eyes and hands" for the headquarters staff.

External DNS Infrastructure

Let's begin with the company's external DNS infrastructure. The company has more than 100 domain names registered through various registrars. Of these, only one, company.com, is widely used (in email addresses and for the company's web properties, for example). The others are mainly registered to protect the company's trade names and to map "www.tradename.com" to the address of the company's web server.

To support the resolution of domain names in the company's external namespace, we recommended a system of four name servers: a primary and three secondaries.

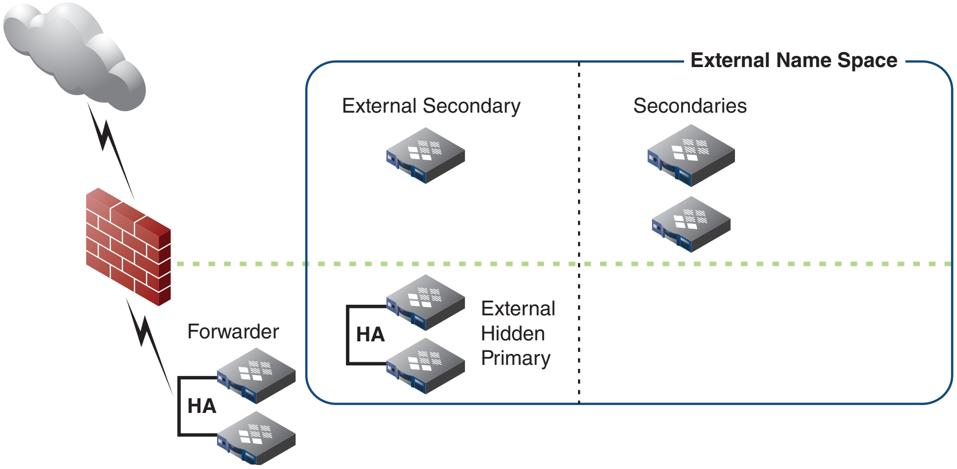

Primary

The primary is behind the company's firewall, and therefore protected from the Internet. Firewall rules permit the secondary name servers to query and transfer zones from the primary, but no other external traffic is allowed. Access control lists on the name server duplicate these restrictions, providing an additional layer of protection. The primary is "hidden," meaning that it doesn't appear in the NS records for the external zones it hosts, nor in the NS records delegating these zones from their parents. This prevents name servers on the Internet from attempting to query it (which would be fruitless, anyway, because of the firewall- and name server-based access controls).

The primary is also a high availability pair. This provides redundancy in the event of hardware failure, to ensure that the company always has a "seat of administration" for the external namespace (a point from which to make changes to zone data). While implementing a high availability pair on a general-purpose operating system such as Unix or Windows is decidedly non-trivial, appliances make it easy.

Secondaries

Two of the three secondaries are on the company's DMZ network, sandwiched between an external and an internal firewall. The other is at a co-location facility with multiple, high-bandwidth connections to the Internet.

The secondary at the co-location facility provides an offsite source of name resolution for the company's customers and partners should the company's connection to the Internet fail. Moreover, even when the company's link to the Internet is up, the colo-based secondary will receive most queries for data in the company's external zones. This is because the round-trip time between most name servers on the Internet and the centrally located, colo-based secondary is lower than their round-trip time to the DMZ-based secondaries. This both minimizes the use of the company's Internet connection for name resolution traffic and provides faster name resolution to the Internet.

All of these name servers—including the primary—restrict zone transfers and have recursion disabled. The secondaries deny all zone transfer requests and provide only non-recursive name service. Restricting them to non-recursive name service helps protect the secondaries, particularly from denial of service attacks.

The appliances are particularly well suited to use externally, as they're hardened against attack: they don't run unnecessary network services and they only listen on the required ports. They're also easy to upgrade, which is especially important since it's imperative that any name server directly exposed to the Internet be upgraded promptly when a vulnerability is announced.

