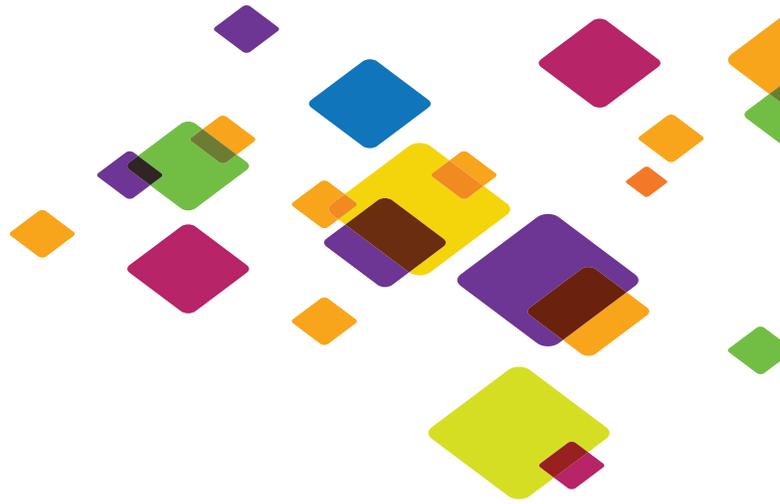




WHITE PAPER

How to Get the Most out of DNS in an Active Directory Environment



Active Directory Domain Names in DNS

Microsoft leverages a relationship between directory service domains and related DNS zones. This allows clients trying to access resources for a given Microsoft domain to leverage DNS by searching for resources in the matching DNS zone. Additionally, Microsoft leverages special subzones (often called underscore zones) to specifically hold the service records for the matching domains.

Active Directory DNS Objects

DNS service records are an Internet Engineering Task Force (IETF) record type standardized in RFC 2782. Microsoft has taken full advantage of this record type and has implemented it as the core method for publishing services. The SRV record type allows for the specification of service type, transport protocol, port, weight, and priority.

Below is an Example SRV record showing that server dc02 has the LDAP service available for the example.net domain.

```
_ldap._tcp.west._sites.DomainDnsZones.example.net. 600 IN SRV 0 100 389 dc02.example.net.
```

Infoblox Enhances Microsoft Active Directory.

Infoblox adds functionality, including highly available and secure DNS, to the Microsoft Active Directory services. The symbiotic relationship between Infoblox and Microsoft AD creates a robust and secure environment, rendering the best possible name-resolution system without creating interdependence. This lack of interdependence is the primary benefit of this particular combination, allowing each of these two core services, Infoblox and Microsoft AD, to be controlled within individual architectures. Each can each be customized, managed, upgraded, and patched independently of the other. By running DNS services on Infoblox, Active Directory is able to gain a number of immediate and important benefits including:

Gain of Resources on Domain Controllers by Removing the DNS Service.

Removing DNS from the domain controllers allows the server to focus on managing the domains. The removal of DNS also allows the domain controllers to function optimally regardless of the DNS load within the environment. Repeatedly, domain controllers respond substantially faster once the DNS authoritative (ADI) functions are adjusted to function as the DNS forwarder to Infoblox. The ability to turn the DNS service off completely means the domain controllers can allocate all available resources to the functioning of Active Directory and its related services.

Remove Interdependence between the Services Being Hosted on the Same Server.

Once interdependence is removed, if one service is being overloaded or even attacked, there will be no effect on the other services. As these two services, DNS and AD, are not interdependent by nature, they can function separately with no risk of issues spreading from one set of services to the other. The separation also allows administrators to resolve one service's issues with little or no concern for the other service.

Independent Patching, Upgrading, and Management

Finally, the separation of these mission-critical services onto separate hardware platforms allows for each service to be maintained without worrying about interruptions and impacts to the other service. In a change-control world, this has a huge impact on the ability to properly maintain and improve the services provided to the enterprise.

Infoblox Advantages

Security

The Infoblox solution is the perfect foundation for extending security with specific features like DHCP Fingerprinting for access policy enforcement, rogue-device detection, and Media Access Control (MAC) spoofing alerts. DNS is also enhanced with Infoblox DNS Firewall for prevention, detection, and mitigation of malware-infected devices. The DNS Firewall has a trusted feed of malicious site addresses that is regularly updated. For zero-day attacks, Infoblox DNS Firewall has been integrated with the FireEye NX Series appliance, making its detection of risk actionable by tying FireEye findings into the feed supporting the DNS Firewall. For externally facing DNS, Infoblox offers an Advanced DNS Protection appliance that can continue to respond to DNS queries even during an attack. Finally, the discovery within Infoblox Network Insight provides the visibility necessary to ensure there are no risks lurking in terms of unmanaged networks and devices. These products that supplement the core DDI make Infoblox an exceptionally secure alternative to using native Microsoft DNS and DHCP.

