



RIPE NCC
RIPE NETWORK COORDINATION CENTRE

Policy routing in K-root servers

Traditional configuration




- Just one interface
- Used for both management and service traffic
- Simple configuration
- Disadvantages
 - Shared fate
 - Exposed management
 - Inaccurate statistics



Single-interface device

Traffic separation



Traffic separation



- Deploy servers with dual network interfaces
- Use one (usually the first one) for management
- Use the second or subsequent ones for service

Advantages of traffic separation



- Service doesn't affect management
 - traffic floods (DoS)
- Different ACLs and security policies
- Different interface settings (MTU, ARP, bandwidth)
- Service traffic statistics are accurate
- Management interface on private network or exposed only through a VPN

Linux advanced routing



- Linux kernel can maintain many routing tables
 - give tables symbolic names in `/etc/iproute2/rt_tables`
- Without any rules, the “main” table is used
- Create rules with “ip rule” command
 - direct Linux to look up alternate tables
 - rules can examine source address, destination address and some other attributes of packets

K-root configuration (interfaces)



```
# ip addr show
```

```
1: lo: <LOOPBACK,UP,LOWER_UP> ...  
   inet 127.0.0.1/8 scope host lo  
   inet6 ::1/128 scope host  
2: em1: <BROADCAST,MULTICAST,UP,LOWER_UP> ...  
   inet 212.36.221.254/30 brd 212.36.221.255 scope global em1  
   inet6 2a00:1590:27:1::2/64 scope global  
3: em2: <BROADCAST,MULTICAST,UP,LOWER_UP> ...  
   inet 194.126.19.38/30 brd 194.126.19.39 scope global em2  
   inet6 2a00:1590:27:2::2/64 scope global  
4: dummy0: <BROADCAST,NOARP,UP,LOWER_UP> ...  
   inet 193.0.14.129/23 brd 193.0.15.255 scope global dummy0  
   inet6 2001:7fd::1/32 scope global
```

K-root configuration



```
# echo 10 anycast > /etc/iproute2/rt_tables
```

```
# ip rule show
```

```
0: from all lookup local
```

```
32764: from 194.126.19.38 lookup anycast
```

```
32765: from 193.0.14.0/24 lookup anycast
```

```
32766: from all lookup main
```

```
32767: from all lookup default
```

