Toward A Dynamic and Automated IP Environment

Abdulrahman Abdulbaki, Mohammed Saad Al-Qahtani

EXPEC Computer Center, Saudi Aramco, Dhahran, KSA
abduax1c@aramco.com and qahtms06@aramco.com

Abstract

The growth and complexity of data center networks in the context of virtualization, high performance, supercomputing and cloud computing services has led to an enormous demand for dynamic and automated IP address environment. Apart from the conventional challenges in business continuity, reliability, security, scalability and management flexibility, the need for a modern dynamic IP Address Management (IPAM) environment has become more pronounced in large-scale energy-sector data centers, foreseeing the exponential growth of IP devices combined with the complexity introduced by IT services (e.g. VoIP, Cloud computing, Server and Desktop Virtualization). Traditional methods of managing IP addresses manually with spreadsheets and static IP assignment cannot keep pace with the growing demands of virtualization and automation. This paper explains a novel design approach suitable for heterogeneous large-scale data centers, aiming to make corporate IPAM services run simpler, more efficient, and secure while utilizing automation as a core building block for corporate network IP services.

Keywords: *Domain Name Serve, Dynamic Host Configuration Protocol, IP Address Management, Identity, Policy and Audit, Network Infrastructure System.*

1. Introduction

Leading global Oil & Gas companies in the energy sector, utilizes a variety of systems and servers host numerous upstream applications to serve business needs. Additionally, data centers are fully populated with High Performance Computing (HPC) systems and clusters, which consist of tens of thousands of compute nodes hosting Geo-Steering, Seismic, Simulation, Exploration and variety of other commercial and in-house applications. They consume massive number of computing nodes that require constant communication between each other, to storage appliances, and other systems utilizing IP addresses in the form of Dynamic Host Configuration Protocol (DHCP) and Domain Name Server (DNS) infrastructure. In such large enterprises, rapid deployment for thousands of cluster nodes in one single installation is typical, and therefore, the traditional method of managing IP networks statically via manual spreadsheet is inefficient and could not keep with the needs of the rapid installation requirements. Thus, a solution for dynamic IP address assignment and IP address management is crucial and obligatory to optimally serve the massive growth and rapid deployment.

2. Conventional Data Center IP Address Management Standards

The main infrastructure of a major Oil & Gas enterprise would typically consist, but not limited, of three distinct environments distributed geographically over multiple data centers which fall in the following three categories.

- **Microsoft -** Dedicated for Windows platforms namely Microsoft Windows workstations and servers hosting a variety of core business applications.

- **High Performance Computing (HPC) -** Dedicated for Geological and Exploration applications that serve the core business. This category contains tens of thousands of compute node clusters, data storage appliances, application servers and other support systems.

- **General Infrastructure** - Built mainly for general purpose applications and other common services.

IP address management and Name Resolution are critical services needed for accurate communication among all the systems in the above environments. DNS & DHCP services were established in the early 90s and partially upgraded during Y2K. At that time they were unique in nature and were installed and distributed across multiple platforms. In such large organization there are three types of DHCP & DNS infrastructure.

- **Microsoft DNS service (MS-DNS) -** This infrastructure was established in the 90s and upgraded in mid-2000 during the deployment of Microsoft Active Directory. This platform was mainly used to provide DNS & DHCP services for Microsoft Windows servers and workstations, VOIP devices, Wireless networks, and a variety of other IP devices that required DHCP & NDS infrastructure.

- **Network Information System (NIS) -** This legacy infrastructure was established in the mid-80s mainly as a DNS service for name resolution, specifically, in the HPC environment. Due to the large number of compute nodes and constant communication amongst them, it was critical to have a DNS service delivering quick response time. It is not out of the ordinary to have thousands of servers querying DNS service simultaneously. To meet these continuous and large number of DNS requests with acceptable response time, huge amount of NIS servers were installed in Master and Slave mode, acting as the DNS infrastructure for this environment.

- **Identity, Policy and Audit (IPA)** – This newly infrastructure was installed for name resolution, authentication and authorization services to control access to Linux platform in both General and HPC environments.

There are three infrastructure services that required three separate databases (MS DNS, NIS DNS and IPA DNS) in the environment. Integration amongst these databases is essential to facilitate continuous and correct communication between the systems across all environments. Merging the databases together from these platforms is not a feasible and practical option due to major differences in data structures and schema. Synchronizing data is a more suitable option, however, this option introduces huge operational overhead and adds enormous complexity to the infrastructure. Each environment requires only a subset of the data from each platform. Data in NIS is mainly used for name resolution and authentication purposes in the HPC environment. Data in MS DNS is mainly used as a name resolution service for MS Windows and other general platforms. Whereas, data in IPA is primarily needed for user authentication and authorization for Linux in both General and HPC environments.