Database and IPAM

At its core, the Infoblox DNS solution is a real-time distributed semantic database. This allows Infoblox to provide solutions to the largest organizations in the world and offer real-time data availability. The Infoblox database is designed with the sole purpose of providing reliable and secure DNS, DHCP, and IPAM. Therefore, it delivers superior data synchronization when compared to other solutions in the industry. That core database provides the building block to a true IPAM solution that allows for the overlay of both protocol and enterprise data. It takes IPAM past the IP address itself and requires the data to be referenced and then cross referenced against any and all fields, providing visual insight and analysis to the network hosts. The Infoblox Grid database's asynchronous nature results in virtually no delays in propagating new data to other members and servers.

Auditing

Infoblox provides complete administration auditing. All administrative actions are logged to both an audit log and optionally written to SysLog. This allows administrators to provide a complete audit history on a per-administrator basis or on a per-protocol-object basis. An organization can now provide all the needed data for various standards-based auditors. Additionally, this enables organizations to clearly identify changes in the environment when troubleshooting identified issues, thus decreasing support times.



Manageability

Infoblox Grid technology adds single-pane-of-glass management for the entire DDI architecture. There is no need to manage individual servers, and therefore, focus can be given to the tasks associated with DDI. From day-to-day data entry to the deployment of new environments, Infoblox creates a platform in which data is deployed centrally to the entire architecture, providing an accurate view of all networks and hosts. For the most complex environments, Infoblox offers Multi-Grid Master, which allows the centralized management of both DNS and DHCP, including both IPV4 and IPV6 networks, in up to 50 Grids, each Grid with up to 250 members for a total of 12,500 members. This level of centralized management of even the most complex environment is unmatched in the industry.

Navigation

The Infoblox interface is optimized for DNS and DHCP tasks. By placing a strong importance on navigation and ease of use, Infoblox customers save time on common tasks associated with DNS management and maintenance. Functions such as Smart Folders and bookmarks allow the users to execute common tasks against data organized to match the organization's architecture. Additionally, Infoblox has created a tasks dashboard that allows users to perform daily tasks without the need for GUI navigation.

Troubleshooting

Infoblox understands that adds, moves, and changes may be the most common tasks performed, but troubleshooting can be the most important. With built-in packet capture, Infoblox provides administrators the ability to look at the traffic being received and returned from any Grid member at any given time. Infoblox also provides detailed logging of all services on all members of a grid for additional troubleshooting resources.

Infoblox Network Insight

Infoblox Network Insight adds additional value by incorporating information about networked hosts and integrating infrastructure device data with IP address management. The collection and correlation of this data provides unprecedented visibility, helping administrators easily gather the necessary information, analyze it, then take the appropriate actions. This enables administrators to better manage their networks, validate designs, and effectively provision, troubleshoot, and deliver network services.

DNSSEC

DNSSEC by Infoblox offers central configuration of all DNSSEC parameters and enforces standards by configuring DNSSEC parameters at a Grid level (default key type and size and validity period, based on NIST-800-81 and RFC 4641 standards and including NSEC and NSEC3 support). Configuring a secondary and/or recursive name server for DNSSEC can be accomplished with a single click, including making it possible to send DNSSEC records as a secondary and enabling validation of DNSSEC for an external zone and easy importing of trust anchors.

Centralized Visibility of Attacks

Through comprehensive reports, Advanced DNS Protection delivers a centralized view of attacks that have happened on the network and provides the intelligence for taking action. These reports include details like number of events by category, rule, severity, member-trend analysis, and time-based analysis. They can be accessed through the Infoblox Reporting Server.

