Containers and IPv6

Docker Container

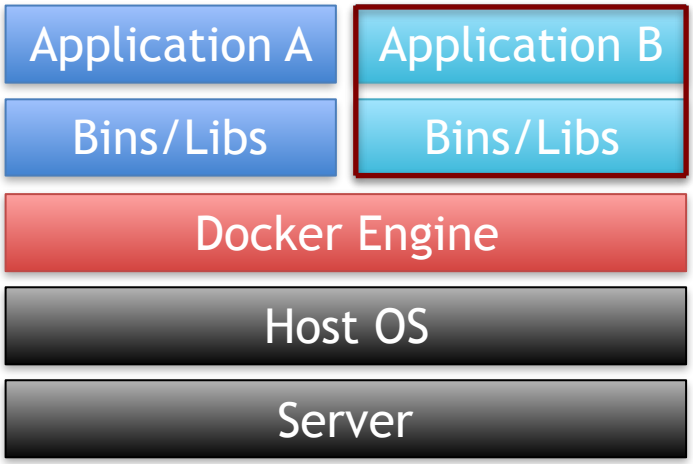

Virtual Machine

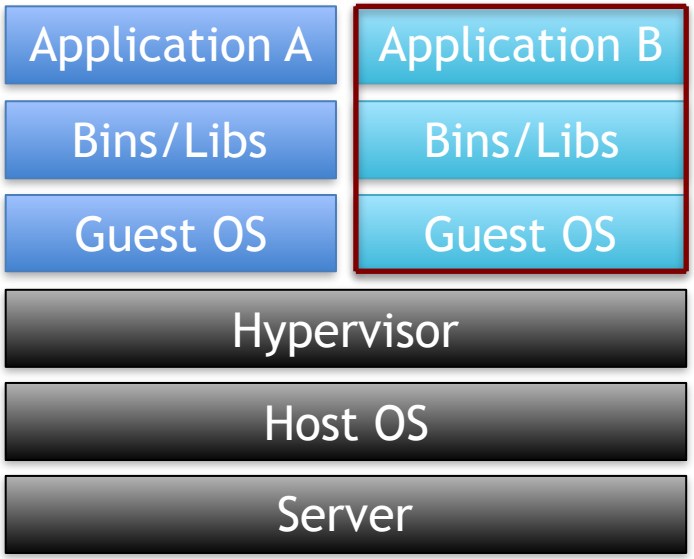

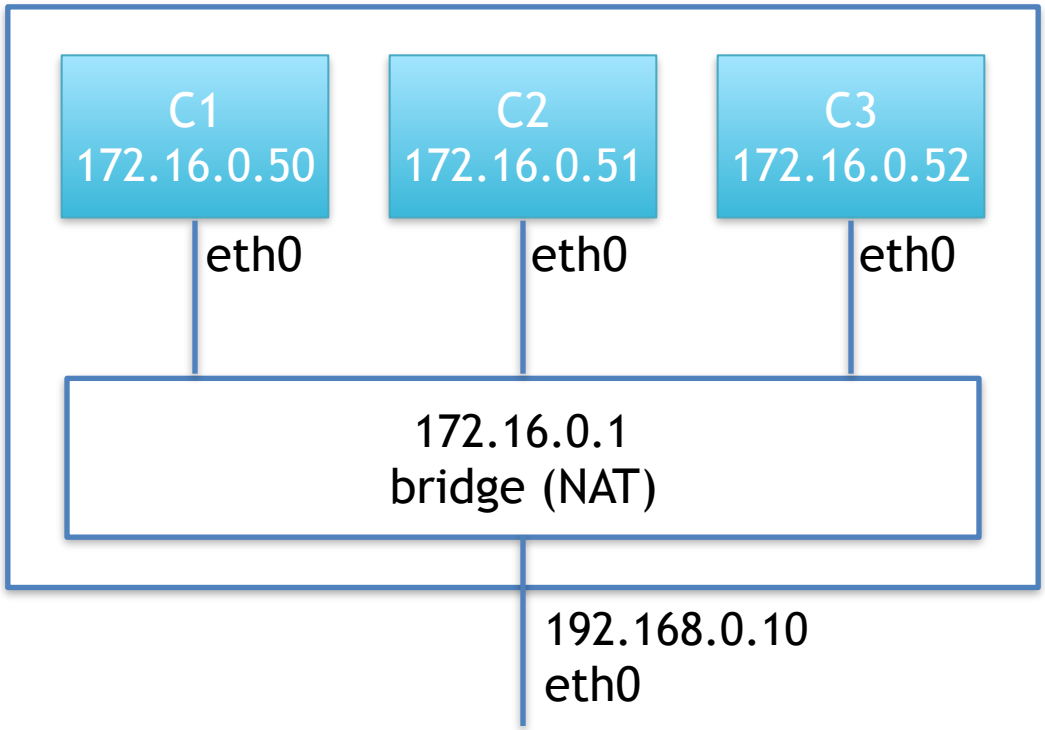
C1
172.16.0.50
C2
172.16.0.51
C3
172.16.0.52
eth0
eth0
eth0
172.16.0.1
bridge (NAT)
192.168.0.10
eth0

