

# IPv6 in Basic Networking



**apachecloudstack**<sup>™</sup>  
open source cloud computing



**pcextreme**

# Who am I?

- Wido den Hollander (1986)
- Co-owner and CTO @ PCextreme B.V.
  - Dutch hosting company
- CloudStack VP
  - Integrated Ceph into CloudStack
  - Work on the KVM Hypervisor code
  - Integrated IPv6 in Basic Networking



# Who is PCextreme B.V.?

- Dutch hosting company since 2004
- 25 employees
- 50.000 customers
- >100.000 shared hosting packages
- Running a *public cloud* on Apache CloudStack with IPv6 and Ceph storage
- CloudStack zones in Amsterdam, Barcelona, Miami, Tokyo and Los Angeles



# Hello IPv6!

Did you dare to touch it yet?



# IPv6

```
if (ipv4addresses.length < 5%) {  
    System.out.println("IPv4 is almost exhausted");  
}
```



# IPv4 vs IPv6

IPv4	IPv6
32-bit (4.294.967.296 addresses)	128-bit (3.402823669×10 <sup>38</sup> addresses)
/24 (255 addresses)	/64 (1.844674407×10 <sup>19</sup> addresses)
Broadcast	Multicast
ARP	Neighbor Discovery
DHCP	StateLess Address AutoConfiguration
NAT	Routed Addresses
iptables	ip6tables
ICMP is a nice to have	ICMP(v6) is <b>mandatory</b>
Virtual Routers (VRRP)	Router Advertisements
169.254.0.0/16	fe80::/10
0.0.0.0/0	::/0



# 2001:db8::/32

Reserved for example and documentation use in  
RFC3849



# No more NAT

With IPv6 each host on the Internet will get a publicly routed address

This means that there is no more need for NAT

Routers will become true **routers** again instead of *packet translators*





# Firewalling

NAT != Firewall

With IPv6 you can still use stateful firewalling

ip6tables with *ESTABLISHED,RELATED*



# Link-Local

fe80::/10

Mandatory for each IPv6 host and is used for communication in that (Layer 2) network segment

Router Advertisements, Gateways, Neighbor Discovery, Duplicate Address Detection all work over Link-Local



# Link-Local

```
root@ubuntu1604:~# ip -6 addr show dev ens3
2: ens3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 state UP qlen 1000
    inet6 2001:db8:100:0:45b:2cff:fe00:bb/64 scope global
        valid_lft 86391sec preferred_lft 14391sec
    inet6 fe80::45b:2cff:fe00:bb/64 scope link
        valid_lft forever preferred_lft forever
root@ubuntu1604:~#
```



# Link-Local

```
root@ubuntu1604:~# ip -6 route show
2001:db8:100::/64 dev ens3 proto kernel metric 256 expires 86385sec
fe80::/64 dev ens3 proto kernel metric 256 pref medium
default via fe80::5054:ff:fe59:c2b3 dev ens3 proto ra metric 1024
expires 165sec hoplimit 64 pref medium
root@ubuntu1604:~#
```



# Link-Local

```
root@ubuntu1604:~# ping6 -c 3 fe80::5054:ff:fe59:c2b3%ens3
PING fe80::5054:ff:fe59:c2b3%ens3(fe80::5054:ff:fe59:c2b3) 56 data bytes
64 bytes from fe80::5054:ff:fe59:c2b3: icmp_seq=1 ttl=64 time=0.719 ms
64 bytes from fe80::5054:ff:fe59:c2b3: icmp_seq=2 ttl=64 time=0.644 ms
64 bytes from fe80::5054:ff:fe59:c2b3: icmp_seq=3 ttl=64 time=0.942 ms

--- fe80::5054:ff:fe59:c2b3%ens3 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2004ms
rtt min/avg/max/mdev = 0.644/0.768/0.942/0.128 ms
root@ubuntu1604:~#
```



# Router Advertisements

Using multicast routers advertise themselves on the network

They advertise the prefix/subnet and their *Link-Local Address*

*(And lifetime, priority, DNS servers, DNS domain)*



# Router Advertisements

*Hello! I am a router in this network*

*My address is fe80::5054:ff:fe59:c2b3*

*The subnet in this network is 2001:db8:100::/64*

*Oh, if you wanted to know, the DNS servers you can use are  
2001:db8:200:1::53 and 2001:db8:200:2::53*