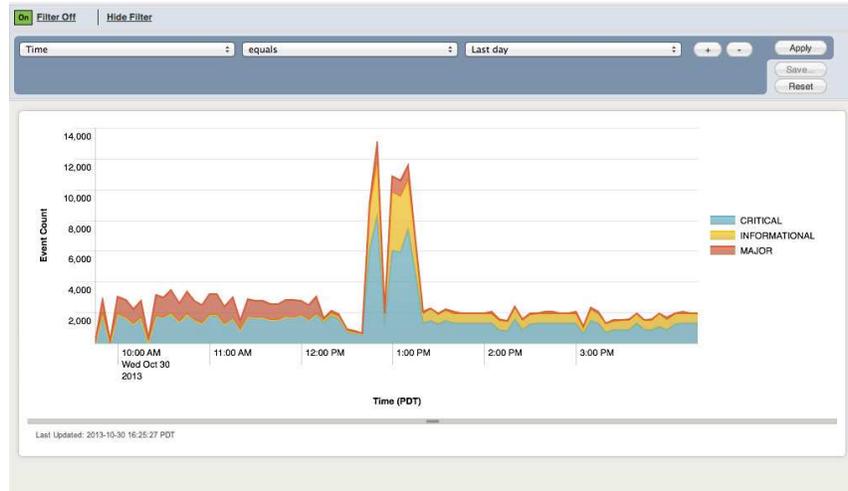


Figure 1.: Report on attack events

Tunable for Unique Needs

Every enterprise has different DNS traffic-flow patterns, and they can vary based on seasonality, time of day, or geography. Advanced DNS Protection provides tunable traffic thresholds that are configurable, making it possible to fine-tune protection parameters based on an organization's unique DNS traffic-flow patterns. This enables responding to good traffic without issues while blocking or dropping malicious traffic.

Why is this significant for Microsoft shops? The following graph was generated during a DDoS simulation. The same attack was launched against a BIND server, Microsoft DNS, and the Infoblox PT Appliance. The BIND server was able to satisfy half of the valid requests; the Infoblox Advanced DNS Protection Appliance was able to continue to support 100 percent of the valid DNS requests while dropping the invalid requests. Microsoft had the least favorable results, simply collapsing under the threat and not responding to any valid DNS requests.

Discovery Phase

Successful deployments begin with fully understanding the current DNS environment and its architecture as well as the needs and requirements of the infrastructure. The Microsoft and Active Directory design and implementation must be fully understood. Specific identification of all current Microsoft DNS servers is necessary, and then all the DNS configuration and zone files are collected from each server. Microsoft provides tools such as DnsCmd to accomplish this task. During the collection process, special attention must be paid to ensuring the collection of all information regarding inter-system relationships such as default forwarders, conditional forward zones, stub zones, and primary-secondary and zone transfer rules. This is often an iterative process as data is gathered and references to other servers and environments are found. The process goes back and forth until all data needed to ensure a simple and easy cutover has been discovered and documented.

For Active Directory integration, a list of all the domain controllers in service is assembled in order to keep track of them during the cutover, ensuring every single domain controller is able to update information about itself.

Additionally, while each infrastructure is different, there are some other environmental factors that can affect the overall architecture as well as the project itself. These factors, which can shape decisions regarding process and architecture commonly include:

- Number of physical sites
- Network connectivity between the sites
- Disaster-recovery requirements and procedures
- Critical services to the organization
- Compliance standards and any related segmentation requirements
- User distribution across geographical sites
- Use and requirements of any advanced or special DHCP options
- For Anycast, the routing protocols being used
- Amount and types of existing IPAM data to be added either now or in a future project
- Security and logging standards
- Change-control procedures

Planning Phase

Due to the robust nature of the DNS protocol, all customer architectures are unique. Below are considerations that make each environment different and are the things that need to be analyzed in the discovery and planning stages of any project.