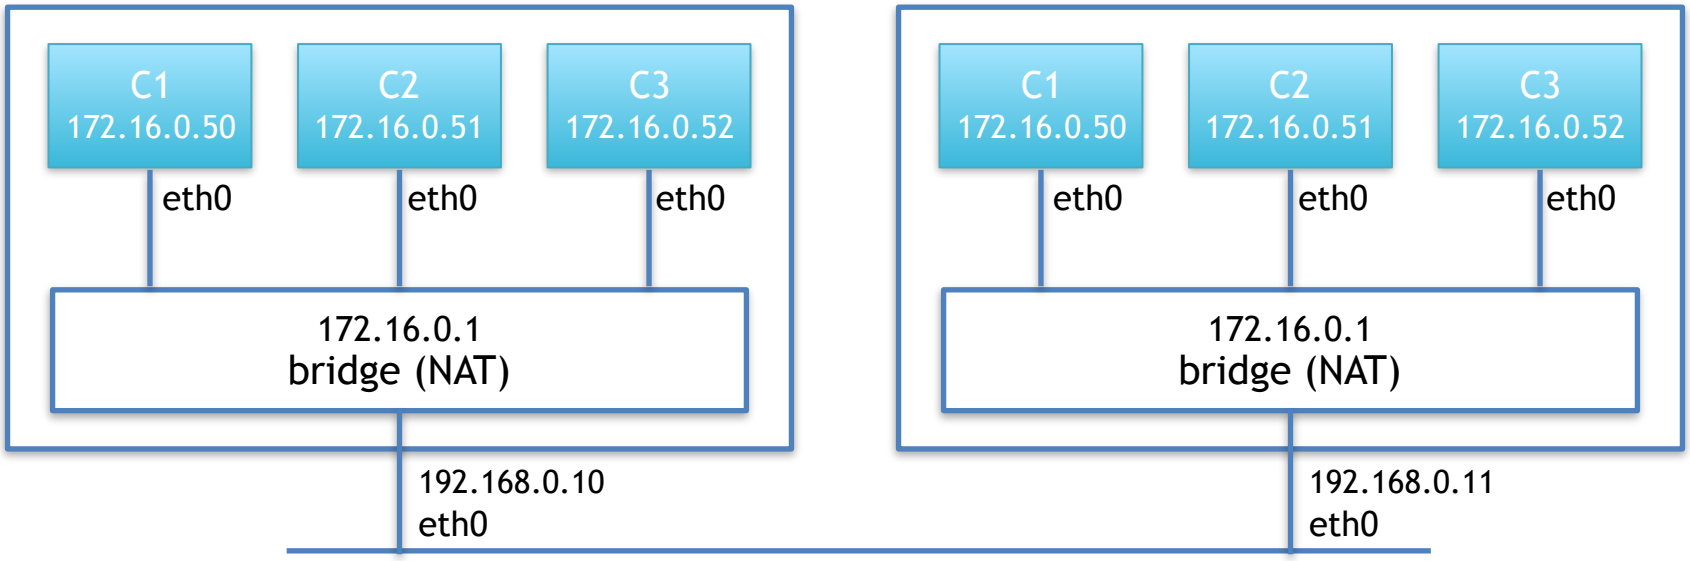
C1
172.16.0.50
C2
172.16.0.51
C3
172.16.0.52
eth0
eth0
eth0
172.16.0.1
bridge (NAT)
192.168.0.10
eth0
C1
172.16.0.50
C2
172.16.0.51
C3
172.16.0.52
eth0
eth0
eth0
172.16.0.1
bridge (NAT)
192.168.0.11
eth0