To ensure complete, accurate and successful synchronization for the required data sets, multiple in-house scripts would have been developed to extract the required subset of data from one platform, restructure it, reformat it, and send it to the next platform. These scripts are performed and applied on all platforms to complete the synchronization cycle among the systems. To avoid overlapping and data duplication, data entry and updates should be initiated from one platform, which in that case is NIS, by using UNIX VI editor to update a file called "Host File". This manual update includes IP addresses, host names, and other related information, such as group names and numbers which are needed for user authentication and authorization in HPC and Linux general environments.

For solely Name Resolution purposes, customized scripts are developed to extract a subset of data from NIS master server. These are stored in a file as a place holder called “Host Map” on another NIS server. The data is reformatted to the correct DNS structure record type, and sent it as a “DNS Record” to yet another file called “DNS Records” and placed in the same server. Once it is stored successfully as “DNS Record” type, the data is pushed via File Transfer Protocol (FTP) from NIS to the target location Microsoft DNS to be used by Microsoft DNS service. Another custom set of scripts are developed to push the data from the “DNS Records” file to the DNS service in IPA server.

For authentication and authorization purposes, more scripts are developed to extract other needed subset data in the “Host Map” file and pushed to the Authentication and Authorization process and services of IPA server. Figure 1. “Legacy DNS data workflow” shows the legacy DNS & DHCP platforms and data workflow involved in those systems.

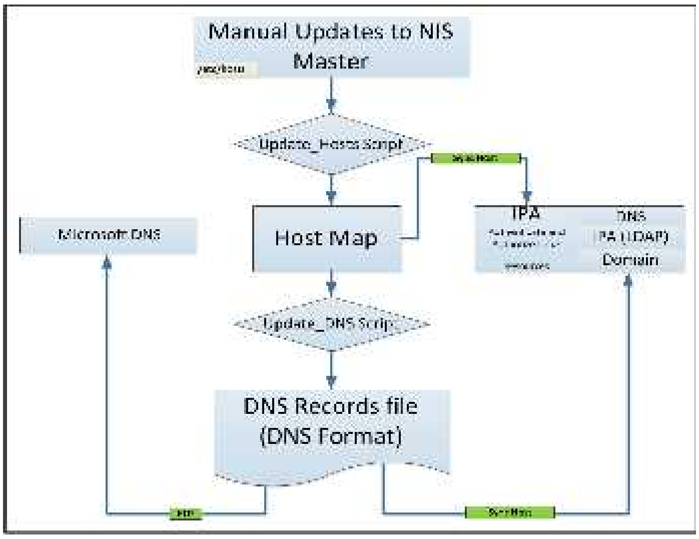

Figure 1. Legacy DNS Data Workflow

While the diagram above presents the complexity in the workflow among all the platforms, the legacy infrastructure lacks an automated method for IP address management. To keep track of the assigned IP addresses and network addresses, dedicated spreadsheets are developed to manually note all the updates for all the networks. Any update such as adding, modifying, or removing a network from the environment requires updating the spreadsheet manually. This process is cumbersome, slow, leaves room for human errors, and causes version control issues and confusions.

As indicated in the diagram above, data workflow and IP management processes are very complex and have introduced operational, management, security and availability challenges.

1. Complex - Multiple customized scripts are developed and executed to extract and push the needed data set from one platform to another. Each script could be written in different format and structure based on the target platform.

2. Slow – Sequential process for data entry and updates. Data entry and any updates started in NIS. Then, an additional time is need for scheduled scripts to extract the data and push the necessary updates to other platforms.

3. Platform security flaws - Server administrators and operational workforce could have access to view and modify the host file.

4. Transmission security weakness - FTP is used to push the updates to the target platforms.
5. Operational overhead - Large number of master and slave NIS servers are installed in the environment that needed to be managed and maintain.

6. Human errors - Susceptible to human errors as result of the direct updates to a flat host text file via VI editor. The updates are saved with any data integrity and error validation, which leaves room for errors.

7. Difficult – Different script formats and version control management tracking issues for the in-house scripts.

8. Risks - Inheritance to any modifications in NIS host file and Microsoft Windows operating systems vulnerabilities.

9. Inadequate IP addresses offering - Splitting DHCP scope amongst multiple servers requires all of them to be available. If one or more goes down, the remaining servers cannot fulfil the requests for IP address assignment, resulting in a shortage.

