

ORACLE®

# IPv6: What You Need to Know

(with Linux & Exadata references)

Dan Norris  
Platform Integration MAA Team

Thanks to Kevin Reardon

# Program Agenda

- 1 IPv6 Basics & Syntax
- 2 OEDA Configuration
- 3 Linux Commands

# IPv6 Basics and Syntax

# The IPv6 Address Format

**2001 : DB8 : 5555 : B77B : 0000 : 0000 : C058 : 6301**

- The IPv6 Address consists of 8 fields
- A field is four Hexadecimal numbers (0-F)
  - The first **64 bits (16 bytes)** are the **Routing Prefix**
  - The last **64 bits (16 bytes)** are the **Interface ID**
  - There is no Subnet Mask!
- Routing Prefix most significant bits are encoded
  - Type of packet - Unicast, Multicast, Tunnel, etc.
  - Where on the Planet it is going to or came from

# Notation: Double Colon

- **At the end** - 2001:DB8:5678:9ABC::
  - 2001:DB8:5678:9ABC:0000:0000:0000:0000
- **In the middle** - 2001:DB8:5678:9ABC::00A1:0001
  - 2001:DB8:5678:9ABC:0000:0000:00A1:0001
- **At the front** - ::1234:5678:9ABC:DEF0
  - 0000:0000:0000:0000:1234:5678:9ABC:DEF0
- The longest run of zeros is chosen to double colon
- Double colon is used only once
  - ::DB8::5678 makes no sense.

# Leading Zeros

- Leading Zeros in a field can be left out (following are equivalent)
  - 2001:DB8:5678:9ABC::00A1:0001
  - 2001:DB8:5678:9ABC::A1:1
- Dropping Leading Zeros and using a Double Colon can be done at the same time

# Netmask

- IPv4: 255.255.248.0 (aka /21), 255.255.255.0 (aka /24), etc
- IPv6: easy - operationally, it should *\*always\** be /64 (except when it isn't)
  - Always written that way too
  - see RFC 6177 - <https://tools.ietf.org/html/rfc6177>



# Link Local

- IPv4 equivalent is 169.254.0.0/16
- Primarily used for zero configuration
  - Every Network Interface will configure with this address
  - RAC HAIP uses the IPv4 version
    - Note - RAC does not yet support IPv6 on private network.
- Routers will not forward any traffic using this address
- Address starts with FE80 ::
  - FE80 :: {Interface Identifier = MAC}
- Still have problems with multiple interfaces on different subnets

# Reserved networks

- Common IPv4 equivalents are 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16
  - See RFC 5735 - <https://tools.ietf.org/html/rfc5735>
- IPv6 - FC00::/7
  - See RFC 4193 - <https://tools.ietf.org/html/rfc4193>
  - Not quite the same, but as close as you can get
- With 2000 IPv6 IPs per square meter, why use non-routable IPs?
  - roughly 42 octillion IPv6 IPs per person on the planet

# Documentation IP Ranges

- Networks reserved for using in documentation
- Ever get a wallet with a mock SS card?
- IPv4 - 192.0.2.0/24, 198.51.100.0/24, 203.0.113.0/24
  - See RFC 5737 - <https://tools.ietf.org/html/rfc5737>
- IPv6 - 2001:DB8::/32
  - See RFC 3849 - <https://tools.ietf.org/html/rfc3849>

# Other IPv6 Unicast Addresses

- Loop Back

- 0:0:0:0:0:0:0:1 or ::1/128

- Just like the 127.0.0.0/24 address in IPv4

- Is never assigned to any physical interface (lo interface)

- Unspecified - all zeros

- 0:0:0:0:0:0:0:0 or ::0/128

- Used before DHCP configuration

# HTTP and scp

- Colon conflicts with oh, so much
  - URI's and URL's
- To solve: enclose IPv6 address in square brackets
  - `[2001:DB8:5678:9ABC::A1:1]`
  - `http://[2001:DB8:5678:9ABC::A1:1]:8888/`
  - `scp file.zip root@[2001:DB8:5678:9ABC::A1:1]:/path/file.zip`

# Case Does Not Matter

- The case of the HEX numbers does not matter!
- An “E” is the same as an “e”
- Mixing case does not matter
  - 2001:DB8:5678:9ABC::A1:1
  - 2001:db8:5678:9abc::a1:1
  - 2001:Db8:5678:9AbC::a1:1
- Wikipedia says lower case is preferred (it never lies)
- Keep in mind, **numbers** ARE case sensitive 😊

# Words With IP Addresses

- 2001:db8::face:feed
  - 2001:db8::f00d:2ea1
  - 2001:db8::1ce:1ce:babe
  - 2001:db8::dead:beef
  - 2001:db8::bad:c0de
- 
- <http://sophiedogg.com/funny-ipv6-words/>

