


Domain Name System

Proper use reduces intranet administration costs



ORIGIN

The Internet continues expanding. Its progress seems unstoppable; for years now the rate of growth has been increasing. To easily continue using a more widely spread Internet, and to keep a grip on it, DNS is vital. This article explains the benefit of DNS for business networks and the technological and administrative conditions necessary for the optimal deployment of this technology. The method described here is particularly important for organizations with many employees spread over multiple locations.

Domain Name System

Proper use reduces intranet administration costs

Network addresses, such as 192.68.44.134, are difficult for people to remember. The need for associating names with network addresses has been recognized almost from the start of the Internet. Initially, a list of the names and network addresses of all computer systems was maintained in a central file, known as the *hosts* file. System administrators needed the discipline to regularly pick up the latest version. This method of working was no longer practical once the Internet starting rapidly expanding. System administrators needed to pick up an increasingly large file increasingly often. Also, the whole Internet was dependent on a single central authority who made changes. This authority also had no way of verifying changes.

By about 1983 it was clear that the hosts file had to be replaced by another mechanism. The successor had to offer the same functions, but also be distributed, consistent, reliable, and autonomous. These four characteristics are brought together in DNS.

1. *Distributed*: the system is hierarchical and allows the delegation of authorities to multiple administrators.
2. *Consistent*: the same answer is given when the same request is made at different places.
3. *Reliable*: redundant data can be held at different places; changes propagate automatically.
4. *Autonomous*: administrators can make changes independently of others.

Advantages for the Intranet

Internet technology is increasingly used on internal company networks, where it is called an *intranet*. The characteristics of DNS bring a number of advantages for intranets. An intranet often has the same characteristics as the Internet, such as rapid growth and increasing geographical dispersal. This means that there is a greater need for distributed management of names and network addresses. DNS can adapt itself to the growth and dispersal because of its distributed and autonomous characteristics.

A reliable translation of names to network addresses is a requirement for the reliability of intranet applications (see figure 1). Because these names are used for applications, the reliability of an application is dependent as much on a reliable translation from name to address as on a reliable network connection (provided by the routers). If the translation from name to network address fails, then the application is, in effect, disconnected from the network. DNS is very suitable for performing the translation because of its reliability and consistency.

A well-structured naming scheme makes maintaining an intranet easier. For example, it makes it possible to make

DNS technology is better for larger networks and business critical applications

changes to the infrastructure without disrupting the services. So the network address of a web server or mail server can be changed without having to inform the users, because the name will not change. DNS can, therefore, reduce intranet management costs.

There are various services which can perform name translation. The most well known, other than DNS, are Network Information name Service (NIS) and Windows Internet Naming Service (WINS). The NIS and WINS services were developed for administrating Local Area Networks (LANs) and are not sufficiently robust for larger networks. It is therefore better to use DNS technology for larger networks and for business critical applications.

SAP	E-mail	WWW
Names		
Network connections		

Figure 1: Applications, names and network connections.

DNS building blocks

What does DNS consist of? DNS is built from three components: *servers*, *resolvers*, and the *content*. Servers and resolvers form the DNS infrastructure. The content consists of what are called *domains*.

- A DNS server is called a *name server*, and its job is to store names or to get them from other name servers. Responses that the name server gets from other servers are temporarily stored in cache to eliminate unnecessary network traffic.
- Resolvers are the DNS clients whose job it is to query the name server (figure 2). A resolver can generally directly query up to three name servers.
- The DNS content can fulfil a number of functions. The best known are the translation of names to network addresses and mail routing. DNS can indicate where e-mail must be delivered. Mail routing can be made more robust by including alternative routes in the DNS. Since DNS is distributed, it is also necessary to store where a domain can be found. This is done with the help of *DNS meta-information*.

The content is held in a tree structure (figure 3), in which the highest level in the hierarchy is called the root. This hierarchical tree structure has different functions. Firstly, the name gives a rough indication of the type and location of the organization. Names that end with *.nl*, for example, are related to the Netherlands. Secondly, it is possible to automatically navigate through the DNS

tree; the DNS meta-information is used for this. Thirdly, the structure allows responsibility for a branch to be delegated to multiple parties. Within the branch *com*, for example, responsibilities can be delegated to *ibm* and to *origin-it*. This creates the domains *ibm.com* and *origin-it.com*.

The content of each domain must be maintained. That is the responsibility of the hostmaster. There is, therefore, a hostmaster for each of the *root*, *com*, and *origin-it.com* domains.

Each hostmaster must manage the names, network addresses, mail routing and DNS meta-information for the domain (figure 3).

