



IPV6 + OPENSTACK NEUTRON

About Me

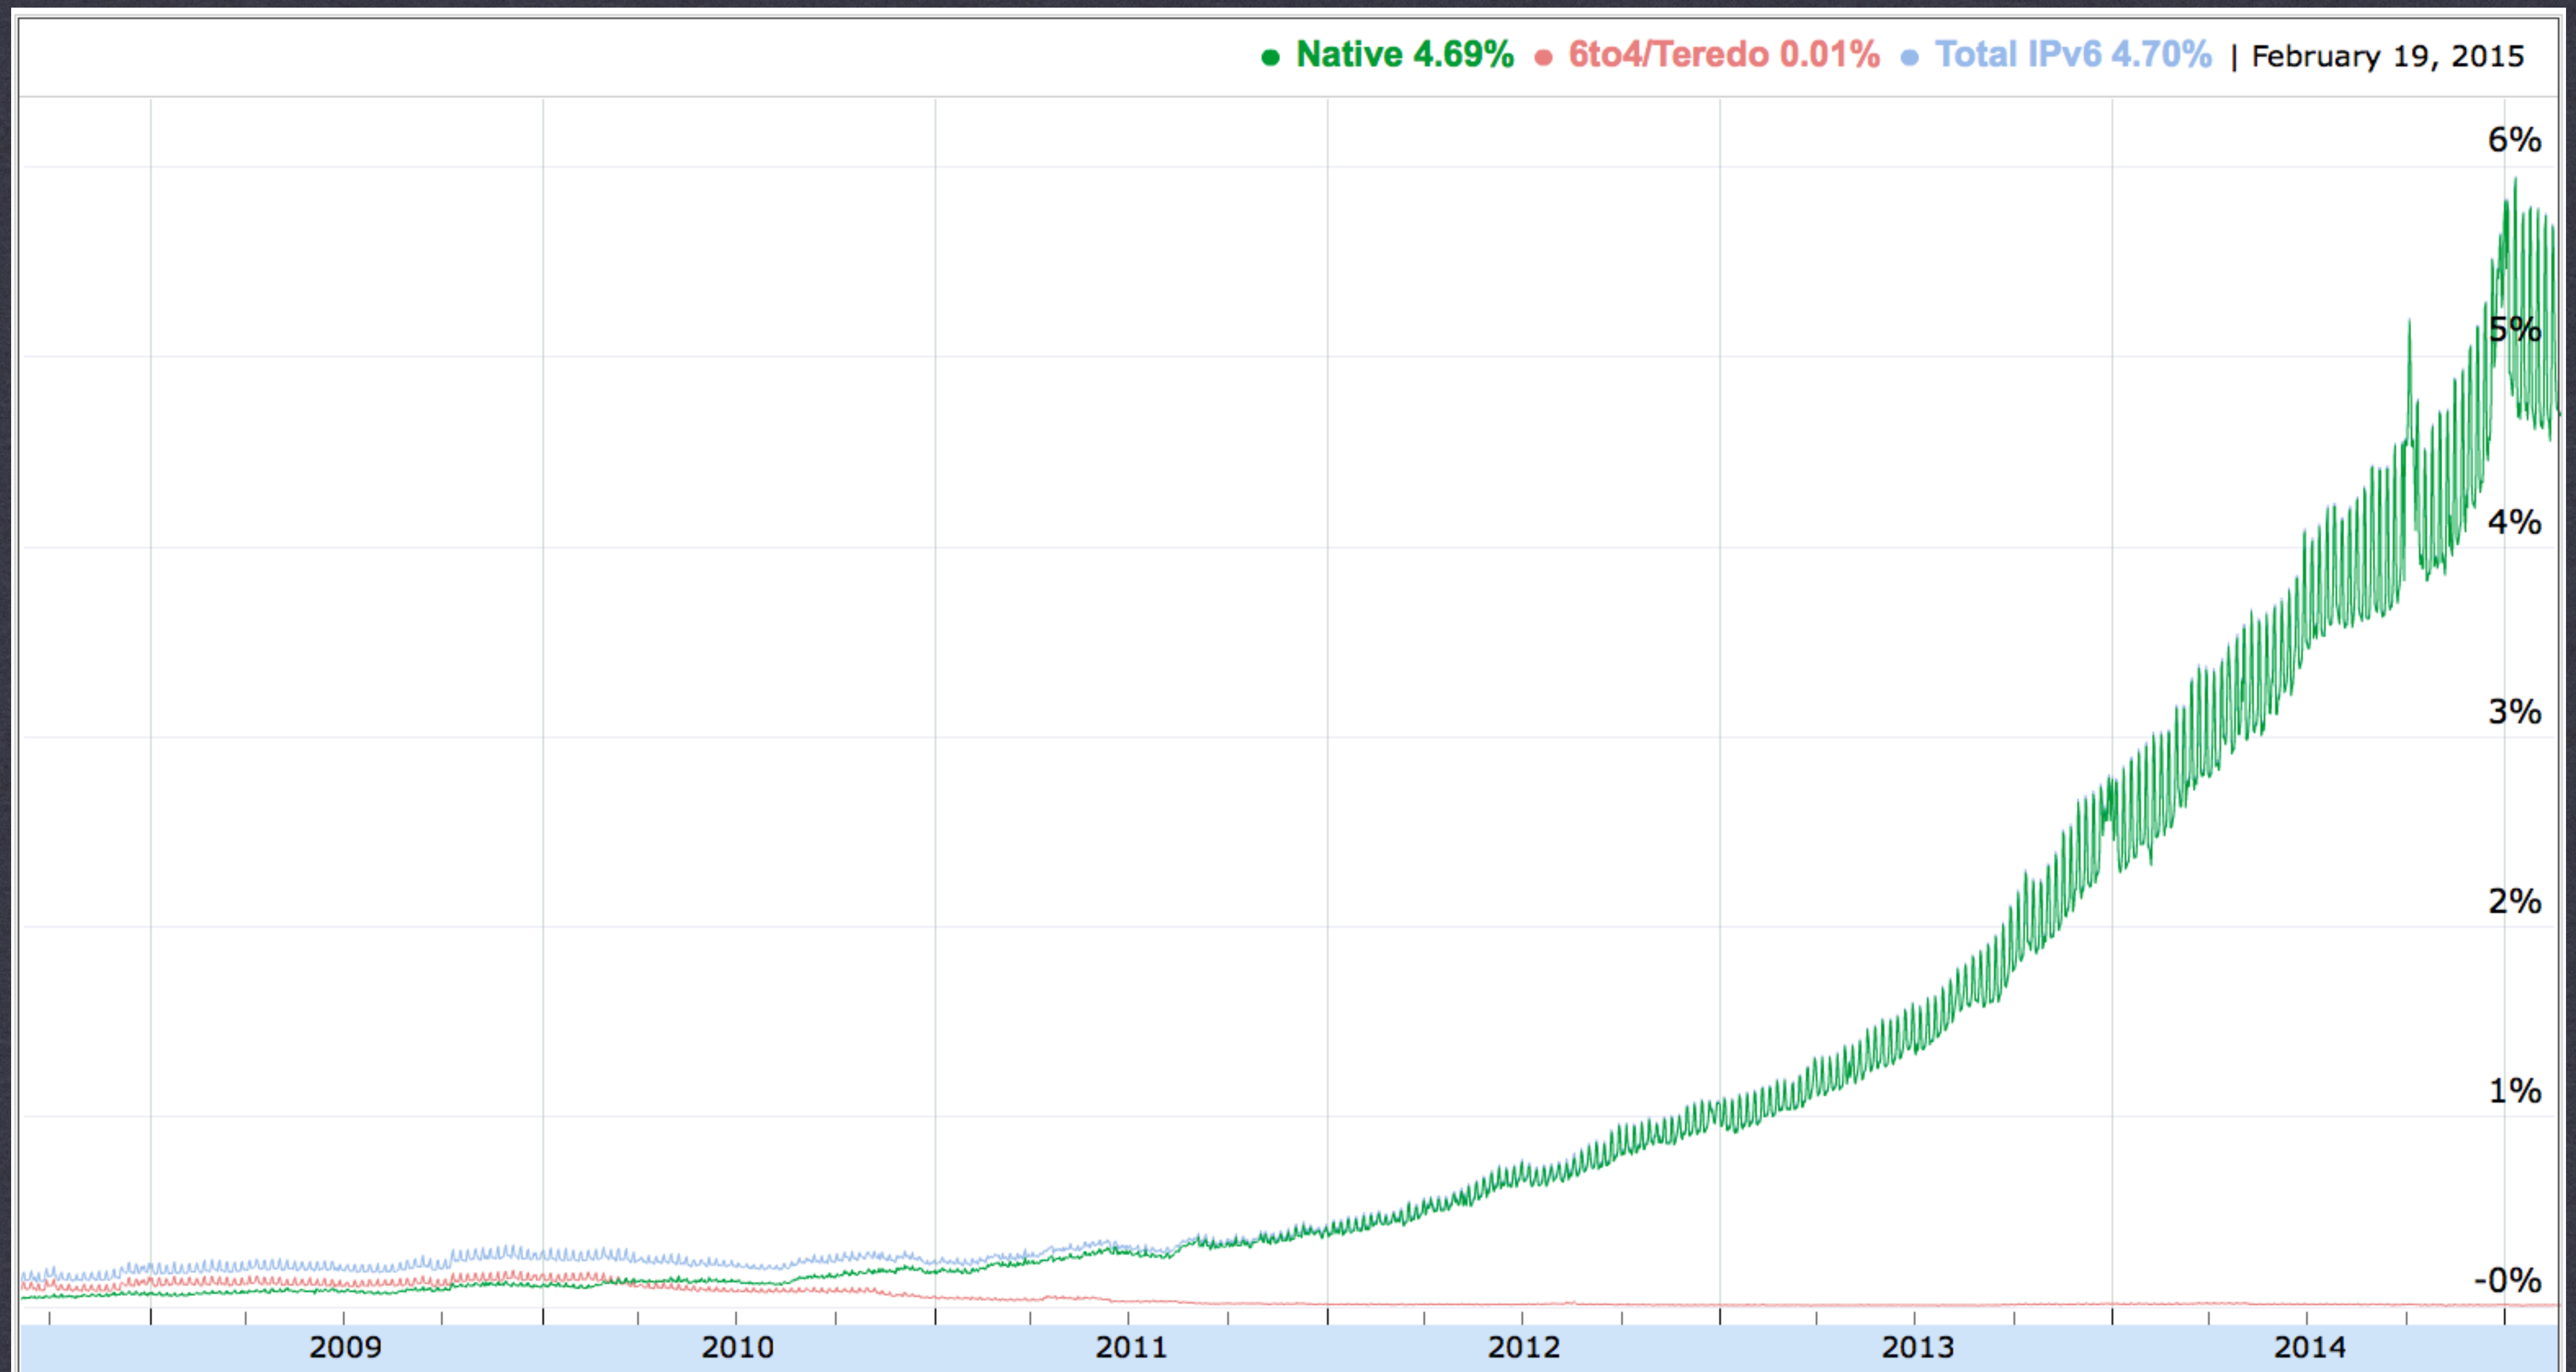
- * CTO at Akanda Inc
- * OpenStack Technical Committee
- * Core Reviewer for Neutron
- * Neutron PTL during Havana/Icehouse

mark@akanda.io

@gtwmm



IPV6




GOOGLE IPV6 TRAFFIC

[HTTPS://WWW.GOOGLE.COM/INTL/EN/IPV6/STATISTICS.HTML](https://www.google.com/intl/en/ipv6/statistics.html)

