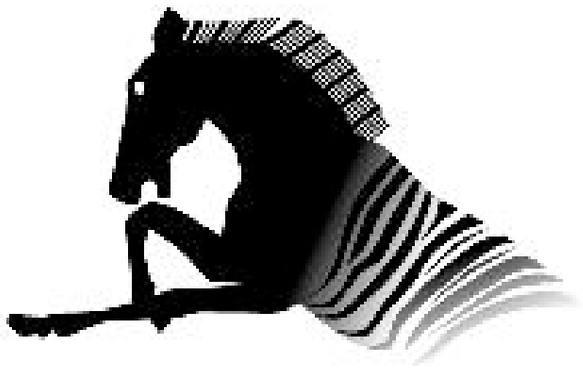

Back to the Future:
BSD at the Edge of the Enterprise

Russell Sutherland
CNS, University of Toronto
russell.sutherland@utoronto.ca

BSDCan 2006

UNIVERSITY of TORONTO

URL for tree huggers

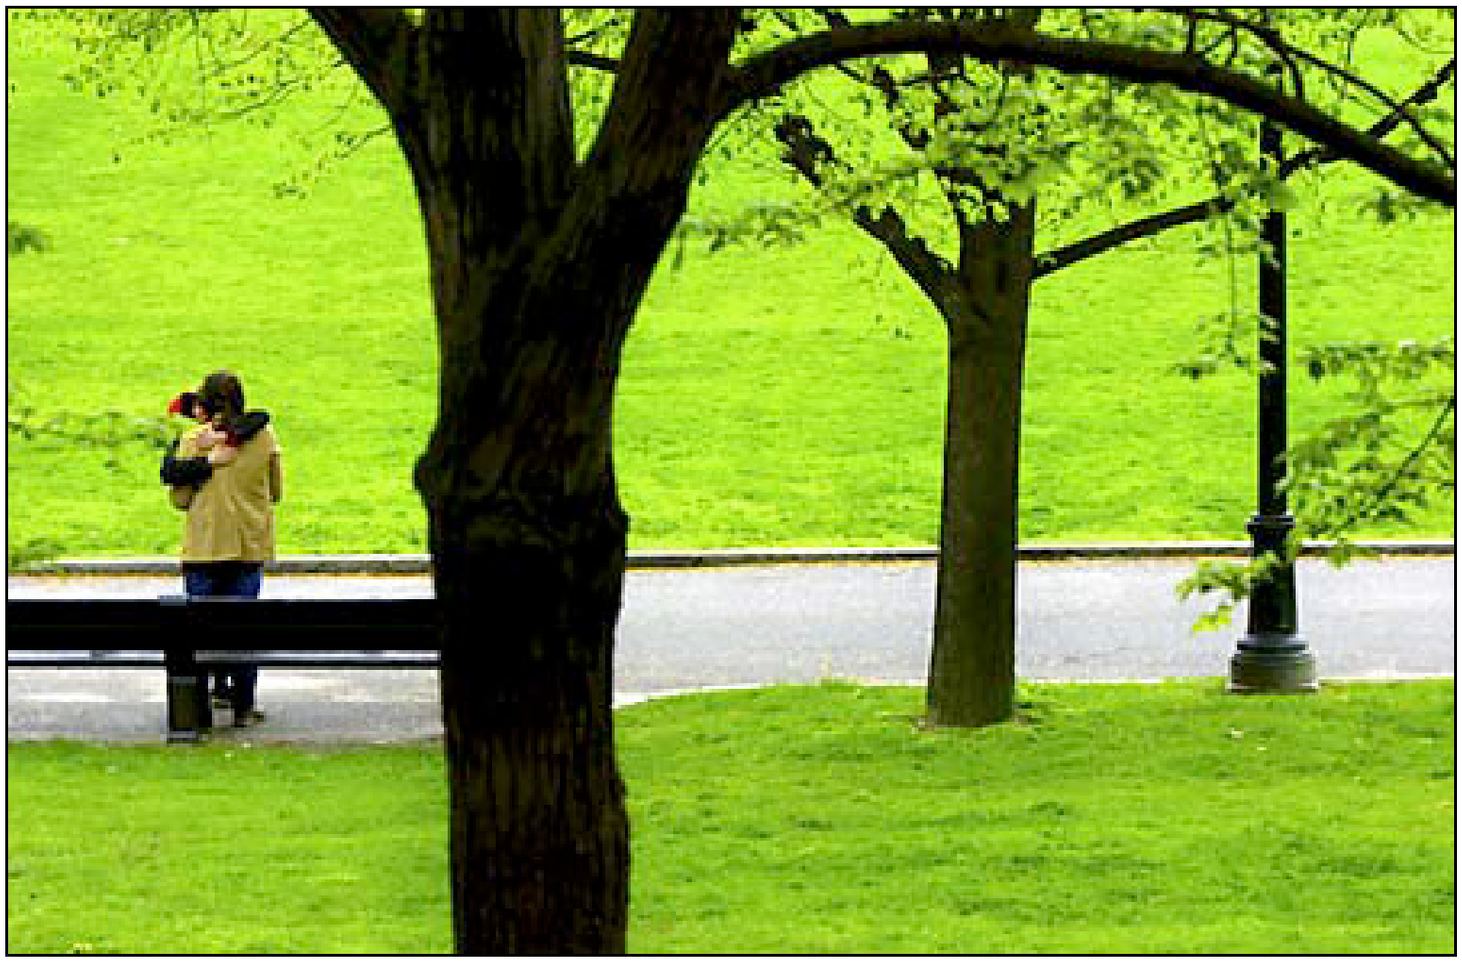

- http://madhaus.cns.utoronto.ca/~russ/bsdcan.pdf
- http://madhaus.cns.utoronto.ca/~russ/bsdcan.html

Routing Chronology I

- 1984
  - BSD ships with routed (RIPv1)
- 1986
  - Fuzz Ball PDP-11 NSFNet Routers
- 1988
  - Dedicated Routers: Proteon, ACC, Cisco

Routing Chronology II

- 1992
  - Gated Consortium Formed
- 1996
  - GNU Zebra
- 2003
  - Quagga forked from Zebra

Routing Chronology III

- 2003
  - XORP Project
- 2004
  - OpenBSD 3.5 ships with bgpd
- 2005
  - OpenBGPD Project
- 2005
  - OpenBSD 3.7 ships with ospfd

Quagga Architecture

- Modular Design
- One process per protocol
  - bgpd, ospfd, ripd, etc...
- One main controlling process
- Extensible