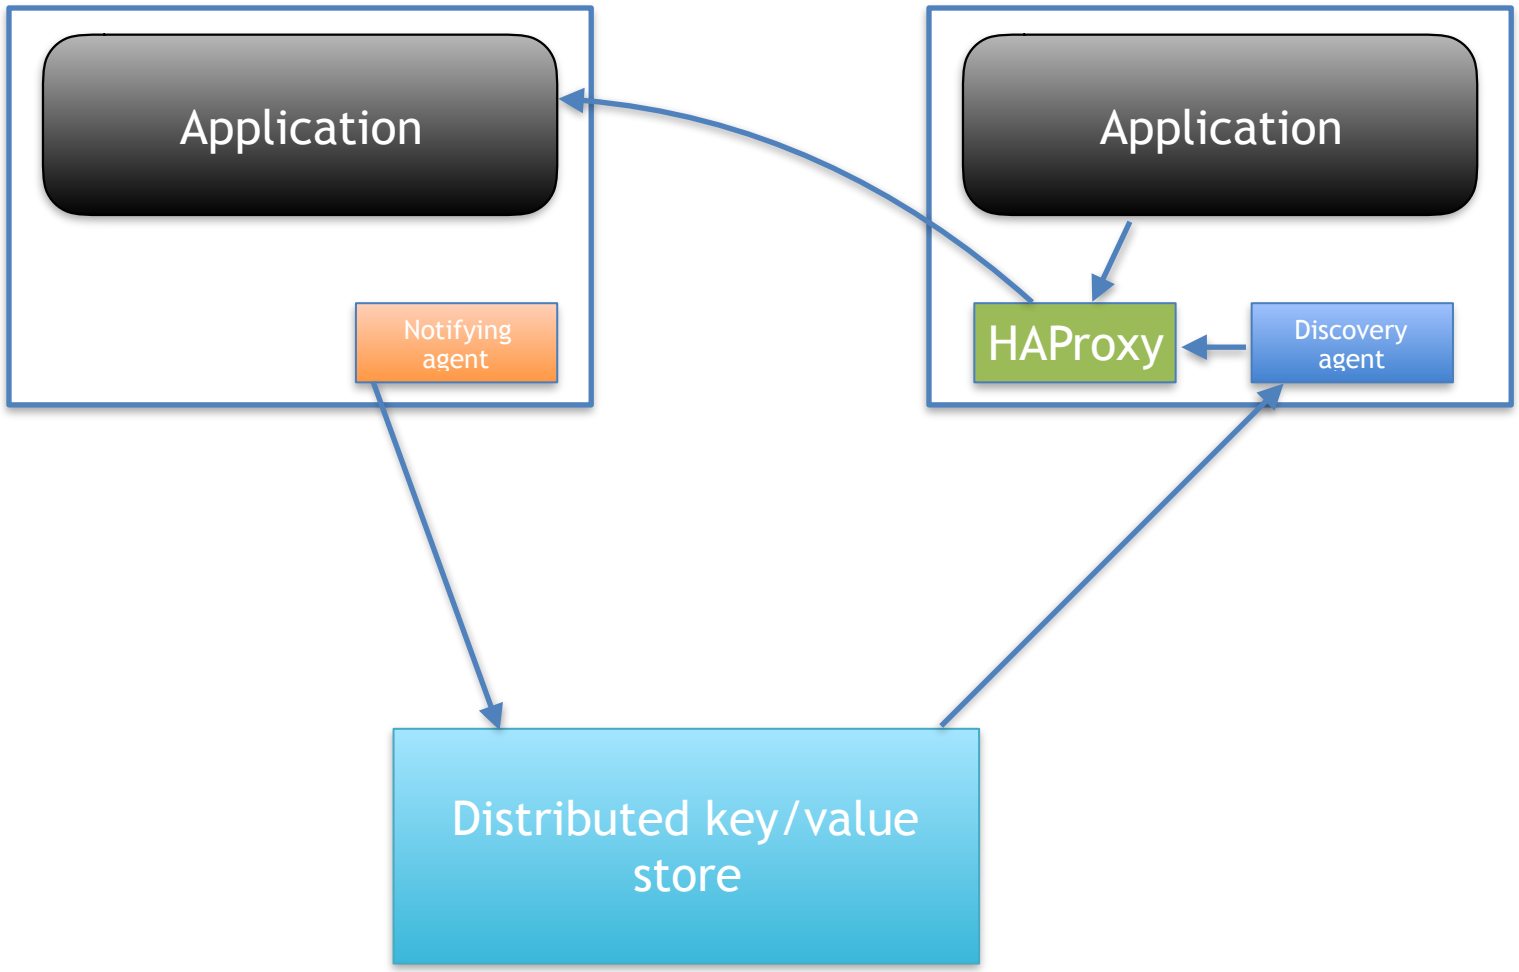
Application
Application
Notifying agent
HAProxy
Discovery agent
Distributed key/value store

