

OpenStack Cascading

...

Zhiyuan Cai

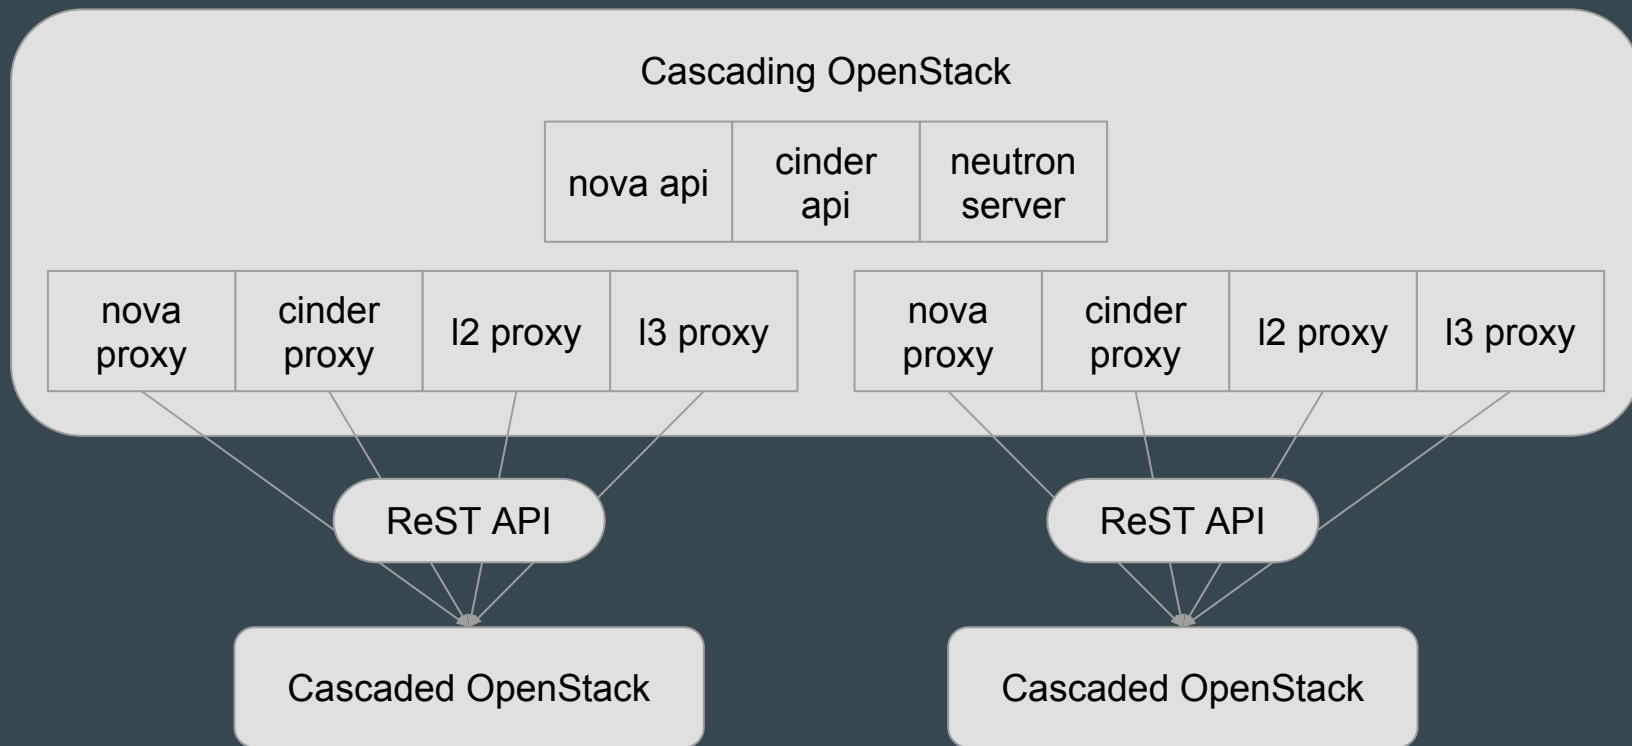
Agenda

- Motivation
- Architecture
- Mapping
- Case Study
 - Create Network
 - Create Router
 - Create Server
- L2 Population
- Cross Site L2 Networking
- Cross Site L3 Networking

Motivation

- Multi-site resource management with unified OpenStack ReST API
 - Global view of resources across sites
 - Large scale cloud working as one single OpenStack
- Multi-site networking
 - L2 connection between servers in different sites
 - L3 connection between servers in different sites