3

12.6%

United States of America



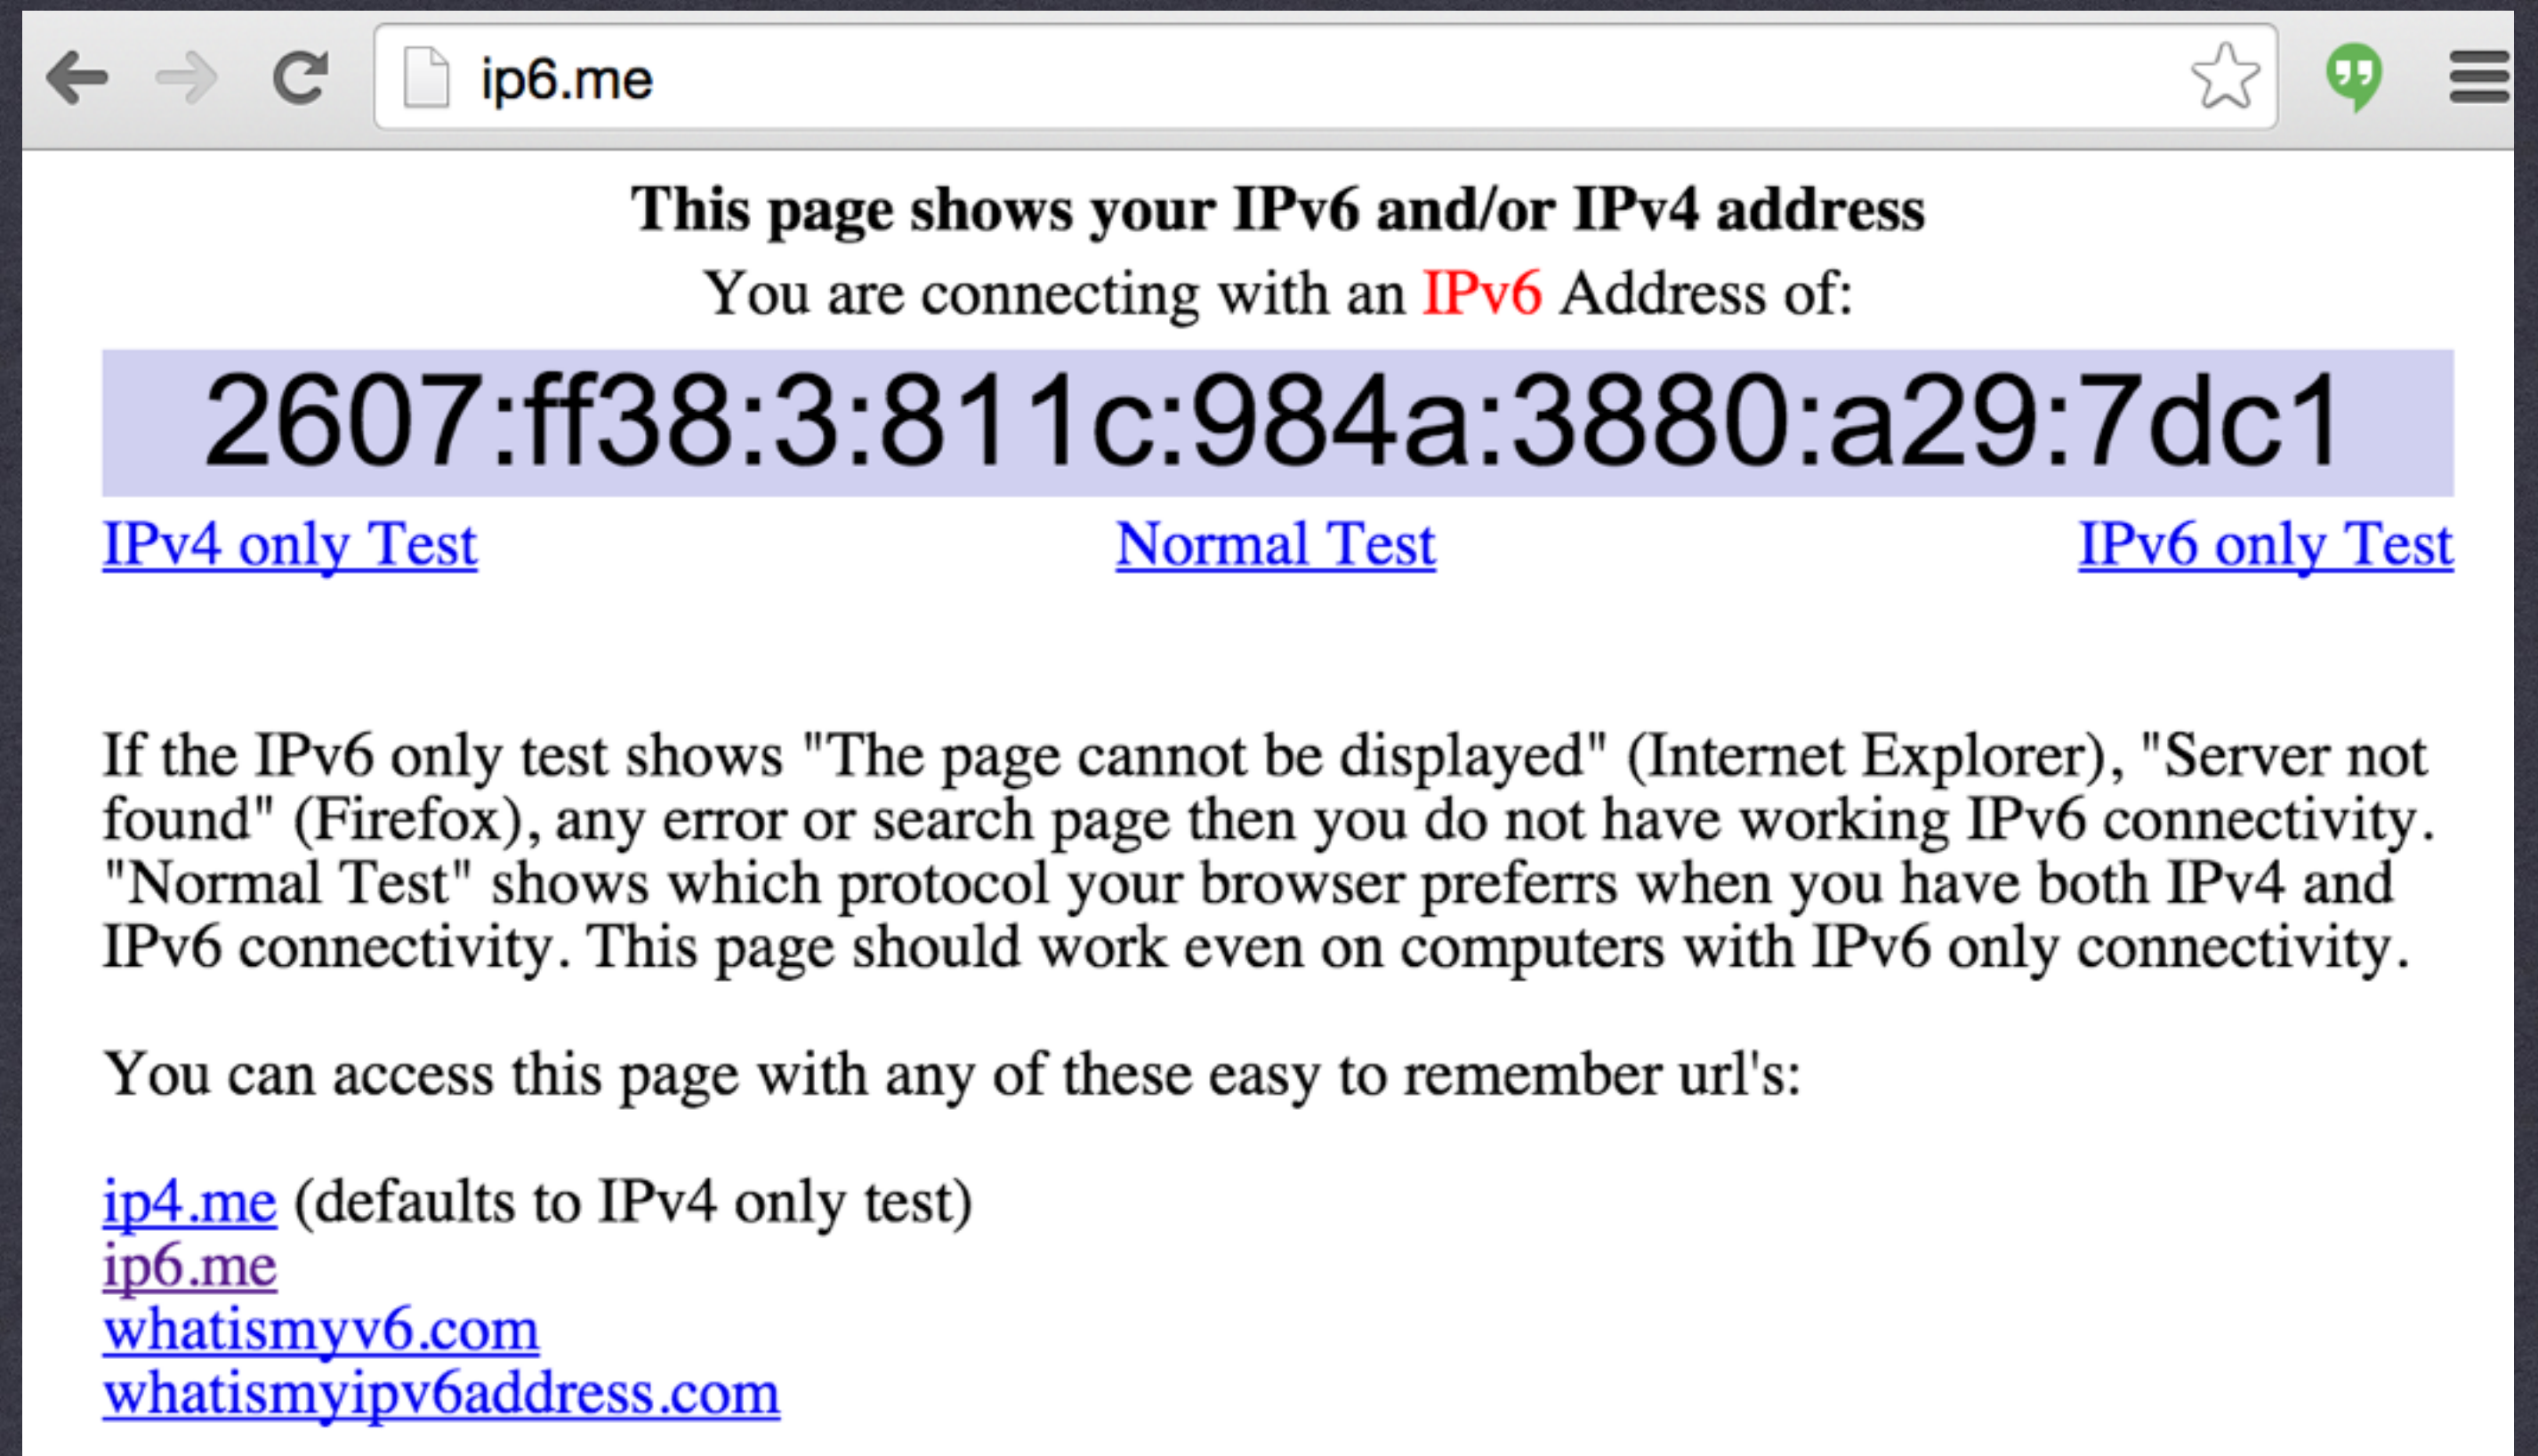
Source: Akamai State of the Internet Report

AKAMAI

[HTTP://WWW.STATEOFTHEINTERNET.COM/TRENDS-VISUALIZATIONS-IPV6-ADOPTION-IPV4-EXHAUSTION-GLOBAL-HEAT-MAP-NETWORK-COUNTRY-GROWTH-DATA.HTML#COUNTRIES](http://www.stateoftheinternet.com/trends-visualizations-ipv6-adoption-ipv4-exhaustion-global-heat-map-network-country-growth-data.html#countries)

DID YOU KNOW?