Excerpt from open letter from Nephos6 to AWS

(http://www.nephos6.com/aws-lack-of-ipv6-support-is-expensive-open-letter/)

"We dealt with your lack of native IPv6 support because you do a lot of things well however, adding an ELB just to meet our IPv6 needs is becoming a pain, a compounding pain:

- *We lose the opportunity to take advantage of improved user experience over IPv6*
- *We take the performance hit on translation*
- *We add a point of failure*
- *And not only it represents 15% of our infrastructure expenses but it is going up with our increasing user base."*

- From Docker 1.5 IPv6 is supported
- IPv6 desirable to avoid port forwarding
- One deployment scenario sees containers used for short lived jobs and destroyed – very high number of containers over time

- Configuration highly static
- Multiple ways to configure IP prefixes:
  - Assign new /64 prefix to Docker
  - Assign longer prefix from /64 assigned to host
  - Use addresses from the host subnet with NDP proxy

- Limited operational experience to understand which are best deployment scenarios

Assign prefix to Docker for use by containers

```
docker –d --ipv6 –fixed-cidr-v6="2001:db8:1::/64"
```

Configure proxy for NDP

```
sysctl net.ipv6.conf.eth0.proxy_ndp=1
ip -6 neigh add proxy 2001:db8::abcd dev eth0
```