#	Information	Impact and Consideration
Layout of Current DNS Environment		
1	Number of existing DNS servers	Map this to the number of Infoblox members. If the relationship is not one-to-one, figure out how to merge functionality and datasets.
2	Number of sites as it relates to DNS	Define and understand the DNS need for each site, such as which site is more critical and requires local survivability, or which site is used as DNS disaster recovery and can house the Grid Master candidate(s).
3	Number of domain controllers	Track this for the cutover. Each domain controller has at least one SRV that needs to be verified before and after cutover. The number of domain controllers is a factor in the technical resources required during the cutover.
4	Number of existing DNS views	Map the existing DNS views to the number of available Grid members. More views could mean needing more members, or larger members that can serve multiple views (more traffic). Sometimes, views also need to work with security policies as they cross security boundaries such as firewalls or router ACLs, and those relationships need to be identified and addressed.
5	Number of existing DNS zones	Not only is the number directly related to the amount of data and resources, but the relationship of which zone is served by which name servers or members is important for the creation of name server (NS) groups on the Grid. Identify how many NS groups need to be created prior to data import, and which zone will be imported into which group.
6	Number of existing DNS records	This directly relates to the size of the final database on the Grid. Ensure that each member in the Grid is sized appropriately to bear the load.
7	Number of subzones or child zones	Ensure that all original data is extracted. When parent-zone data is extracted, data from the child zone is not automatically taken. It has to be identified and extracted separately.
8	Activity Directory domains and forest architecture	AD domains are tightly integrated with DNS zones. Understand the relationships, identify all the domains, forests, and trust relationships, and plan the data migration and cutover accordingly.
9	Forward-mapping and reverse-mapping zones relationship	Forward-mapping and reverse-mapping zones can end up with a N-to-M relationship with multiple mismatching or conflicting versions. This is quite common in the world of reverse-mapping zones, and it should be identified early on to determine which "version" of the reverse-mapping zone to import, or sometimes, whether to discard reverse-mapping zone data altogether and rebuild based on forward-mapping zone data.
10	DNS Service IP addresses before and after cutover	Determine all DNS service IP addresses in use, and map it to DNS service IP addresses after the cutover. Then decide which IP addresses can be assumed by Infoblox members, and which ones cannot. This directly impacts the level of effort during the cutover and any cleanup or follow-up tasks that need to be performed.
Relationship with Other Name Servers		
11	Default forwarder settings	Identify all current default forwarder settings. Determine whether not there is a need to standardize across all name servers, or some sites need to forward specifically to other servers. Determine if there are dedicated forwarders in place, and if not, how to implement with given Infoblox grid members.
12	Conditional forwarding rules	Identify all conditional forwarding rules (DNS forward zones). Ensure that new Infoblox Grid members can still forward to these name servers.
13	Stub zone configuration	Identify all stub zones in use, and make sure they are transitioned to the Grid intact.
14	Zone transfer from other name servers	Identify all the zones in which the current DNS environment is relying on external entities. Notify the external entities of the migration if necessary, usually when there is an IP address change.
15	Zone transfer to other name servers	Identify all the zones in which external entities are relying on the current DNS environment. Notify the external entities of the migration if necessary, usually when there is an IP address change.
16	Zone delegation to other name servers	Identify all zone delegation to external entities from the current DNS environment. Notify the external entities of the migration if necessary, usually when there is an IP address change.
17	DNS access control list (ACL)	Identify entities that can perform queries, recursion, and zone transfers. These rules need to be reviewed and usually created manually on the Grid.

Data Accuracy		
18	Accuracy of existing DNS data	Analyze collected DNS data and look for correctness of the records (contains illegal characters?) and consistency of data across servers (do server A and server B disagree on an A record?) These conflicts and discrepancies require discussion with the current DNS team to resolve.
19	Reverse-mapping zone data accuracy	Often the reverse-mapping zones contain dangling pointer records. Identify the amount of inaccurate data in reverse-mapping zones (based on studies from forward-mapping zones) and resolve with the current the DNS team. Sometimes the solution is to discard reverse-mapping zone data and rebuild from forward-mapping zone data during import.
20	Current use and accuracy of dynamic DNS (DDNS)	Identify whether or not DDNS is enabled for each zone and whether or not all clients are performing updates. Identify servers other than domain controllers that need to update DNS dynamically. Explore future DDNS update policy (Check TXT Only versus ISC Mode).
Architectural Considerations		
21	Frequency and number of DDNS updates	DDNS updates are sent to DNS primary. Ensure that the DNS primary member can handle the amount of traffic from DDNS.
22	TTL values	Short TTLs result in high traffic and high load; long TTLs result in slow propagation and response to change. After identifying all current TTLs in use, discuss implementing a standard TTL on the Grid with the DNS team, and only allow exceptions where needed.
23	Start of authority (SOA) timer values	SOA timer values impact how frequently the secondary servers communicate with the primary servers for zone updates and zone transfers, as well as how long a secondary server can hold on to the secondary zone. These need to be considered to strike the best balance between performance and functionality.
24	Environment factors identified in discovery	I FORGOT WHAT THIS WAS FOR like network connectivity
25	Other DDNS needs to be identified	Identify any other devices or subnets that need to update DNS dynamically, and determine how they will be to updating DNS. For example, decide to allow updates from a /24 subnet for servers and use TSIG for some servers. Everything else needs to become a new operational procedure to add into DDI first before the machine can come online.
Cutover Procedures & Resource Planning		
25	Timing and duration of cutover	Work with change management to establish the ideal time for change to take place. Using information gathered, estimate length of cutover.
26	Change freeze and data re-import prior to cutover	Identify whether or not there is a need to re-import the DNS data prior to cutover. If the dataset has changed between the time data was first extracted and when the cutover will take place, a fresh new data extraction should be performed to ensure the latest data is collected, and prior to the cutover, the procedures developed during data import and testing should be used to re-import the dataset.
27	Operational support to make necessary changes and testing	Coordinate to have teams update Microsoft servers, DHCP servers, routers, and firewalls during the cutover, as well as test DNS functionalities before and after cutover, sometimes from multiple sites.

