

YAHOO!

OPENSTACK CLOUD AT YAHOO SCALE: HOW TO AVOID DISASTER


Anders Beitnes

Yahoo Paranoids

Yahoo Scale

- 800 million + monthly users
- 10x traffic spikes
- News , finance, millions/day
- Mail 110+ million daily users, 170+ petabyte data, 30 billion emails sent daily.
- Mobile photo uploads +50% quarter over quarter
- We're hiring many developers, to experiment, build, and deliver.

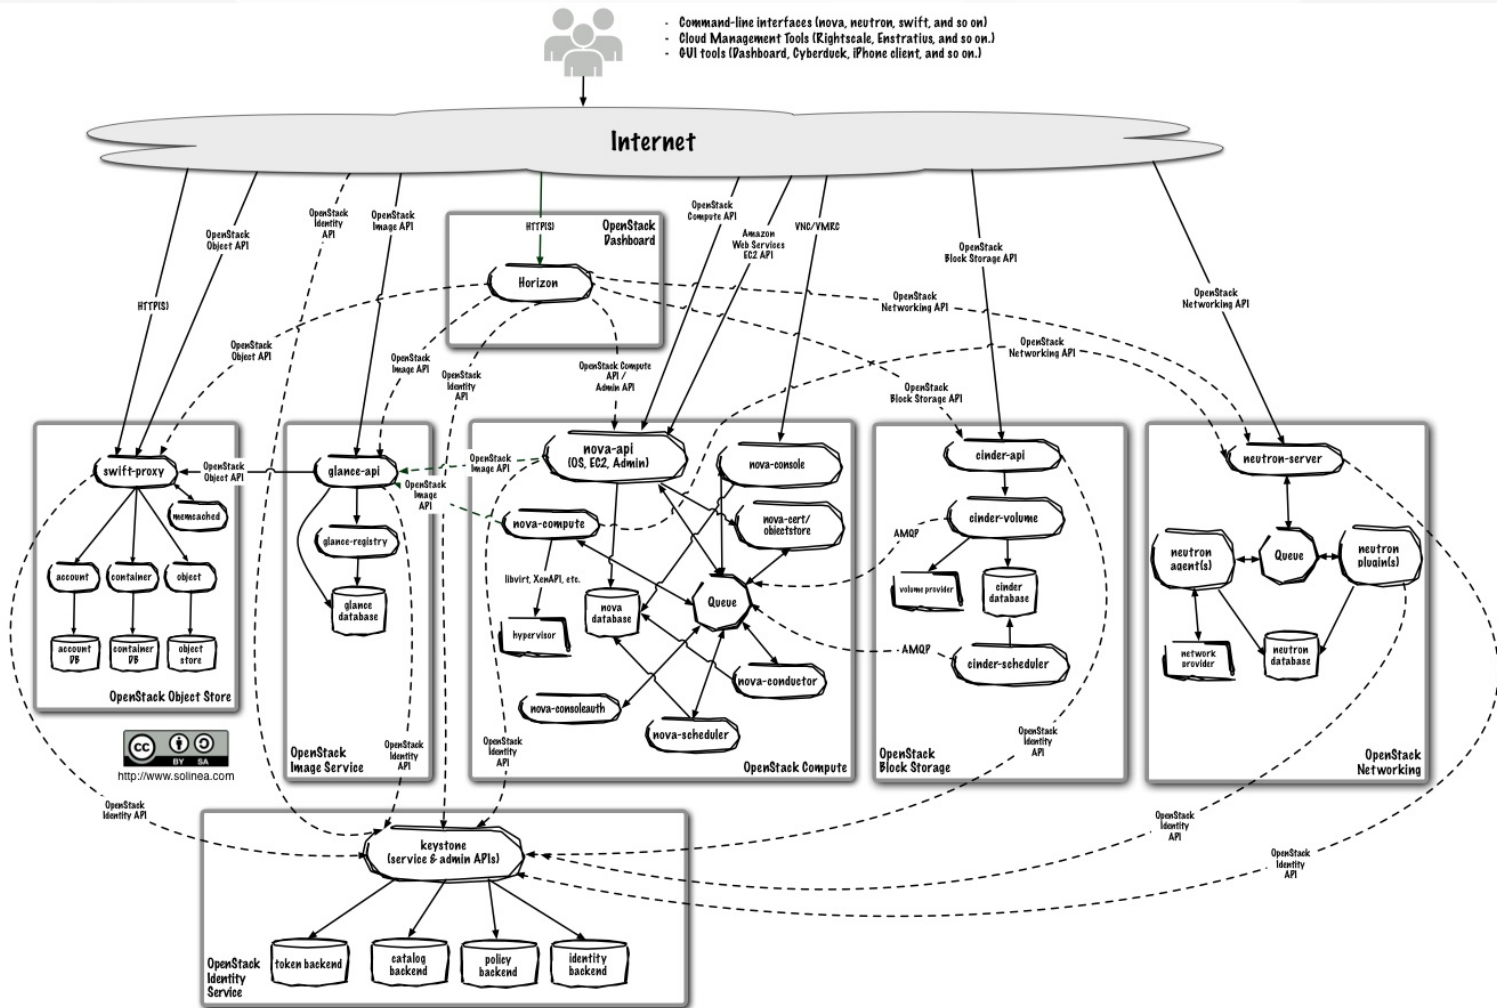
OpenStack Reach



Slice and Dice

- Multiple Private Clouds
- The production clusters are in separate security domains.
- We control what enters and leaves the VMs.
- Employee access only.
- We own the VMs and the data they contain, and can inspect it.
- We manage more than VMs

Don't Panic!