Architecture



Mapping

- URL Mapping
 - Proxy needs to know where to send the ReST request
- Resource Mapping
 - Resource has data record in both cascading layer and cascaded layer
 - Need to map resource in cascading layer to resource in cascaded layer
 - OpenStack doesn't allow creating one resource while specifying its id at the same time
 - Direct id mapping is not usable

URL Mapping

- URL information is stored in Keystone
- URL information is sorted by region
- Configure proxy process with proper region
- Proxy process retrieves URL from Keystone with configured region

service_id	url	region_id
d18c3c97c1ae43fc9fb0da3118e02a2b	http://10.1.16.233:9696/	CascadedC
d18c3c97c1ae43fc9fb0da3118e02a2b	http://10.1.16.170:9696/	CascadedJ

ml2_conf.ini for l2 proxy

```
[agent]
cascading_os_region_name =
CascadingC
os_region_name =
CascadedC
```

Resource Mapping

- Use “name” field for the mapping

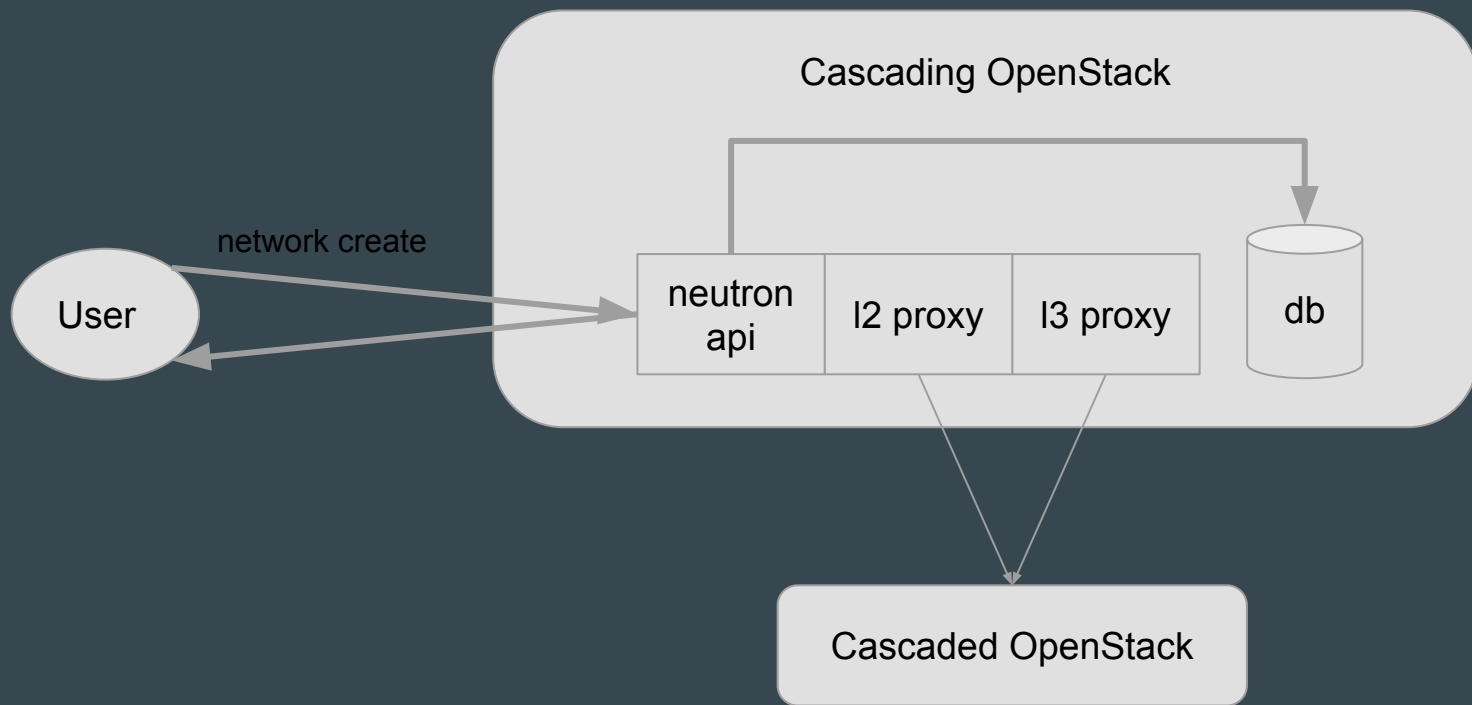
Server record stored in cascading layer

ID	Name	Status
8b934887-4dfe-4b11-aed6-7e2849fc6583	DB-C	ACTIVE

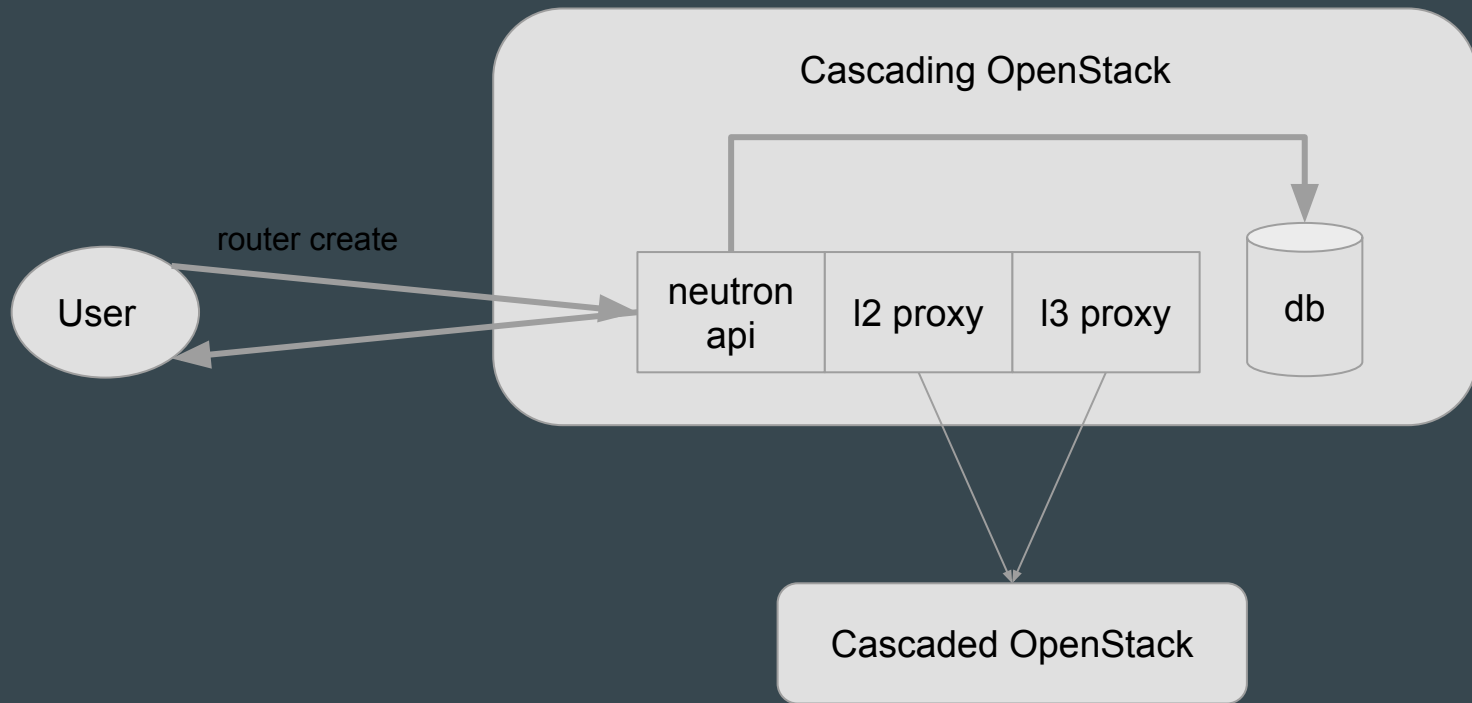
Server record stored in cascaded layer

ID	Name	Status
0910ae31-0b46-47c0-99c9-5e331b3462ba	DB-C@8b934887-4dfe-4b11-aed6-7e2849fc6583	ACTIVE