All of these need to be taken into account when designing the migration and transition to the new DNS environment.



Migration Phase

#	Task	Description
Configuration of Grid Before Data Import		
1	Default forwarder configuration	
2	Setting default TTL and SOA values	
3	Query and allowed recursion ACL configuration	
4	DNS view mapping and ordering	
5	Name server group creation	
Data Migration		
6	Service IP addresses	Determine which Microsoft DNS service IP addresses can be assumed by Infoblox members. Service IP addresses that cannot be assumed may result in consolidation of two or more Microsoft DNS name servers into a single Infoblox DNS member, and data merging may be necessary.
7	Zone-to-member mapping	For each of the Infoblox DNS members, determine which DNS zones will be served on that particular member, and create more name server groups if necessary.
8	A-zone import source and target identification	For each DNS zone identified to be migrated, determine which Microsoft name server to import the zone data from, and decide what DNS view, zone type, and name server group to import the zone into.
9	Reverse-mapping of zone data	Determine how to handle reverse-mapping zone data: import as is, or recreate based on records in the forward-mapping zone data.
10	DDNS entries	Determine how to handle DNS entries created as a result of dynamic DNS: import as is, or skip and let it repopulate later.
11	Data import	Use Infoblox provided tools to perform data migration to the Grid. Record the exact steps to take and any data transformation performed (such as fixing unreadable characters in zone data prior to import).
12	Verification	Verify all data imported, both by human verification and using tools and scripts to compare before-and-after datasets.
Configuration of Grid after Data Import		
13	Allow zone transfer and notify list	For each DNS zone imported, update allow zone transfer ACL and notify list as needed.
14	Verification of all domain controller records	Verify that all A, pointer record (PTR), and SRV records from each domain controller are successfully imported. If any record was imported as a host record, it needs to be reconfigured as separate A and PTR record pairs.
15	Integrated Active Directory list	For each AD-integrated zone, check that domain controllers are listed under the Active Directory section under zone configuration (to allow each domain controllers to update DNS dynamically).
16	Allow update ACL	For each zone, update the allow update list to allow all other devices (non-domain controllers) to update DNS dynamically.
17	Adjustment of zone settings	Make zone setting adjustments if necessary, such as special SOA timer values for certain zones.
Verification and Testing		
18	Verification correctness of DNS data imported:	Verify of correctness of data imported <ul style="list-style-type: none"> • Check TTL and SOA values. • Check NS records generated. • Check correctness of authoritative data. • Check reverse-mapping zone data.
19	Verification the following zone settings and behaviors, if applicable:	Verify of zone settings and behaviors <ul style="list-style-type: none"> • Forward zones • Stub zones • All sub-zones for correct post-cutover configuration • All delegated zones
20	Recursive lookup	Test and verify recursive lookup behavior, if enabled.
21	Zone transfer	Test and verify zone transfer behavior, if enabled.
22	DNS views	Test and verify DNS views behavior, if enabled.

This phase focuses on migrating the Microsoft data to the Infoblox platform. The migration phase is as critical as the other phases because this is when the real data gets moved from the existing environment over to the new environment. This is where the bulk of the work is completed, enabling a simple and seamless cutover.

During lab testing, identify which of the tools Infoblox offers are best suited for the specific data migration. It may be necessary to repeat the data migration process one more time closer to the cutover time depending on change-control and change-freeze policies. In these cases, get the estimated time needed for executing the data extraction, import, and verification steps, and use these to update the migration and cutover plans as necessary.

Cutover Phase

After migrated data has been fully verified and tested, the Infoblox Grid is now ready to take over DNS service from Microsoft. Below are some general guidelines, listed in the order they should be carried out, on how to execute a cutover, and a sample cutover plan is included below.