Loneliness can be a good thing

- Who should the Infrastructure nodes talk to?
- Hypervisors should listen to their parents.
- Who should VMs talk to?

Who should talk to the Corporate Infrastructure nodes?

YAHOO!
openstack
DASHBOARD


Select your region

User Name
IwantItNow

Region

- ✓ OpenHouse US (West)
- OpenHouse Asia
- OpenHouse Europe
- OpenHouse US (East)

Sign In



Launch Instance

Details Options User-Data

Instance Source
Image

Image
Select Image

Instance Name

Flavor
medium

Instance Count
1

Specify the details for launching an instance.
The chart below shows the resources used by this project in relation to the project's quotas.

Flavor Details

Name	medium
VCPUs	2
Root Disk	25 GB
Ephemeral Disk	0 GB
Total Disk	25 GB
RAM	4,096 MB


Project Quotas

Number of Instances (3) 2 Available

Number of VCPUs (4) 2 Available

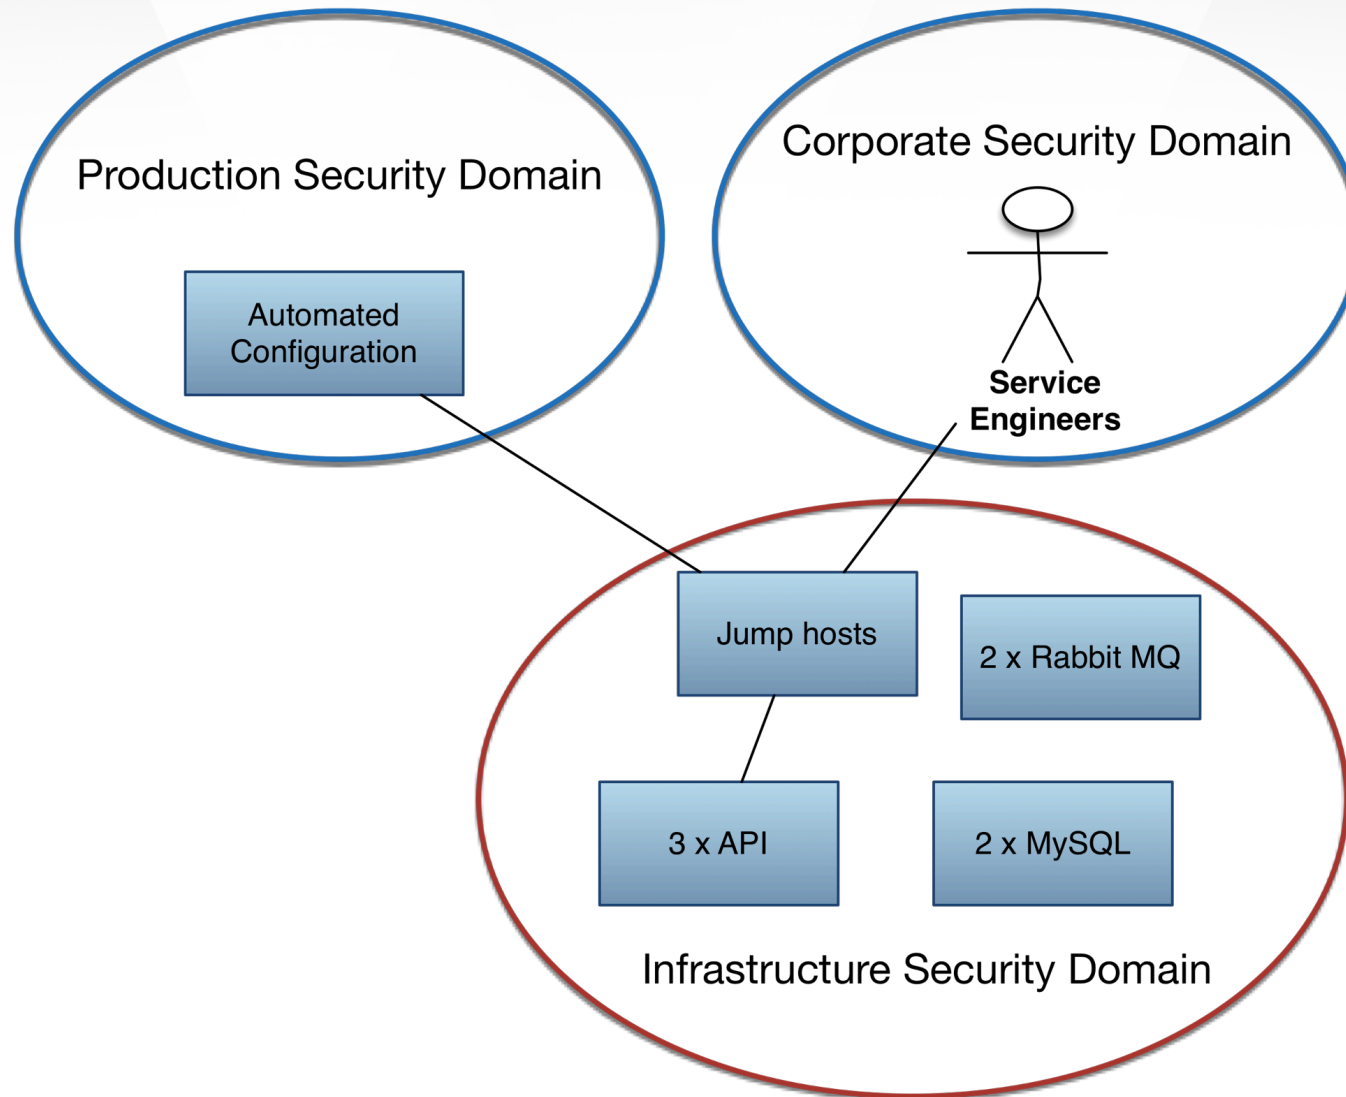
Total RAM (6,656 MB) 5,632 MB Available

Cancel Launch



while true; do ssh ihatetowait.corp.ne1.yahoo.com; sleep 5; done

Who should talk to the Production Infrastructure nodes?



Who should the Infrastructure nodes talk to?

Infrastructure Nodes

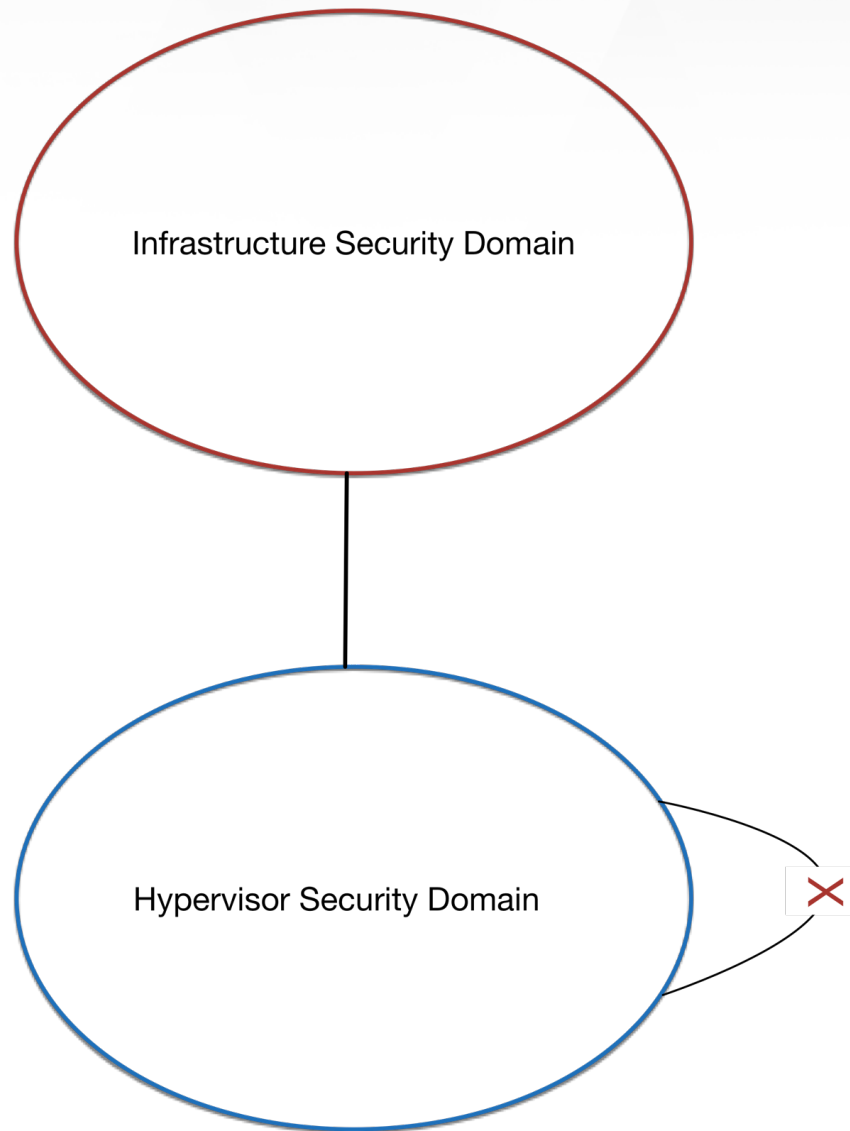
nova-compute
hypervisor

DNS IP address
assignment

Operations
database

ETC.