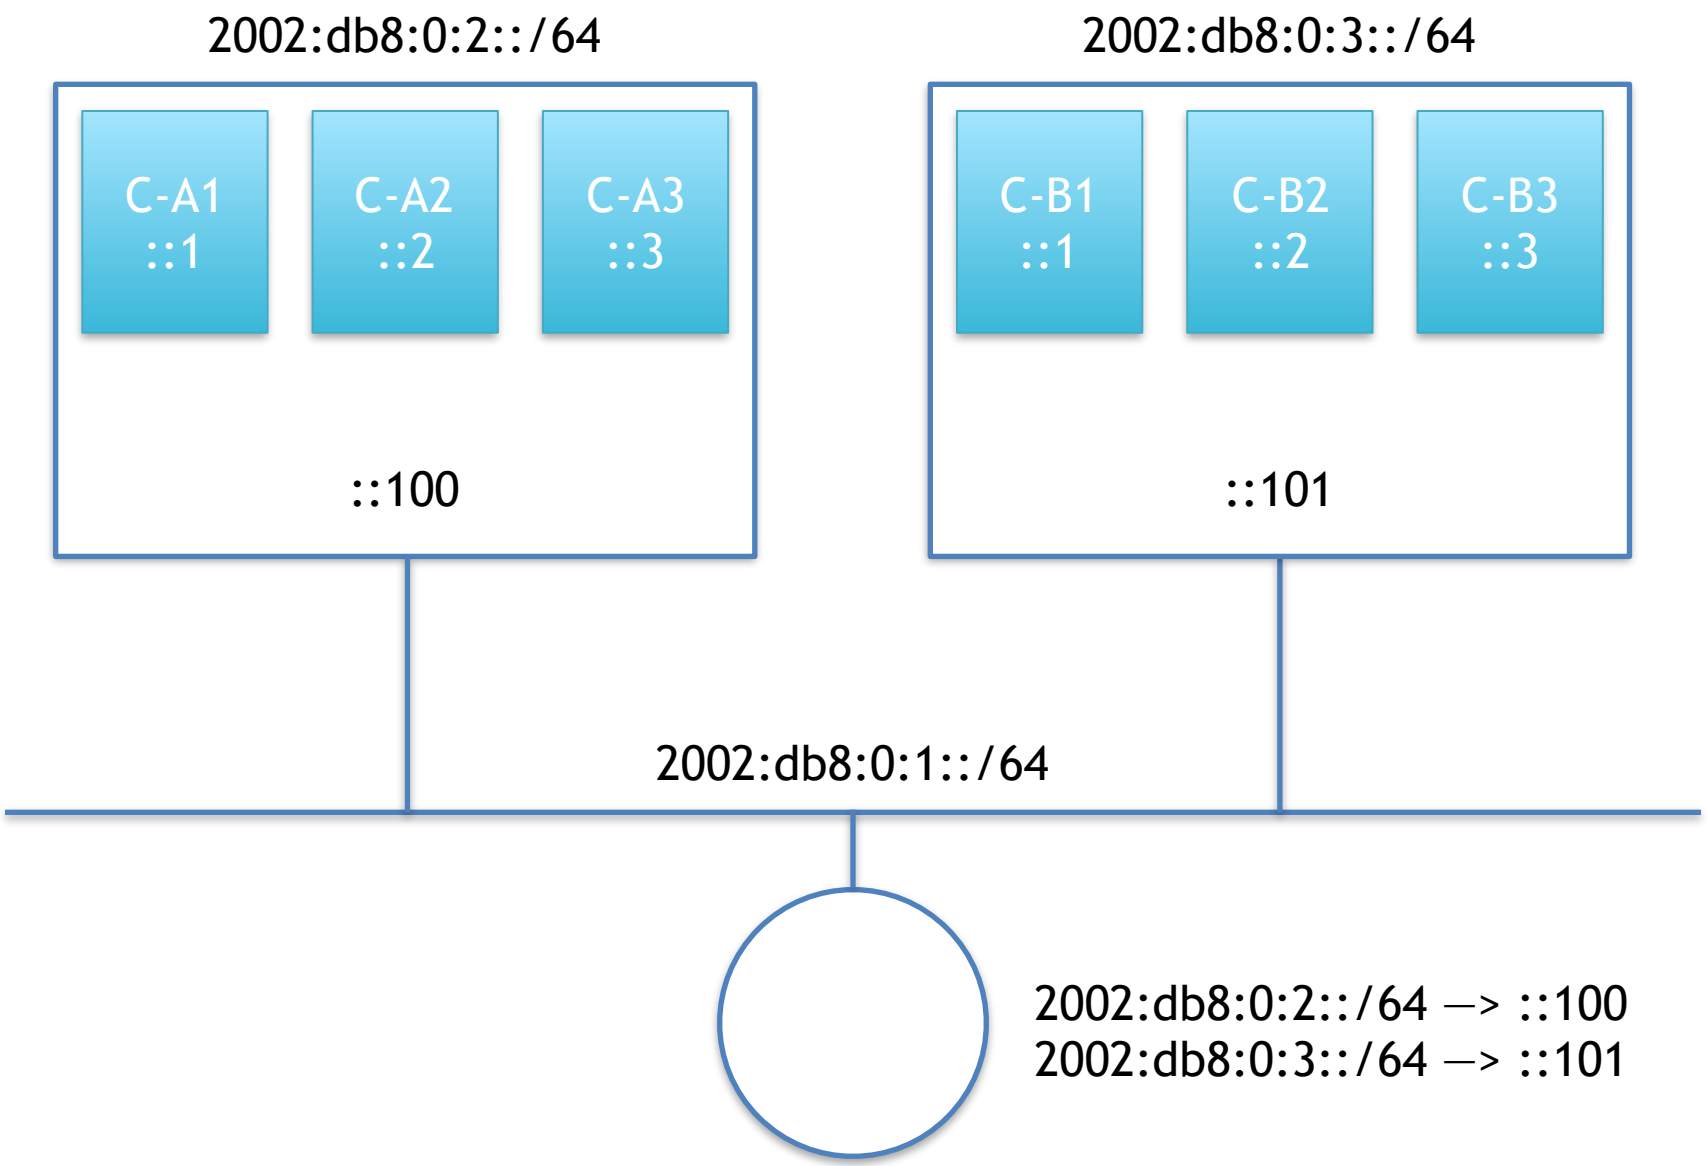
2002:db8:0:2::/64
2002:db8:0:3::/64
C-A1
::1
C-A2
::2
C-A3
::3
C-B1
::1
C-B2
::2
C-B3
::3
::100
::101
2002:db8:0:1::/64
2002:db8:0:2::/64 —> ::100
2002:db8:0:3::/64 —> ::101