Forwarder

The final component of the company's external DNS infrastructure is the forwarder. The forwarder is responsible for handling internal name servers' queries for Internet domain names—queries they can't process themselves, as they lack connectivity to the Internet. Like the primary, the forwarder is located inside the firewall. Firewall rules allow the forwarder to send queries to Internet name servers and to receive responses from those name servers, but don't allow unsolicited DNS messages from Internet name servers to the forwarder. An access control list on the forwarder reinforces this restriction, denying queries from addresses outside of the company's address ranges. The forwarder is at headquarters, close to the company's connection to the Internet, to minimize delay and speed its resolution of Internet domain names.

The forwarder is also implemented as a high availability pair of appliances. High availability is an especially useful feature with forwarders, since many name servers employ very primitive forwarding algorithms. Most will accept a list of addresses to forward queries to, but wait a fixed amount of time after querying a forwarder before trying the next in the list. If the first forwarder in the list fails, the name server will nonetheless query it first each time it consults a forwarder, then wait for a timeout, then try the next forwarder. This can slow resolution of Internet domain names by several seconds per query. With high availability, there's only a single, virtual address to configure as the forwarder, and there's a very high likelihood that the name server at that address is responding.

Headquarters DNS Infrastructure

Next, let's describe the internal DNS infrastructure at headquarters. These are name servers whose primary role is to answer queries from internal stub resolvers, such as desktop PCs and production servers. They're authoritative for the zone company.com, which contains data about all computers at headquarters, as well as printers, routers, hubs, and other network resources.

Primary

The core of the headquarters DNS infrastructure is the primary name server for company.com. Like the external primary, the internal primary is hidden, but for a different reason: operational flexibility. The primary isn't listed in NS records for the zone, nor does it answer direct queries from any resolvers. Its role is simply to serve zone transfers to the zone's secondaries, which means that its availability has no impact on resolution of company.com domain names. The administrators can bring it down for maintenance or upgrade it without degrading name resolution.

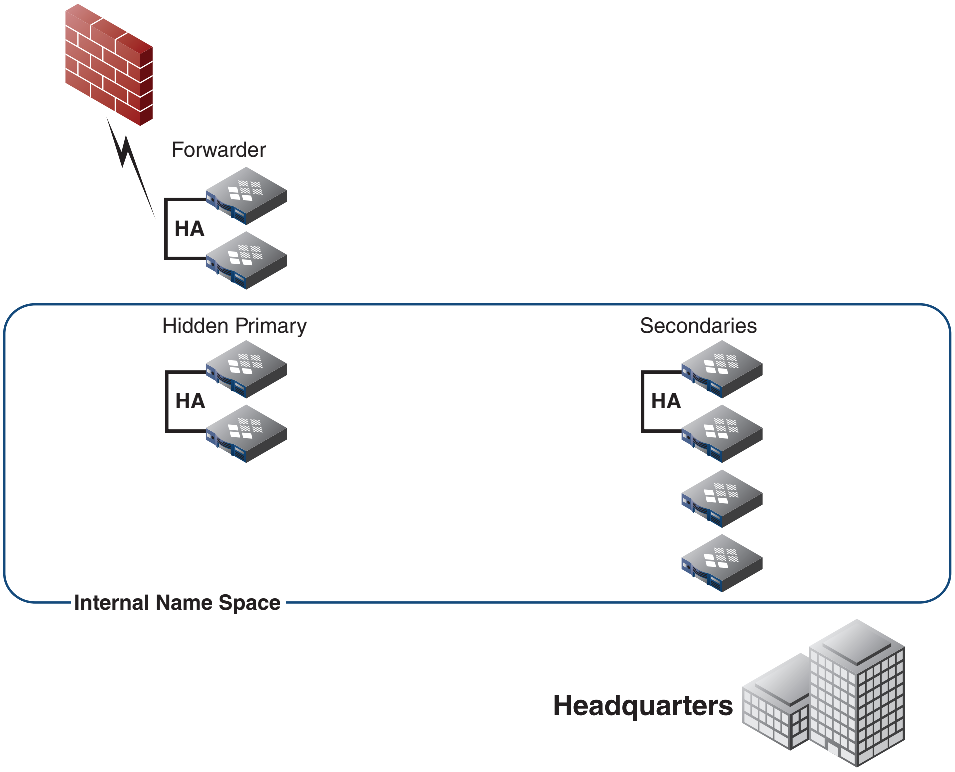

The primary is implemented as a high availability pair to provide redundancy for the seat of administration. Though running hidden provides operational flexibility, the primary is still critical to the administration of company.com.