SCALE WIFI



The screenshot shows a web browser window with the address bar containing "ip6.me". The page content includes:

- Header: "This page shows your IPv6 and/or IPv4 address"
- Text: "You are connecting with an **IPv6** Address of:"
- Large text box containing the IPv6 address: **2607:ff38:3:811c:984a:3880:a29:7dc1**
- Three links: [IPv4 only Test](#), [Normal Test](#), and [IPv6 only Test](#)
- Explanatory text: "If the IPv6 only test shows 'The page cannot be displayed' (Internet Explorer), 'Server not found' (Firefox), any error or search page then you do not have working IPv6 connectivity. 'Normal Test' shows which protocol your browser prefers when you have both IPv4 and IPv6 connectivity. This page should work even on computers with IPv6 only connectivity."
- Text: "You can access this page with any of these easy to remember url's:"
- List of alternative URLs: [ip4.me](#) (defaults to IPv4 only test), [ip6.me](#), [whatismyv6.com](#), and [whatismyipv6address.com](#)

Differences from v4

- * Huge Address Space — 128-bit

340,282,366,920,938,463,463,374,607,431,768,211,456

- * End-to-end Connectivity

- * Configuration

- * Routing

- * ICMP + Multicast

Auto Configuration

- * Link Local Addressing
- * SLAAC
 - * Based on EUI-64 Address
 - * Prefix Announced via RA
 - * Requires a /64 available for each tenant

Auto Configuration (con't)

- * DHCPv6
 - * Stateless
 - * RA + DHCP for extra information
 - * Stateful
 - * Nearly the same as v4 DHCP today

Changes for Routing

- * Fixed Header Size
- * No checksum
 - * Lower/Upper layers provide
 - * No packet fragmentation
- * Router Advertisements

ICMPv6

- * Neighbor Discovery
- * Router Advertisements
- * Some messages require multicast
 - * ie ff02::1
- * Path MTU Discovery

Challenges

- * New Hardware Deployment
- * Software Updates
- * Application Updates for
 - * IP, DNS, etc
- * Addresses not really human friendly
 - * 2607:f295:6050::101:ca75
 - * fe80::893:33f1:b5c7:67ee
 - * fe80::feed:cafe