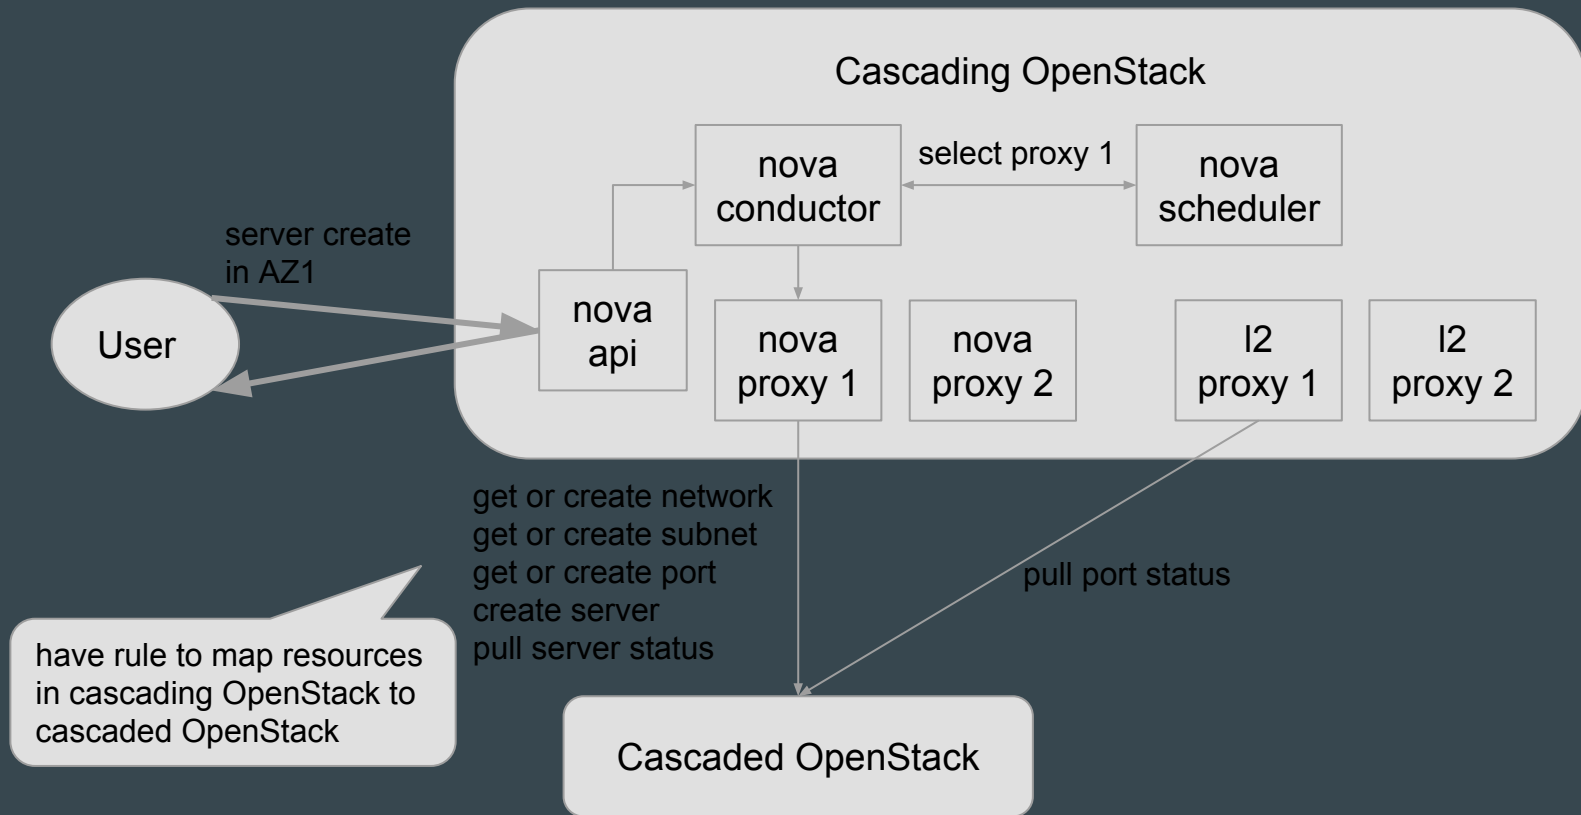
Case Study - create network



Case Study - create router




Case Study - create server




L2 Population

- Work in tunnel type networks(VXlan, GRE)
- When one server port is up, Neutron server will send forwarding table entry to other l2 agents which serving servers in the same network
- Help to decrease broadcast traffics
 - ARP request is intercepted and answered locally using the entry, no broadcast packet needed
 - Broadcast packets are only sent to related l2 agents