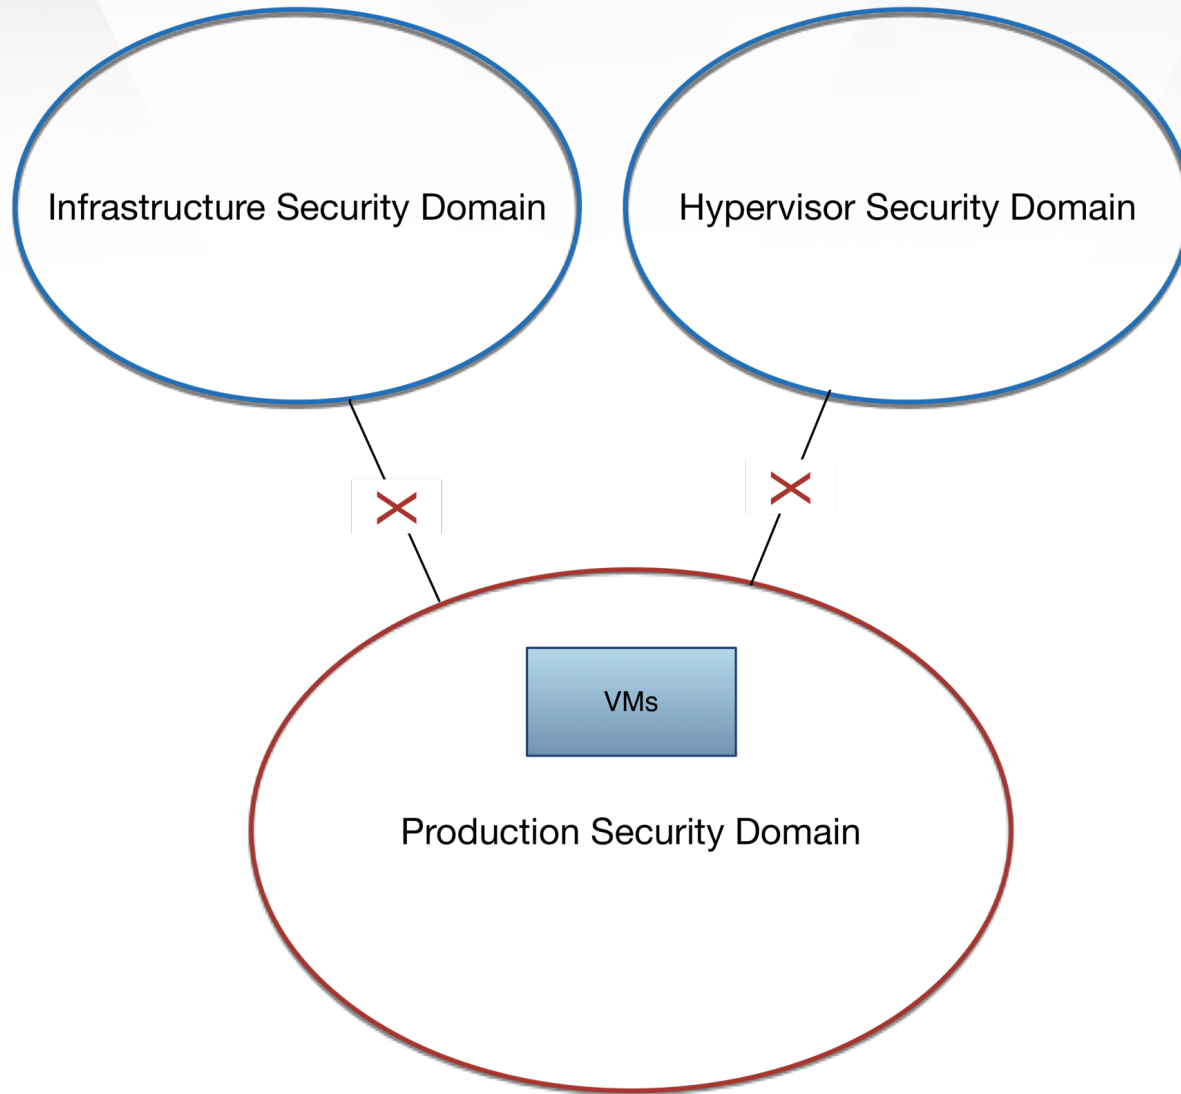
Who should the Hypervisors talk to?