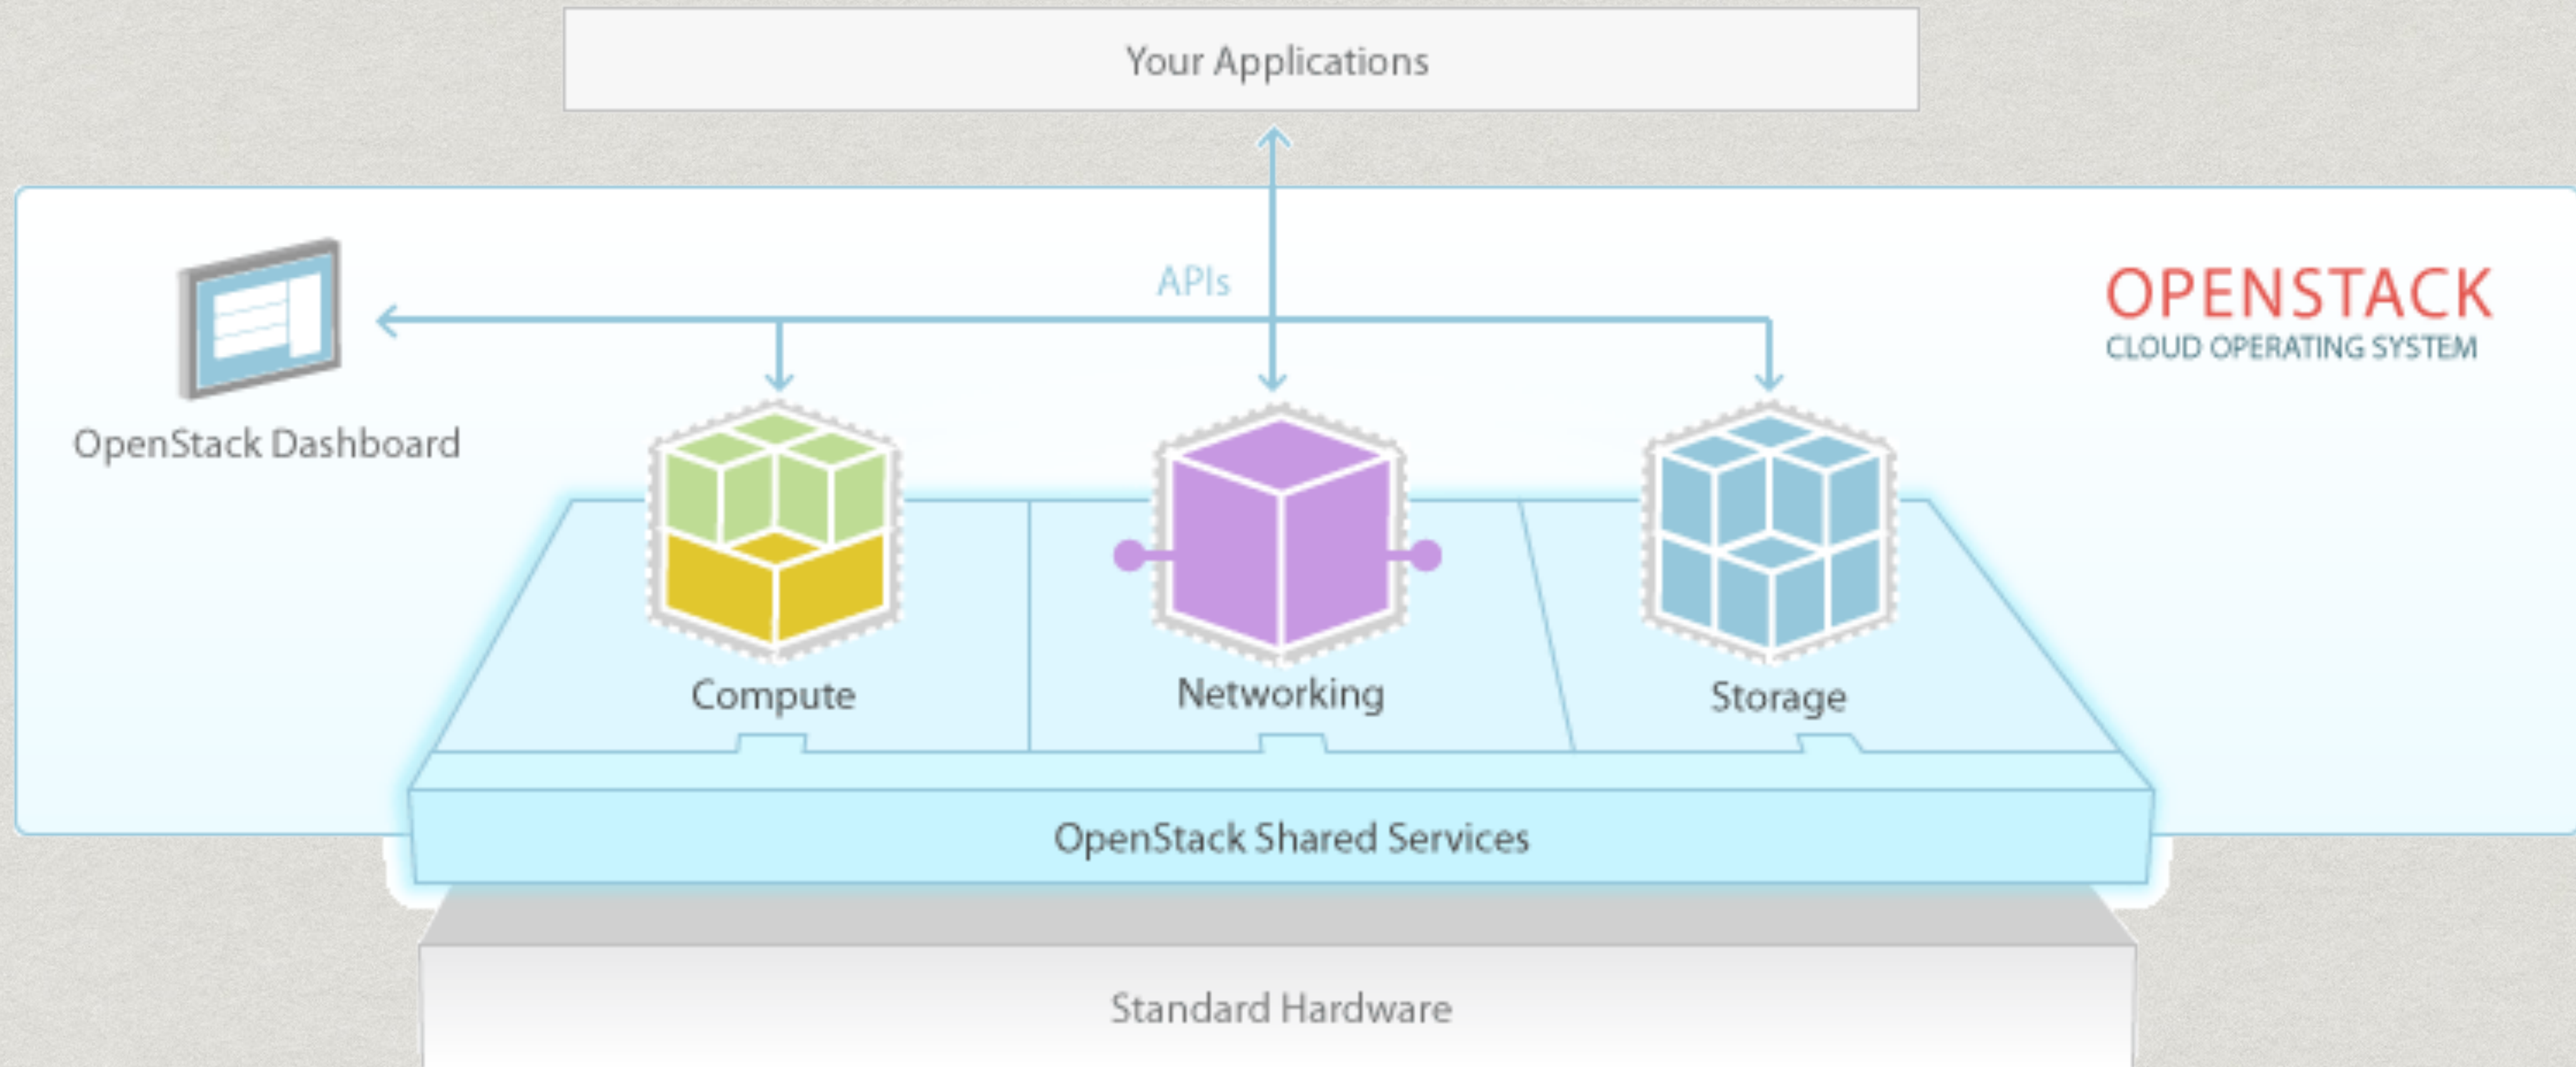
IPV6 IS A MUST

OPENSTACK NEUTRON

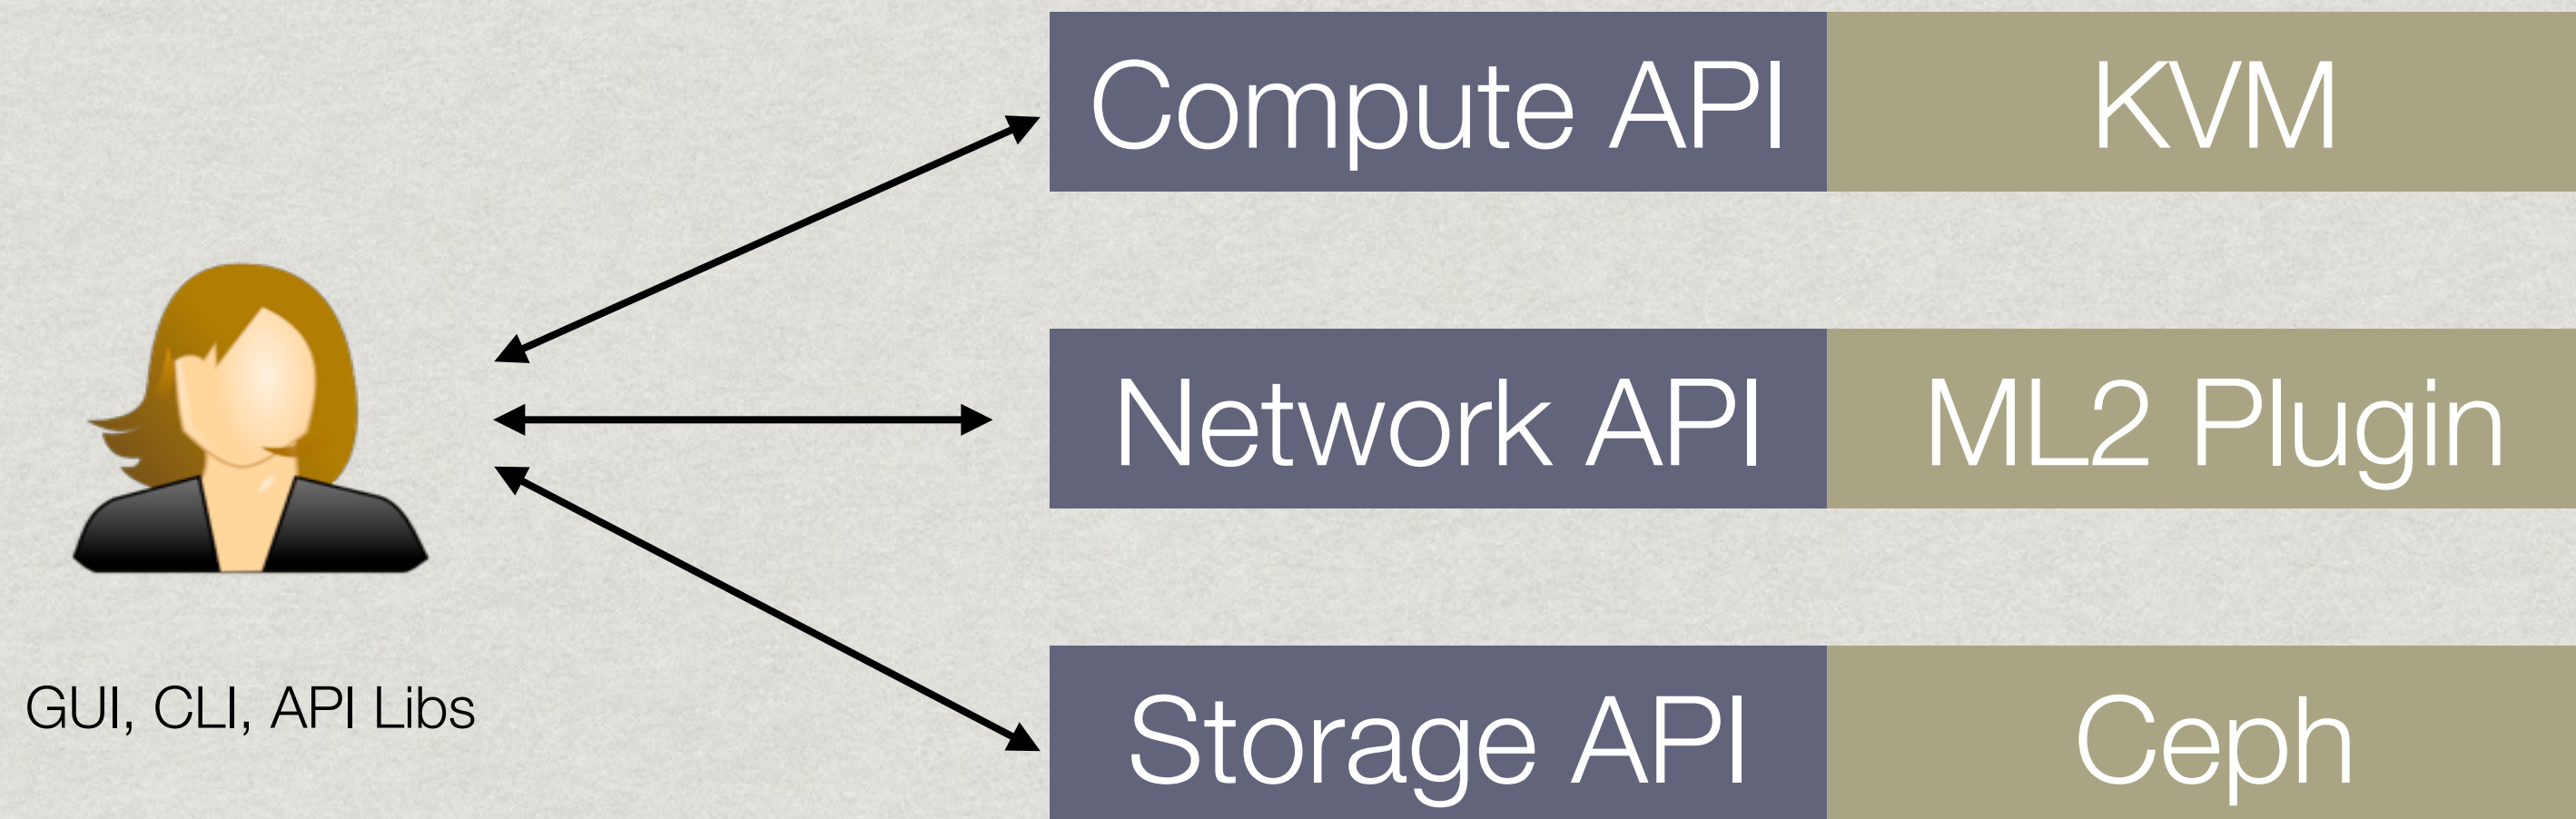
OpenStack

- * Large ecosystem of projects
- * Common Projects:
 - * Compute, Identity, Networking, Storage projects
- * Released every 6 months