Without L2 Population



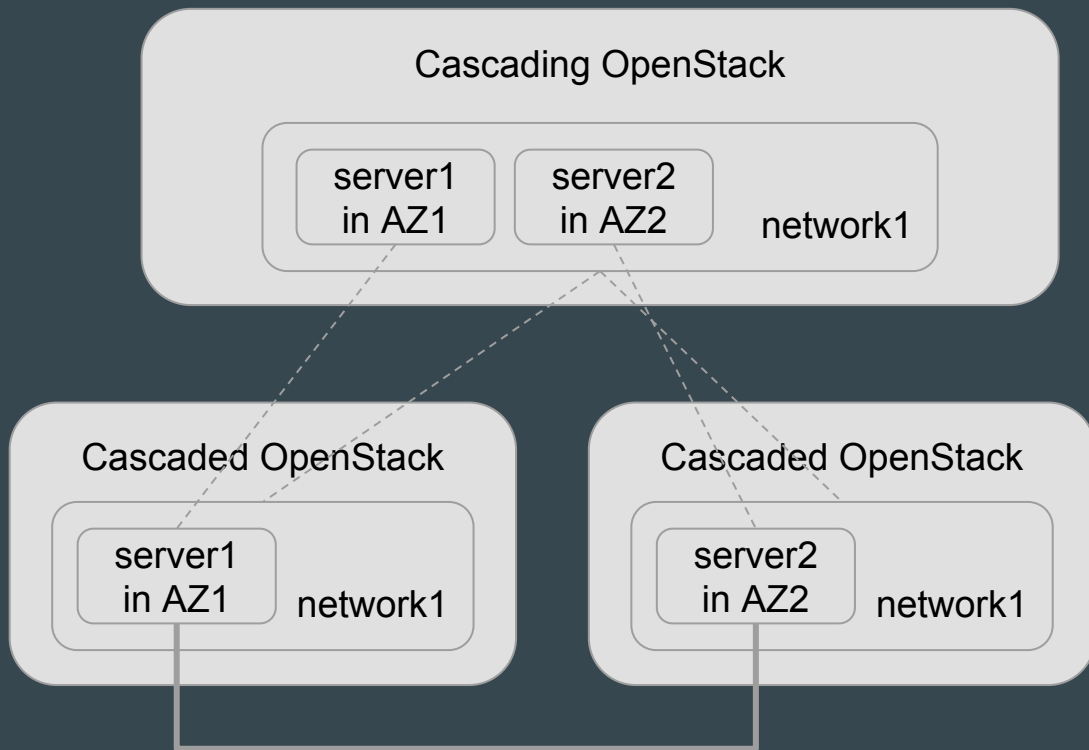
With L2 Population - broadcast



With L2 Population - unicast



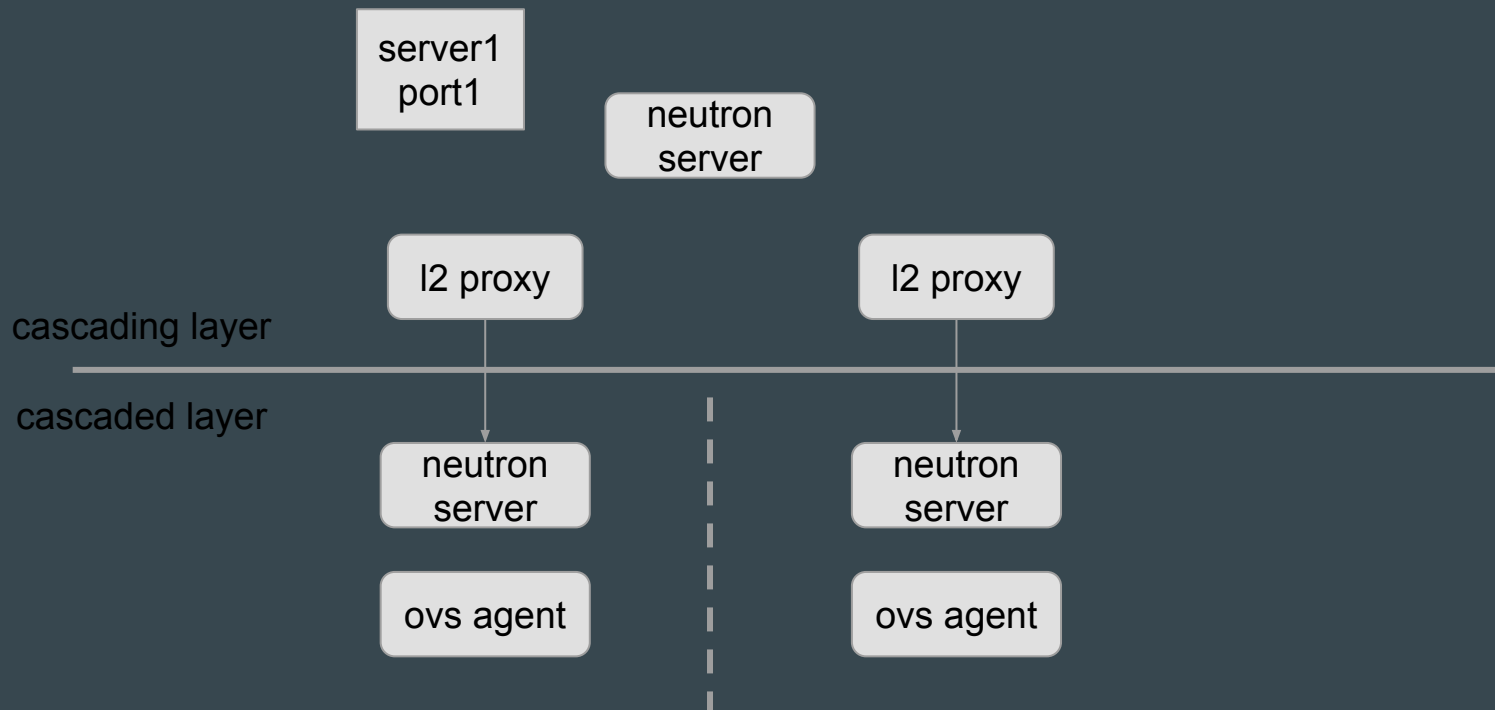
Cross Site L2 Networking - what to do



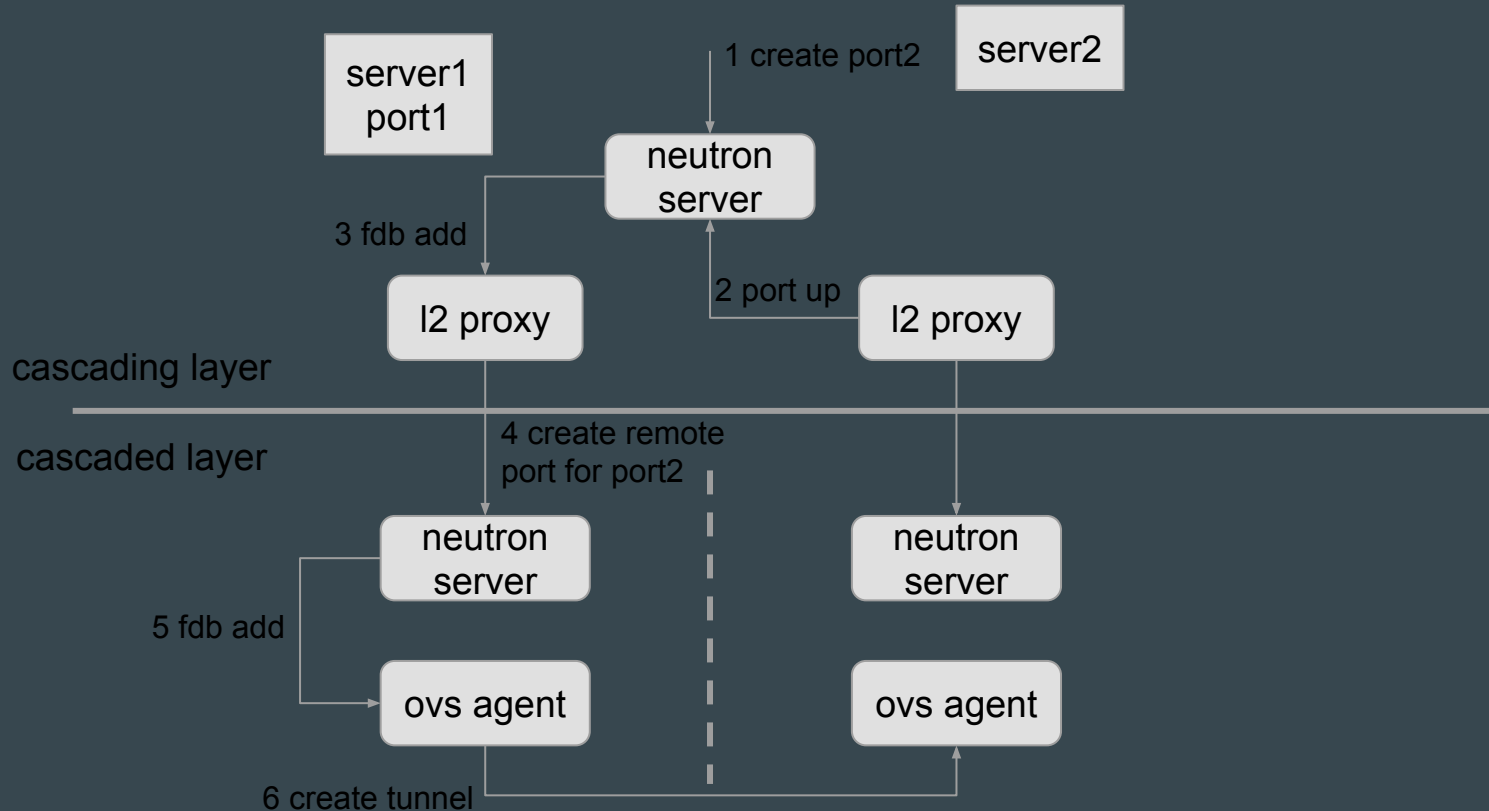
Cross Site L2 Networking - how to do

- Introduce “remote port”
 - exact the same mac, ip information with the original port
- In cascading layer, utilize l2 population to create “remote port” in other sites
- In cascaded layer, utilize l2 population to populate port information to l2 agents
- L2 agents across sites build VXlan tunnel for connection

