WORLD METEOROLOGICAL ORGANIZATION

COMMISSION FOR BASIC SYSTEMS

OPAG ON INFORMATION SYSTEMS & SERVICES

Expert Team on Communication Techniques and Structure

Geneva, 30 Apr – 03 May 2012

ET-CTS/Doc. 3.2.1 (1)
(12 IV 2012)

ITEM 3

ENGLISH only

Adoption of IPv6 at ECMWF

(Submitted by ECMWF)

Summary and purpose of document

This paper presents the situation regarding IPv6 at ECMWF. In particular, one very important application (the dissemination of products) in now IPv6 ready.

ACTION PROPOSED: The meeting is invited to review the document.

TABLE OF CONTENT

1. Introduction

It is now well-known that IPv4 addresses are rapidly being depleted. It is inevitable that sites, especially newly created organizations, will be assigned IPv6 networks by their Internet Service Providers. This is especially true in areas such as the Asia/Pacific region.

As ECMWF produces Mid-Range Weather Forecasts for many organizations throughout the world it is a clear business case that ECMWF needs to be prepared for native IPv6 delivery of products to these destinations.

As a secondary target, and as more and more sites are using IPv6 addresses, there is also the need to allow native IPv6 web access to this web sites.

2. The DMZ Architecture

The Internet DMZ architecture at ECMWF is split into two main areas: DMZ frontend and DMZ backend. These areas are isolated from each other and completely separate Internet and Local Area Network (LAN) network traffic.

Internet traffic is allowed only to/from the external network interfaces of hosts installed on the DMZ, and in a similar manner traffic to/from the LAN is only allowed from the internal interfaces of hosts installed on the DMZ.

Network packets are never allowed to traverse from the LAN to the Internet. All network connections terminate at or are initiated from a DMZ host.

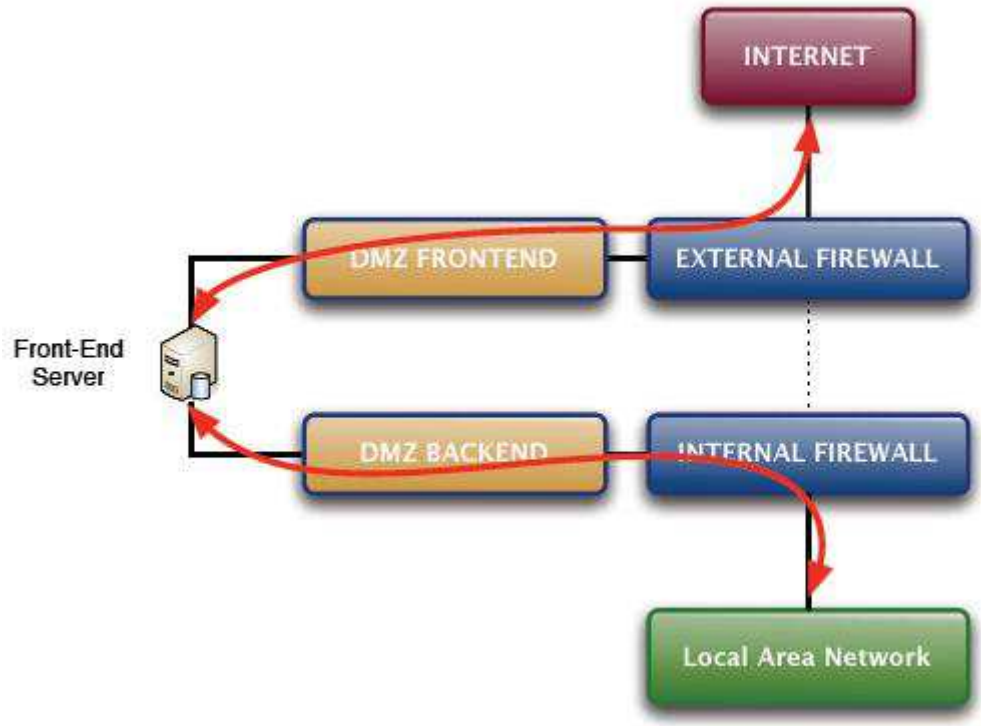

Figure 1: DMZ Architecture

With this architecture, the delivery of ECMWF products to external sites using IPv6 only requires allowing IPv6 data flows initiating or terminating on the external interfaces of DMZ hosts. Data flows between the LAN and the DMZ as well as flows within the LAN itself can still use IPv4.

3. The IPv6 product delivery roadmap

3.1. Obtaining native IPv6 connectivity

Due to the fast depletion of IPv4 addresses Internet Service Providers (ISPs) are very keen to provide native IPv6 connectivity to customers. Obtaining an IPv6 /48 network from the UK Educational and Research Network (UKERNA/JANET) and setting up native connectivity to ECMWF was a quick process delegated mainly to the ISP.

IPv6 addresses were assigned to the Internet facing routers, currently Vyatta[i], and the routing configuration was modified.