OpenStack



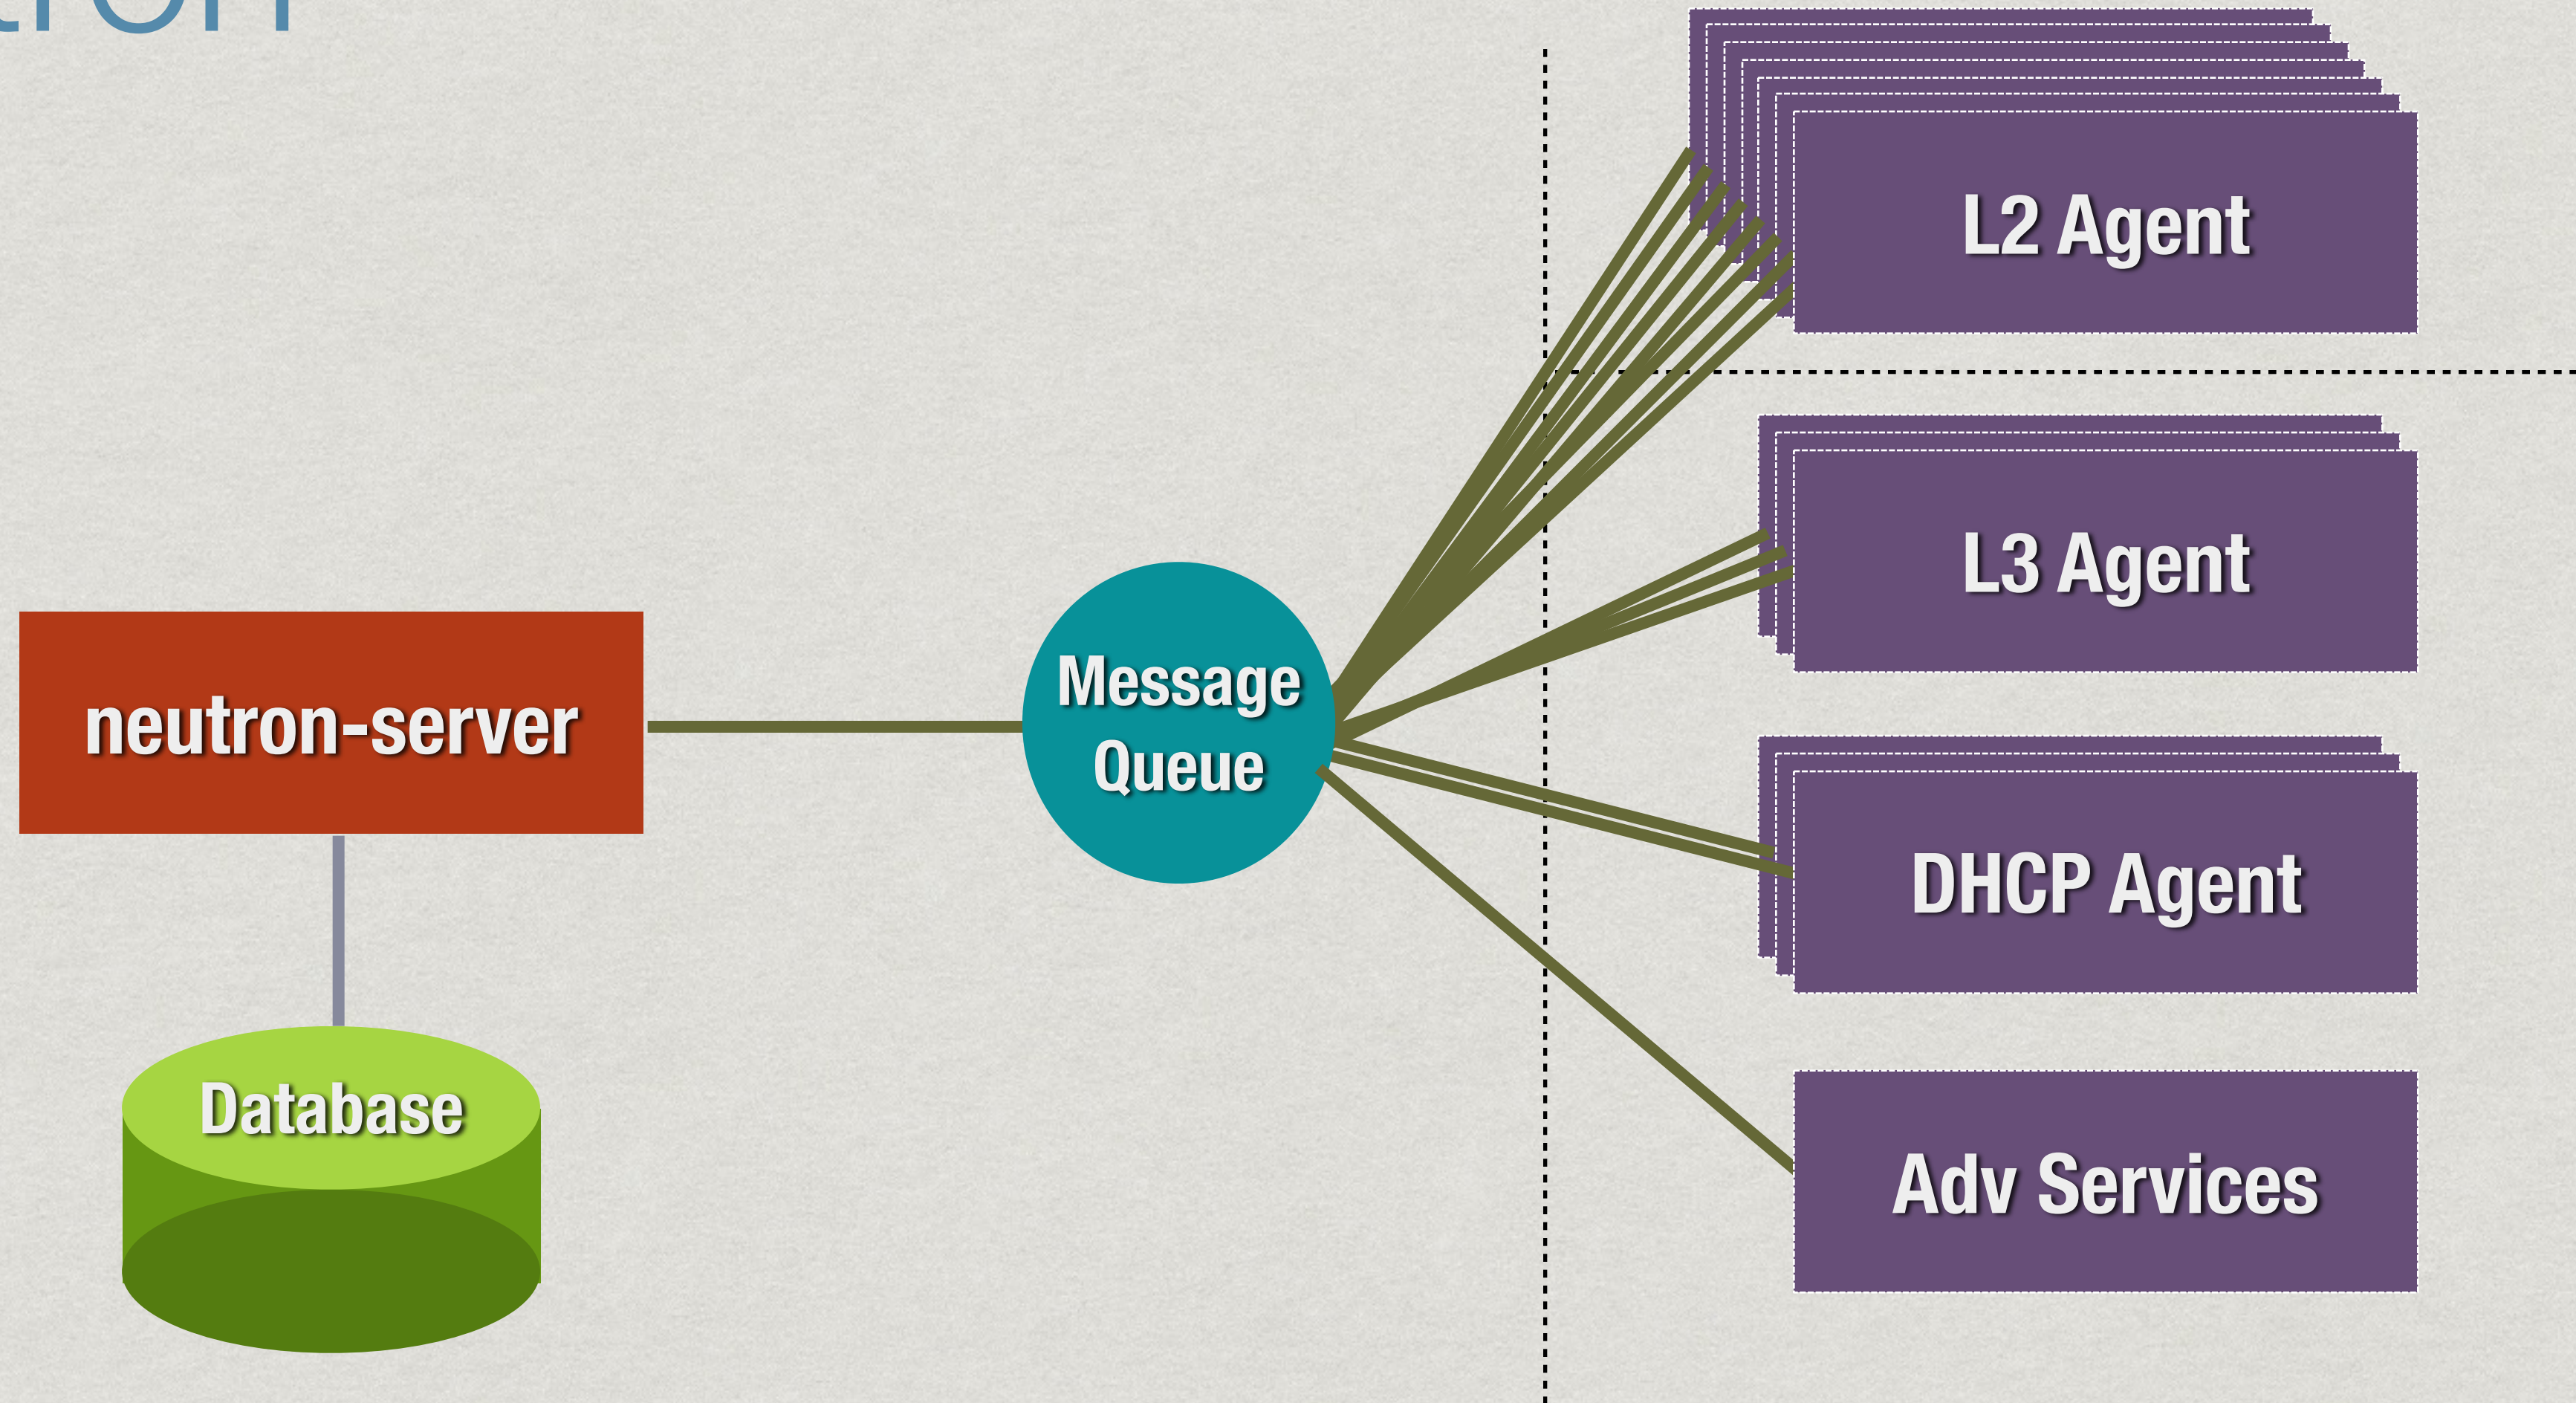
OpenStack



What is Neutron?

- * Unified API
- * Small Core
- * Pluggable Architecture
- * Extensible

Neutron



IPv6 in Neutron

- * IPv6 IPAM
- * Security Groups, L2 Pre-Pop
- * Early support starting in Grizzly
- * Enhanced support in Juno

```
+-----+-----+
+-----+-----+
| admin_state_up      | True
| allowed_address_pairs |
| binding:vif_details | {"port_filter": true}
| device_id           | 2ff37d12-c0fe-4cf5-a73a-7c4846b9df6a
| device_owner        | compute:iad-1
| fixed_ips            | {"subnet_id": "7e294b5b-9ee2-4dd0-8c7a-1438e8652926",
"ip_address": "10.10.10.46"}
|                       | {"subnet_id": "a4abc1d2-e51c-4bc9-9ac8-a071412ca6da",
"ip_address": "2607:f298:6050:eb8f:f816:3eff:fe39:b947"}
| id                   | f8f71519-9b78-4f7c-b83f-ad0a858981db
| mac_address          | fa:16:3e:39:b9:47
| name                 |
| network_id           | 7c7d0ba8-a74e-4b2f-b0a3-58ce5c7c090b
| port_security_enabled | True
| security_groups      | d05819c5-7eb6-48b2-87bc-f2bc381f0fe3
| status               | ACTIVE
```

IPv6 Tenant Benefits

- * No NAT
- * No Floating IP Required

```
$ ssh 2607:f298:6050:eb8f::a1f:ea75:ca75
```