1. Initiate change freeze on Microsoft name servers.
2. Extract and import data into Infoblox again as needed.
3. Re-verify any newly imported data with methods perfected in the initial migration.
4. Enable and restart necessary services on Infoblox.
5. Change IP addresses of Microsoft servers and Infoblox members as needed to assume service IP addresses.
6. On all Microsoft name servers, change default global forwarding settings to forward all DNS traffic to service IP addresses on the Infoblox authoritative DNS Grid members.
7. On all domain controllers, change client TCP/IP DNS settings to use service IP addresses on Infoblox first and themselves second.
8. On each domain controller where Infoblox can assume the Service IP address:
 - 8.1 Disable or remove DNS service.
 - 8.2 Re-IP the Microsoft Server off of the service IP and replace it with an Infoblox DNS Grid member.
 - 8.3 Restart net logon services, which triggers dynamic updates to register SRV resource records in DNS.
 - 8.4 Use Microsoft's IPConfig tool to register DNS, which updates the A and PTR records in DNS of the domain controllers.
9. On Microsoft name servers not replaced by Infoblox that need to forward DNS until all static clients can be configured to point directly to Infoblox:
 - 9.1 Remove all DNS zones and conditional forwarders.
 - 9.2 On Infoblox remove SRV, A, and PTR records associated with the domain controller.
 - 9.3 Restart net logon services, which triggers dynamic updates to register SRV resource records in DNS.
 - 9.4 On Infoblox, verify that all SRV records are repopulated by the domain controller.
 - 9.5 Use Microsoft's IPConfig tool to register DNS, which updates the A and PTR records in DNS of the domain controllers.
 - 9.6 On Infoblox, verify that domain controller A and PTR records are repopulated by the domain controller.



Cutover Diagram

The diagram in Figure 2 shows the example environment immediately after the cutover phase and assumes the service IP addresses of only Microsoft servers A and B can be taken over. Now all of the AD servers (A,B,C) point to Infoblox Grid members for DNS service. Clients are updated to use the new DNS servers, or their DHCP ranges are updated, and the clients receive new information on DHCP renew.

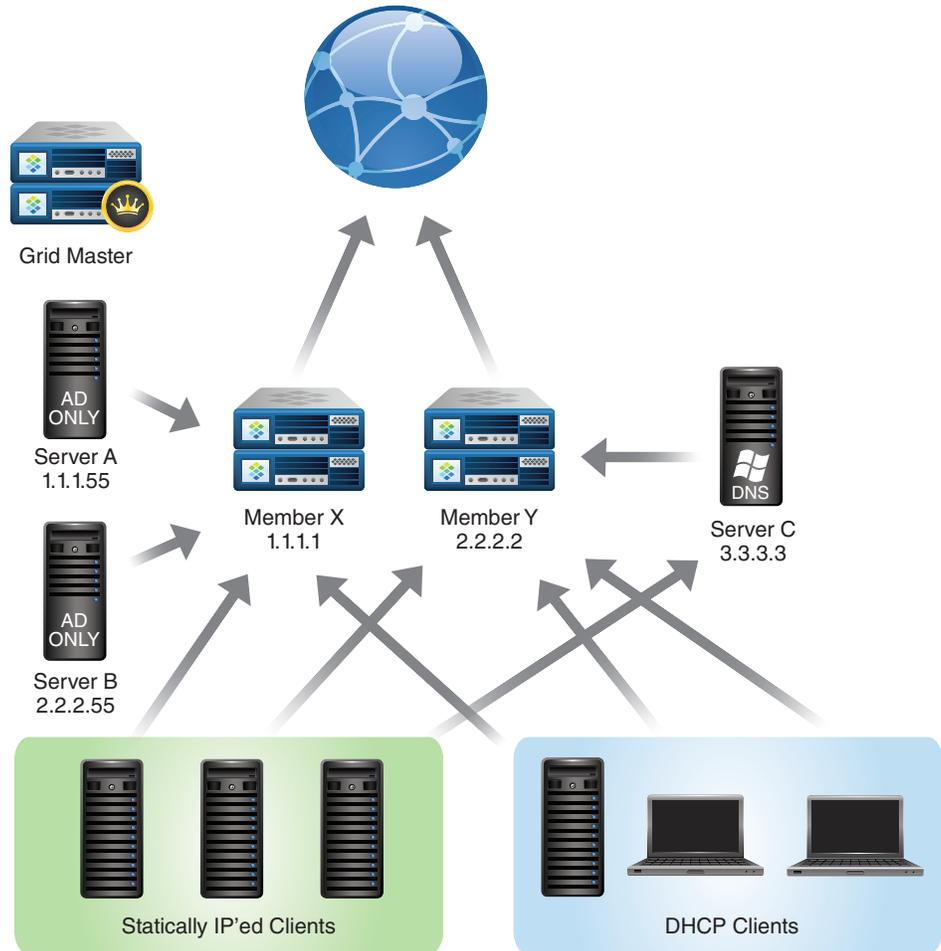


Figure 2: Initial cutover is complete with Server C still running Microsoft DNS.

