

# Welcome

## Case Study – IPv6 Challenges for Cloud Service Providers

**IPv6 Business  
Conference**

Organized by  
SWISS  
**IPv6**  
COUNCIL

| 2016  
**June, 16**

Manuel Schweizer  
CEO

**cloudscale.ch**

# Agenda


1. Initial Situation
2. Getting Started
3. Status Quo
4. Lookout
5. Questions

# Agenda

- 1. Initial Situation**
2. Getting Started
3. Status Quo
4. Lookout
5. Questions

# Background Information

- Founded 2014
- Swiss IaaS Provider
- Linux Cloud Servers (VMs)
- Focus on Simplicity



**„For Developers Who Care“**

# Lineup

- Ubuntu 14.04 LTS
- OpenStack Kilo
- Brocade VDX
- **No IPv6**



**BROCADE**®

The Brocade logo, featuring the word "BROCADE" in a bold, black, uppercase, sans-serif font, followed by a red stylized "B" logo with a registered trademark symbol.

cloudscale.ch

The CloudScale logo, featuring the word "cloudscale" in a blue, lowercase, sans-serif font, followed by ".ch" in a green, lowercase, sans-serif font.

# Agenda

1. Initial Situation
- 2. Getting Started**
3. Status Quo
4. Lookout
5. Questions

# Goals

- „Do it right“
- Ease of Operation
- Scalable Design & Implementation
- Support Different Operating Systems
- Release by Mid June

# Brainstorming

- Address Allocation
  - 1 IPv6 address per VM
  - /64 per customer
  - /64 per virtual server
- Address Modes
  - Static
  - SLAAC
  - DHCPv6



# Brainstorming

- Address Allocation
  - 1 IPv6 address per VM
  - /64 per customer
  - **/64 per virtual server**
- Address Modes
  - Static
  - SLAAC
  - **DHCPv6**

# First Attempt

- /64 per VM
- DHCPv6

Let's see where that got us...

# Prefix Allocation

## **Problem**

OpenStack allocates random IPv6 address

## **Solution**

Assign a /64 per VM in OpenStack

# Scalability

## Problem

No. of unicast routes on (our) L3 ToR switches

Maximum IPv4 unicast routes	12,000
Maximum IPv6 unicast routes	3,000

Source: brocade.com

## Solution

Create a /48 that contains all those /64s

# Default Gateway

## Problem

No routing information in DHCPv6 (yet)

DHCPv6 Draft: <https://tools.ietf.org/html/draft-droms-dhc-dhcpv6-default-router-00>

Best Practice: <https://tools.ietf.org/html/rfc4861> (ND for IPv6)

## Solution

Router advertisements (RA) on L3 ToR switches

# Router Advertisement (RA)

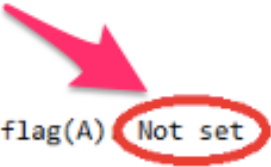
## Problem

/48 prefix information with „A“ flag set

## Solution

Advertise with no-autoconfig => „A“ flag not set

```
4 ICMPv6 Option (Prefix information : [REDACTED]::/64)
  Type: Prefix information (3)
  Length: 4 (32 bytes)
  Prefix Length: 64
  4 Flag: 0x80
    1... .... = On-link flag(L): Set
    .0.. .... = Autonomous address-configuration flag(A) Not set
    ..0. .... = Router address flag(R): Not set
    ...0 0000 = Reserved: 0
```



Source: insinuator.net

# TAC Case

```
ipv6 nd prefix 2a06:c01:abcd::/48 infinite infinite no-autoconfig
```

## Process

2016-04-05: Case opened

2016-04-06: Escalation

2016-04-07: Confirmation

2016-04-12: Workaround (do not shut/no shut!)

2016-04-18: CCE patch required

2016-06-07: CCE patch available

# SLAAC

## Problem

VM is still doing SLAAC

## Solution

Send RA with „M“ flag set

1-bit "Managed address configuration" flag. When set, it indicates that addresses are available via Dynamic Host Configuration Protocol [[DHCPv6](#)].

If the M flag is set, the O flag is redundant and can be ignored because DHCPv6 will return all available configuration information.

Source: RFC4861



# DNS Resolvers

## Problem

VM ignores DNS resolvers received by DHCPv6

## Solution

Send RA with „O“ flag set

1-bit "Other configuration" flag. When set, it indicates that other configuration information is available via DHCPv6. Examples of such information are DNS-related information or information on other servers within the network.

