

Preparing CMS for an IP Version 6 World

David Boyes
MVMRUG
January 27, 2006

Agenda

- Quick Overview of IPv6
- Why IPv6 Matters
- Current State of IPv6 Support in z/VM
- Current State of IPv6 Support in Other IBM OSes
- What Shall We Do?
- A Proposal: A IPv6 Gateway for the VM TCPIP Stack
- A Proposal: A Parallel IPv6 Stack for z/VM
- Construction of the IPv6 Gateway Appliance
- Construction of the Parallel IPv6 Stack
- Outstanding Difficulties

z/VM will have usable IPv6 support.....

What Is IPv6?

- **Internet Protocol version 6 (IPv6)** is a major architectural change in the basic protocol suite of the Internet
- Changes occur in almost every part of the protocol:
 - Addressing and Addressing Formats/Notation
 - Network Segmentation and Routing
 - Host Configuration and Path Discovery
 - Security and Authentication
 - DNS Behavior
 - Socket APIs

Addressing

- The biggest visible change from IPv4 to IPv6 is the length of network addresses.

IPv6 addresses are 128 bits long (as defined by RFC 2373 and RFC 2374), typically written as 32 hex digits. This gives a total of 2128 or 3.4×10^{38} addresses, enough to provide an IPv6 48-bit range of network addresses **for every 10 square meters on earth.**
- IPv6 addresses are usually composed of two logical parts: a 64-bit network prefix, and a 64-bit host-addressing part, which is often automatically generated from the interface MAC address.

Address Notation

- IPv6 addresses are normally written as eight groups of four hexadecimal digits. These are often followed by a slash and the prefix length which indicates a range of IPv6 addresses.

For example,

2001:0db8:85a3:08d3:1319:8a2e:0370:7334

is a valid IPv6 address.

- If a four-digit group is 0000, the zeros may be omitted.
 - Example: 2001:0db8:85a3:0000:1319:8a2e:0370:7344 can be shortened as 2001:0db8:85a3::1319:8a2e:0370:7344.

Address Notation

- Having more than one double-colon abbreviation in an address is invalid because it allows for ambiguity in the number of 0 groups omitted from different parts of the address.
 - For example, consider the following four valid IPv6 addresses:
2001:0000:0000:0000:0000:25de:0000:cade
2001:0000:0000:0000:25de:0000:0000:cade
2001:0000:0000:25de:0000:0000:0000:cade
2001:0000:25de:0000:0000:0000:0000:cade
 - If multiple omissions of 0 groups were allowed, these four could all be abbreviated to the same ambiguous address: 2001::25de::cade.
- Leading zeros in a group can be omitted.
 - Thus 2001:0DB8:02de::0e13 may be shortened to 2001:DB8:2de::e13.

Network Notation

- A network is denoted by the first address in the network and the size in bits of the prefix, separated with a slash.

For example, 2001:1234:5678:9ABC::/64 stands for the network with addresses

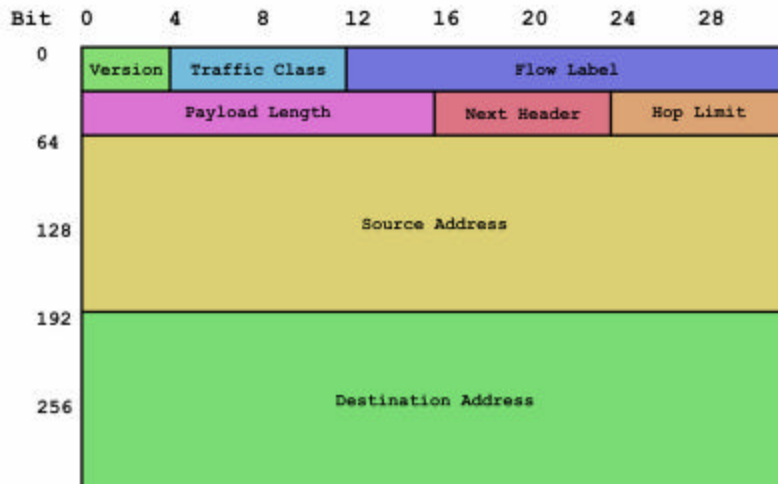
2001:1234:5678:9ABC:0000:0000:0000:0000 through
2001:1234:5678:9ABC:FFFF:FFFF:FFFF:FFFF

- Network segments must be sized in powers of 2 (just as today).

IPv4 Addresses in IPv6

- IPv4 addresses are easily converted to IPv6 format. Three methods exist to do this:
 - Standard IPv6 hexadecimal,
 - Hybrid IPv4 compatible address
 - IPv4 mapped address.

IPv6 Packet Header



Autoconfiguration and Path Discovery

- Basic support for automatic address assignment is built-in to the protocol
- DHCP still used for more extensive configuration parameters
- All systems capable of acting as routers are required to announce themselves
- All end systems must be able to discover the presence of routers and desired MTU size as a basic part of participating in a IPv6 network

Routing Protocols

- Existing IPv4 OSPF routing protocol extended for IPv6 addresses
- New IDRP for inter-domain routing (replaces BGP)

Security and Authentication

- Support for encryption is **mandatory**
- Support for endpoint authentication is **mandatory**
- **Spec includes basic authentication and encryption methods that MUST be implemented, but does not limit use of additional methods that can be negotiated between peers**

IPv6 and the DNS

- IPv6 addresses are represented in the DNS by *AAAA records* (so-called quad-A records). This scheme is defined in [RFC 3596](#).
- The AAAA scheme was effectively standardized on in August 2002 by RFC 3363
- Works pretty much the same way it did before.