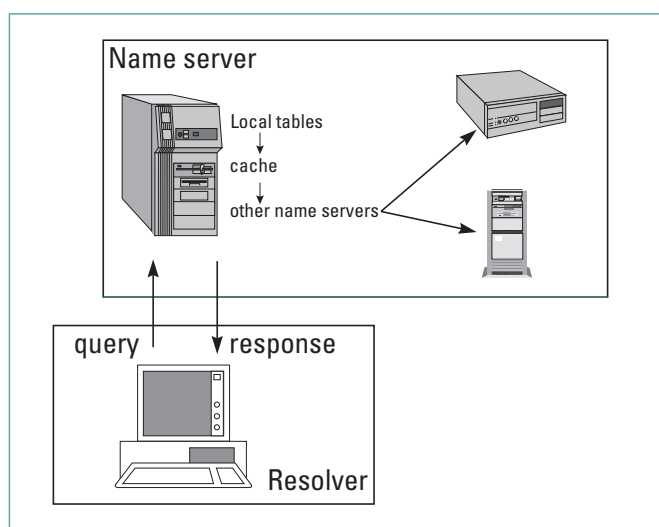


Figure 2: Name servers and resolvers

DNS architecture

Without a clear architecture for name servers, resolvers and the content, DNS cannot be effectively deployed. One possible architecture is looked at here. For the *name server* architecture a distinction is made between the backbone name servers and the LAN servers. A backbone name server has information about the whole intranet. Within the backbone a division is made into *master* name servers and *regional* name servers. The master backbone name server passes data on to the regional name servers. The LAN name server only has information about the LAN. LAN name servers ask the regional backbone servers for DNS content that is not known locally. To increase reliability the backbone name servers can be geographically dispersed (figure 4).

Resolvers first approach the name server on the LAN. In case of problems they can fallback to the backbone servers. This arrangement minimizes usage of bandwidth on the Wide Area Network (WAN), because the LAN name server holds information in cache. With this set-up the resolver gets a response as quickly as possible. Furthermore, this gives extra robustness in two ways. Translating local names is still possible if the connection to the WAN is down, and translation of all names is still possible if the local name server is down, provided that the resolvers can fallback to using the backbone name servers.




Figure 3: Part of the Internet naming hierarchy

Naming convention

To make optimal use of DNS a naming convention is vital. A naming convention defines the architecture of the content. It describes the structure and rules for the names that may be used within the company. A practical naming convention ensures that names need to change as little as possible, and that they have a logical structure. At the same time, the naming convention must also have both a technical and a human dimension. The technical dimension gives the relationship between the application and the infrastructure (see figure 1). The human dimension is concerned with ensuring that the end users and administrators can apply the convention in practice.

Bearing in mind that organizations frequently reorganize, it is advisable to include as few organizational indications in the domain as possible. The geographic spread of a business is often the most stable factor. A convention that recommends <location>.<country>.company.com is logical. A machine in Chicago would then have the name webserver.chi.us.company.com. The convention can also make a link to the network infrastructure. A LAN is connected to the domain <location>.<country>.company.com with the associated network addresses.

In addition to domains, the naming convention must indicate the names of the applications. DNS offers the possibility to separate, in a logical sense, the services from the infrastructure. The e-mail service in Chicago, for example, is called mail.chi.us.company.com. The mail server could have another name, for example, mailserver.chi.us.company.com. This offers the manager of the e-mail service great flexibility. If the mail server in Chicago must undergo maintenance, then a simple change to the DNS takes care that the e-mail service mail.chi.us.company.com is temporarily moved to, for example, mailserver.was.us.company.com. This change is transparent to end users: e-mail carries on being delivered without delays.

DNS implementation

Assuming that the DNS name server, resolvers and content are to be correctly implemented, then a combination of a centrally administered backbone and LAN name servers gives the most effective solution. The configuration of the resolvers can be distributed per LAN, for

corporate level, with the LAN manager remaining responsible for the LAN names and addresses. Therefore, a DNS domain is partly centrally managed, and partly locally managed.

A combination of centrally managed backbone and LAN name servers is the most effective.

example, by using Dynamic Host Configuration Protocol (DHCP) servers. For the implementation of the content, a naming convention is needed that states, among other things, which domain names are allowed. The hostmaster must maintain the names, network addresses, mail routing, and DNS meta-information for each domain. LAN managers often have no knowledge of mail routing and DNS meta-information, and therefore cannot adequately fill this role. This means that mail routing and DNS meta-information must be administered at a

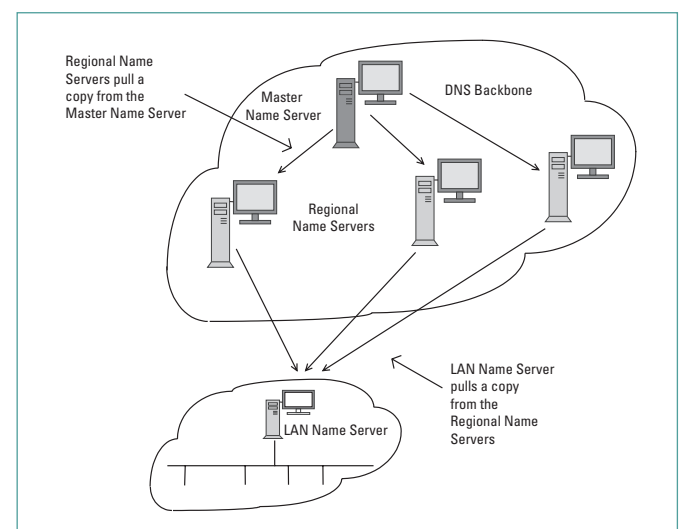


Figure 4: Name server architecture

Tool

The tool used for managing the content must allow for both central and local LAN management. The right tool lets the LAN manager concentrate on the LAN names and network addresses without worrying about mail routing and meta-information. A good tool will automatically check the syntax and semantics of changes made by the LAN manager. In this way, decentralized administration with central control is possible. This allows a company to implement a solid layer of names which can support business critical applications. There are numerous tools on the market that claim to make DNS administration simpler. Many of these, however,