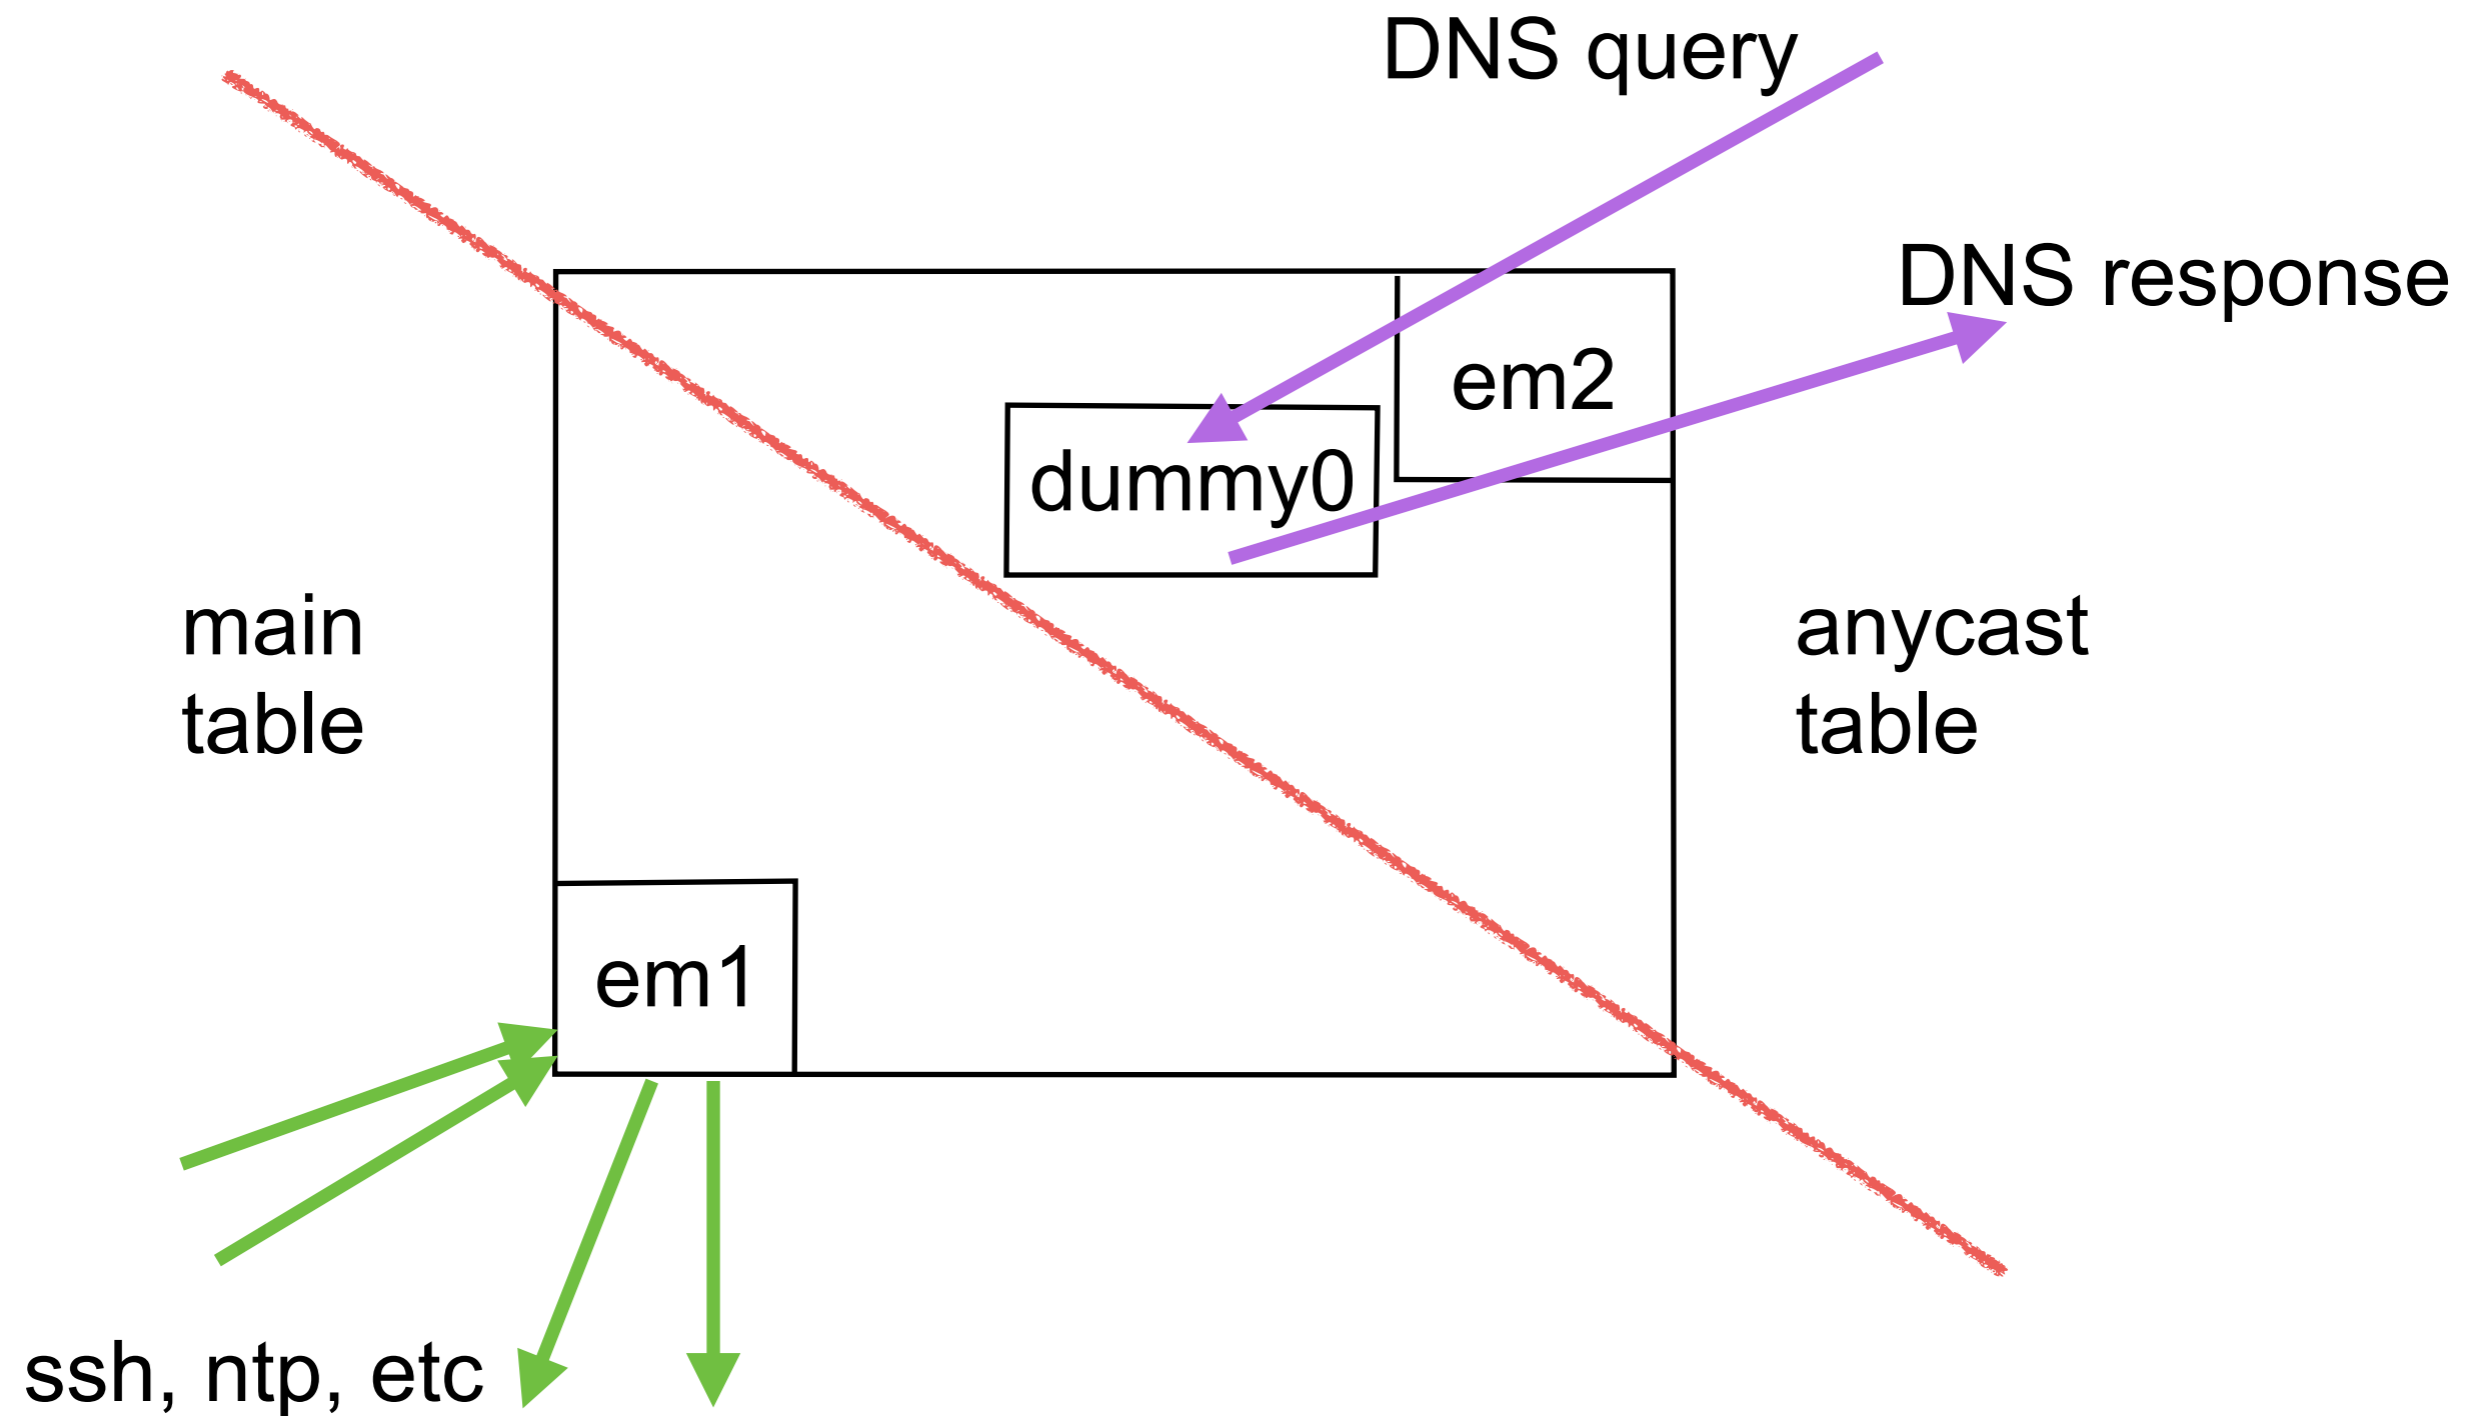
```
# ip route show | grep default
```

```
default via 212.36.221.253 dev em1
```

```
# ip route show table anycast | grep default
```

```
default via 194.126.19.37 dev em2
```


K-root configuration



Things to watch out for



- Debugging becomes trickier
 - you have to be aware of the multiple routing tables and rules
- If there is no match in the anycast table, Linux will look up the main table
 - a query will come into one interface, and its response will go out of another

The important of default



- Queries often originate from unannounced address space
- The anycast table needs a default route
 - without the default, responses will go out of the management interface
 - we avoid this by having a blackhole route in the anycast table
 - this can cause queries to time out



Questions



anandb@ripe.net
@aabdn