In many cases the number of Infoblox appliances will be smaller than the number of Microsoft servers being replaced. In Figure 2, two Infoblox appliances have been installed so reassignments and reconfigurations will have to account for that, and all dynamically assigned clients are using Infoblox for DNS services. Statically configured clients that pointed to the DNS services that Infoblox assumed are also now served by Infoblox DNS. Clients that were statically pointing to Server C must be identified and reconfigured.

Example Cutover Plan

Once all data has been migrated and verified, the cutover is simply a matter of assuming the service IPs where the plan calls for it. The Microsoft servers are reconfigured to either shut off DNS (where Infoblox assumes service IPs) or become global forwarders to Infoblox where the Microsoft DNS has been not shut off (for whatever reason). Next, start with configuring global forwarding, configure DNS clients, and back up all zone data. In the second pass the Microsoft configuration is changed to either OFF or Forwarding, and the Active Directory SRV records are repopulated to ensure the DCs can update their own records properly on Infoblox.

#	Infoblox Task	Microsoft Task	Notes
1		DNS change freeze and maintenance window starts.	
2		Export zone data.	Use DnsCmd.
3	Import zones as primary to Grid.		
4	Enable DNS update.		
5	Restart DNS and DHCP services.		
6		Change IP address of server A to 1.1.1.55.	
7	Change member X IP address to 1.1.1.1.		
8		Update server A default forwarders to 1.1.1.1 and 2.2.2.2.	
9		Change IP address of Server B to 2.2.2.55.	
10	Change member Y IP address to 2.2.2.2.		
11		Update server B default forwarders to 1.1.1.1 and 2.2.2.2.	
12		Update server C default forwarders to 1.1.1.1 and 2.2.2.2.	
13		Update server C client TCP/IP DNS configuration to use 1.1.1.1 and 2.2.2.2.	
14			Default DNS query path now flows through Infoblox.
15		Update server A client TCP/IP DNS configuration to use 1.1.1.1 and 2.2.2.2.	
16		Disable or remove DNS service on server A.	
17		On server A, restart net logon service and reregister DNS.	
18	Verify appropriate A, PTR, and SRV records were created by server A.		Server A cutover is complete.
19		Update server B client TCP/IP DNS configuration to use 1.1.1.1 and 2.2.2.2	
20		Disable or remove DNS service on server B.	
21		On server B, restart net logon service and re-register DNS.	
22	Verify appropriate A, PTR, and SRV records were created by server B.		Server B cutover complete.
23		On server C, remove all DNS zones.	
24		On server C, restart net logon service and re-register DNS.	
25	Verify appropriate A, PTR, and SRV records were created by server C.		Server C cutover complete.
26		Restart net logon services and re-register DNS on all other domain controllers.	
27	Verify appropriate resource records were created by domain controllers.		Cutover complete.



Post-cutover Phase Follow-up

The post-cutover phase consists of the tasks that are completed after the cutover has been deemed a success and additionally includes tasks that need time to allow Infoblox to be fully populated with hosts on the network that may not have been present during the cutover. These tasks are either documented to complete after a given period of time, or can be scoped into the project for an engineer to come back and complete at a given period after the cutover is complete.

- On Infoblox, change DDNS update mode if necessary.
- For statically configured clients, update DNS settings to use Infoblox service IP address(es).
- On Microsoft name servers still forwarding DNS:
 - Capture all clients still using this name server and reconfigure each client to use service IP address(es) on Infoblox.
 - Once all clients have been re-pointed, disable DNS service on those Microsoft servers to regain system resources.

Final State Diagram

Figure 3 shows the example environment after the follow-up phase. This shows that all statically pointed servers using Microsoft server C have been reconfigured to point to Infoblox for DNS, and all three Microsoft servers have had DNS turned off and are dedicated to supporting Active Directory services.