do make administration more difficult. For example, the LAN manager must still know about DNS *internals*, such as the definitions for mail routing, to be able to do his job. Some tools put the emphasis more on configuration management than DNS management. Many businesses use the tools without first sorting out the configuration management process. An organization that does so is saddled with another, often incomplete, configuration management database. There is a tendency to neglect building a decent DNS infrastructure of name servers and resolvers.

The connection with the Internet

A company must, of course, also make names known on the Internet, for example, for e-mail and web servers. It is not wise to 'publish' the complete intranet contents on the Internet. An intranet's DNS content can give competitors more information than you might first realize. Based on the number of machines included in the research department in a particular town, for example, it is possible to get an indication of the number of staff working on a particular project. The solution is to set up a shadow DNS structure which only contains that information necessary for communicating with the Internet. This is called a *split DNS*.

Abuse

The domain name is a company's calling card on the world-wide web. Misuse is made of this by people

who register minor variations on names, banking on users making typing mistakes. This has the greatest consequences for e-mail. An example is the domain shell.com from Shell Oil Company. It is expected that people will forget the second l and type in shel.com. A web pirate has therefore registered shel.com. An e-mail sent to user@shel.com is then sent to the pirate's e-mail machine. In this way, even internal e-mail can be sent outside the organization as the result of a typing mistake, if the appropriate measures are not put in place. In addition to variations, the names of products are also popular with pirates. There are a few technical measures that can be taken to limit the damage as far as possible. Administrative measures are the most effective: companies must treat domain names with the same care they apply to trademarks and patents.

DNS management

For DNS to work to its full potential, an organization must fill a number of roles.

- The *naming authority* is a central role which sets up the naming convention, guards it, and continually keeps it up to date. Each time a new service or location is added to the network, additions are made to the naming convention. The naming authority must have a corrective function if the convention is not being adhered to.
- A *hostmaster* is required for each domain, whose role is to assign names to the locally known network addresses. Considering that the network

addresses are spread over the entire organization, assigning names to network addresses must also be done in a distributed manner. The role of hostmaster must be filled for each LAN. By using the right tool, the hostmaster need not be concerned with complex issues such as mail routing.

- The *backbone administrator* implements the naming convention determined by the naming authority. Using the right tools allows the administration of mail routing and DNS meta-information. The backbone administrator also manages the backbone and the LAN name servers.

A glimpse into the future

DNS technology is still developing. Within this article we have just scratched the surface. The first example of a new development is that Microsoft is embracing DNS technology in Windows 2000.

Windows 2000 uses DNS for finding services on servers. It is also going to use dynamic DNS, which effectively gives the desktop write access to DNS. This brings dangers with it, because an incorrectly configured desktop

could, for example, claim the name of a corporate web server or mail server. Therefore, dynamic DNS must be combined with a proper method for authentication and verification. Another new development is secure DNS, which guarantees the authenticity of DNS responses. The name server adds a signature to each response, which the resolver can verify. A prerequisite for secure DNS is a public key infrastructure.

Conclusions

Companies invest a lot of time and money in network connections (routers, bandwidth, etc) and in applications, such as SAP. The poor relation here is often the names in between. The necessary organizational roles are often not filled, and there is lack of a decent infrastructure. This increases the operational costs of the

intranet, because, for example, changes can not be implemented without involving end users. A good DNS service brings with it a clear divide between LAN management and WAN management. This simplifies the daily administration, increases service availability, and prepares the intranet for the future.

Anton Holleman

Origin provides a full spectrum of information technology services for global corporations and other complex "extended enterprises". With a staff of over 17,000, Origin conducts operations in 32 countries around the world and is headquartered in Eindhoven, the Netherlands. Origin's customers include over one hundred of the world's Fortune 500 firms.

Origin is a member of the Royal Philips Electronics group of companies.

Origin Managed Services is responsible for all of Origin's network services, system management and managed application services, providing a focus on global synergy and efficiency to the best possible benefit of its customers. To achieve this, it uses standardized processes, technology and global management centers in North America, South America, Europe, and Asia Pacific, together with national data centers worldwide, to deliver 7 x 24 hour services.

**Anton Holleman works for Origin Managed Services.
E-mail: Anton.Holleman@nl.origin-it.com
This article first appeared in Dutch in Open Computing 3, April 1999.**

Origin BV
Origin Managed Services
Eindhoven
The Netherlands

dns-services@origin-it.com