JUNO

RELEASED OCTOBER 2014



"Amicalola Falls" by Sean Morgan

CC BY-ND 2.0

<https://www.flickr.com/photos/seanm1025/3646862123>

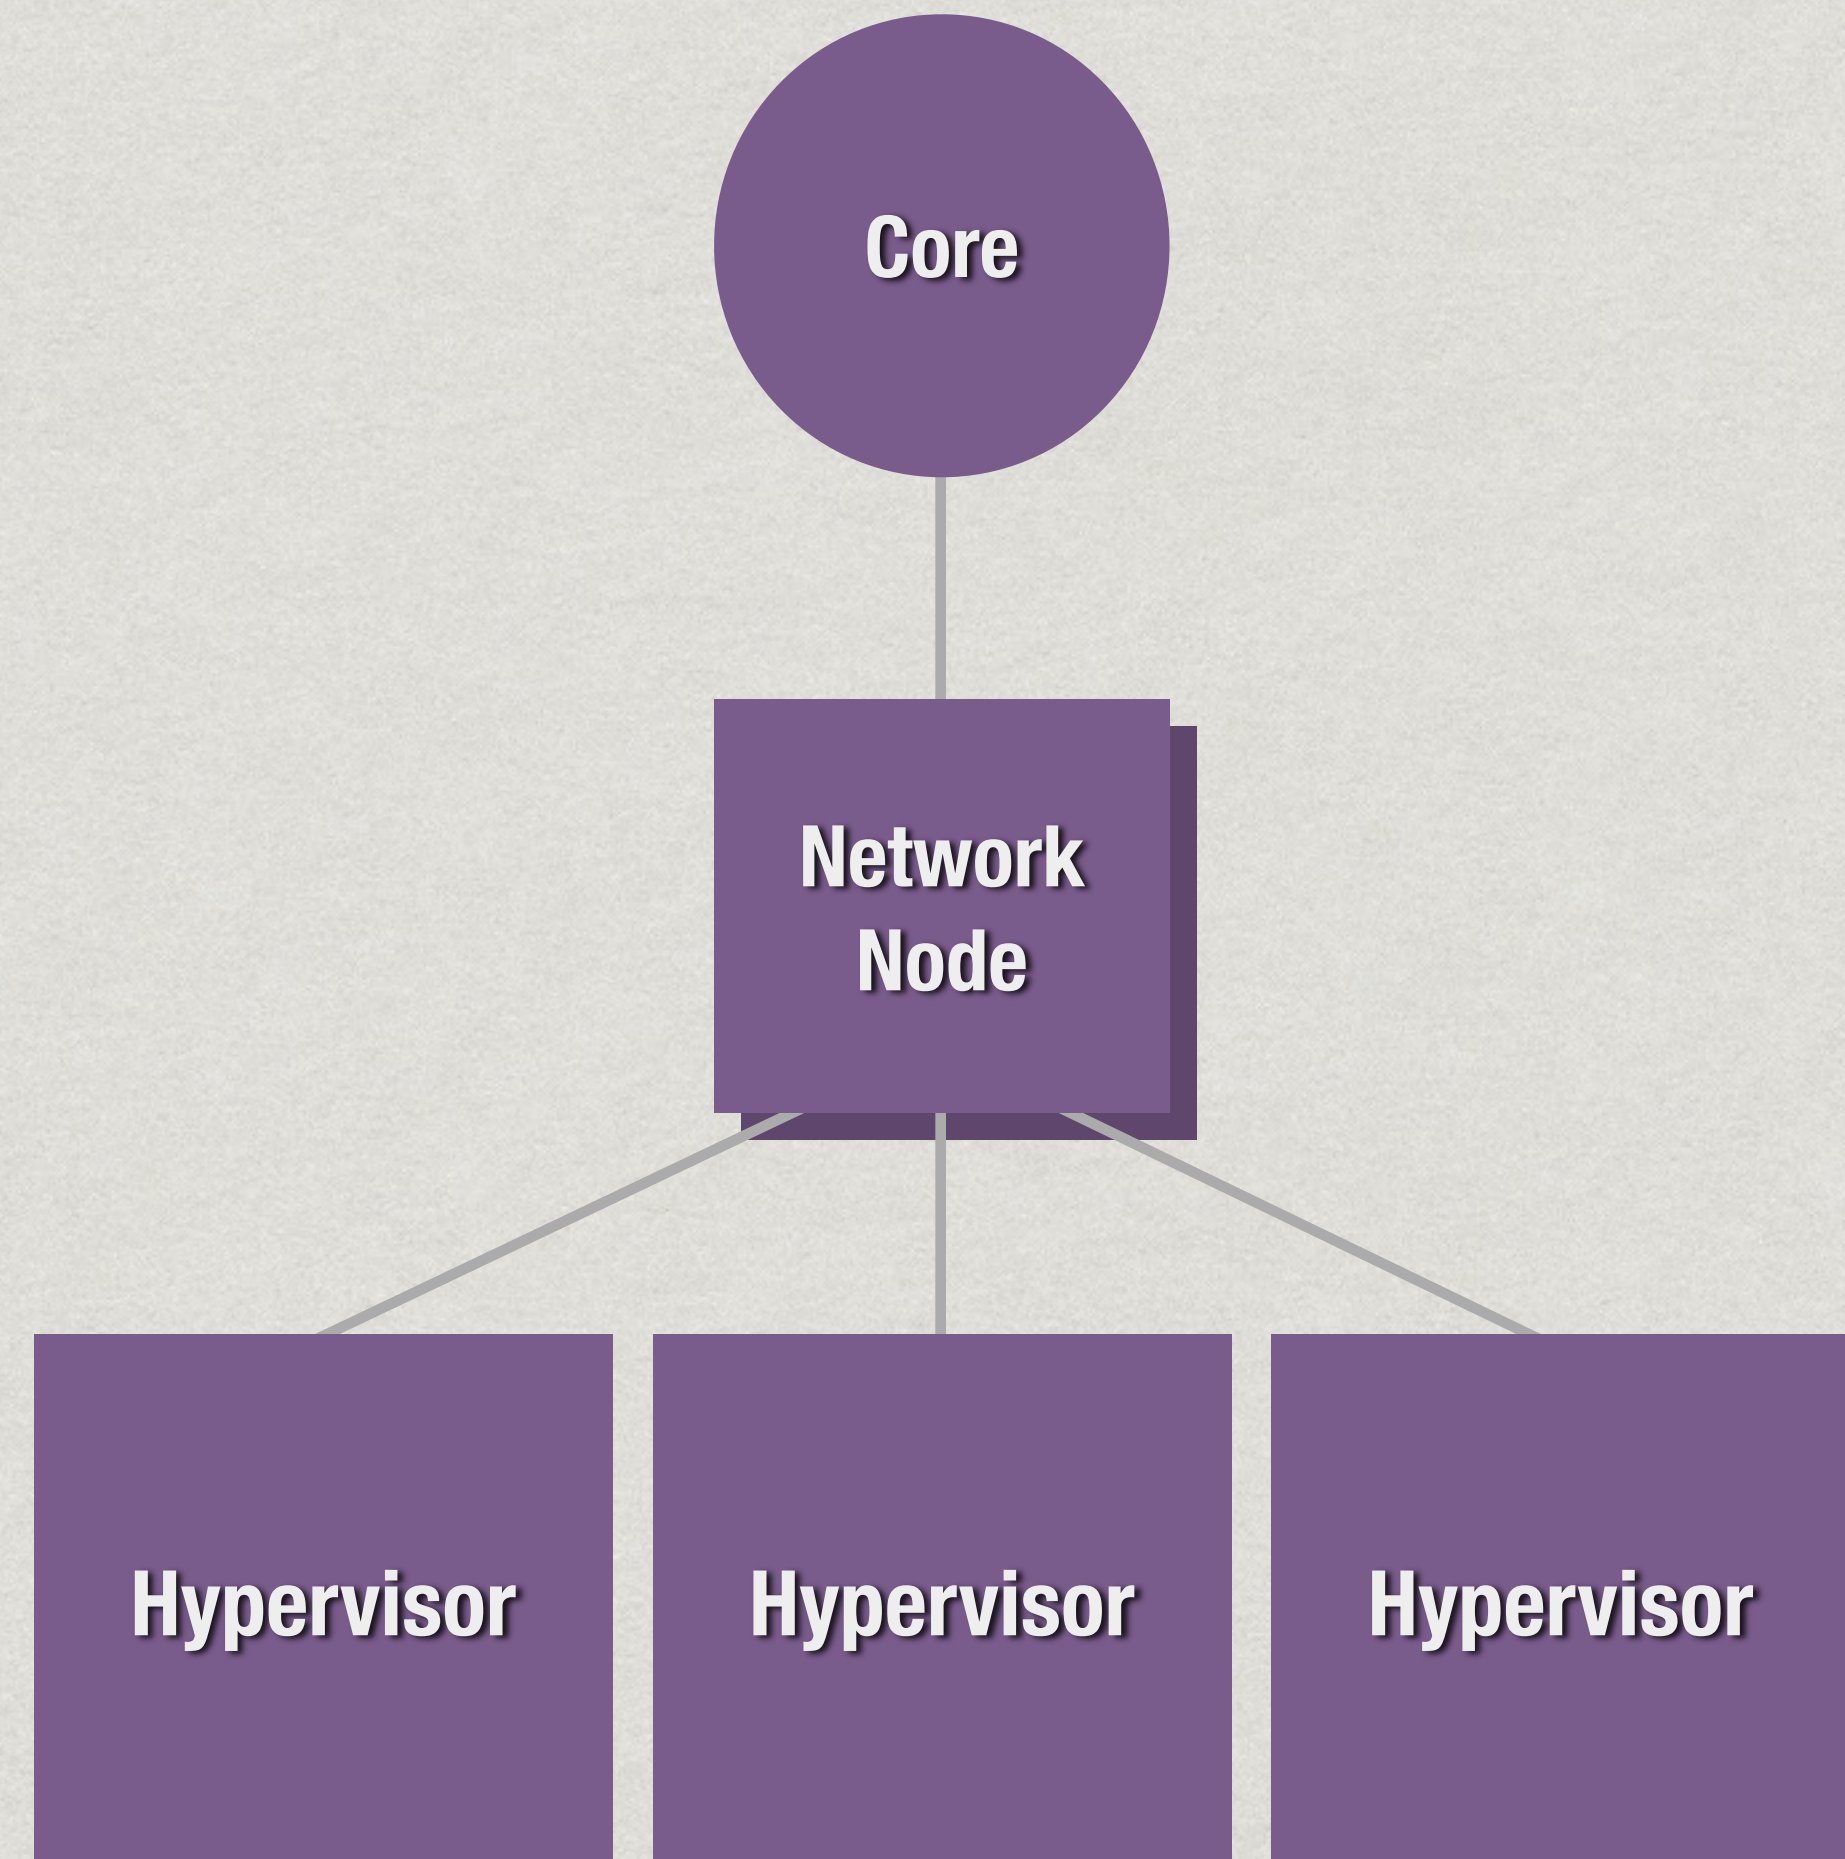
Configuration

- * SLAAC
 - * via radvd
- * DHCPv6
 - * stateless
 - * stateful
 - * via dnsmasq

Field	Value
allocation_pools	{"start": "fdd6:a1fa:cfa8:fc34::2", "end":
cidr	fdd6:a1fa:cfa8:fc34::/64
dns_nameservers	
enable_dhcp	False
gateway_ip	fdd6:a1fa:cfa8:fc34::1
host_routes	
id	b1cbd9ee-bf96-4e67-8e49-b9d6c00c070a
ip_version	6
ipv6_address_mode	slaac
ipv6_ra_mode	slaac
name	
network_id	8f74ddf8-aca3-4c47-b663-e5b2c169bf30
tenant_id	f0edc959fca740bf81eccc2b67b87b7b

Routing

- * L3 Agent
 - * Provider Support
 - * RA within project
 - * Static External Gateway
- * Provider Network



IPv6 Security

- * End to end connectivity
 - * NAT as a security blanket
- * DHCP Spoofing
- * RA Spoofing

MANAGEMENT LAYER

IPv6 Readiness

- * API Layer
- * DB Layer
- * Client Libraries
- * Messaging
- * Overlay Tunnels*

Linux Tunnels

- * Kernel 3.9+
- * iproute(2)
- * Works

```
# ip link add vtep1 type vxlan id 2 local 2607:f298:6050::feed  
remote 2607:f298:6050::beef
```

Open vSwitch Tunnels

This command will silently fail:

```
# ovs-vsctl add-port br1 tun1 -- set interface tun1 type=vxlan
options:remote_ip=2607:f298:6050::face
```

However, in the logs you'll get this error:

```
2015-02-21T05:44:38.213Z|00031|socket_util|
ERR|"2607:f298:6050::face" is not a valid IP address
```

KILO

APRIL 2015



Kilo

- * v4 and v6 DHCP Options
- * Multiple Prefixes per Subnet
- * Support External Prefix Delegation
- * Link Local Default External Gateway

Linux Routing RA+Forwarding

- * `net.ipv6.conf.eth0.accept_ra=1`
- * `net.ipv6.conf.all.forwarding=1`

Linux Routing RA+Forwarding

- * `net.ipv6.conf.eth0.accept_ra=2`
- * `net.ipv6.conf.all.forwarding=1`

CHALLENGES

Metadata Service

- * Cloud Init
 - * Default EC2 address is IPv4: 169.254.169.254
- * Config Drive is your friend

IPv6 Privacy Extension

- * SLAAC address is identifiable
- * Generate and use temporary address for a limited time
- * Incompatible with Spoofing Prevention
 - * Do we care? private network, shared network, public network

Routing

- * Prefix Delegation
 - * Your routers and Neutron need to know this information
- * Static Routing
- * Dynamic Routing
 - * OSPFv3
 - * BGP