# DNS Terminology

- IPv4

- "A" records - maps name to IPv4 address

- dm01adm01 -> 198.51.100.2

- "PTR" records - aka "reverse DNS records" - maps IPv4 address to name

- 198.51.100.2 -> dm01adm01

- IPv6

- "AAAA" records (aka "quad A") - maps name to IPv6 address

- dm02adm01 -> 2001:DB8:5678:9ABC::00A1:0001

- "PTR" records - same as IPv4

- 2001:DB8:5678:9ABC::00A1:0001 -> dm02adm01



# Real World IPv6

```
$ host www.oracle.com
```

```
www.oracle.com is an alias for ipv6-www.oracle.com.edgekey.net.
```

```
ipv6-www.oracle.com.edgekey.net is an alias for  
e11582.dscx.akamaiedge.net.
```

```
e11582.dscx.akamaiedge.net has address 23.211.116.177
```

```
e11582.dscx.akamaiedge.net has IPv6 address 2001:668:108:986::2d3e
```

```
e11582.dscx.akamaiedge.net has IPv6 address 2001:668:108:999::2d3e
```

```
$
```

```
$ host www.google.com
```

```
www.google.com has address 216.58.219.36
```

```
www.google.com has IPv6 address 2607:f8b0:4007:808::2004
```

```
$
```

# Real World IPv6

```
# ip addr list eth0

2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP
qlen 1000

    link/ether 00:0c:29:39:22:56 brd ff:ff:ff:ff:ff:ff
    inet 172.16.85.138/24 brd 172.16.85.255 scope global eth0
    inet6 2001:db8::3/64 scope global
        valid_lft forever preferred_lft forever
    inet6 fe80::20c:29ff:fe39:2256/64 scope link
        valid_lft forever preferred_lft forever

[root@ol6-7-64bit ~]#
```

# IPv6 @ Oracle

- Oracle corporate VPN
  - Issues IPv4 and IPv6 IPs when connecting via VPN
- Oracle internal test labs (some, not all)
  - IPv6 subnet for every IPv4 subnet

# OEDA IPv6 Configuration

# OEDA IPv6 Changes

- Same OEDA for all configs (IPv6 or IPv4) – Sept 2015 or later
- IP input fields have tooltips with good information
- Do not use "subnet mask" pulldown for IPv6
  - It will grey out once you enter proper IPv6 syntax in the address field
- Use /64 for all router IP address designations

# OEDA Valid Configurations

- Valid configurations (Private IB is always IPv4):
  - Mgmt IPv4, Client IPv4 (aka "the old way")
  - Mgmt IPv6, Client IPv4
  - Mgmt IPv4, Client IPv6
  - Mgmt IPv6, Client IPv6
- Invalid configurations:
  - Mgmt IPv4+IPv6 (one or the other – not both)
  - Client IPv4+IPv6 (one or the other – not both)

# OEDA IPv6 Changes

Name : **Admin**  Bonded

Subnet Mask : 255.255.255.0  Non Bonded

Gateway :

Network Format : Enter the gateway IP address for admin network either IPv4 with a Subnet from above or IPv6 including a /CIDR

Name : **Client**  Bonded

# OEDA IPv6 Changes

## Subnet 1

Name : **Admin**  Bonded

Subnet Mask :   Non Bonded

Gateway : 2001:0db8:418:1ea5::1/64

Admin Network Format :  1/10 Gbit Copper Base-T  10 Gbit Optical

## Subnet 2

Name : **Client**  Bonded

Subnet Mask :   Non Bonded

Gateway : 2001:0db8:418:1ea4::1/64

Client Network Format :  1/10 Gbit Copper Base-T  10 Gbit Optical

## Subnet 3

Name : **Private**  Bonded

Subnet Mask : 255.255.248.0  Non Bonded

Private Network Format :  InfiniBand



# OEDA IPv6 Changes

## Administration Network

Starting IP Address for Pool :

**[Valid network range : 2001:db8:418:1ea5:0:0:0:2 - 2001:db8:418:1ea5:ffff:ffff:ffff:fffe]**

# OEDA IPv6 Changes

## Client Ethernet Network

Starting IP Address for Pool :

**[Valid network range : 2001:db8:418:1ea4:0:0:0:2 - 2001:db8:418:1ea4:ffff:ffff:ffff:fffe]**

Pool Size :

Ending IP Address for Pool :

Is the default gateway for database servers

Defines the hostname for the database servers

# OEDA IPv6 Changes

IB is still IPv4 only

## InfiniBand Network

Starting IP Address for Pool :

Pool Size :