Why Should You Care?

- IPv6 provides vastly more addresses for networked devices, allowing, for example, each cell phone and mobile electronic device to have its own address.
 - **IPv4** supports 4.3×10^9 addresses, which is inadequate to give one to every living person.
 - **IPv6** supports 3.4×10^{38} addresses, or 5×10^{28} addresses for each of the roughly 6.5 billion people alive today.
- The US government and many other national governments have specified that all agencies must deploy IPv6 (US Feds by 2008).

State of IPv6 Support in z/VM

- As of z/VM 4.4, the IUCV application socket interface has been enhanced to support IPv6 addresses.
- The stack can receive and send native IPv6 packets
- No provided CMS applications utilize IPv6
- No routing software included with the package can receive or process IDRP or other IPv6-aware routing protocols
- Limited testing so far indicates that resource consumption in the current stack is measurably higher using IPv6 APIs.

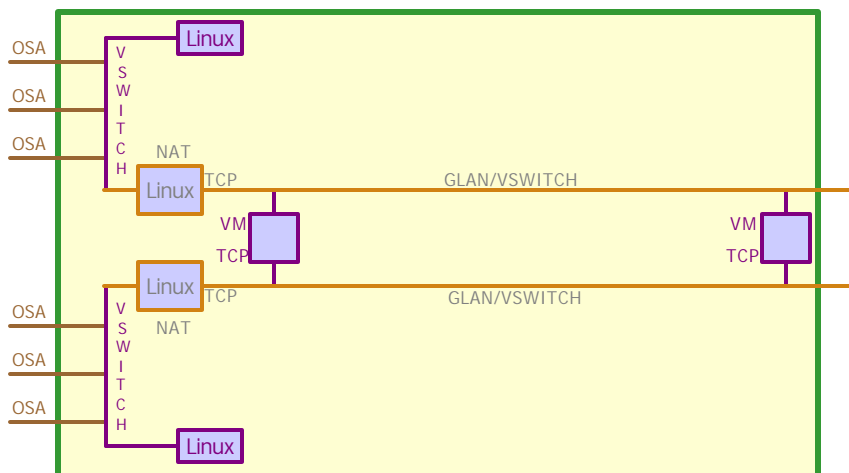
What Shall We Do?

- Option 1: Bridge IPv6 traffic to IPv4 via a Proxy Server
- Option 2: Provide A Native IPv6 Stack and Applications

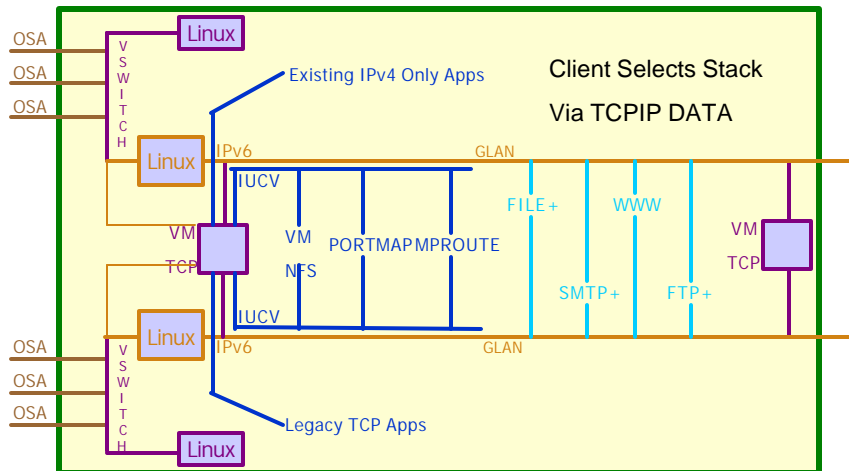
Observations

- Linux IP stack already has a rich IPv6 native implementation
- IPv6-to-IPv4 gateway code already implemented for Linux (NAT-PT, proxy code)
- Existing difficulties with the current stack lie in two (mostly undocumented) areas:
 - Terminal management
 - NFS support

IPv6 Proxy



Parallel IPv6 Stack



© 2006 SNA

Solving Terminal Handling

- VMPassthru (aka PVM) has all the necessary CP and communications method interfaces already implemented... and it accepts multiplexing traffic over a IUCV or CTC connection
- The Linux IUCV driver or CTC driver is easily connected to this protocol (documented in the PVM manual), and a minor modification to the telnet server allows mapping incoming TCP connections to PVM sessions.
- Benefits:
 - Direct VTAM or other legacy comm integration
 - All terminal handling moved out of the TCPIP stack
 - Session scripting and automation
 - Pseudo-session manager

© 2006 SNA

Solving NFS

- Documentation of the IUCV interface
- Not much else we can do – once IBM cooperates with this, many things become possible

Summary

- IPv6 offers dramatically improved capabilities for the future of networked devices and systems
- IPv6 is a deployable and concrete reality, and z/VM will need to have the tools and applications to participate to remain viable
- We have a opportunity to adopt a clean break with the old VM TCP/IP stack at the point of migration, and VM's ability to separate communications from machine makes this easy
- We can provide a immediate alternative solution that employs best of breed technologies without a long development cycle or large IBM expenditure

z/VM will have usable IPv6 support.....

.... when IBM works with us to make it happen.

Q&A/Contact Us

David Boyes
Sine Nomine Associates
www.sinenomine.net
info@sinenomine.net