Secondaries

The headquarters site also runs two company.com secondaries. The site's resolvers are configured to query the closer secondary first, then the other secondary. This provides redundancy and spreads the resolution load between the two name servers. Normally, however, resolvers would experience a timeout of several seconds if their preferred name server failed. Consequently, both secondaries are implemented as high availability pairs, since the business applications and executives that rely on these name servers tend to misbehave when name resolution slows or fails. This allows the administrators to provide continuous name service even while upgrading the name servers: an upgrade of a high availability pair upgrades the passive member of the pair first, then forces the active member into the passive state and upgrades it.

Regional DNS Infrastructure

Each region manages a subdomain of company.com, which we'll call region.company.com. The regional subdomains contain data about all of the computers in the region, including computers at branches in the region.

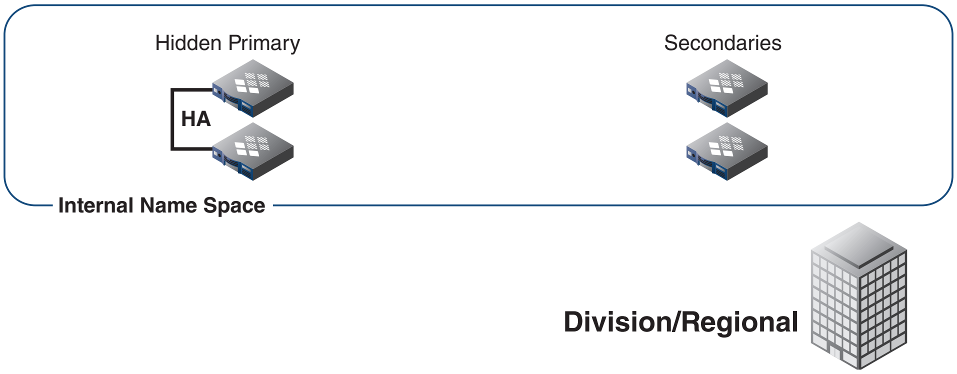

Primary

Each regional subdomain is managed on a primary name server at the regional office. The primary is hidden to provide operational flexibility, just as the company.com primary is. The primary is also a high availability pair, to provide redundancy for the seat of administration. The regional subdomains are also dynamically updated by DHCP servers at the regional offices and at branch offices. The use of high availability ensures that the primary will always be accessible to process the updates.

Secondaries

For cost reasons, the two secondaries at each regional office are individual appliances rather than high availability pairs. The regional office's resolvers are configured to query the closer secondary first, then the other secondary. This provides redundancy and spreads the resolution load between the two name servers. Should one secondary fail, roughly half of the office's resolvers will experience a degradation of name service as queries to their preferred name server time out, but this is considered acceptable.

The appliances' graphical user interface and its support for fine-grained authorization makes them particularly well suited to use at the regional offices. The headquarters IT staff can retain overall control of the appliances' configurations, ensuring that they use the right forwarder and apply the necessary access control lists, for example. They can create administrative users for IT staff within the regional office and allow them to edit only locally relevant zone and DHCP data. The GUI prevents local IT staff from accidentally or deliberately making changes they shouldn't, and provides a simple user interface for managing zone and DHCP data without specialized expertise in DNS or DHCP.

Branch DNS Infrastructure

Each branch office in a region runs a single secondary name server (secondary for the zone corresponding to the region that branch is part of). These are run as stealth secondaries, which means that they are not listed in the zone's NS records. This prevents other name servers in the company from querying them and unnecessarily congesting the branch's connection to the corporate network.