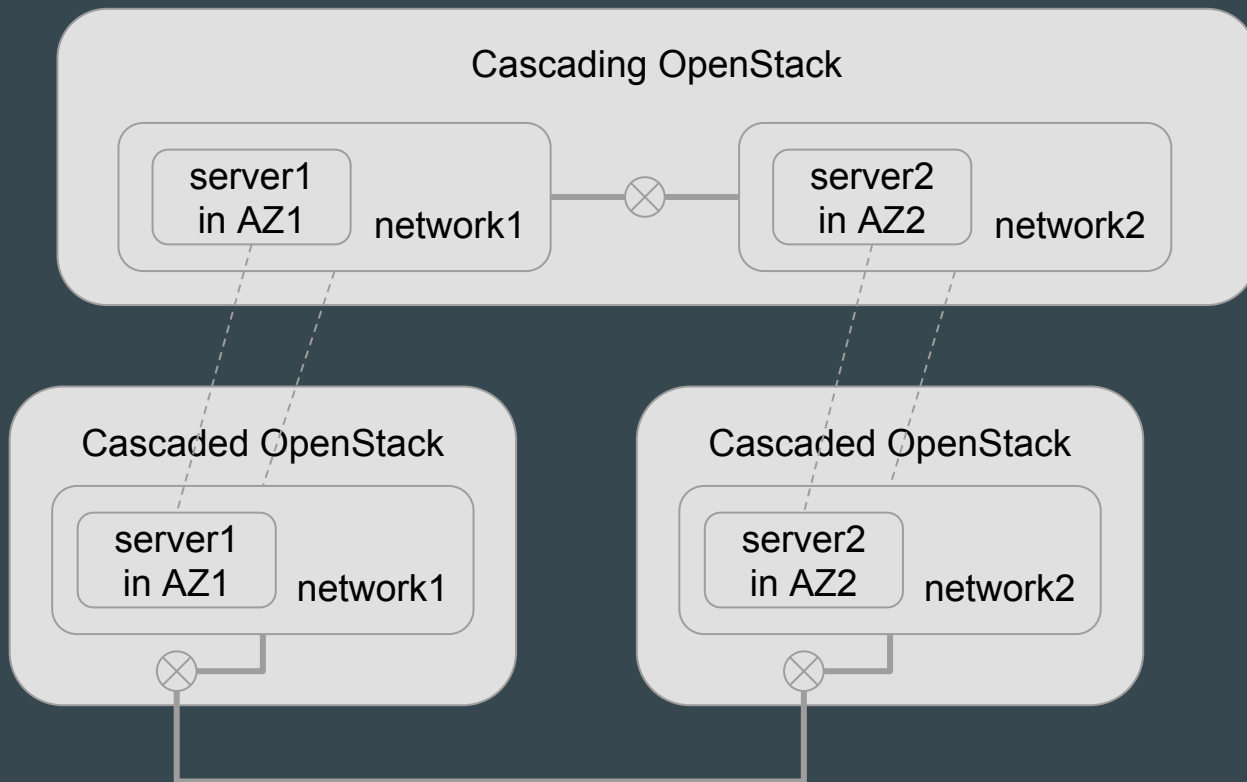
Cross Site L2 Networking - detail



Cross Site L2 Networking - detail



Cross Site L3 Networking - what to do

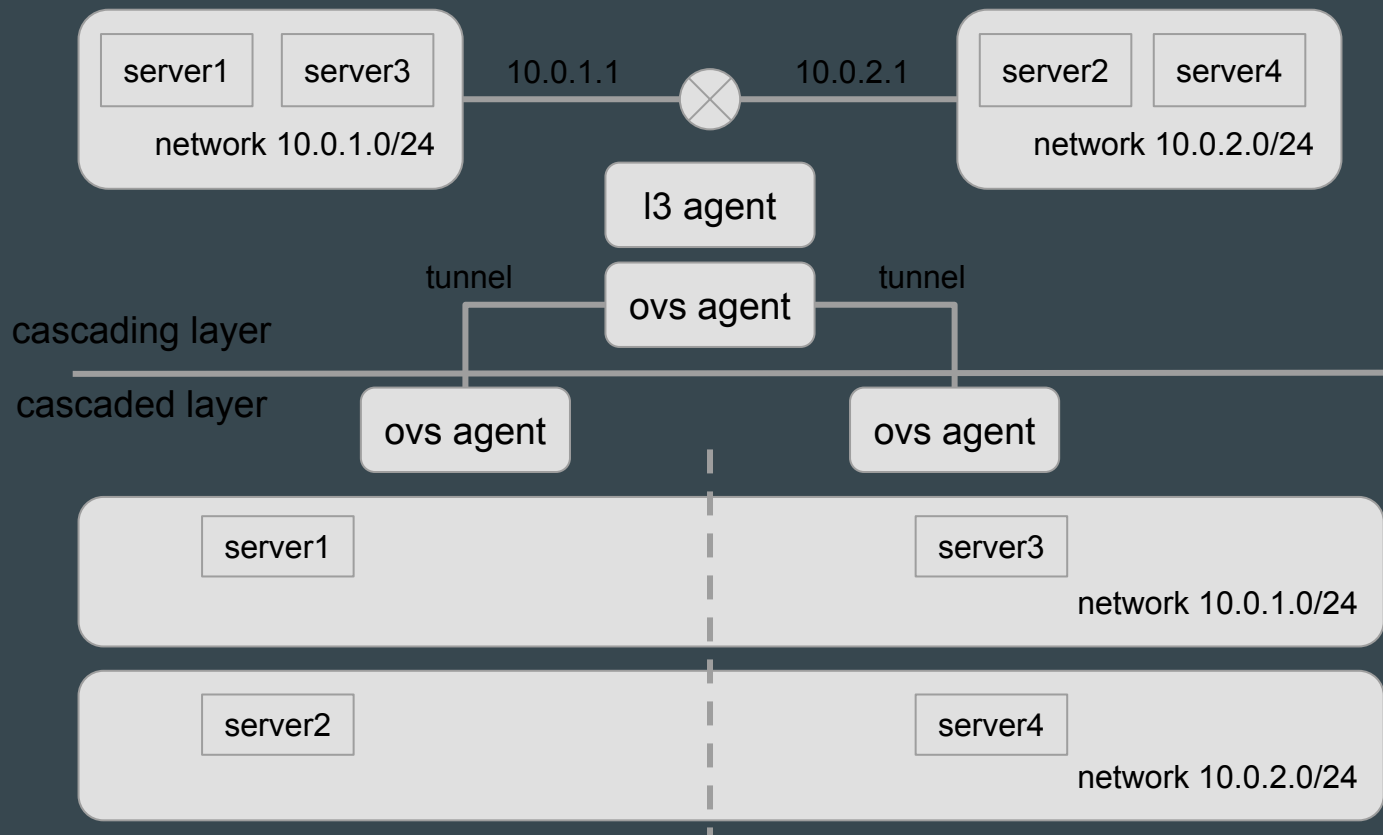


Cross Site L3 Networking - how to do

- Central router
 - Utilize cross site l2 networking
 - L2 agents connect to the cascading layer
 - No router in cascaded layer
- Distributed router
 - Routers distributed in cascaded layer
 - Utilize external network to connect routers
 - Configure routes to route packets across sites


Cross Site L3 Networking - central

networks CAN
cross sites



Cross Site L3 Networking - distributed

networks CANNOT
cross sites



Q&A