# StateLess Address AutoConfiguration

Routers advertise the subnet for a network

Using a combination of the subnet and it's MAC  
the host generates a address

This is *Autoconfiguration* **without** a database





# StateLess Address AutoConfiguration

```
root@ubuntu1604:~# ip addr show dev ens3
2: ens3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc
pfifo_fast state UP group default qlen 1000
    link/ether 06:5b:2c:00:00:bb brd ff:ff:ff:ff:ff:ff
    inet6 2001:db8:100:0:45b:2cff:fe00:bb/64 scope global
        valid_lft 86368sec preferred_lft 14368sec
    inet6 fe80::45b:2cff:fe00:bb/64 scope link
        valid_lft forever preferred_lft forever
root@ubuntu1604:~#
```



# Linux and BSD

Make sure Privacy Extensions are disabled

```
net.ipv6.conf.all.use_tempaddr = 0  
net.ipv6.conf.default.use_tempaddr = 0
```



# Windows

## Use MAC as identifier

```
netsh interface ipv6 set privacy state=disabled store=persistent  
netsh interface ipv6 set global randomizeidentifiers=disabled store=persistent
```



# Basic Networking

Directly connected to the network/internet



# Address selection

We know the subnet and MAC!



# Obtaining the address

The Management Server calculates the IPv6 Address the Instance will obtain using the subnet and MAC address



# Obtaining the address

```
mysql> select * from vlan \G
***** 1. row *****
      id: 1
      uuid: 5e8b1050-ab7c-4fcc-8c67-cac1528cdfd5
      vlan_id: vlan://untagged
      vlan_gateway: 192.168.200.1
      vlan_netmask: 255.255.255.0
      description: 192.168.200.10-192.168.200.250
      vlan_type: DirectAttached
      data_center_id: 1
      network_id: 204
      physical_network_id: 200
      ip6_gateway: 2001:db8:100::1
      ip6_cidr: 2001:db8:100::/64
      ip6_range: NULL
      removed: NULL
      created: 2017-04-11 05:58:55
1 row in set (0.00 sec)

mysql>
```



# Obtaining the address

NIC 1 (Default)	
ID	fc7d4024-4af4-418a-baba-6545494f39a4
Network Name	defaultGuestNetwork
Type	Shared
IP Address	192.168.200.95
Secondary IPs	
Gateway	192.168.200.1
Netmask	255.255.255.0
IPv6 IP Address	2001:db8:100:0:45b:2cff:fe00:bb
IPv6 Gateway	2001:db8:100::1
IPv6 CIDR	2001:db8:100::/64





# Security Grouping

Works just like with IPv4



# Security Grouping

Details   Ingress Rule   Egress rule

Add by:  
 CIDR    Account

Protocol	Start Port	End Port	CIDR	Add
TCP	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="button" value="Add"/>
TCP	22	22	0.0.0.0/0	<input type="button" value="tag"/> <input type="button" value="x"/>
ICMP	-1	-1	0.0.0.0/0	<input type="button" value="tag"/> <input type="button" value="x"/>
TCP	80	80	0.0.0.0/0	<input type="button" value="tag"/> <input type="button" value="x"/>
TCP	22	22	:::0	<input type="button" value="tag"/> <input type="button" value="x"/>
ICMP	-1	-1	:::0	<input type="button" value="tag"/> <input type="button" value="x"/>
TCP	80	80	:::0	<input type="button" value="tag"/> <input type="button" value="x"/>



# DEMO

Live demos are dangerous. They always fail..



# TODO

- System VMs should get IPv6
- Cloud-init doesn't work over IPv6 (config drive?)
- Prefix Delegation
  - Route a subnet to a Instance
- Advanced Networking (?)



# When

Apache CloudStack 4.10



# Thank you!



**pcextreme**

- Twitter: @widodh
- E-Mail: wido@widodh.nl
- Blog: <https://blog.widodh.nl/>
  
- <https://www.pcextreme.com/>