Resolvers at the branch are configured to query the local name server first, and then to try one of the secondaries at the regional office. This makes use of the branch name server's copy of the regional zone data and cache of frequently looked up domain names, but provides redundancy in case the local name server fails.

The appliances' graphical user interface is especially useful for the remote management of name servers at the branches. The GUI gives IT staff at headquarters complete control over the entire appliance, including both service (DNS and DHCP) and platform (i.e., appliance) configuration. The IT staff can easily reconfigure or upgrade the appliances remotely, with no assistance from local staff.

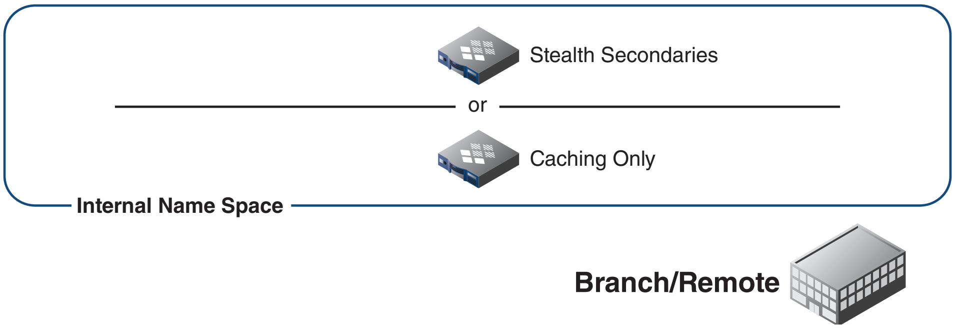

Design Considerations

In designing the system, there are other key considerations that inform the final choices made to develop a best practices architecture.

Forwarding

All internal name servers in the company DNS architecture are configured to forward to the forwarder at headquarters. To avoid creating a bottleneck in name resolution at the forwarder, we've configured all internal name servers not to forward queries that end in company.com, but to resolve these domain names using iterative queries. This ensures that internal name resolution is more efficient and robust, and avoids unnecessarily burdening the forwarders. Only queries for domain names that don't end in company.com—that is, Internet domain names—will be forwarded.

Integrated DHCP

A bulletproof DHCP architecture is just as important to the company as DNS, particularly because the company recently completed a company-wide implementation of Voice over IP telephony. Their VoIP gear relies on DHCP to function—as does most VoIP equipment. Consequently, our design provides redundant DHCP servers for all of the company's sites.

All internal secondary name servers also provide DHCP to their respective sites. At headquarters, the secondaries—and hence the DHCP servers—are run in a high availability configuration. This provides redundant DHCP service to the headquarters networks.

At the regional offices, the individual appliances that serve as secondaries are also peers in a DHCP failover association. This allows them to share a database of DHCP leases. Should either peer fail, the other can assume its responsibilities. The appliances' graphical user interface makes it simple to configure DHCP failover associations, and to assign any lease pool to an association.

At branch offices, the local name server also serves DHCP. (In fact, DHCP is arguably the more critical of a branch office appliance's responsibilities.) The branch DHCP server is a member of a failover association with one of the two DHCP servers at the branch's regional office.

All DHCP servers automatically update the appropriate primary name servers with their DHCP clients' forward and reverse mappings. DHCP clients are not allowed to update DNS data directly. This keeps the namespace synchronized with the state of the network without compromising security, and supports both newer and older versions of Windows.

Disaster Recovery

All appliances are backed up nightly to headquarters. This only requires a simple script, which can run on any computer at the headquarters site that is itself regularly backed up. The backup is sufficient to restore any appliance to its state at the time of the snapshot. For disaster recovery purposes, the backups can be copied to a remote site.

Active Directory Integration

The company has a single Active Directory domain with the same name as the company's parent zone, company.com. To support Active Directory, we must let the Domain Controllers for company.com update records that advertise the services they offer. We created four new zones on the primary name server for company.com to "capture" these updates: _msdcs.company.com, _sites.company.com, _tcp.company.com, and _udp.company.com. The name server is also configured to allow updates from the Domain Controllers to these zones.