Quagga Architecture

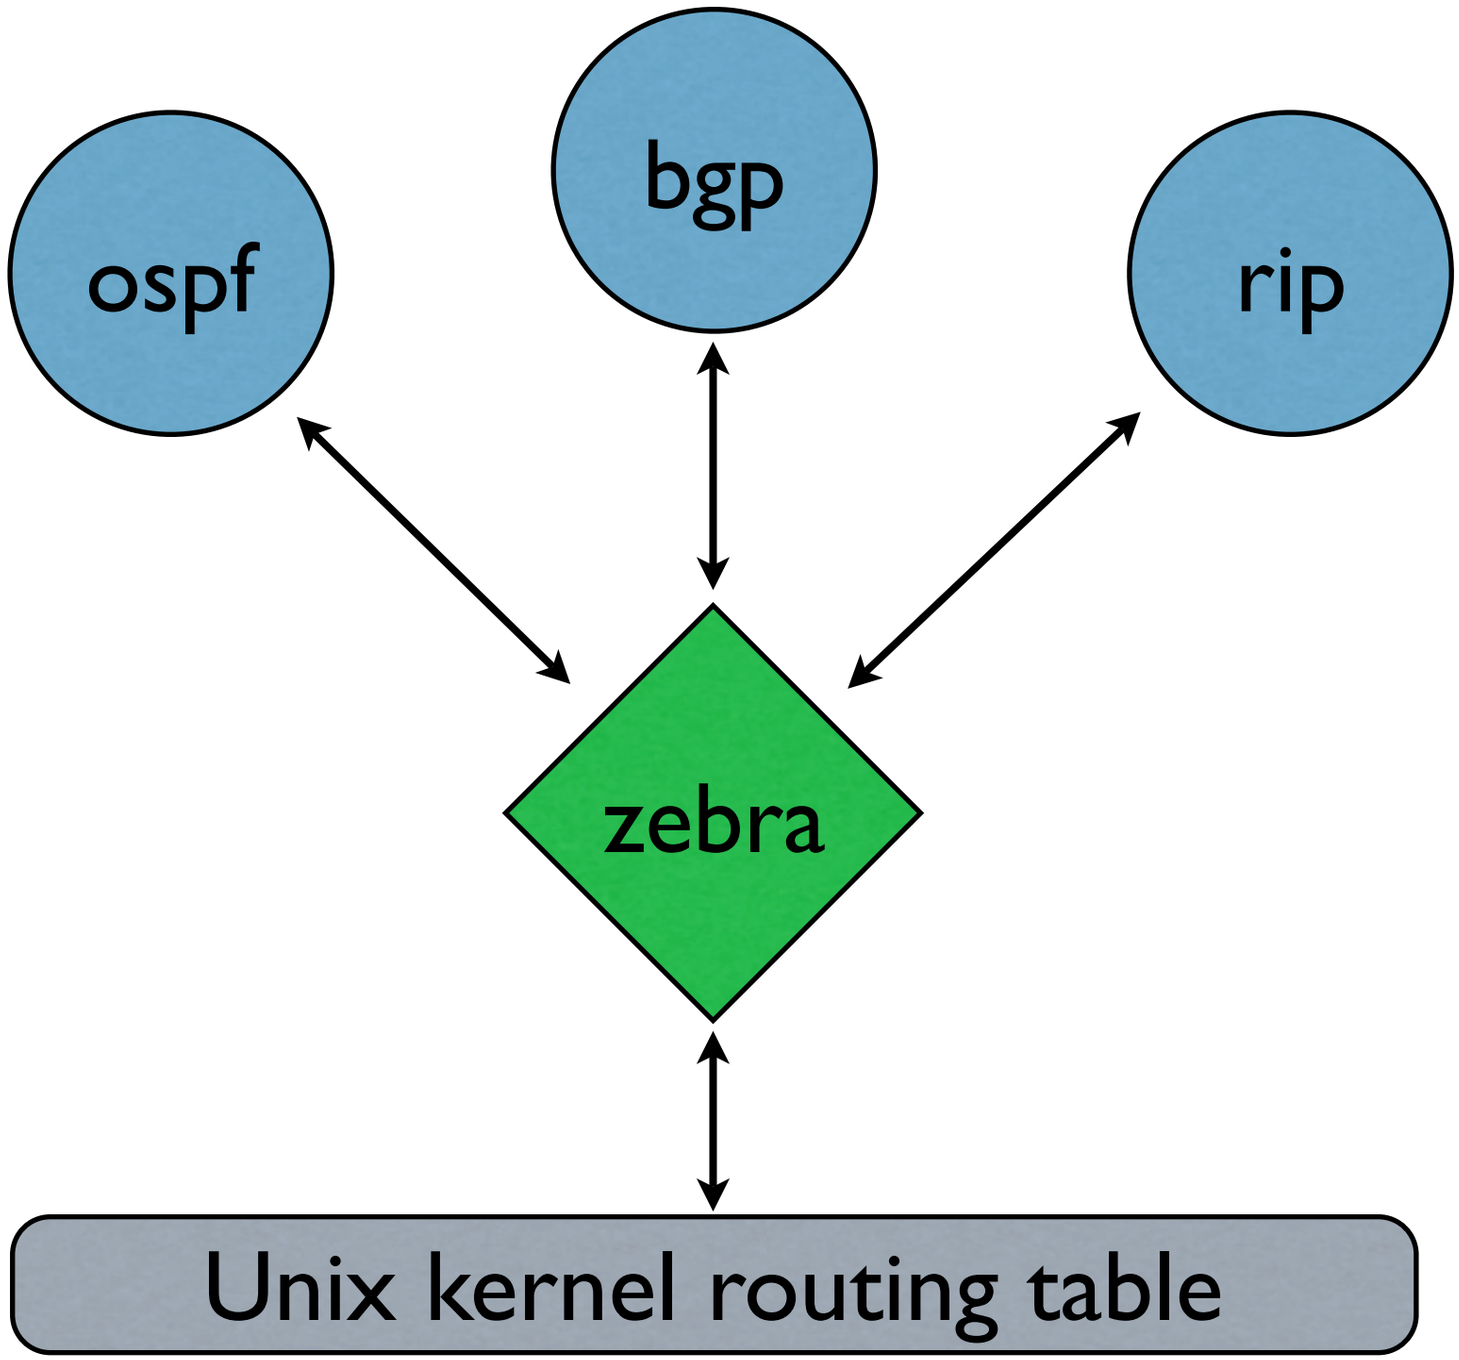

Quagga Routing Protocols

- RIPv1, RIPv2, RIPng
- OSPFv2, OSPFv3, OSPFng
- BGPv4, BGP route server
- Filtering, Route Map
- IPv6
- SNMP data via SMUX

Quagga Supported Platforms

- FreeBSD
  - versions 4.x, 5.x, 6.x
- OpenBSD
  - versions 3.x
- NetBSD
  - version 1.4
- Linux & Solaris

Hardware Specs.

- CPU: Intel, 2.x GHz
- Memory: 512Mb
- Disks: 18 Gb
  - SCSI or IDE
  - RAID 1 (optional)
- Ethernet: 3 x 10/100/1000 Mbps
- e.g. Dell PowerEdge 2650 or IBM x335

Software Configuration

- OS: FreeBSD 5.x, 6.0
- ipfw firewalling and policy based routing
- dummynet traffic shaping
- ssh for secure access
- from ports tree
  - Routing Quagga v 0.96
  - SMNP net-snmp 5.1.x
  - Logging/Mail syslog/qmail 1.03

Scottish Economics

Vendor	Model	List Price [Cdn k]
Cisco	Catalyst 3550	13-30
Foundry	BigIron 4000	12-20
Extreme	Alpine 3800	31-38
Intel 2U	Dell 2650	2-3

Topology Legend

Symbol	Device
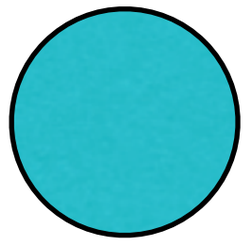	Edge Router
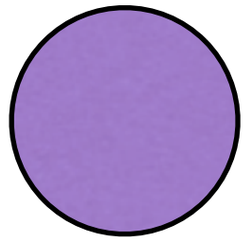	Interior Router
	External ISP Router
	Traffic Shaper