10. Multiples Management Interfaces – Administrators have to deal with multiple application interfaces and portals for the various platforms, leading to confusion and maintenance overhead.

3. Migration Strategy

Due to the above limitations in the legacy environment, completely different solution should be architected and designed to provide an optimal DNS, DHCP and IP Address Management (IPAM) applications as comprehensive set of core network services to manage any large environment. With the novel solution, the new established architecture should have simplicity and security guidelines focusing on four principles.

1. Simple - Keep the architecture and design simple.

2. High Availability - Provide continuous service availability.

3. Security - Ensure robust system to defend against security.

4. Performance - Improve throughput and system response time.

During the initial planning variety of commercial products can be evaluated including Microsoft, Cisco, QIP, Infoblox, and BIND. Although, most of the products conformed to the guidelines and provided the desired DHCP/DNS services, one product should be selected that fulfill the requirements and suites the environment best. This solution is based on a modular architecture accommodating versatile deployment scenarios, agnostic, and independent from the underlying operating system. This provides complete isolation for DNS/DHCP and IPAM services from the operating system.

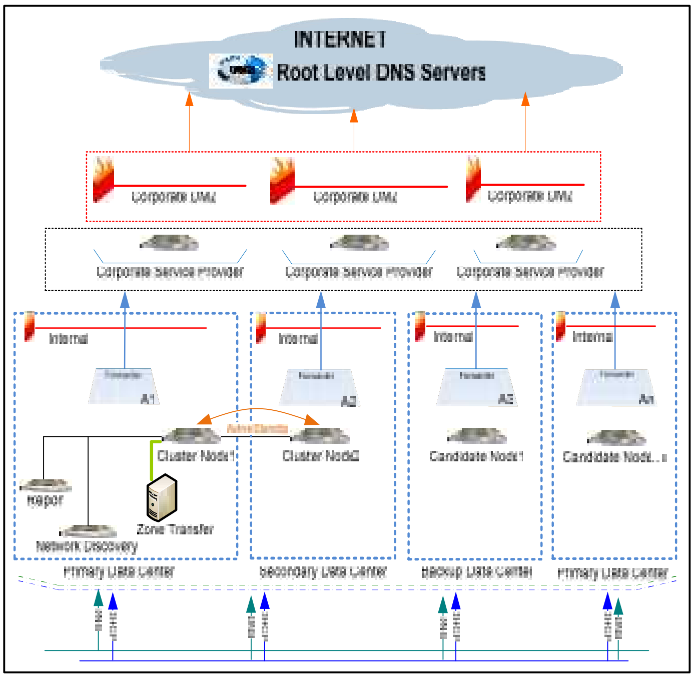

Figure 2. DNS Solution Architecture

"DNS Solution Architecture" shown in Figure 2 provides an overview of the architecture solution of the installed cluster. The components of the selected product are distributed across multiple data centers. They are physically connected on the network, establishing a logical domain forming one virtual cluster. The cluster contains multiple physical DNS & DHCP nodes that shared common elements and characteristics. Managing the configuration for this virtual cluster is executed from a single portal.

As depicted in the example above, the cluster consists of several nodes. Some of the nodes are dedicated for DNS & DHCP services, others are dedicated for network discovery, and the remaining nodes are dedicated for data collection and reporting.

As shown in Figure. 2, Cluster node1 (A1) is installed in the primary data center as an Active node of the cluster. Cluster node2 (A2) is installed in the secondary data center as Passive node in the same cluster. Both nodes share common characteristics and form the cluster identity. Active and passive act as one virtual node for management and operational needs.

More nodes (A3-An) are installed as member nodes in the same cluster that also offer DNS & DHCP services. The Passive node is ready to become Active automatically upon any failure in the Active node. Failure in case of both Active and Passive nodes simultaneously, requires manual intervention to set a candidate node to assume the Active role.

To ensure complete separation between DHCP & DNS data and management data, each node is configured with two network interfaces. LAN interface, which is dedicated for legitimate DHCP/DNS traffic, and Management interface, which is dedicated for management traffic, i.e. Operating System upgrades.

Distributing the nodes across multiple data centers in cluster mode provides persistent and continuous service availability. The new setup utilizes the organization DNS infrastructure as the primary DNS service to the end devices. With this integration, it utilizes the service provider DNS infrastructure as secondary DNS infrastructure in the case of any failure in the primary DNS infrastructure.