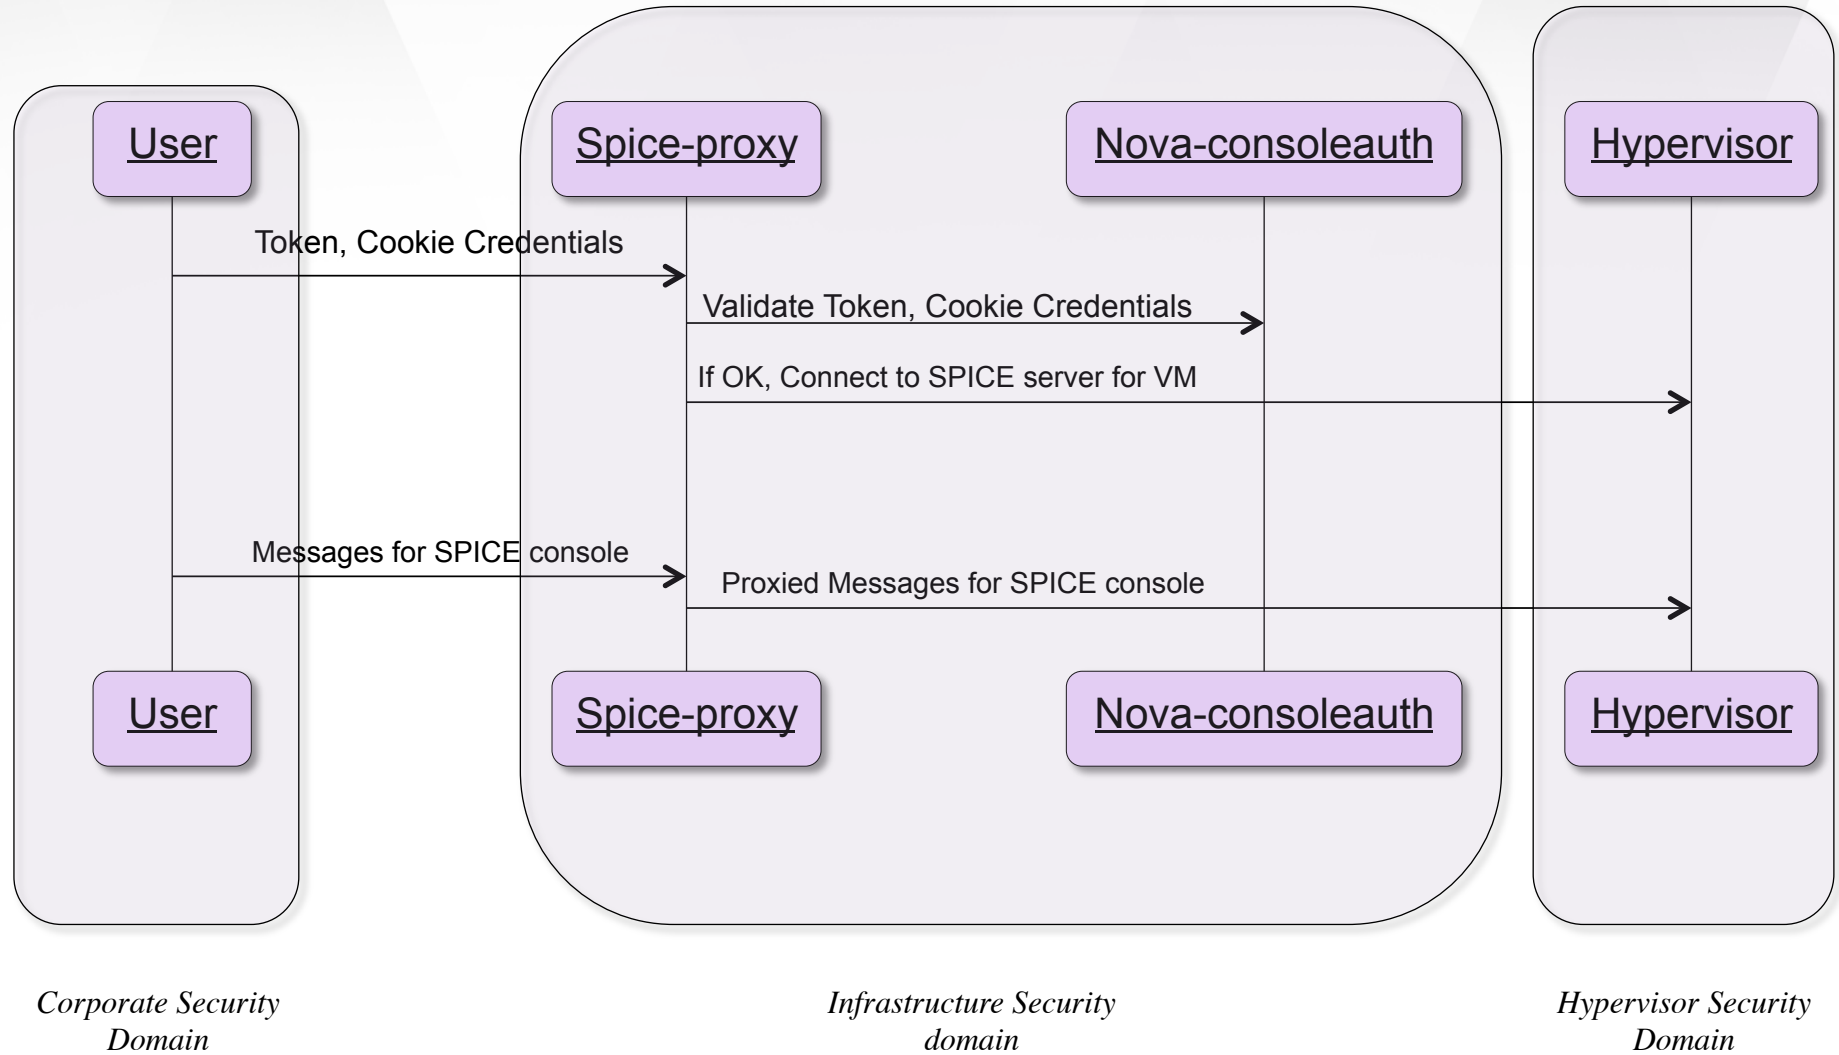
No VM Migration?



Who should the VMs talk to?



Who should talk to the VMs?



One Infrastructure node cluster to rule them all



What about those responsible for the Infrastructure nodes?



Keystone



Keystone policy file

```
{  
  "admin_required": [{"role:admin"}, {"is_admin:1"}],  
  "gsd_admin": [{"rule:admin_required"}, {"role:GSDAdmin"}],  
  "default": [{"rule:admin_required"}],  
  
  "get_role": [{"rule:gsd_admin"}],  
  "get_roles": [{"rule:gsd_admin"}],  
  "get_role_refs": [{"rule:gsd_admin"}],  
  "get_all_projects": [{"rule:gsd_admin"}],  
  "get_project": [{"rule:gsd_admin"}],  
  "get_project_users": [{"rule:gsd_admin"}],  
  "get_user": [{"rule:gsd_admin"}],  
  "get_users": [{"rule:gsd_admin"}],  
  "get_user_roles": [{"rule:gsd_admin"}]  
}
```