Network Topology I

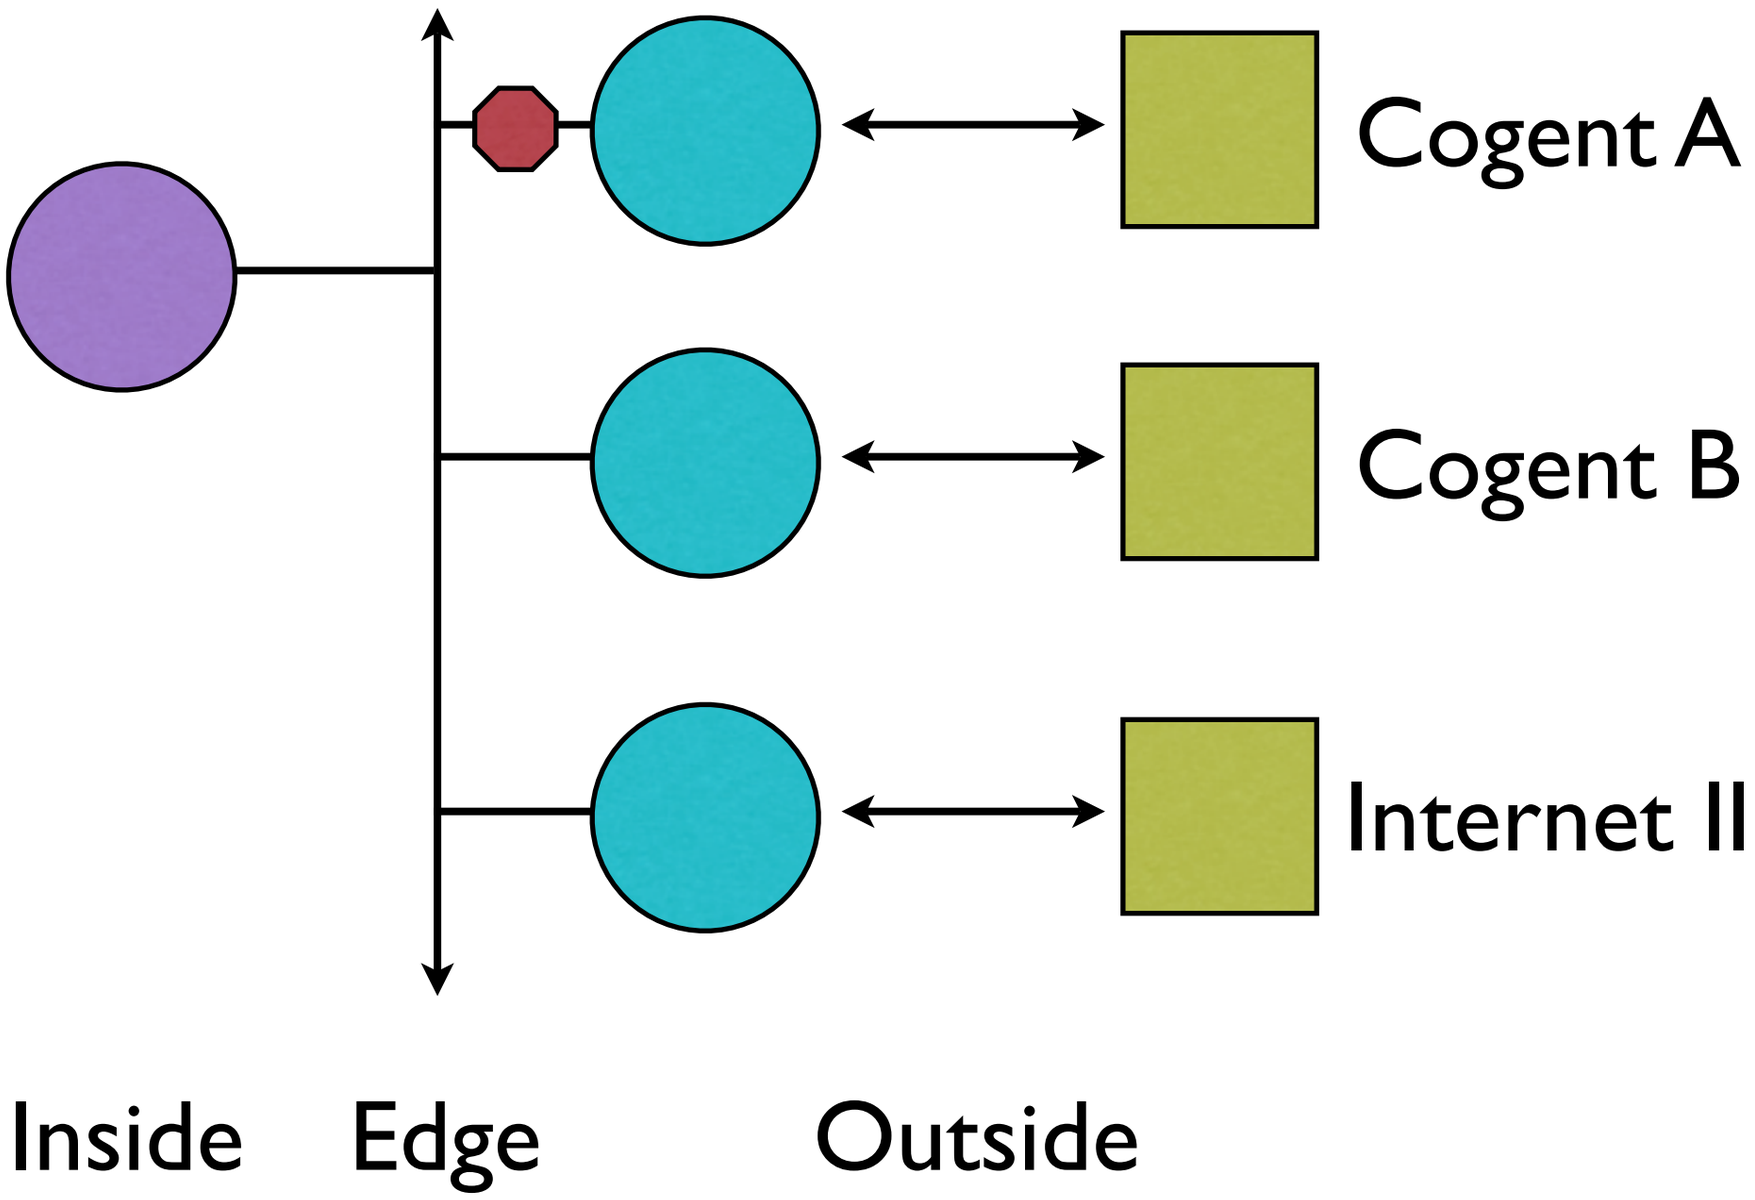

Network Pipe Size 1

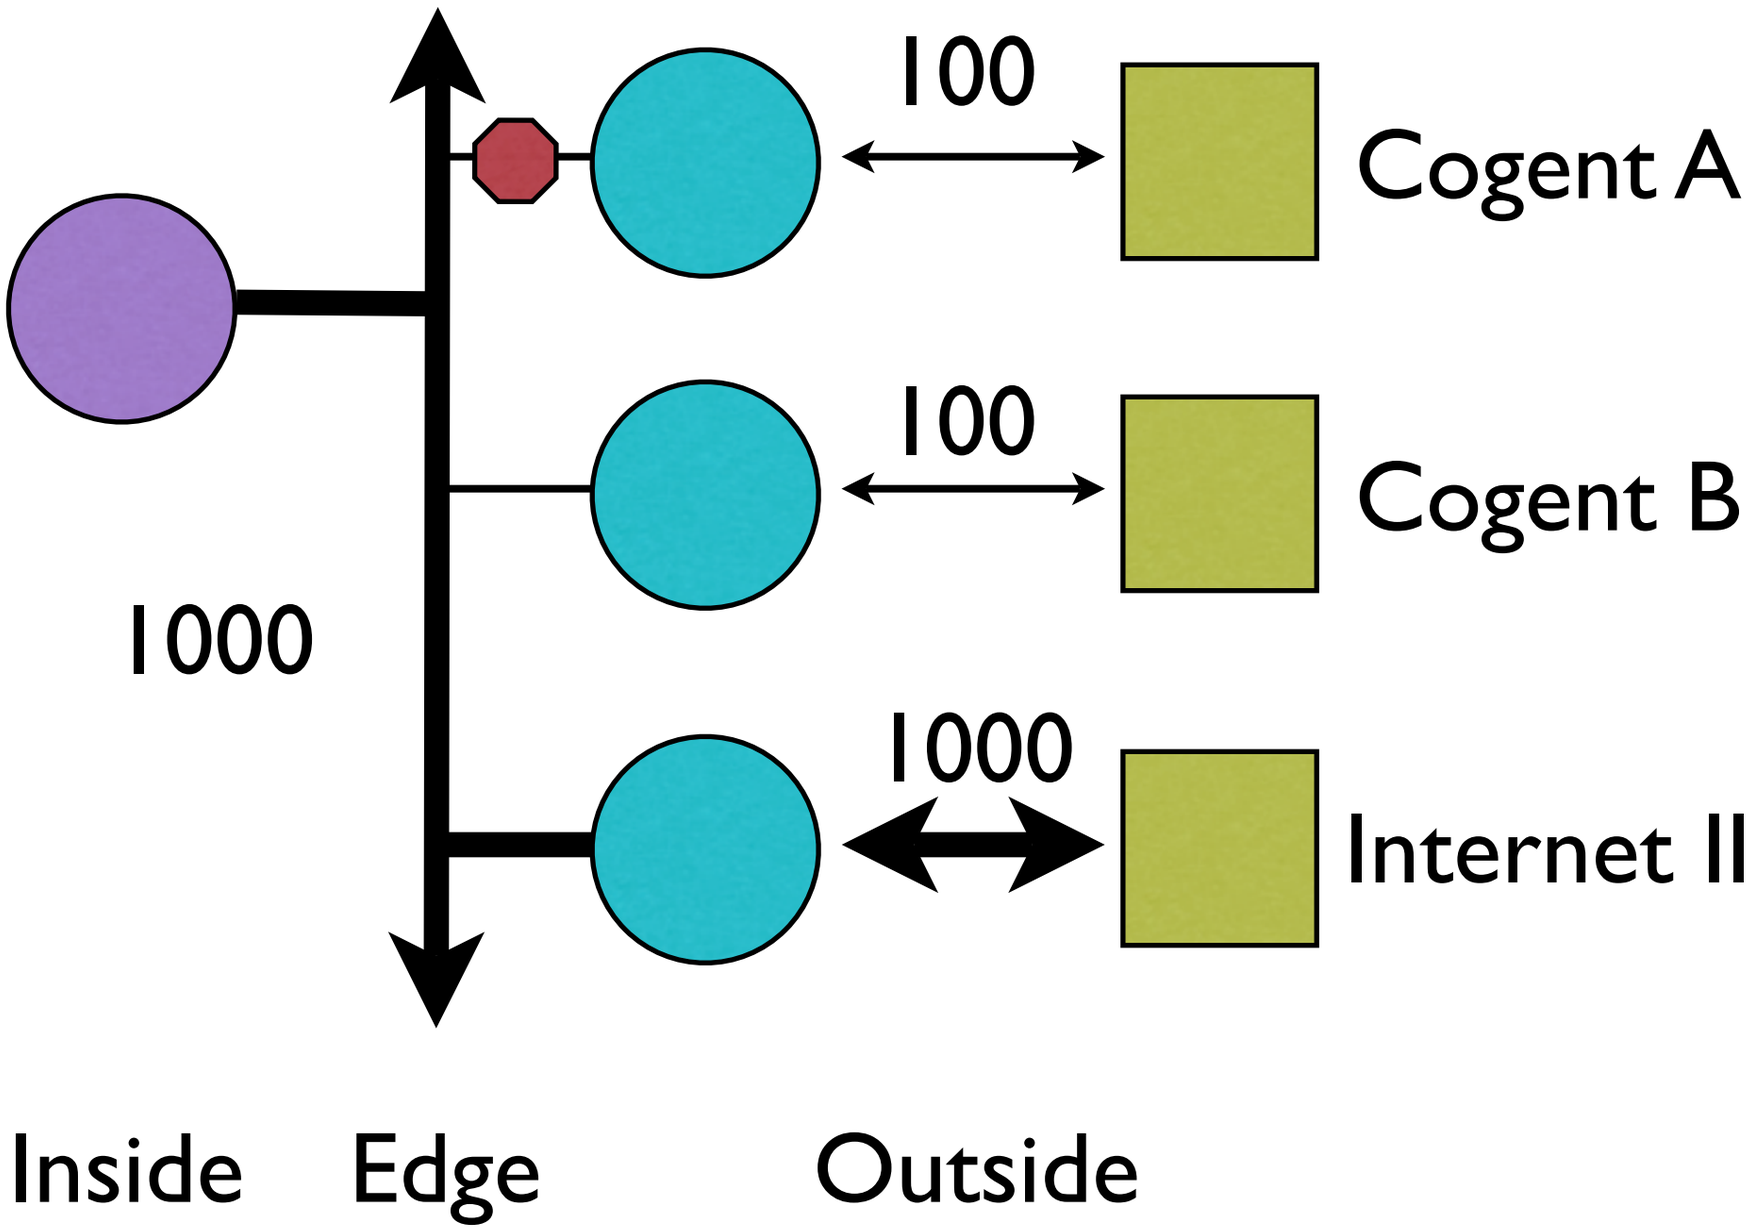

Network Topology II

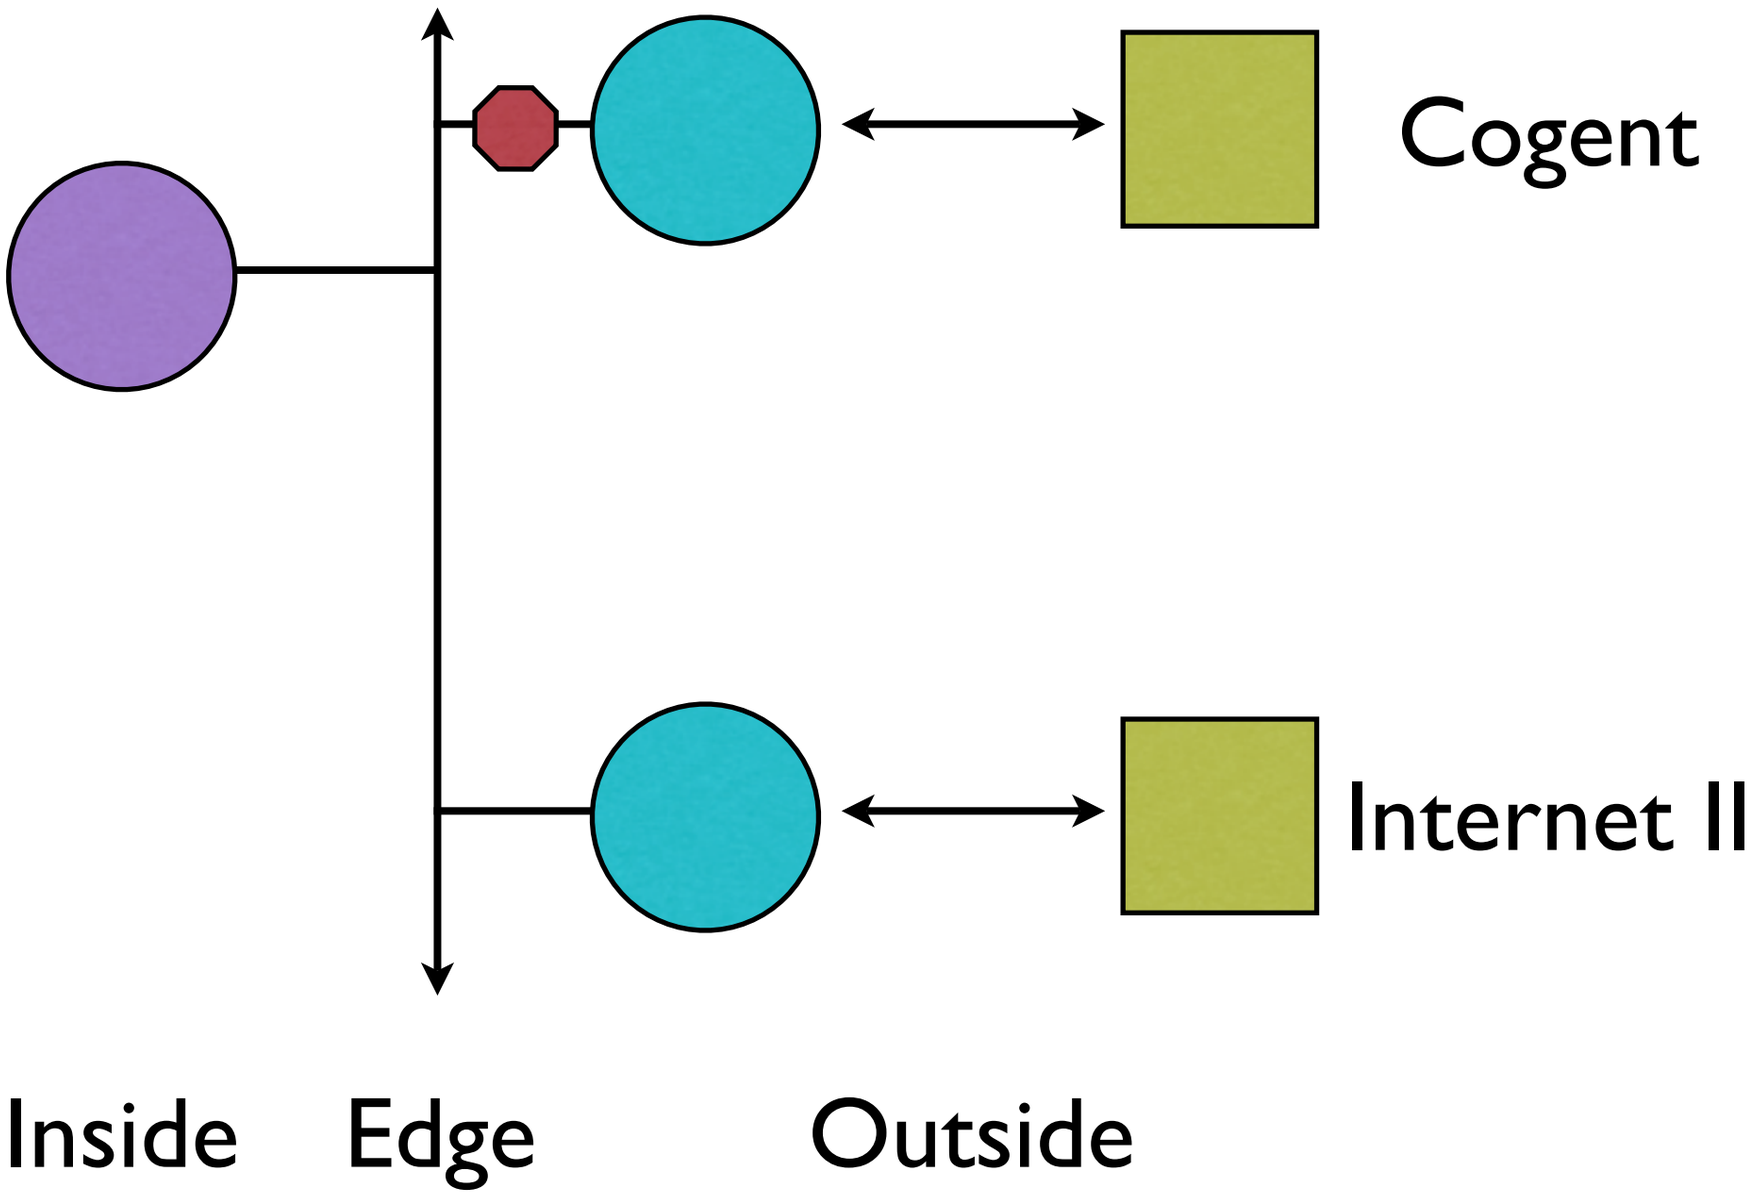

Network Pipe Size II

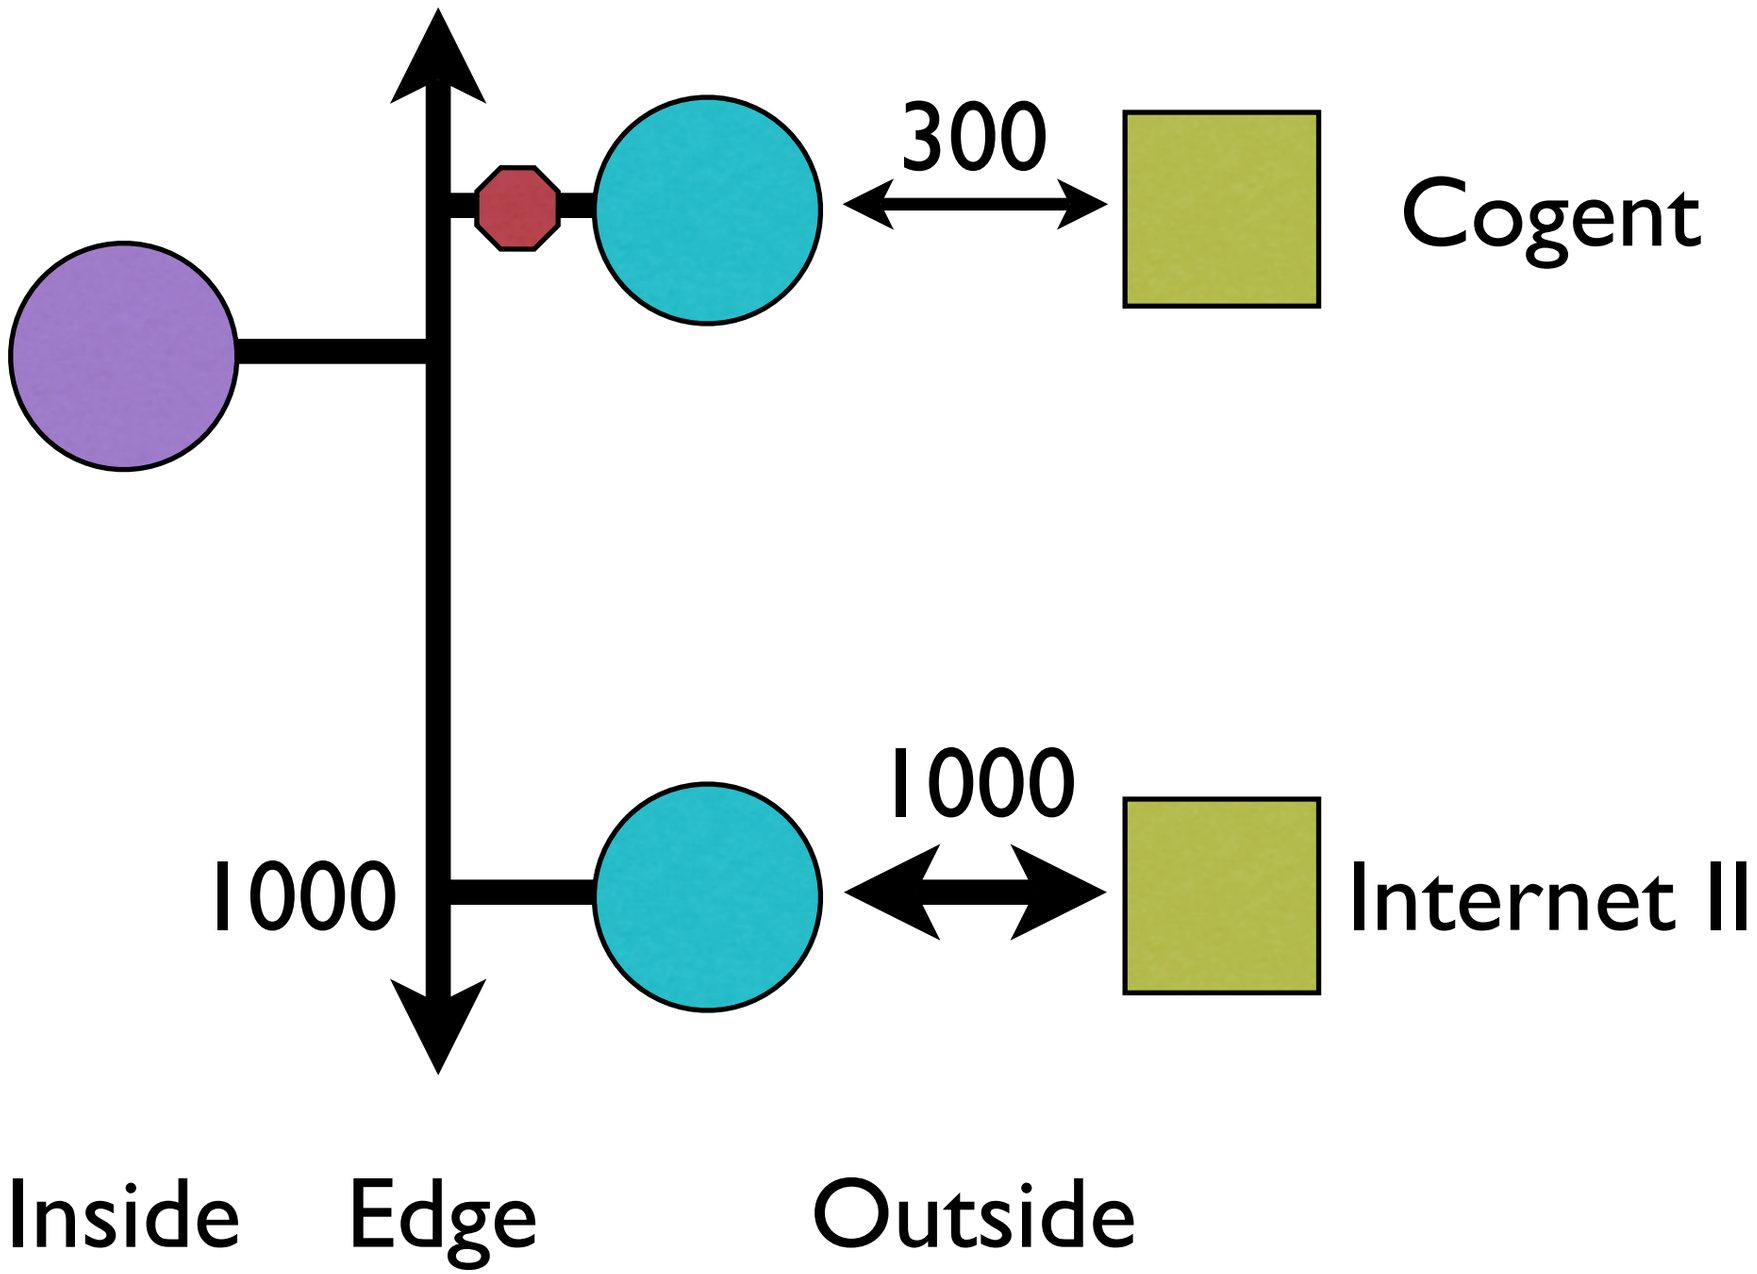

Network Filtering Policies

- KISS
  - UofT is a collection of fiefdoms
- Drop packets in accord with
  - generally accepted ISP norms
  - broad community accord
- Allow everything else
- Do not keep state

Drop Filtering Rules

- Drop the following packets:
  - spoofed [non UofT] source IP addresses
  - non-routable destination addresses
    - 0/8, 10/8,127/8
    - 172.16/12, 192.168/16, etc.
  - Nasty M$ tcp/udp ports
    - 67-69, 135,137,139
    - 161-162, 445, 593, 707, 1433,1434,3127,4444

Network Routing Policies I

- Three distinct [/16] internal networks
  [based on source IP address]
  - ResNet
  - UofT A
  - UofT B

ResNet Routing Policies

- All traffic to and from Cogent A
- Traffic Shaper to mangle P2P
- No Internet II transit !