The new IP Address Management solution automatically routes data workflow among the concerned platforms transparently and efficiently. Rather than modifying the host file in NIS directly, the updates can be initiated in the new DNS infrastructure using the consolidated portal. Traditional DNS zone transfer is performed from DNS system to NIS, and if needed, only one script is developed to extract the needed data from NIS and pushed it to IPA for Linux platform name resolution as shown in Figure 3 "New DNS / NIS/ IPA workflow".

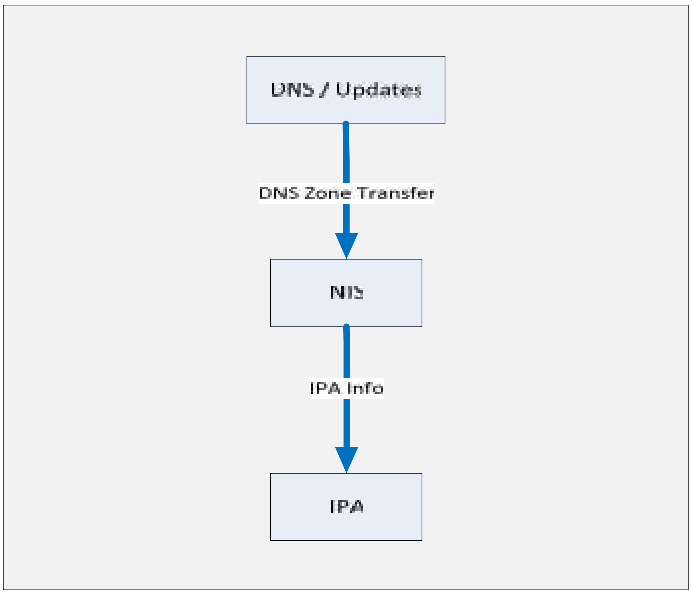

Figure 3. New DNS/NIS/IPA

Figure 3 "New DNS / NIS / IPA" workflow shows the simple and automated data route amongst the components in the new solution, with compression of the complexed workflow in the legacy system.

Simplified management is another important objective for the re-architected solution. In addition to DNS & DHCP services, system management functions and operational tasks including software upgrades and security patches are now conducted from a single portal, utilizing one easy dashboard. The new infrastructure eliminates the complicated processes in the legacy system, where several scripts are developed to move data from one platform to another. With the modern infrastructure there is no longer any need for creating multiple DHCP scopes for the same network in two separate DHCP servers, patch and secure the operating systems for NIS, IPA and Microsoft DNS servers.

Security vulnerabilities are a major concern in hosting the DNS database and service in the legacy system. The new DNS & DHCP infrastructure solution is designed with security in mind. Design considerations are made for provisioning a high level of security to reduce risks in DNS cache poisoning or DNS spoofing, Denial of Service (DoS), and Distributed Denial of Service (DDOS) attacks. Based on the integrated DNS firewall module and the built-in DNS advanced protection feature, the modular DNS is capable of defending against a wide range of external and internal DNS-based threats, exploitation attempts, and DNS hijacking. It detects and mitigates DNS attacks intelligently, while serving only legitimate queries.

The new solution provides Secure Dynamic DNS (DDNS), updates of dynamic A and PTR records upon workstation registration, and login to Active Directory domain. This mechanism is highly scalable as all dynamic DNS records registrations are held by DHCP service, thus offloading potential and manual changes on workstations, and eventually, offloading workstations from the overhead of registering their own records in DNS. The end result is a vast improvement in the security posture, and lower traffic footprint on the underlying network.

DHCP Fingerprinting offers an additional layer of security to identify MAC spoofing. By tracking DHCP fingerprinting and MAC address pairs, the system can determine when the associated fingerprint has changed, which can be used to prevent access to what might be a malicious behavior inside the network. Unauthorized access to the networks can no longer be granted thru an existing client on the network by spoofing its MAC address.

4. IP Address Management

IPAM is the complete solution for planning, tracking, and managing IP address space in the network without any manual intervention. This application is designed to discover the network, and automate DNS management processes, which allows the devices to be populated to the corresponding networks correctly without any chance of human error. With a powerful collection of tools and automation that streamline the workflow and processes, the new IPAM solution automated the process of "on-demand" IP address management requests.

IPAM is built utilizing simple Application Programming Interfaces (APIs) which make it modular and easy to maintain. With secure authentication, APIs can enable rapid and efficient actions, such as; Get, Put, Delete, and Post, data search and exchange between different functions and objects in the DHCP & DNS databases. With enabling the built-in workflow module and applying the API scripts, we save time and effort when deploying large-scale clusters and systems.