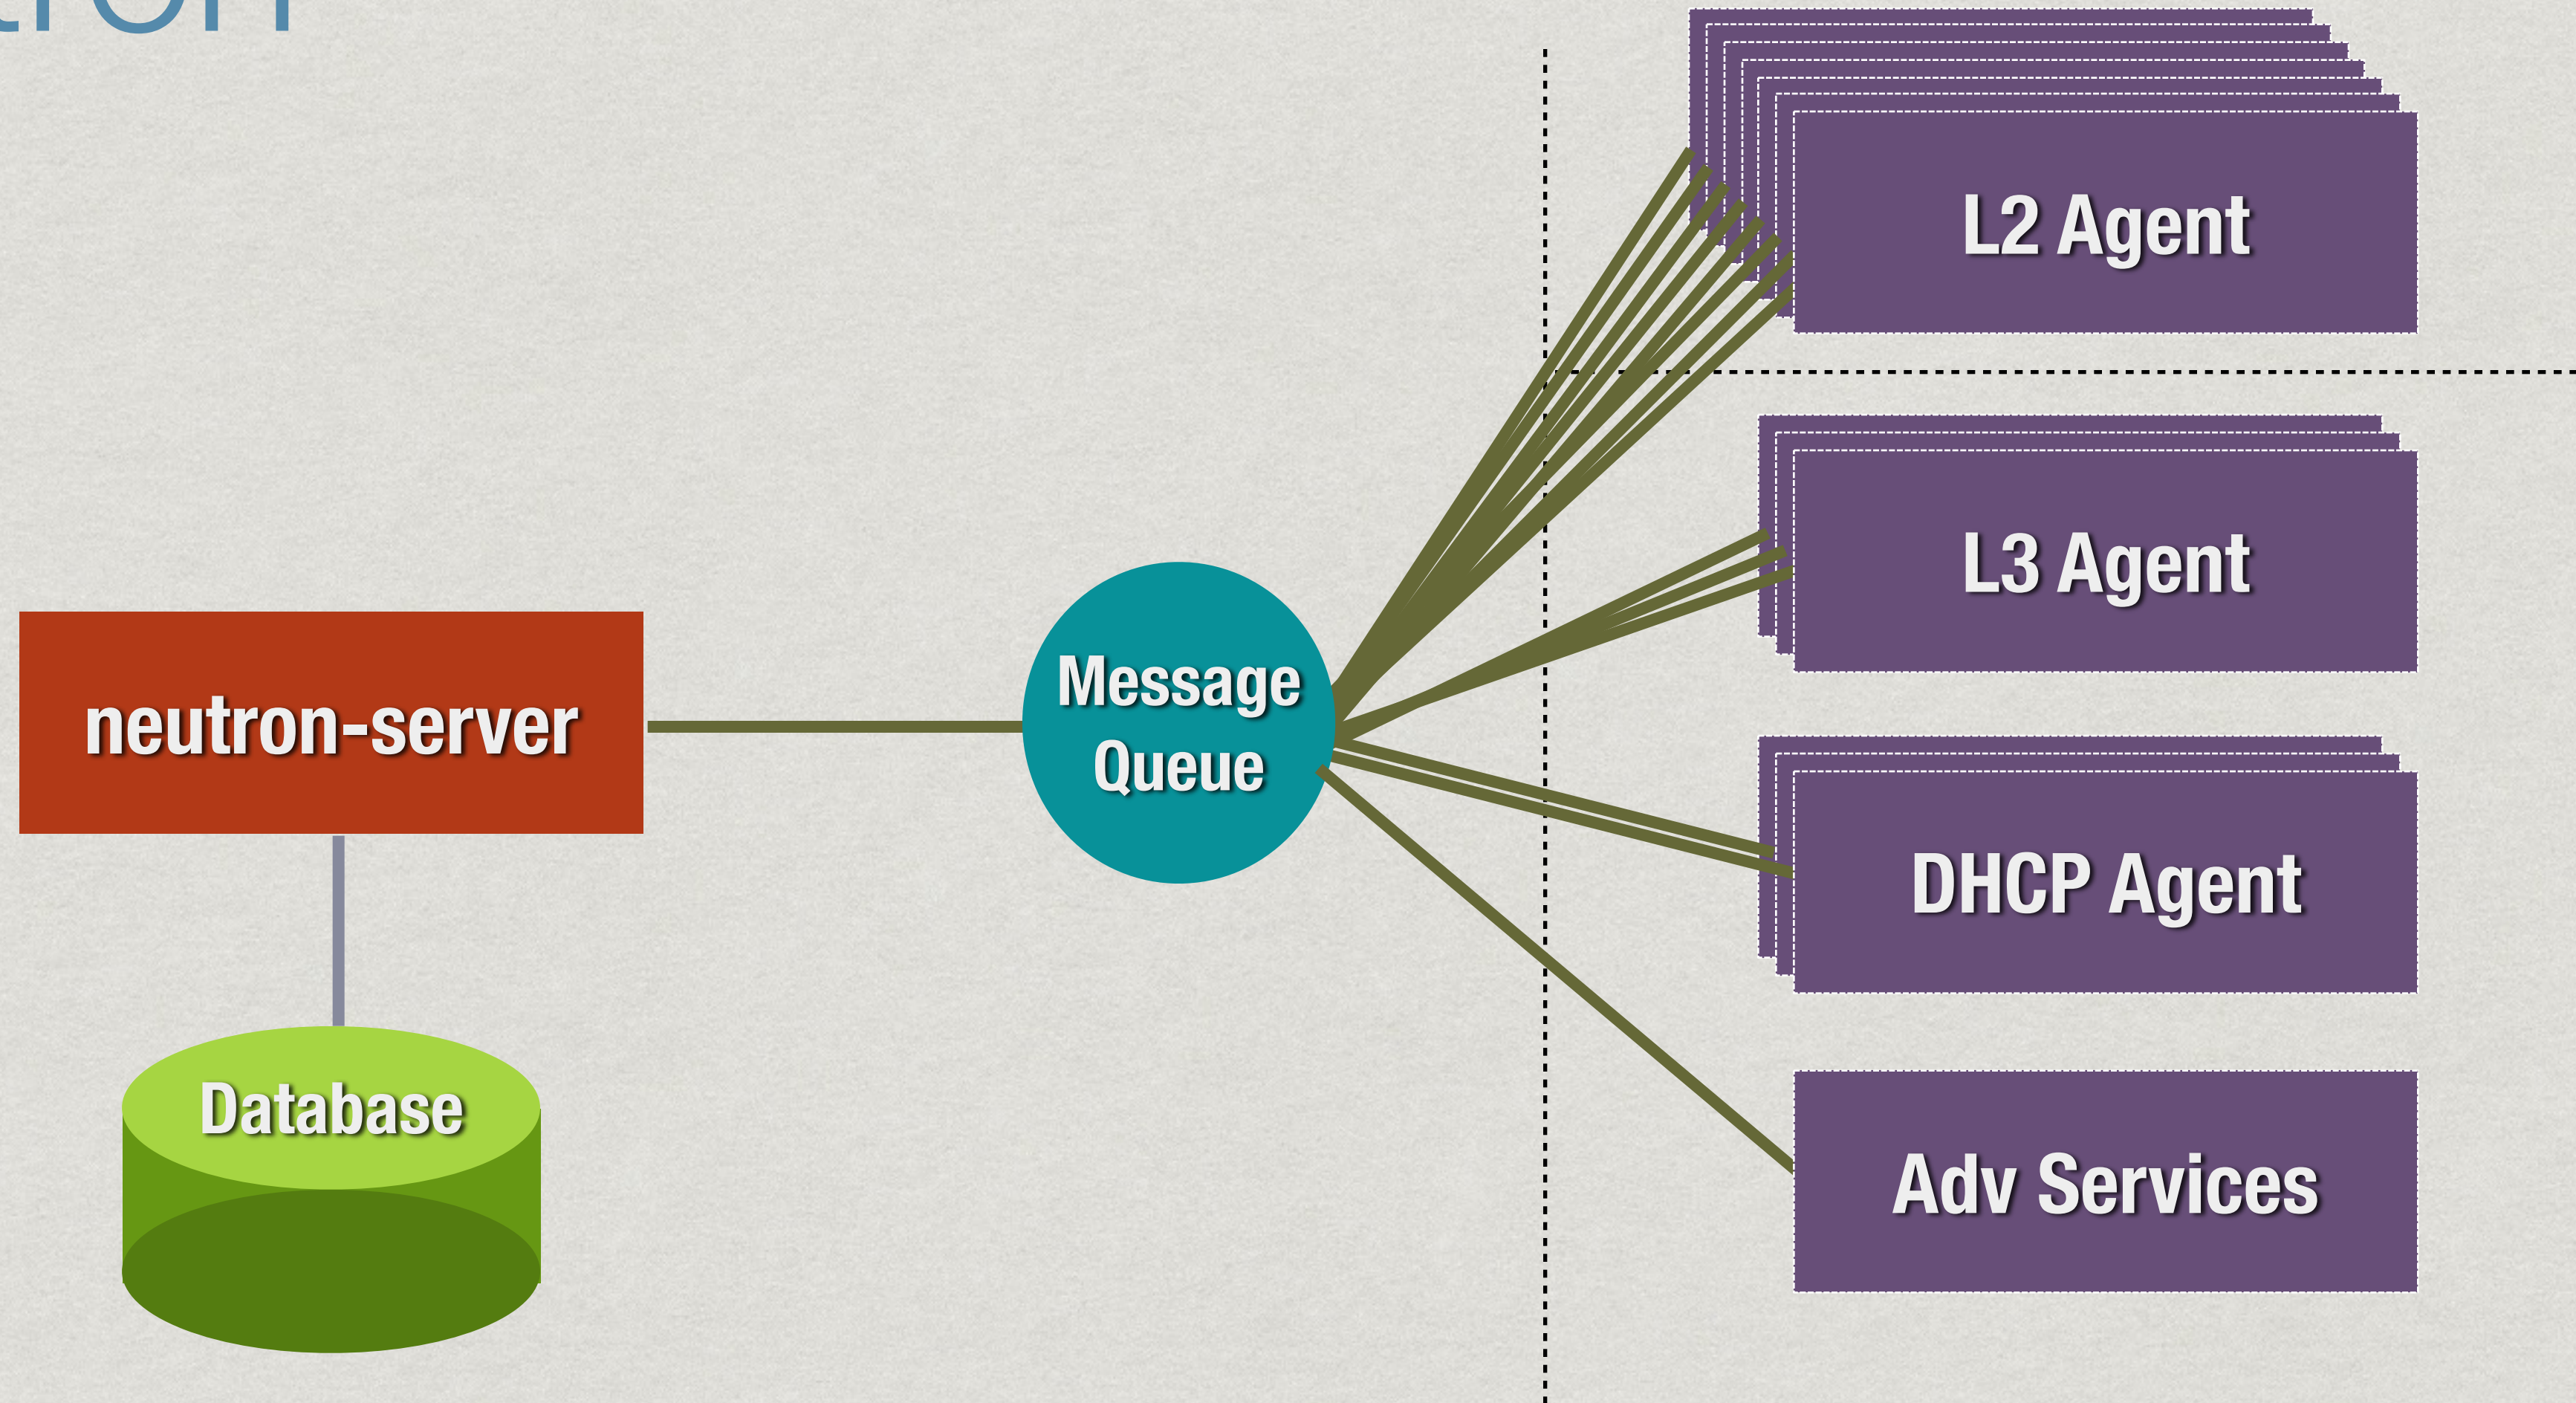
AKANDA

What is Akanda?