ResNet Routing Policies

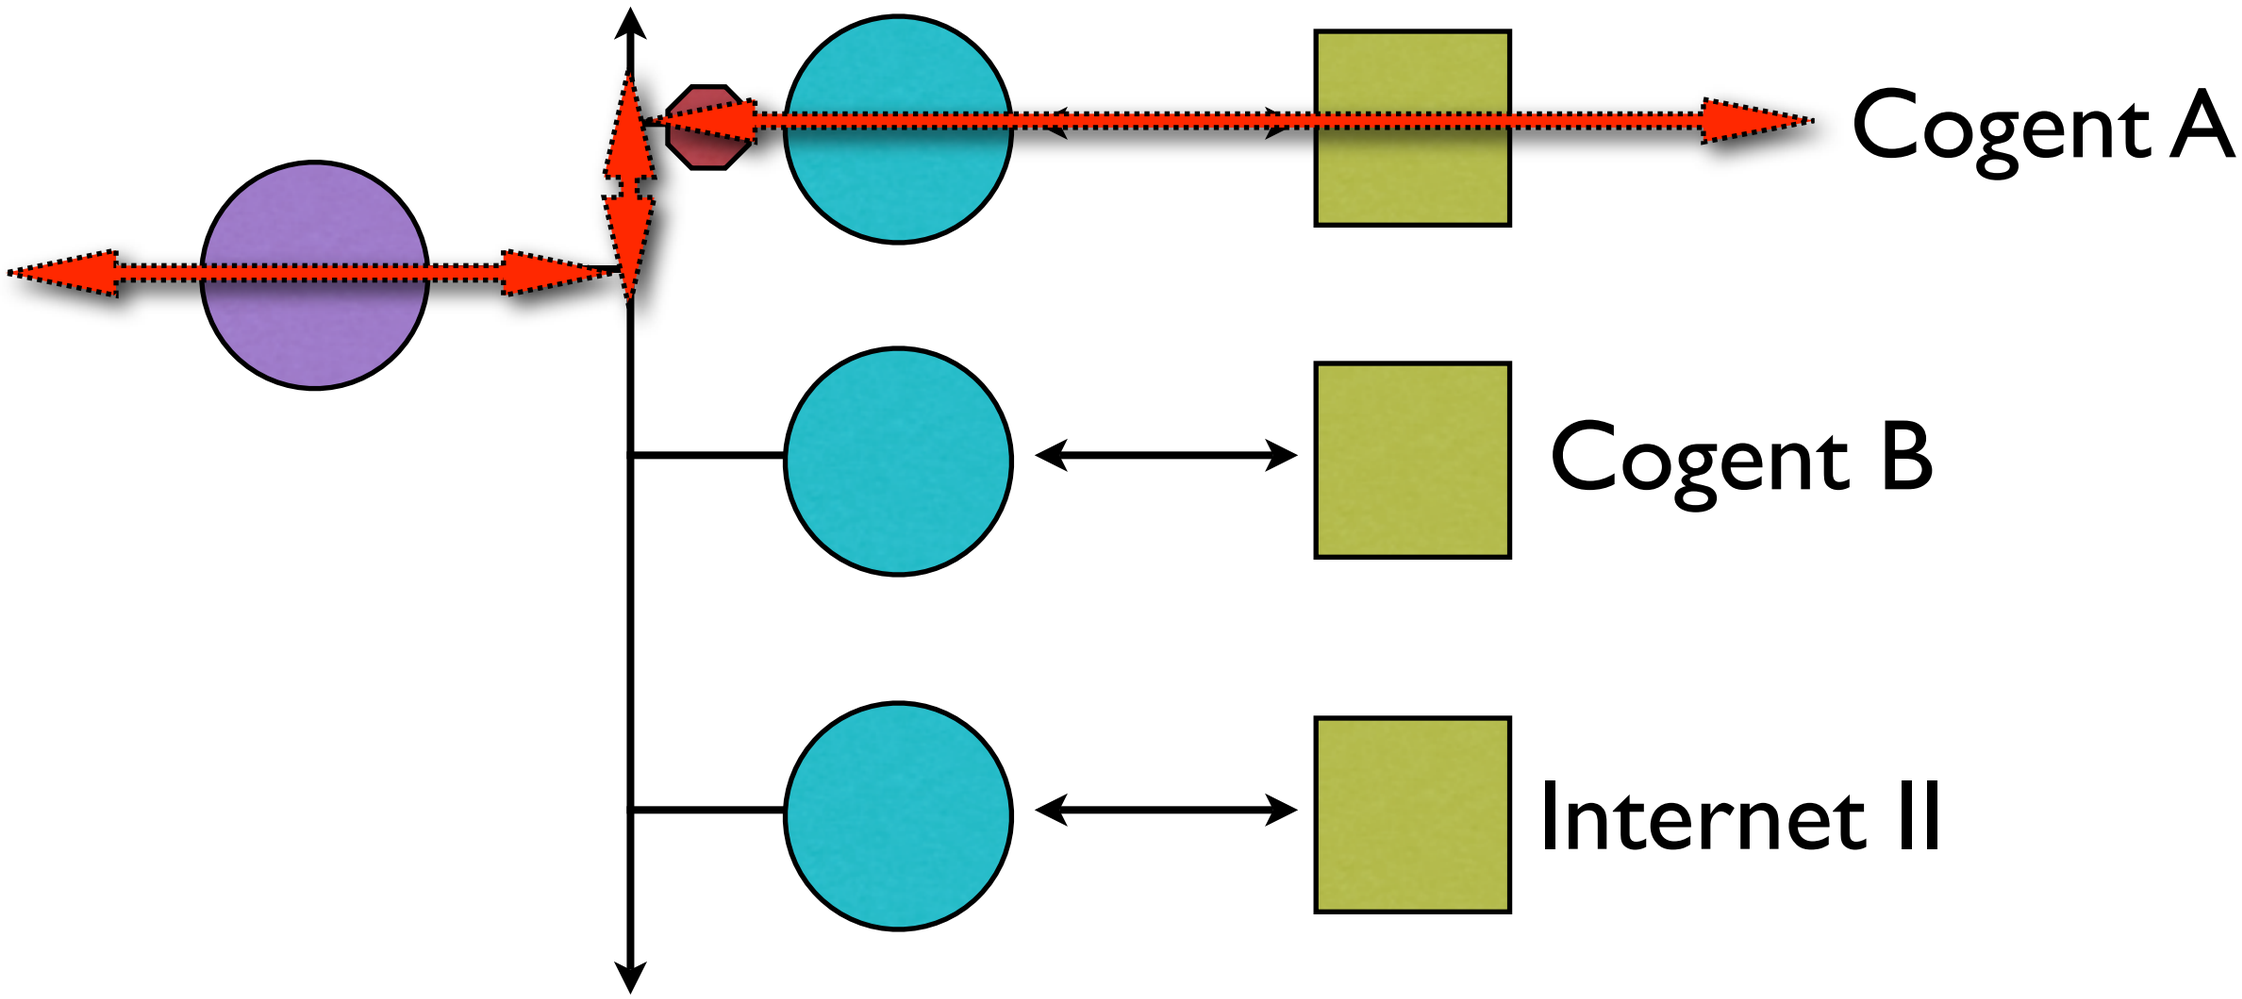

UofT A Routing Policies

- Internet II traffic via the dedicated GTAnet/ Orion link.
- All other Internet traffic via Cogent A [shared with ResNet traffic]