IPAM is developed as a comprehensive solution which is fully equipped with many rich features and tools that eased management and facilitated operations. One of the most used features is the file Export and Import, which allows administrators to easily upload or modify records in bulk with one single push to DNS zones.

One of the biggest challenges in the legacy system is lacking of IP address management. The legacy system is manual, cumbersome, and error prone. It does not offer any automatic mechanism for managing the ever increasing number of IP addresses. Therefore, IP address management is performed by developing static Excel spreadsheets which needs careful and constant updates. This contributes to human errors, IP address overlapping, IP address duplication, subnet scheme mismatch, and version control management challenges.

The new solution provides an integrated IP Address Management service that is capable of discovering new networks, IP devices on the network, and dynamically updates the database with newly discovered networks, eliminating the need for human intervention. Using a repository, this service shows the consumed and free addresses for all networks and can be used as a reference for future network capacity planning and baselining. For large enterprises, requiring rapid deployment with 1,000s of servers in a single installation, it is necessary to have a solution that offers an automated IP Address Management (IPAM) process and dynamic IP address assignment quickly.

IPAM adds another layer of security to the infrastructure utilizing DHCP Fingerprinting feature. It gathers information from DHCP requests and identifies the host types connected to the network. This information can be inspected and labeled as Linux, Windows, Printers, Wireless, IP Phone, and a wide list of devices connected to the network. Special Access Lists (ACLs) of devices are developed for DHCP scopes to deny or permit obtaining IP addresses from DHCP service. For example, a Linux workstation cannot obtain an IP address if it is connected to the Windows network or VLAN. However, it can obtain an IP address when it is connected only to the permitted VLAN.

Oil & Gas enterprise companies, typically, have a huge DNS infrastructure. Rather than operating on one flat view of DNS database, it can be structured reflecting the organization support groups. The new solution, logically, categorizes and organizes the database in a friendly and easy to follow smart folders, reflecting the organizational structure and following a defined standard naming convention.

5. Experimental Results

A key objective in the new solution is improving throughput in DNS response time and performance. During the development and testing phases we conducted stress testing on a thousands of nodes to measure and evaluate the elapsed time for DNS queries and replies. Although there are many local Master and Slaves NIS servers that served the environment, the stress testing results indicated a gain of an average of 1 millisecond for each node reply in the new solution. Having tens of thousands nodes in the environment contribute to a huge difference in the overall performance. The chart in Figure 4 "NIS V.S new DNS Response", shows the results for 30 servers and corresponding DNS response times between NIS (legacy) and IPAM (new solution).

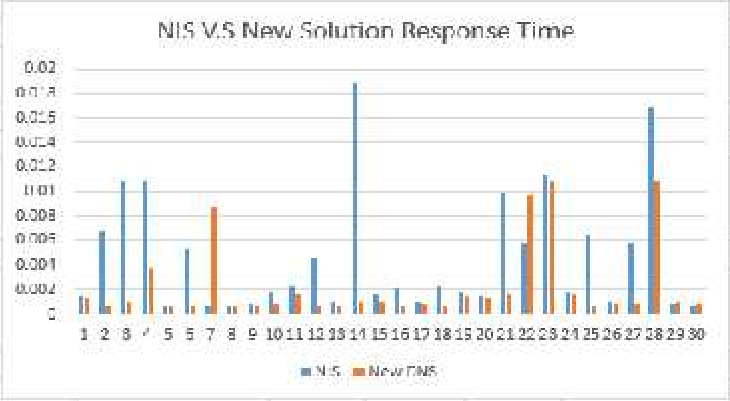

Figure 4. NIS V.S New DNS Response

5. Conclusion

Managing IP address space is becoming an important topic, as more devices and services are deployed on modern networks. These are increasing in number, diversity, and complexity and deployed on networks which are also becoming complex with virtualization and cloud computing. This has resulted in more and larger address spaces that require a dynamic approach to IP address management. The discovery and automation features in IPAM offers an essential tool to effectively and efficiently manage and secure the IP address space for heterogeneous large-scale data centers. The seamless and transparent integration between IPAM and DNS automates and secures DNS management processes, as networks become more difficult to manage.

References

[1] R. Droms (March 1997), IEEE RFC 2131 *"Dynamic Host Configuration Protocol" Bucknell University.*

[2] D. Eastlake, 3rd (September 2000), IETF RFC 2929 "*Domain Name System (DNS) IANA Considerations*"

[3] P. Mockapetris (November 1987). "*Domain Names – Concept and Facilities*".