Impostors and Spies



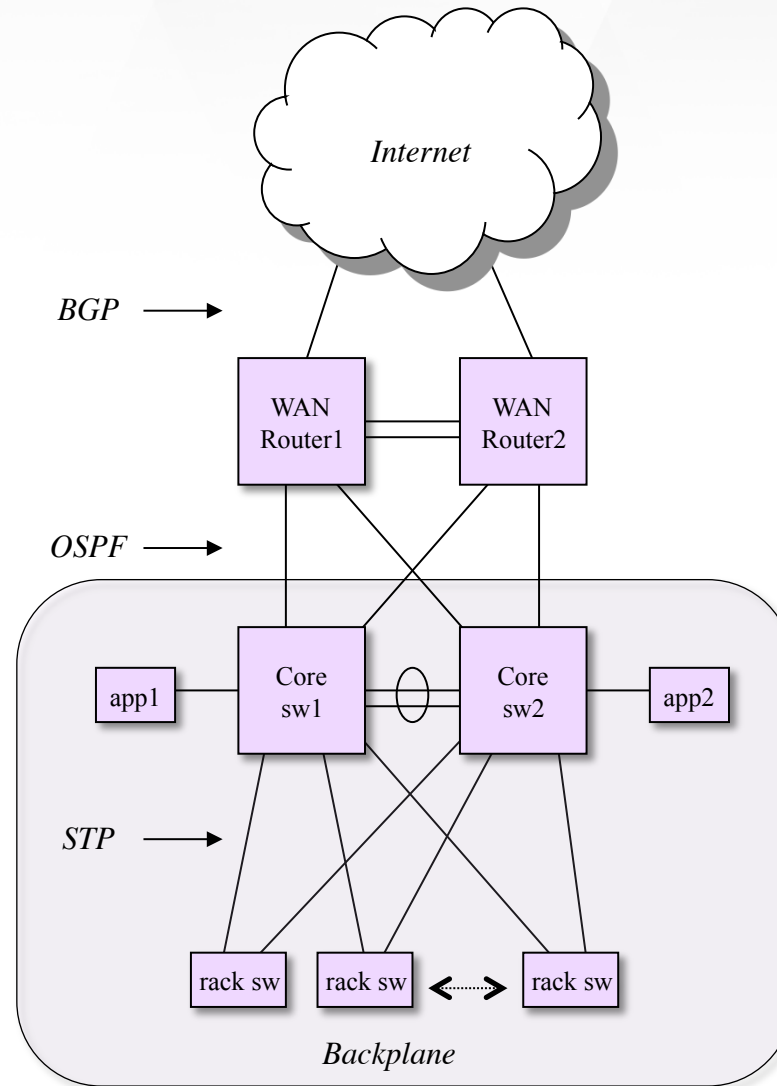
Someone to watch over me



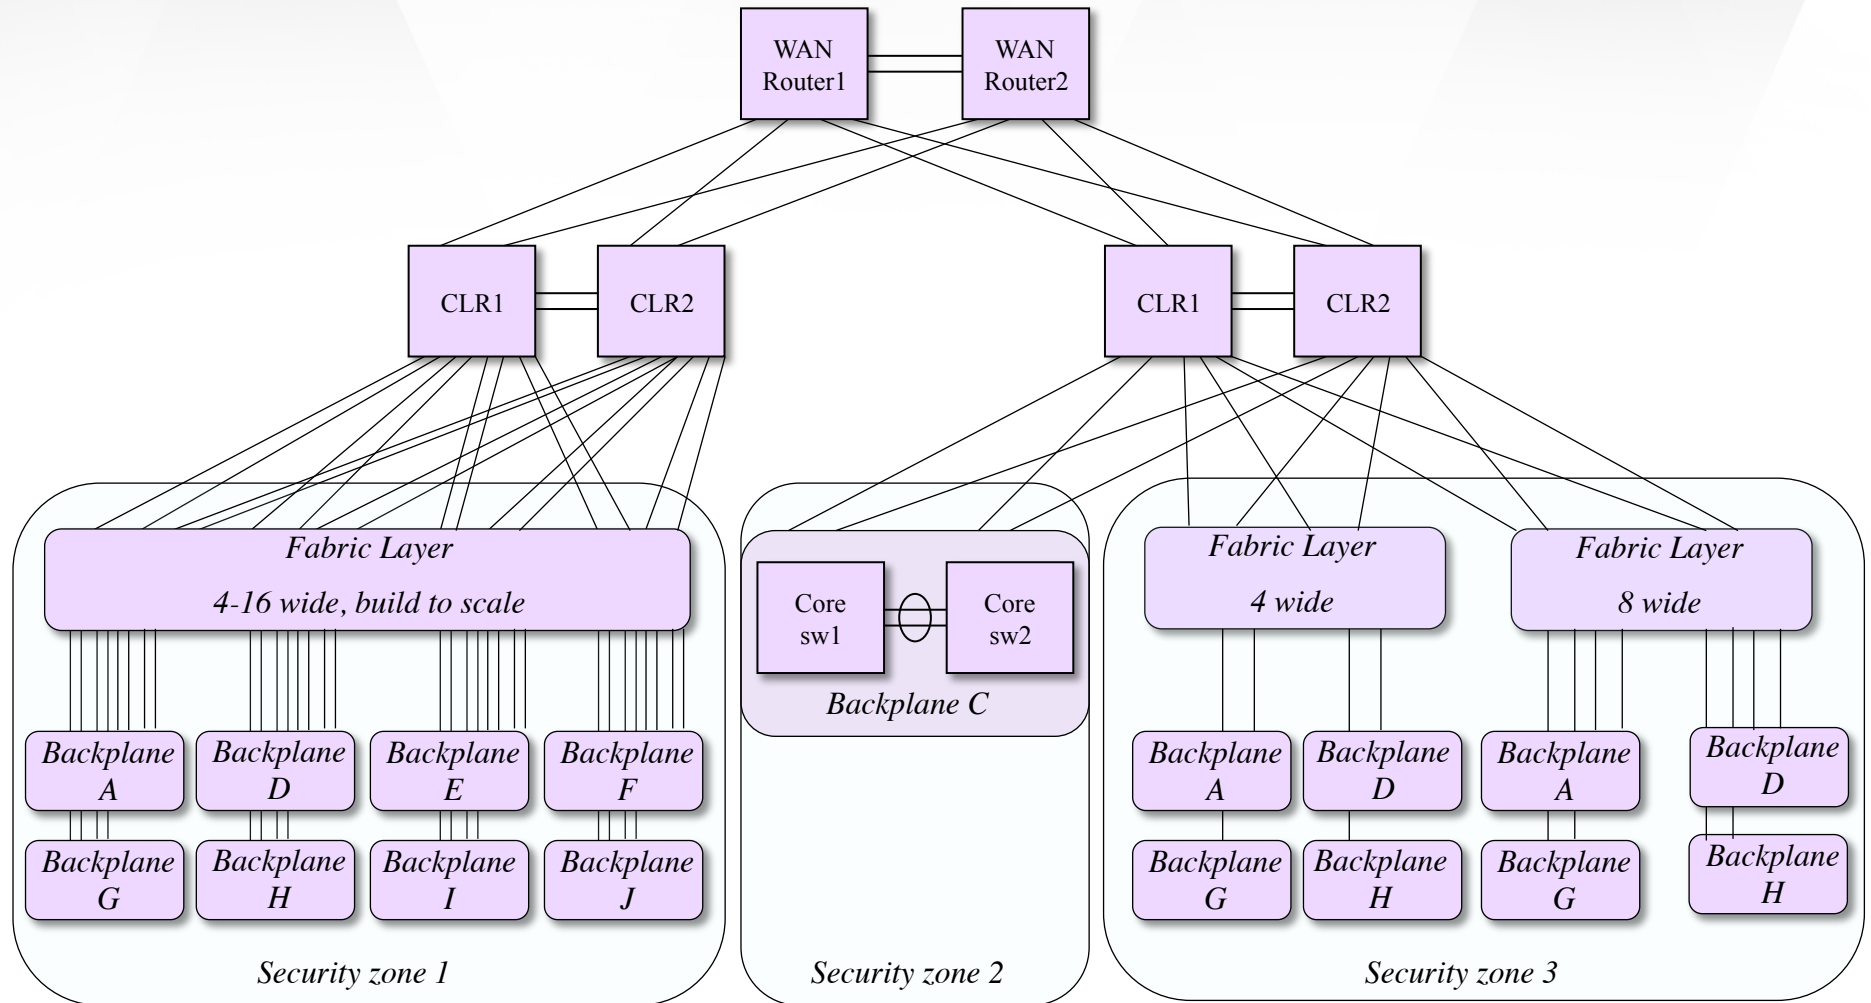
Next generation networks

- What are the issues with Networking at scale?
- L2 networks provide VLAN isolation, but have an associated cost problem
- L3 networks solve the cost issue, but pose a security issue.