UofT A Routing Policy

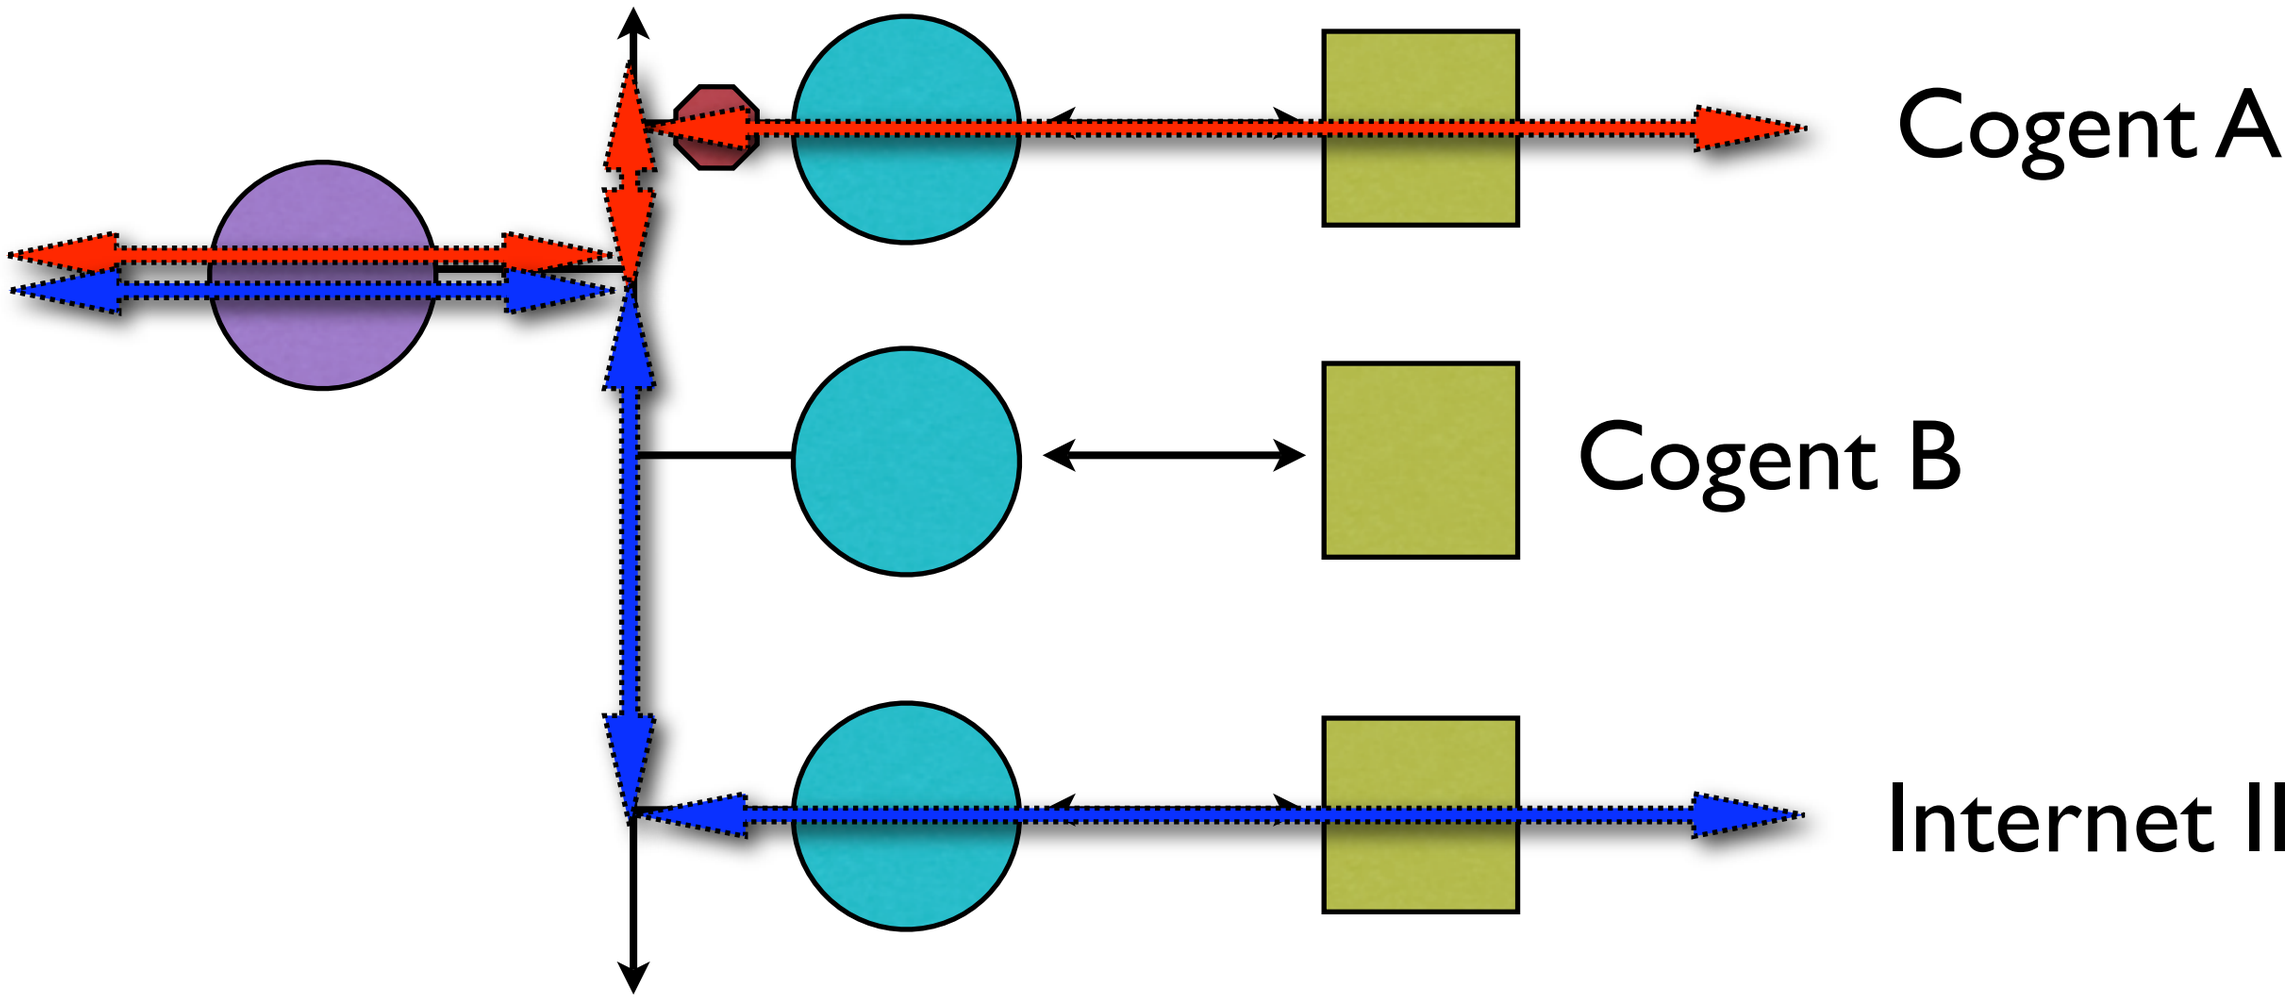

UofT B Routing Policies

- Internet II traffic via the dedicated GTAnet/ Orion link.
- All other Internet traffic via Cogent B [preferred traffic]