Ending IP Address for Pool :



Enable Active Bonding on Compute node Network

# OEDA IPv6 Changes

## Exadata X5-2 Compute Node 8

Rack 1 - Rack Location 26

Admin Name : test01db08.us.oracle.com

Admin IP : 2001:db8:418:1ea5:0:0:0:207

Ilom Name : test01db08-c.us.oracle.com

Ilom IP : 2001:db8:418:1ea5:0:0:0:21d

Priv Name : test01db08-priv1.us.oracle.com

Priv IP : 192.168.10.15

## Exadata Cell Node HC 4TB 1

Rack 1 - Rack Location 2

Admin Name : test01cel01.us.oracle.com

Admin IP : 2001:db8:418:1ea5:0:0:0:208

Ilom Name : test01cel01-c.us.oracle.com

Ilom IP : 2001:db8:418:1ea5:0:0:0:21e

Priv Name : test01cel01-priv1.us.oracle.com

Priv IP : 192.168.10.17

# OEDA IPv6 Changes

## Cluster acs-clu1

SCAN Name : test01-scan

SCAN IP 1 : 2001:db8:418:1ea4:0:0:0:213

2 : 2001:db8:418:1ea4:0:0:0:214

3 : 2001:db8:418:1ea4:0:0:0:215

### Compute Node 1

Client Name : test0101

IP : 2001:db8:418:1ea4::203

VIP Name : test0101-vip

IP : 2001:db8:418:1ea4:0:0:0:204

### Compute Node 2

Client Name : test0102

IP : 2001:db8:418:1ea4:0:0:0:205

VIP Name : test0102-vip

IP : 2001:db8:418:1ea4:0:0:0:206

# Linux Command Changes

# Linux Commands: Changed & Deprecated commands

- Commands changed

IPv4	IPv6
iptables	iptables6
arp	ip neighbor show
ping	ping6

- Deprecated commands (OL5->OL6, not really IPv4 vs IPv6)

OL5	OL6
ifconfig	ip
nslookup	host, dig

# Linux Commands: Flags

- Some utilities have IP version options

IPv4	IPv6
<code>netstat -A inet -r -n</code>	<code>netstat -A inet6 -r -n</code>
<code>traceroute -4</code>	<code>traceroute -6</code>
<code>scp -4</code>	<code>scp -6</code>
<code>ssh -4</code>	<code>ssh -6</code>
<code>host -4</code>	<code>host -6</code>
<code>iperf -4</code>	<code>iperf -6</code>

- If DNS responds with both, the OS (Linux, Windows) tries IPv6 first



# Other Notation

- When scripts use command output

IPv4	IPv6
INET or inet	INET6 or inet6
. (period)	: (colon)

# Configuration Files

`/etc/sysconfig/network`

```
NETWORKING_IPV6=yes
IPV6FORWARDING=no
IPV6INIT=yes
IPV6_AUTOCONF=yes
IPV6_ROUTER=no
IPV6_AUTOTUNNEL=no
```

`/etc/sysconfig/network-scripts/ifcfg-{int}`

```
IPV6INIT=yes
IPV6FORWARDING=no
IPV6ADDR=2001:db8:418:1ea5:4fff:ffff:ffff:9f03
```

# Configuration Files

## /etc/hosts

##IPv6 address	FQDN	Alias
2001:DB8:5678:9ABC::00A1:0001	dm02adm01.us.oracle.com	dm02adm01
198.51.100.2	dm01adm01.us.oracle.com	dm01adm01

Routing Config files (for eth0, for example):

`/etc/sysconfig/network-scripts/rule6-eth0`

`/etc/sysconfig/network-scripts/route6-eth0`

# host

```
host -4 <dnsname>    ## uses IPv4 to DNS
host -6 <dnsname>    ## uses IPv6 to DNS
host <dnsname>       ## uses IPv4 to DNS
```

By default, the host command will search for all records (IPv4 or IPv6)

# traceroute -6

```
traceroute -6 2001:DB8:5678:9ABC::00A1:0001
```

```
traceroute -6 <dnsname>
```

# ip – some examples

```
ip neighbor list
```

```
ip address list ### can also abbreviate: "ip addr list"
```

```
ip route list table all
```

```
ip rule list
```

# About Exadata Support for IPv6

- Minimum Exadata image version: 12.1.2.2.0
- Supported for Ethernet only (not on IB)

# References

- MOS Note: IPv6 support status on Exadata Database Machine (Doc ID 2056895.1)
- MOS Note: Migrating IPv4 to IPv6 on Exadata Database Machine (Doc ID 2056690.1)



## Safe Harbor Statement

The preceding is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

# **Hardware and Software Engineered to Work Together**

ORACLE®