The appliance's graphical user interface makes this configuration especially simple: it provides a section specifically for the configuration of Active Directory support and creates these "underscore zones" automatically.

The company.com secondaries at headquarters are also configured as secondaries for these "underscore zones." They provide redundant sources of these records, which are critical to locating a Domain Controller. In fact, any other internal name server can be configured as a stealth secondary for these zones: They're very small and rarely change, but are indispensable to the correct operation of Active Directory.

Additional Requirements

When choosing an appliance to implement this DNS architecture, the company should look for more than just an appliance that supports the design. This customer, in particular, required the ability to monitor the appliances using an SNMP- and Syslog-based management system. They also needed to integrate an existing, home-grown network management application into the DNS infrastructure. The appliances' support for a Perl-based remote API satisfied this requirement.

Finally, the customer took into account features not yet commonly available in appliances that would simplify administration and extend this platform for the support of DNS and DHCP to other critical network infrastructure protocols.

Traditionally, appliances have been configured and managed one by one. To upgrade a group of appliances, for example, required uploading new firmware to each independently. Changing authorization rights for a particular user on that group of appliances meant visiting each appliance, in turn, and making changes via the GUI or command-line interface.

Once deployed, it makes sense to exploit this appliance infrastructure to support more than just DNS and DHCP. Some appliances promise the extensibility to manage other network infrastructure protocols in the future, so that the company could use the same graphical user interface and in some cases the same appliance platforms to implement protocols such as RADIUS and LDAP into a unified network identity platform.

Conclusion

While many of the principles in this design are agnostic of how DNS is deployed, appliances have provided the ability for companies to deliver this architecture as core network infrastructure—a significant leap in how DNS is deployed and managed. As with other network infrastructure, DNS appliances offer better reliability and security than name servers based on general-purpose operating systems, as well as significant features including high availability, powerful management interfaces, and easy backup and restore.

As companies revisit and revise their DNS architectures in light of the increased importance of DNS as a network service, appliances should be considered as the key delivery mechanism for companies that want to optimize management and cost-effectiveness, enhance the security of their networks, and build a scalable, reliable platform for future network growth.

About the Author

Cricket Liu is the co-author of all of O'Reilly & Associates Nutshell Handbooks on the Domain Name System, *DNS and BIND, DNS on Windows 2000, DNS on Windows Server 2003,* and the *DNS & BIND Cookbook,* and was the principal author of *Managing Internet Information Services.*

Cricket worked for five and a half years at Hewlett-Packard's Corporate Network Services, where he ran hp.com, one of the largest corporate domains in the world, and helped design the HP Internet's security architecture. He later joined HP's consulting organization to found their Internet consulting business.

Cricket left HP in 1997 to start his own company, Acme Byte & Wire, with his friend and co-author Matt Larson. Acme Byte & Wire specialized in consulting and training on the Domain Name System, including both the BIND and Microsoft DNS Server implementations. Acme Byte & Wire's customers included over ten percent of Fortune 100 companies.

Network Solutions acquired Acme Byte & Wire in June of 2000. Subsequently, Network Solutions merged with VeriSign. Cricket became Director of DNS Product Management of the merged company, helping determine which new DNS-related products VeriSign would offer.

Cricket left VeriSign in June, 2001, and is currently vice president of architecture of Infoblox. In that role, he helps guide the development of Infoblox's product strategy and service offerings, and serves as a liaison between Infoblox and the technical community.

CORPORATE HEADQUARTERS:

+1.408.986.4000

+1.866.463.6256

(toll-free, U.S. and Canada)

info@infoblox.com

www.infoblox.com

EMEA HEADQUARTERS:

+32.3.259.04.30

info-emea@infoblox.com

APAC HEADQUARTERS:

+852.3793.3428

sales-apac@infoblox.com

© 2013 Infoblox Inc. All rights reserved. Infoblox-whitepaper-dns-appliance-best-practices-June2013