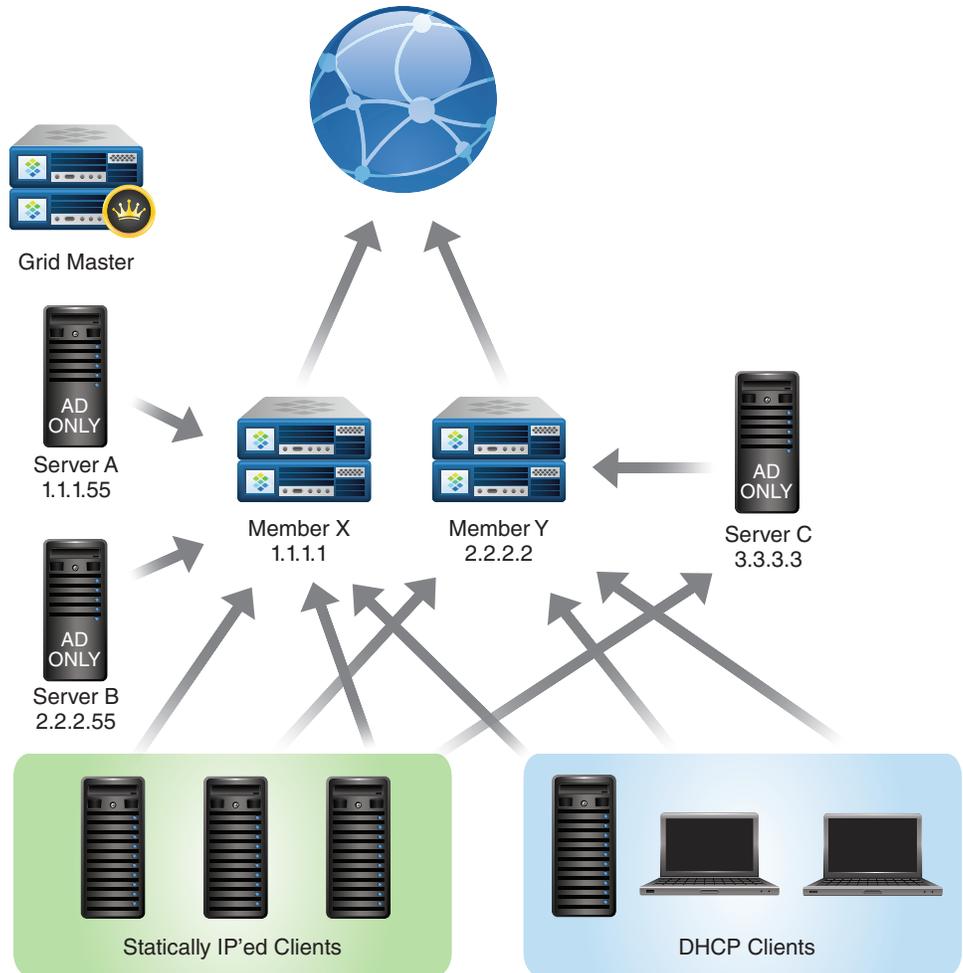


Figure 3: All dynamically assigned clients are using Infoblox for DNS services and that statically configured clients that pointed to the DNS services IPs that Infoblox assumed are also pointing to Infoblox. Because clients were statically pointing at server C in the diagram, and could not assume its service address, server C has been set to operate as a DNS forwarder and must identify and re-point those statically configured clients.

Summary

By enhancing the DNS service provided to Active Directory, IT organizations see immediate benefits in the performance, speed, and accuracy of data across both Active Directory and DNS. A common comment on the follow-up day is that things are running “so much faster.” But many of the advantages are seen over time, as by separating the services network teams are able to manage and maintain the environments separately, leading to more accurate DNS, decreased Active Directory issues, fewer issues getting change-control approvals for managing either service, improved visibility, security of DNS and hence IPAM databases, and dramatically decreased troubleshooting time. These benefits have allowed Infoblox to become the market leader in DDI and, as indicated by Microsoft’s Gold Partner status, a market-leading enhancement of Active Directory.

Contact Infoblox

If you have Microsoft Active Directory and Microsoft DNS and are considering adding an Infoblox solution please contact Infoblox Sales at 1-866-463-6256 or sales@infoblox.com

About Infoblox

Infoblox (NYSE:BLOX) helps customers control their networks. Infoblox solutions help businesses automate complex network control functions to reduce costs and increase security and uptime. Our technology enables automatic discovery, real-time configuration and change management and compliance for network infrastructure, as well as critical network control functions such as DNS, DHCP and IP Address Management (IPAM) for applications and endpoint devices. Infoblox solutions help over 7,100 enterprises and service providers in 25 countries control their networks.



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