Source: RFC4861

# High-Availability

## Problem

VM receives RA with router link-local address

## Solution

Set up HSRP/CARP/VRRP (why not fe80::1?) and disable router link-local advertisement

# TAC Case (2)

`ipv6 vrrp-suppress-interface-ra`

## Process

2016-06-08: Reported

2016-06-09: Case opened

2016-06-12: Escalation

2016-06-13: Confirmation

2016-06-13: Workaround (ra-interval / ra-lifetime)

2016-06-15: CCE patch required

# Scalability

## Problem

~250 IPv6 addresses per DHCP agent per port

## Solution

Back to the roots: One /64 per L2 domain

+ Prefix routing up to /48 per VM

# Security / Stability

## **Problem**

Issues with OpenStack and IPv6

## **Solution**

Upgrade OpenStack from Kilo to Liberty

# Operating Systems

## Problem

Lack of proper support for DHCP dual-stack

## Solution

### Ubuntu 16.04 LTS & Debian 8

```
auto eth0
iface eth0 inet dhcp
iface eth0 inet6 dhcp
# A timing issue prevents the DHCP client from binding to the interface
# before DAD has completed.
# See: https://bugs.launchpad.net/ubuntu/+source/isc-dhcp/+bug/1447715
# Workaround: Wait 2 seconds for DAD to complete.
pre-up sleep 2
```

# Operating Systems (2)

## Solution

### Ubuntu 14.04

```
auto eth0
iface eth0 inet dhcp
iface eth0 inet6 dhcp
# The DHCP agent automatically takes down the link after releasing the
# IPv4 address. To avoid hanging for a long time, take the link back up
# again to allow the IPv6 address to be released immediately.
# See: https://bugs.launchpad.net/ubuntu/+source/ifupdown/+bug/1013597
pre-down ip link set dev eth0 up
# We need to explicitly accept router advertisements because this gets
# disabled by ifupdown.
accept_ra 1
```



Corin Langosch (ipfo) wrote on 2015-10-05:

#26

This is still broken on latest ubuntu server 14.04. No wonder that ipv6 gets no adoption when one needs hours (or dig through bug reports) to get a basic ipv4/ipv6 dual stack setup working.

Source: [bugs.launchpad.net](https://bugs.launchpad.net)

cloudscale.ch

# Operating Systems (3)

## Solution

### CentOS 7 & Fedora 23 (base)

```
DEVICE="eth0"  
BOOTPROTO="dhcp"  
ONBOOT="yes"  
TYPE="Ethernet"  
USERCTL="yes"  
PEERDNS="yes"  
IPV6INIT="yes"  
DHCPV6C="yes"  
PERSISTENT_DHCLIENT="1"
```

## Good News

Many modern OS work out of the box!



# Agenda

1. Initial Situation
2. Getting Started
- 3. Status Quo**
4. Lookout
5. Questions

# Status Quo

- GA of IPv6 at [cloudscale.ch](https://cloudscale.ch)
  - Released 2016-06-15
  - One IPv6 address per VM
  - Up to /48 routed to this IPv6 address

# Demo

Launch of a VM with IPv6 at [cloudscale.ch](https://cloudscale.ch)

# Demo

cloudscale.ch

Balance: 24.00 CHF | manuel@cloudscale.ch

Servers

Billing

Support

Launch Server

my-swiss-cloud-server

Hint: Must only contain letters (a-z), digits (0-9) and hyphens (-) but must not start or end with a hyphen.

Compute Flavor

Flex-2

1 vCPU  
2 GB Memory  
1 CHF per day

Flex-4

2 vCPUs  
4 GB Memory  
1.5 CHF per day

Flex-8

4 vCPUs  
8 GB Memory  
3 CHF per day

Flex-16

6 vCPUs  
16 GB Memory  
6 CHF per day

show all compute flavors

Storage Capacity

50

10 50 100 150 200 250 300 350 400 more

SSD Storage

50 GB Storage  
0.5 CHF per day

Operating System

Ubuntu

14.04 LTS

Debian

Select Version

CentOS

Select Version

Arch

Select Version

show all operating systems

Network Settings

☒ Public Network Interface

☒ Enable IPv4

☒ Enable IPv6

☐ Private Network Interface

Server Location

cloudscale.ch

# Goals Reached?

- „Do it right“
  - /64 on link ✓
  - Routed up to /48
- Ease of Operation ✓
- Scalable Design & Implementation ✓
- Support Different Operating Systems ✓
- Release by Mid June ✓

# Agenda

1. Initial Situation
2. Getting Started
3. Status Quo
- 4. Lookout**
5. Questions

# Lookout

- Reverse DNS
- Automate routed prefixes
- Investigate DHCPv6-PD
  - DHCPv6 Relay Agent Prefix Delegation Notification
- Our Infrastructure
- Security

# Agenda

1. Initial Situation
2. Getting Started
3. Status Quo
4. Lookout
- 5. Questions**



# Questions



# Thank you!

I am looking forward to your feedback:  
**[manuel.schweizer@cloudscale.ch](mailto:manuel.schweizer@cloudscale.ch)**

 cloudscale.ch

# Thanks to all our Sponsors