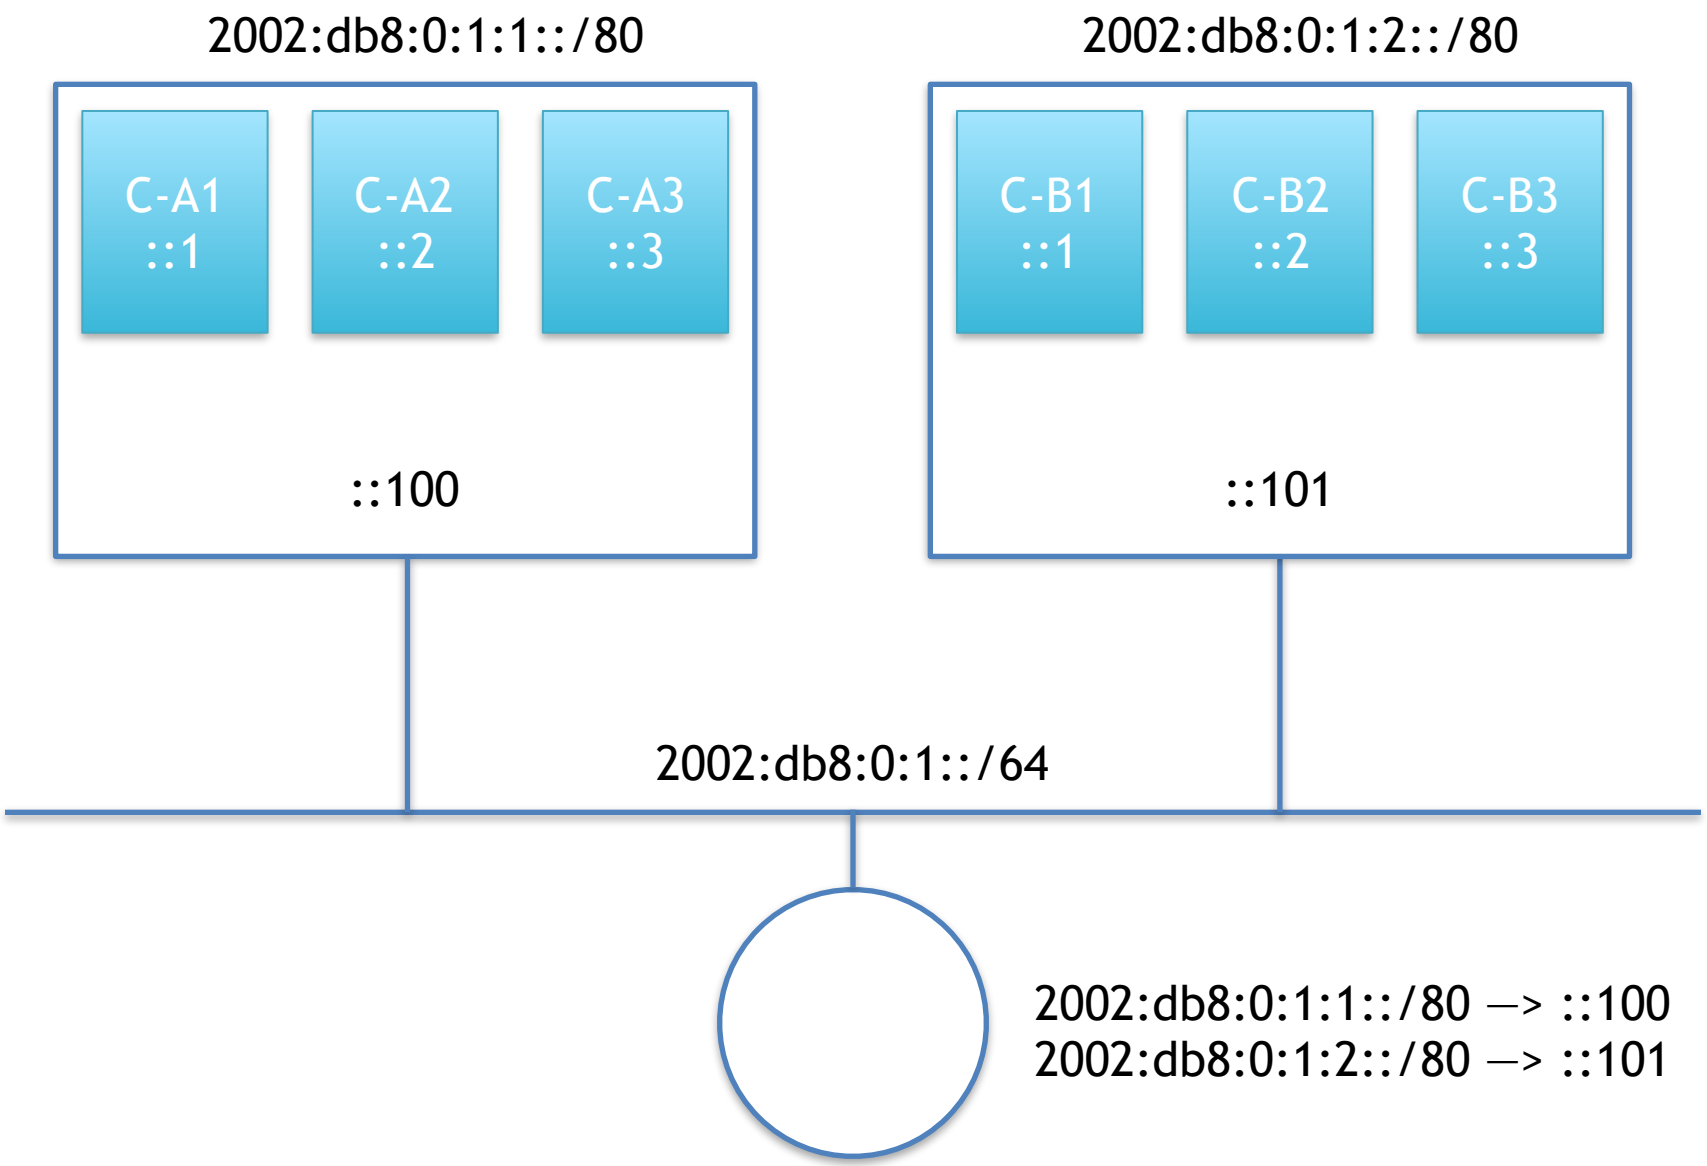
2002:db8:0:1:1::/80
2002:db8:0:1:2::/80
C-A1
::1
C-A2
::2
C-A3
::3
C-B1
::1
C-B2
::2
C-B3
::3
::100
::101
2002:db8:0:1::/64
2002:db8:0:1:1::/80 —> ::100
2002:db8:0:1:2::/80 —> ::101