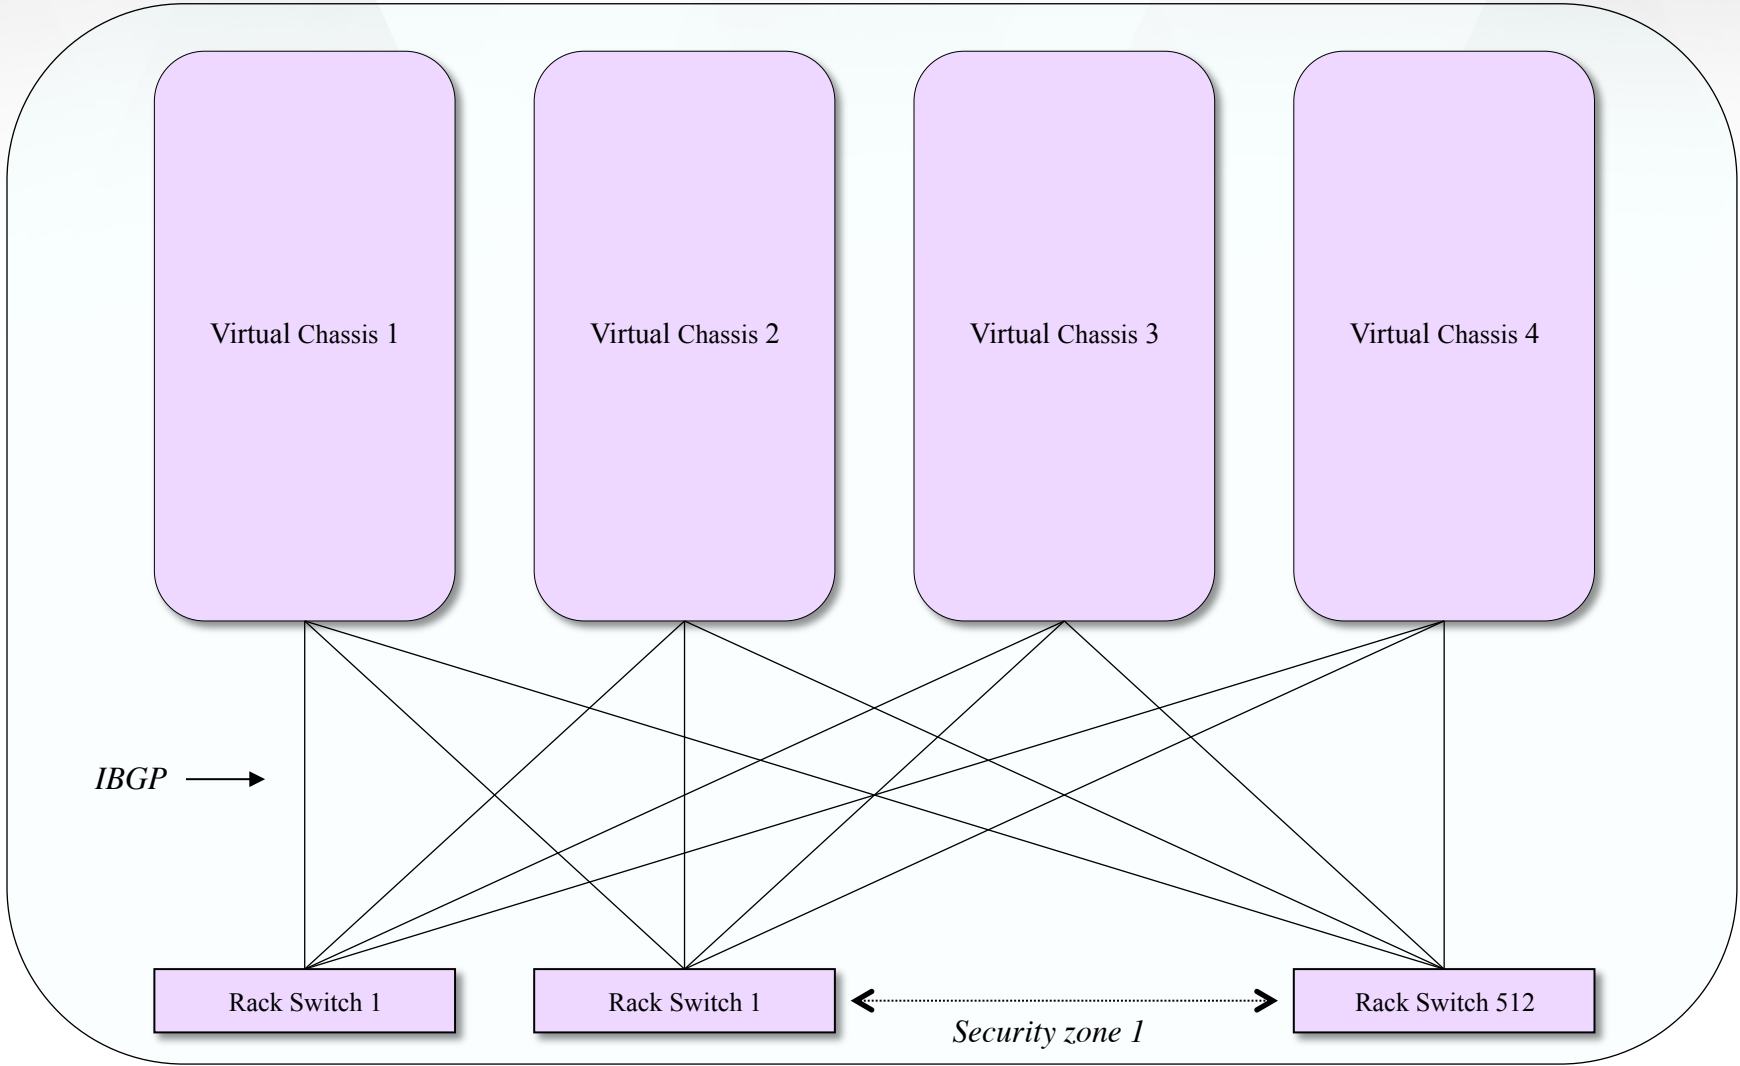
L2 Network



L2 network + Yahoo Scale



Next generation networks at Yahoo Scale





Thanks!

Credits and attributions

- <https://www.flickr.com/photos/andrewgustar/7280161708>
 - › <https://creativecommons.org/licenses/by-nd/2.0/>
- <https://www.flickr.com/photos/arbron/8278577848/>
 - › <https://creativecommons.org/licenses/by-nd/2.0/>
- <https://www.flickr.com/photos/clearlyambiguous/131620637>
 - › <https://creativecommons.org/licenses/by-nd/2.0/>
- <https://www.flickr.com/photos/87913776@N00/460181144>
 - › <https://creativecommons.org/licenses/by-nd/2.0/>
- <https://www.flickr.com/photos/macgodbrad/4869419570>
 - › <https://creativecommons.org/licenses/by-sa/2.0/>
- <https://www.flickr.com/photos/sklathill/2255718951/>
 - › <https://creativecommons.org/licenses/by-sa/2.0/>