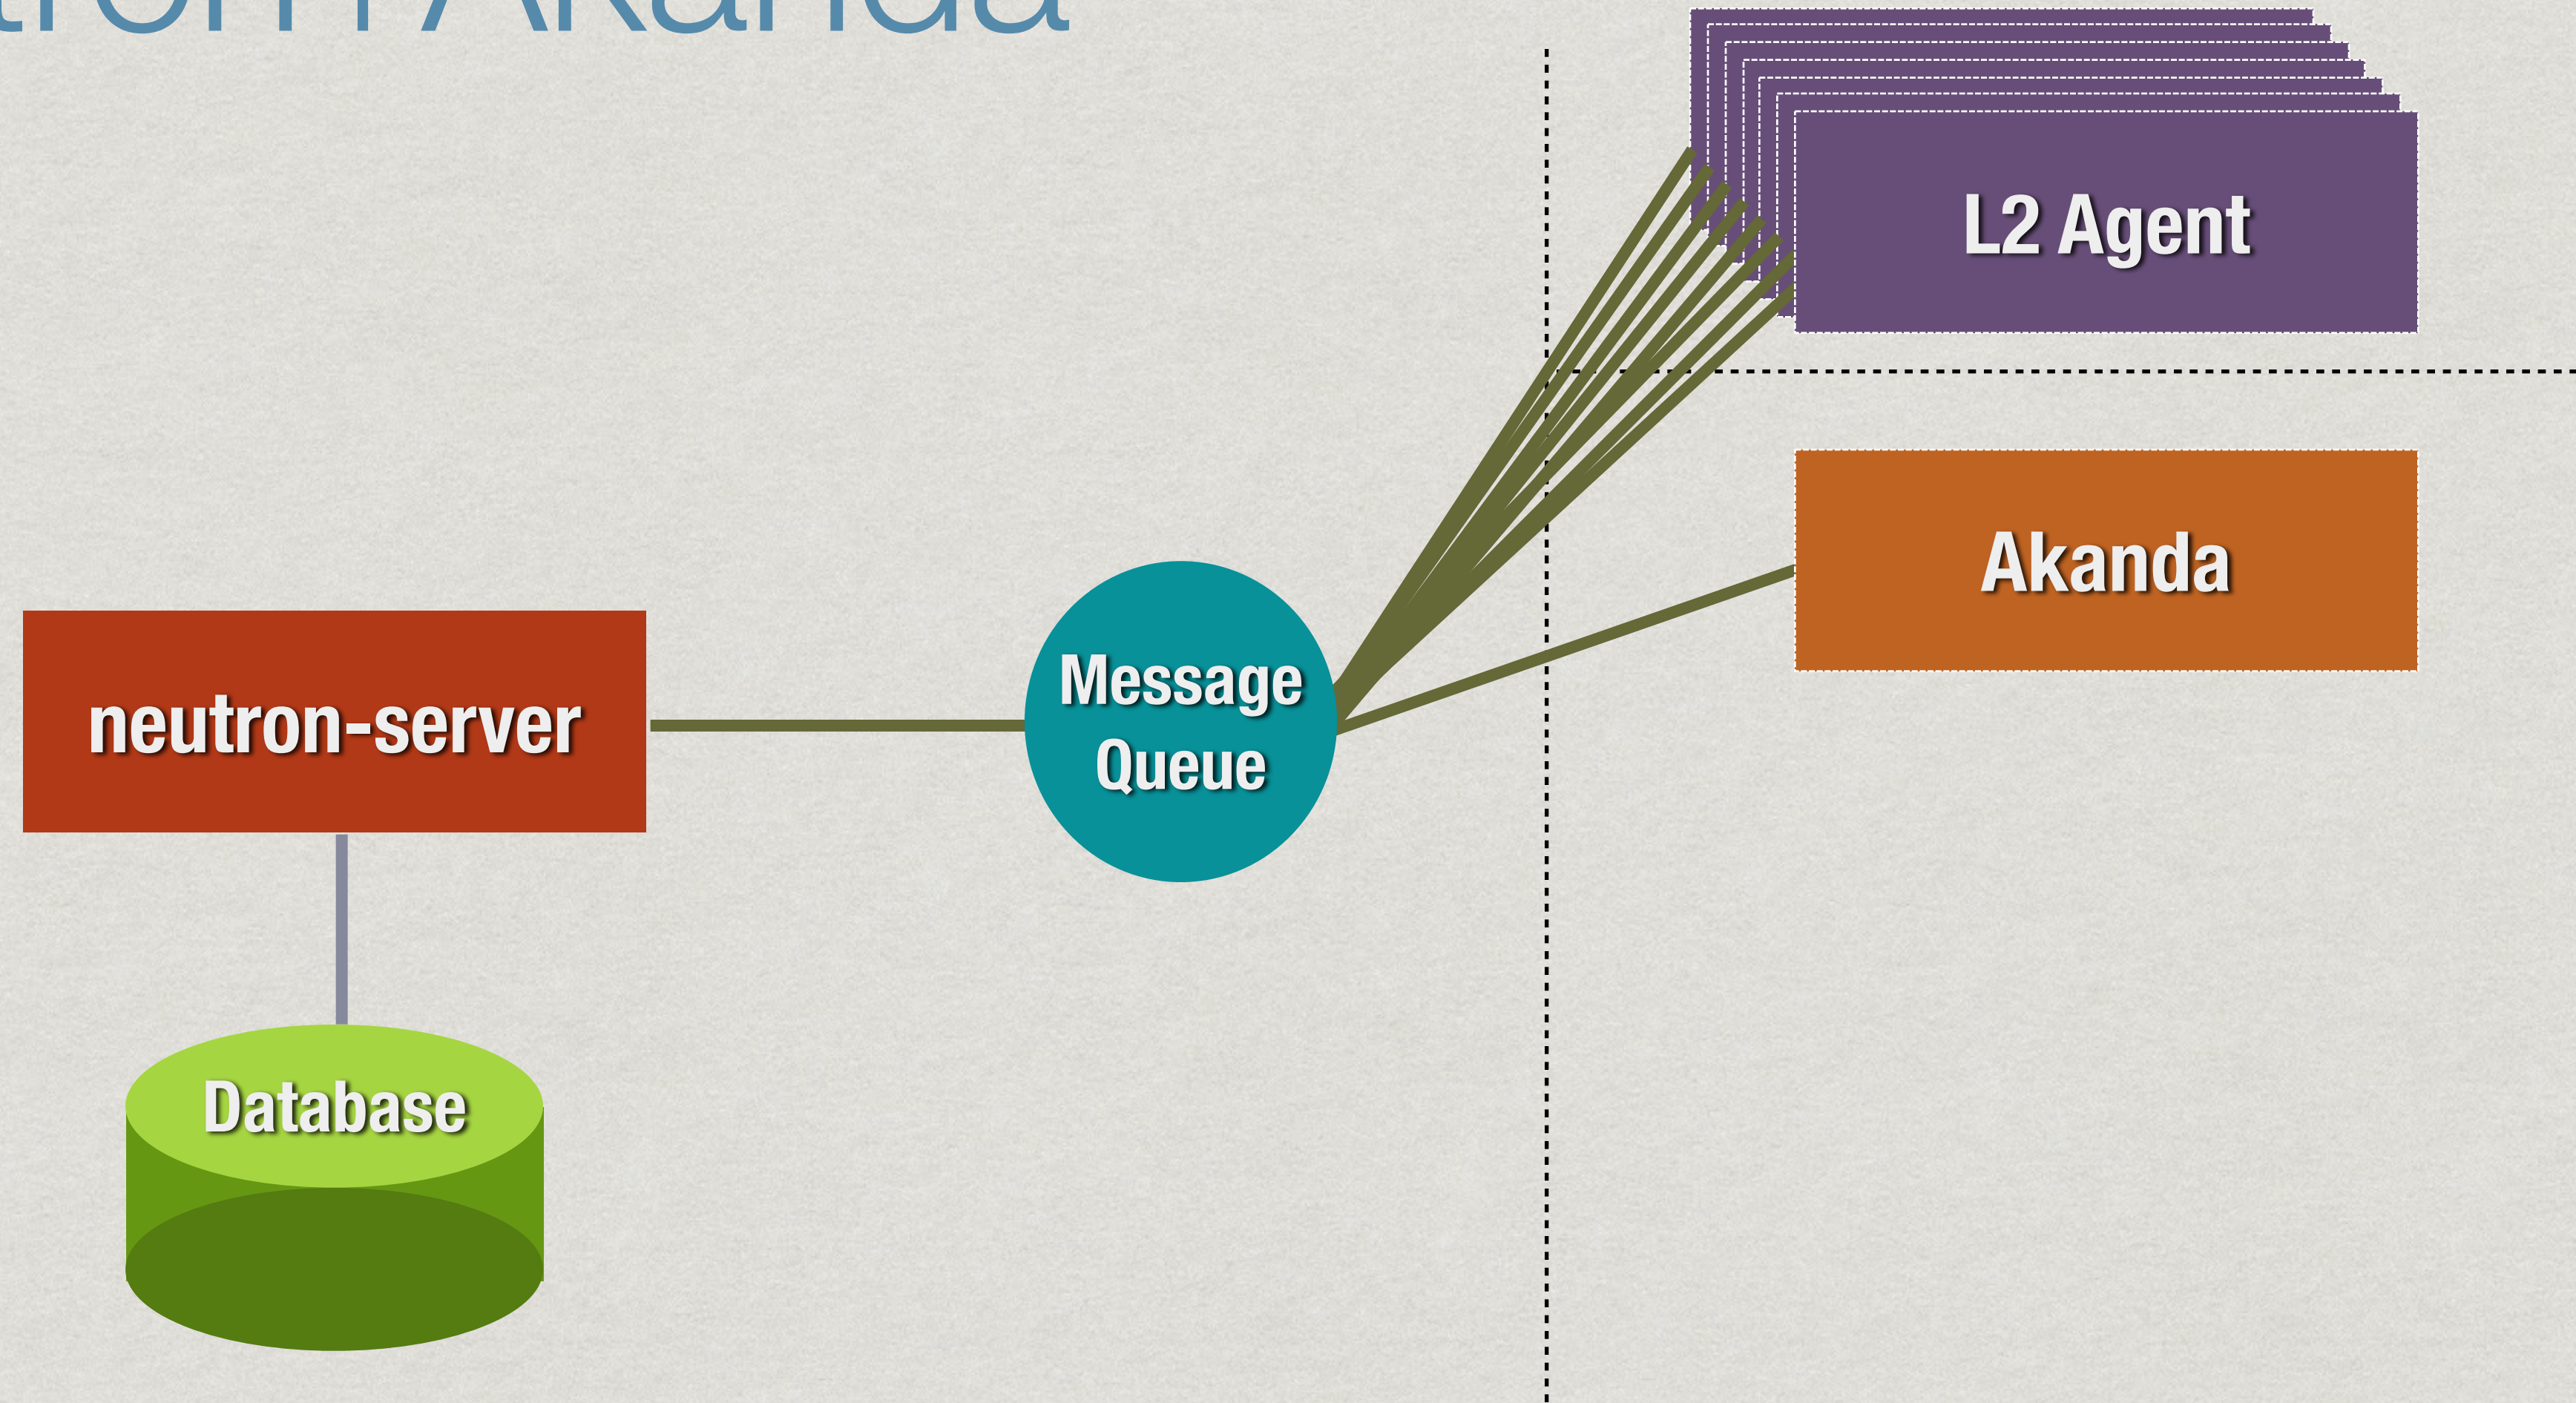
- * Open Source (Apache 2)
- * Dynamic Routing via Service VMs
- * Drop in replacement for Neutron L3, DHCP, Metadata Agents
- * In Production Now
- * Created by DreamHost



Neutron



Neutron+Akanda



THANK YOU

mark@akanda.io

@gtwmm

QUESTIONS?