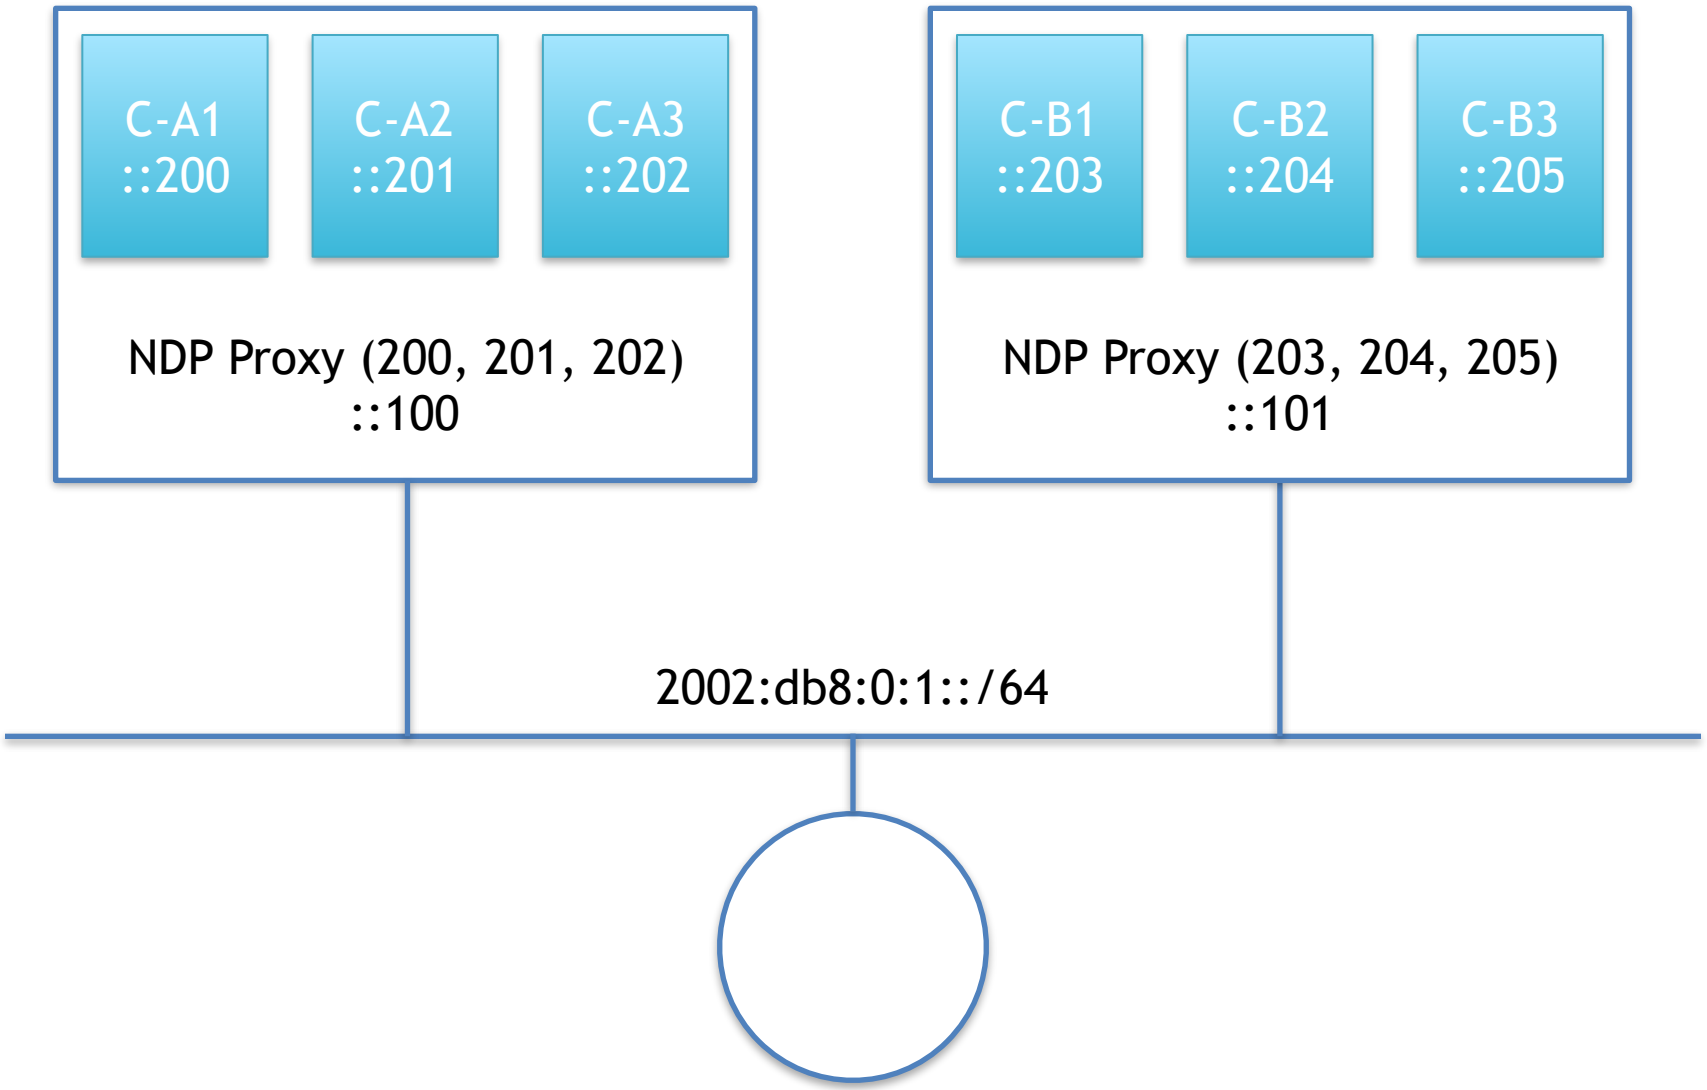
C-A1
::200
C-A2
::201
C-A3
::202
NDP Proxy (200, 201, 202)
::100
C-B1
::203
C-B2
::204
C-B3
::205
NDP Proxy (203, 204, 205)
::101
2002:db8:0:1::/64

- IPv4 -> IPv6 flows for containers desirable

- Highly orchestrated create/destroy lifecycle
- Only some containers require public communication
- SIIT-DC suitable solution (draft-ietf-v6ops-siit-dc-03)

- Add dynamic creation of SIIT-DC entries by
  - DNS lookup for AAAA record referring to container
  - SIIT-DC translation entry removed when container destroyed or at TTL of AAAA
- Or
  - Flagged as required in container configuration
  - SIIT-DC translation entry removed when container destroyed
  - Needs cost model otherwise every container owner sets “public” flag “just in case”