As ECMWF has dual resilient links to the ISP it is the responsibility of ECMWF to advertise the BGP routes of the newly created IPv6 network using the appropriate local preference parameters. The procedure to add these routes was similar to the already existing IPv4 BGP routes. The ISP had to make modifications at their point of presence (PoP) routers to accept these newly advertised routes.

3.2. Agreeing internally on an IPv6 addressing scheme

The IPv6 network allocated to ECMWF by the Internet Service Provider is **2001:630:55::/48**

At ECMWF it was decided to split this network into ***/64*** subnets using a three level hierarchical schema:

The levels are defined using the first four hexadecimal digits of the ***/48*** network. These levels are:

- Super-Role (8 bits): Indicates whether address is internal or external,
- Role(8 bits): Indicates type of device such as network equipment, server, workstations etc, and
- Instance number (16 bits): An integer to allow multiple subnets with the same Super-Role and Role.

The idea behind this schema is to make addresses more human readable and easier to locate for the benefit of network support and staff training at ECMWF. As an example, an address such as

< 2001:630:55 >**:ed22::1**

Indicates that this is an **external** address on a device which is accessible from the Internet ('**e**') and is located on a DMZ network ('**d**'). In addition ECMWF staff would recognize (**'::1')** as being a router or firewall address as gateways at ECMWF are always assigned the first available address in the subnet.

3.3. Enabling IPv6 on Firewalls

Only the external firewalls (these are Checkpoint firewalls and they are IPv6 ready) need to be IPv6 enabled as with this architecture the internal firewalls would continue to carry IPv4 traffic only.

Setting up the firewall's IPv6 address and default IPv6 gateway was straight forward. Once this was done defining and applying the rules relevant to IPv6 traffic was no different from any other IPv4 rule.

To avoid problems related to hybrid IPv4-IPv6 communication models it was decided that all forms of tunnelling IPv6 over any other protocol would be explicitly blocked from the beginning. This is in line with standard security practises but more importantly it also keeps IPv6 network flows easier to understand by network support staff and avoids difficult to solve

problems at the host and application level arising in cases where for example a given host may automatically fall back into tunnelling mode (ie IPv6 over IPv4) to gain connectivity under certain (unpredictable) conditions.

3.4. Enabling IPv6 on hosts

Systems that run product delivery software modules and hosts that run web proxy software are the only systems given a dual-stack with both IPv6 and IPv4 addresses and their respective default routes.
All systems are running recent Linux distribution (SLES 11), and again, IPv6 is available “in the box”.

Again, thanks to the simple DMZ architecture, host configuration settings such as NTP, DNS servers, etc do not need any modification as all these network services are supplied to the host via the internal IPv4 only network.

3.5. Enabling IPv6 on the applications

Delivery of products at ECMWF is achieved via an in-house developed dissemination system which runs distributed across a number of hosts in the LAN, Internet DMZ and RMDCN DMZ.

The application runs under Java and therefore the low level code which interacts with the operating system socket services is part of the standard Java JVM and already was IPv6 enabled. At a higher level, some database fields which hold IP destination information needed to be enlarged to cater for the longer IPv6 addresses. Both IPv4 and IPv6 addresses are kept as strings on the applications database.

In addition a modification to the ftp java library was needed to add Extended Passive support (EPSV) and Extended Port support (EPRT) . These are the standard ftp commands to allow IPv6 data channel sockets.

No other modification was done to the dissemination software.

To provide web access to ECMWF staff, the web proxy currently deployed on the DMZ (a product called Delegate) was re-configured to resolve DNS queries using quad-A records (IPv6 records) as a preference. This enables internal workstations to transparently access external IPv6 web sites without any change to their configuration or to the configuration of the browsers. HTTP connections are initiated from the workstations to the proxy using IPv4 and from the proxy to the Internet using IPv6.

3.6. Testing

For product delivery an IPv6 external host located in France was selected and an ftp server was installed. An end to end product delivery stream was configured on the operational ECMWF dissemination software. Daily delivery of products to this host has been occurring continuously during the past 12 months. The number of daily products sent to this host using IPv6 and standard ftp is approximately 40 per day with a total volume of 0.5GB. So far this dissemination stream has never failed.

For standard web access via IPv6 ECMWF took the opportunity offered by the “*World IPv6 Day*” which took place on 8 June 2011 and where service providers such as Google, Facebook and Yahoo enabled IPv6 only networks for their services. All internal staff at ECMWF could access these IPv6 sites transparently via the ECMWF web proxy from their internal (IPv4 only) workstations.

4. Conclusion

The initial IPv6 deployment at ECMWF has been relatively effortless and has caused zero impact on the operational IPv4 network.

This was possible thanks to the simple architecture defined on our DMZ and by sticking only to the strategic main objectives that ECMWF needed to address rather than the more ambitious alternative of a full deployment providing all hosts and workstations at ECMWF with an IPv4/IPv6 dual-stack.

Future work will involve enabling incoming IPv6 for other web services that are currently under development and, maybe, IPv6 access from mobile devices.

[i] Vyatta (www.vyatta.com) is a software solution for high performance routers running on x86 platform.