UofT B Routing Policies

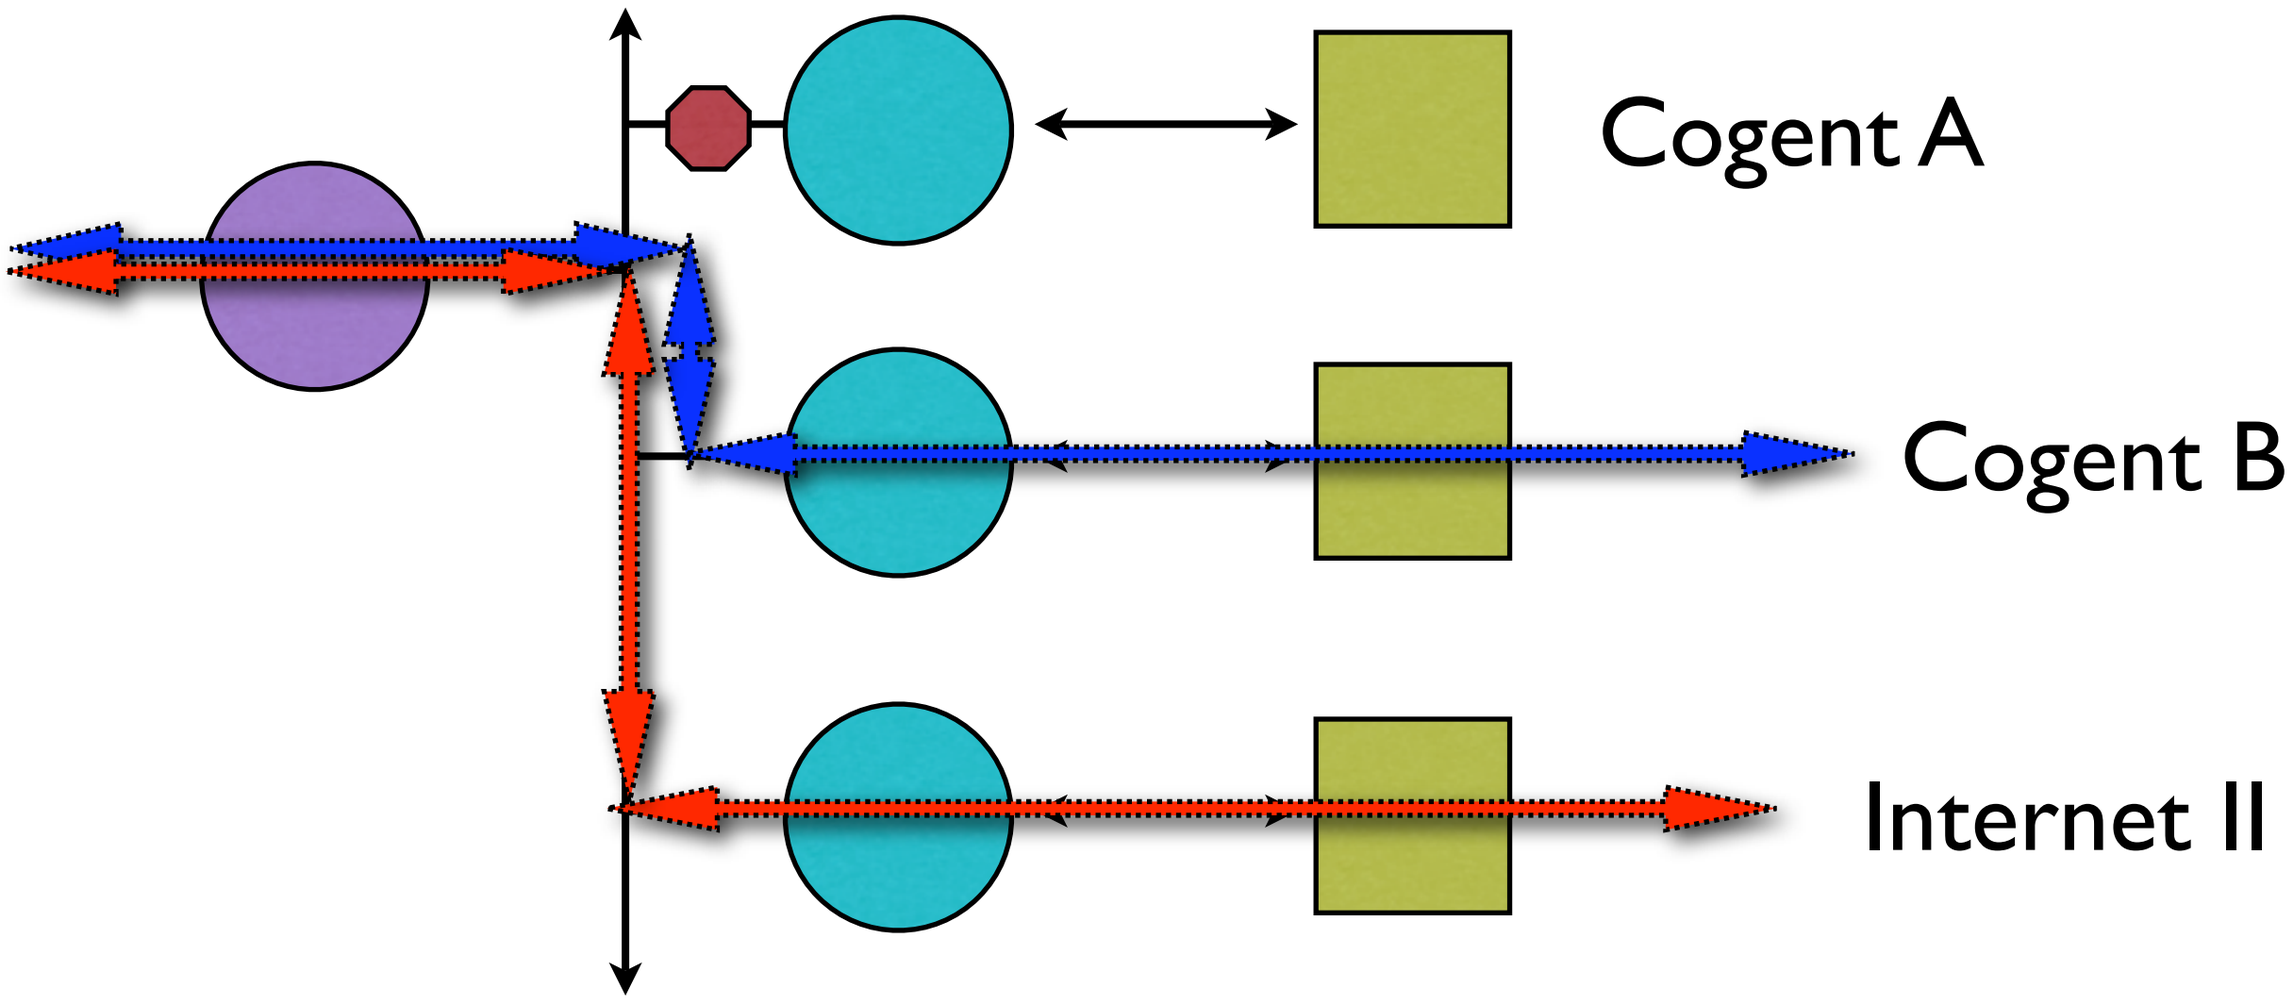

Routing Protocols

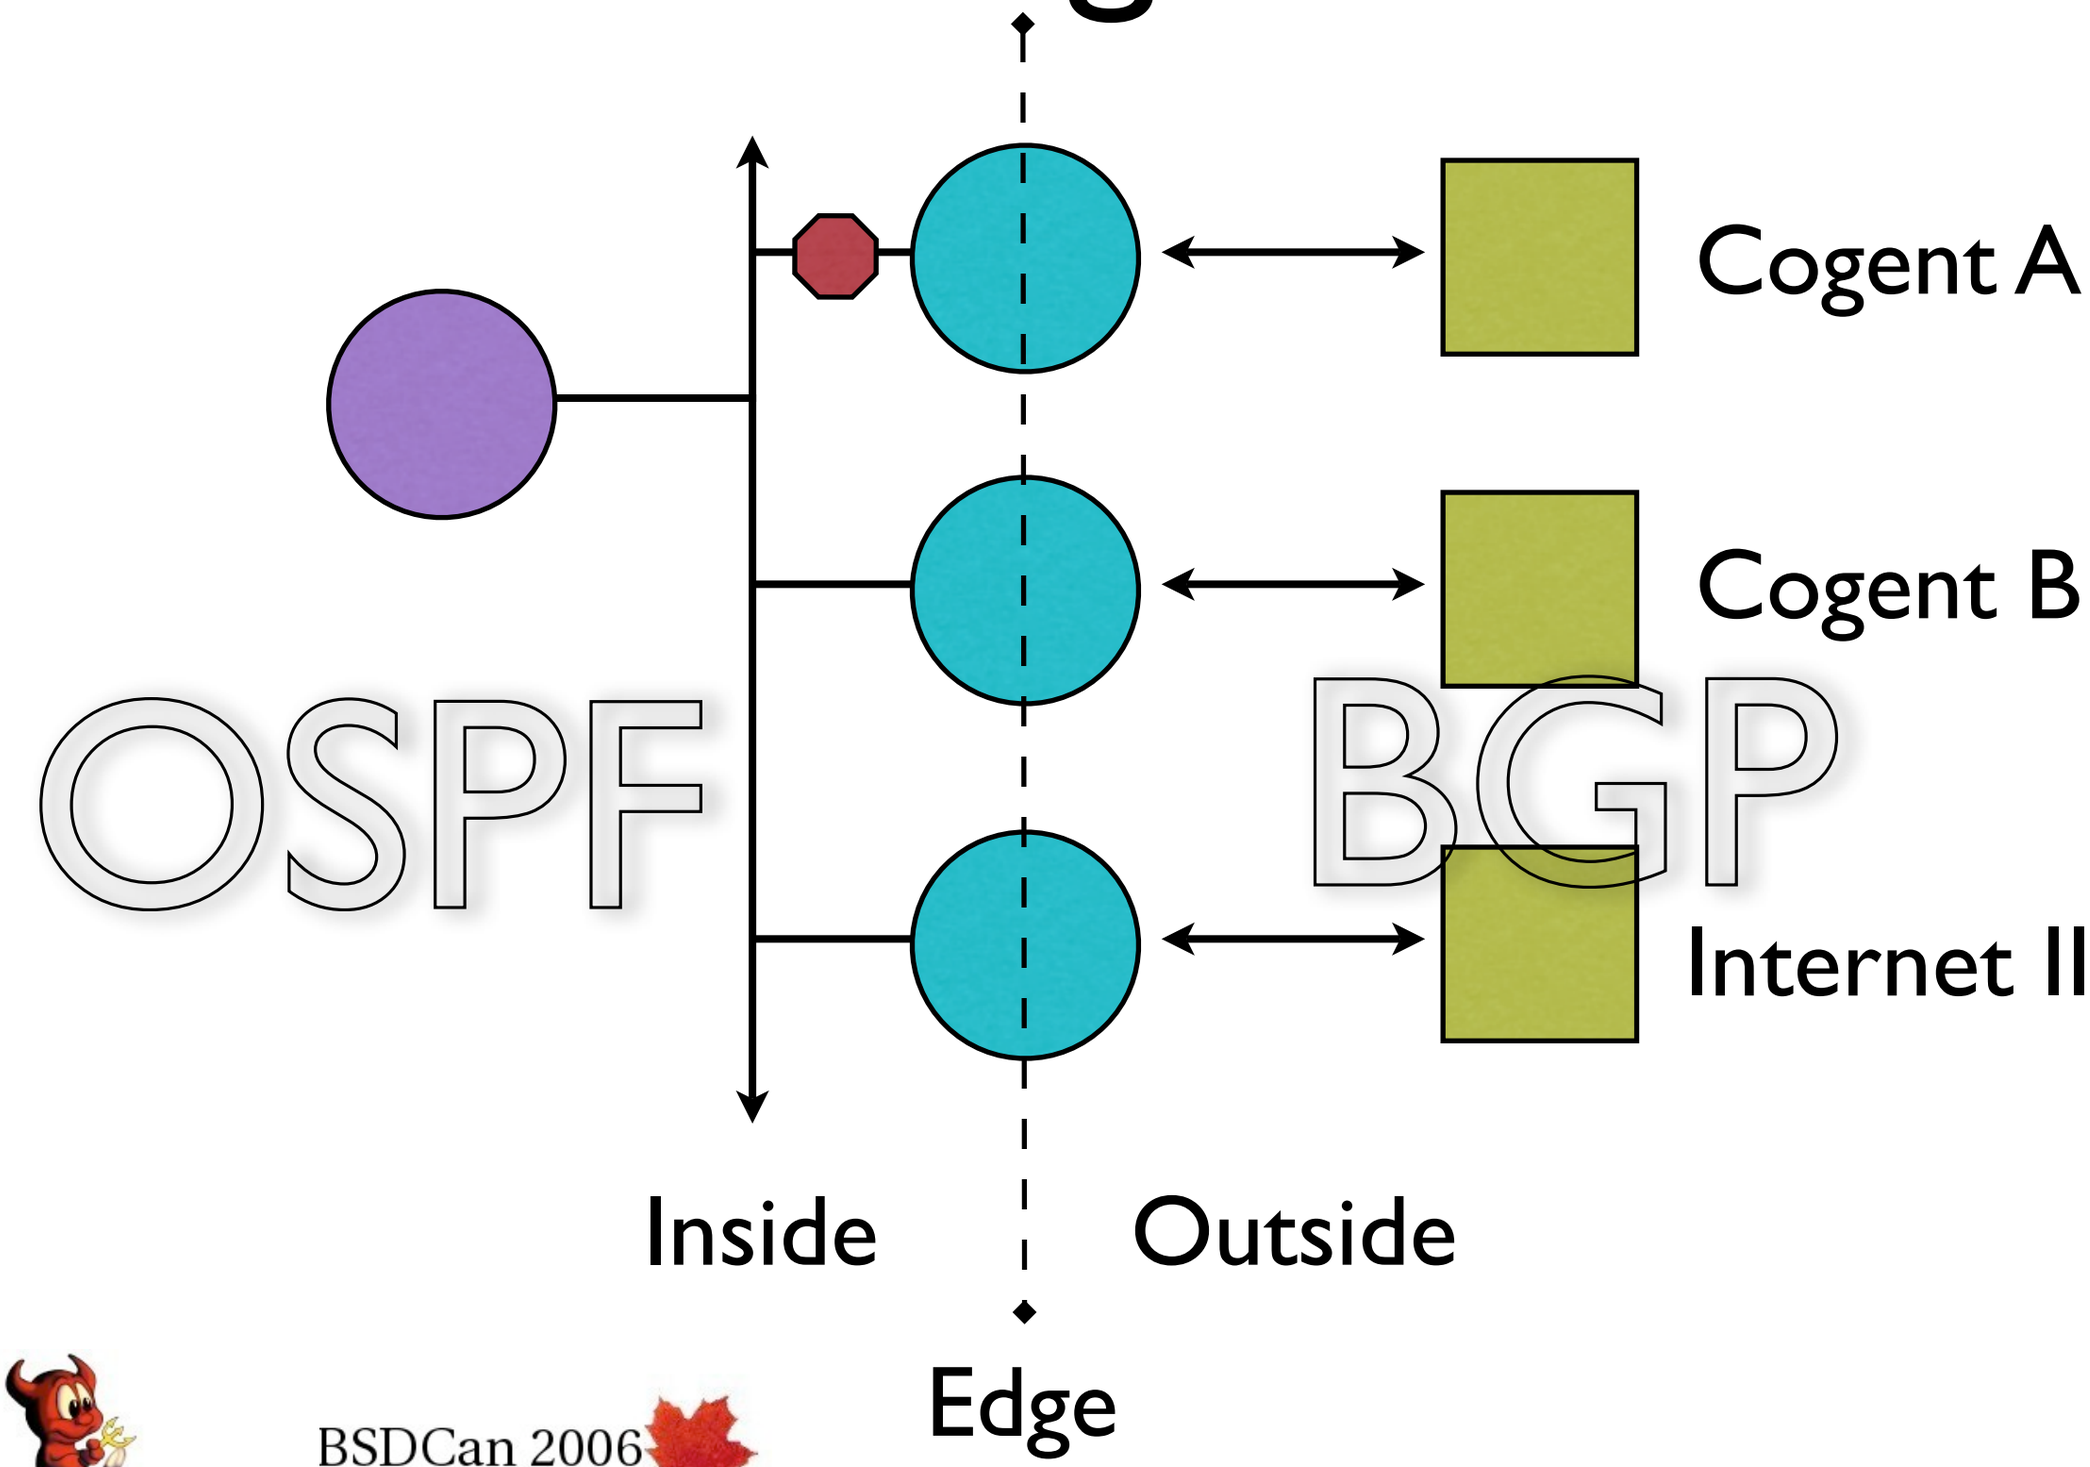

Scottish Router Names

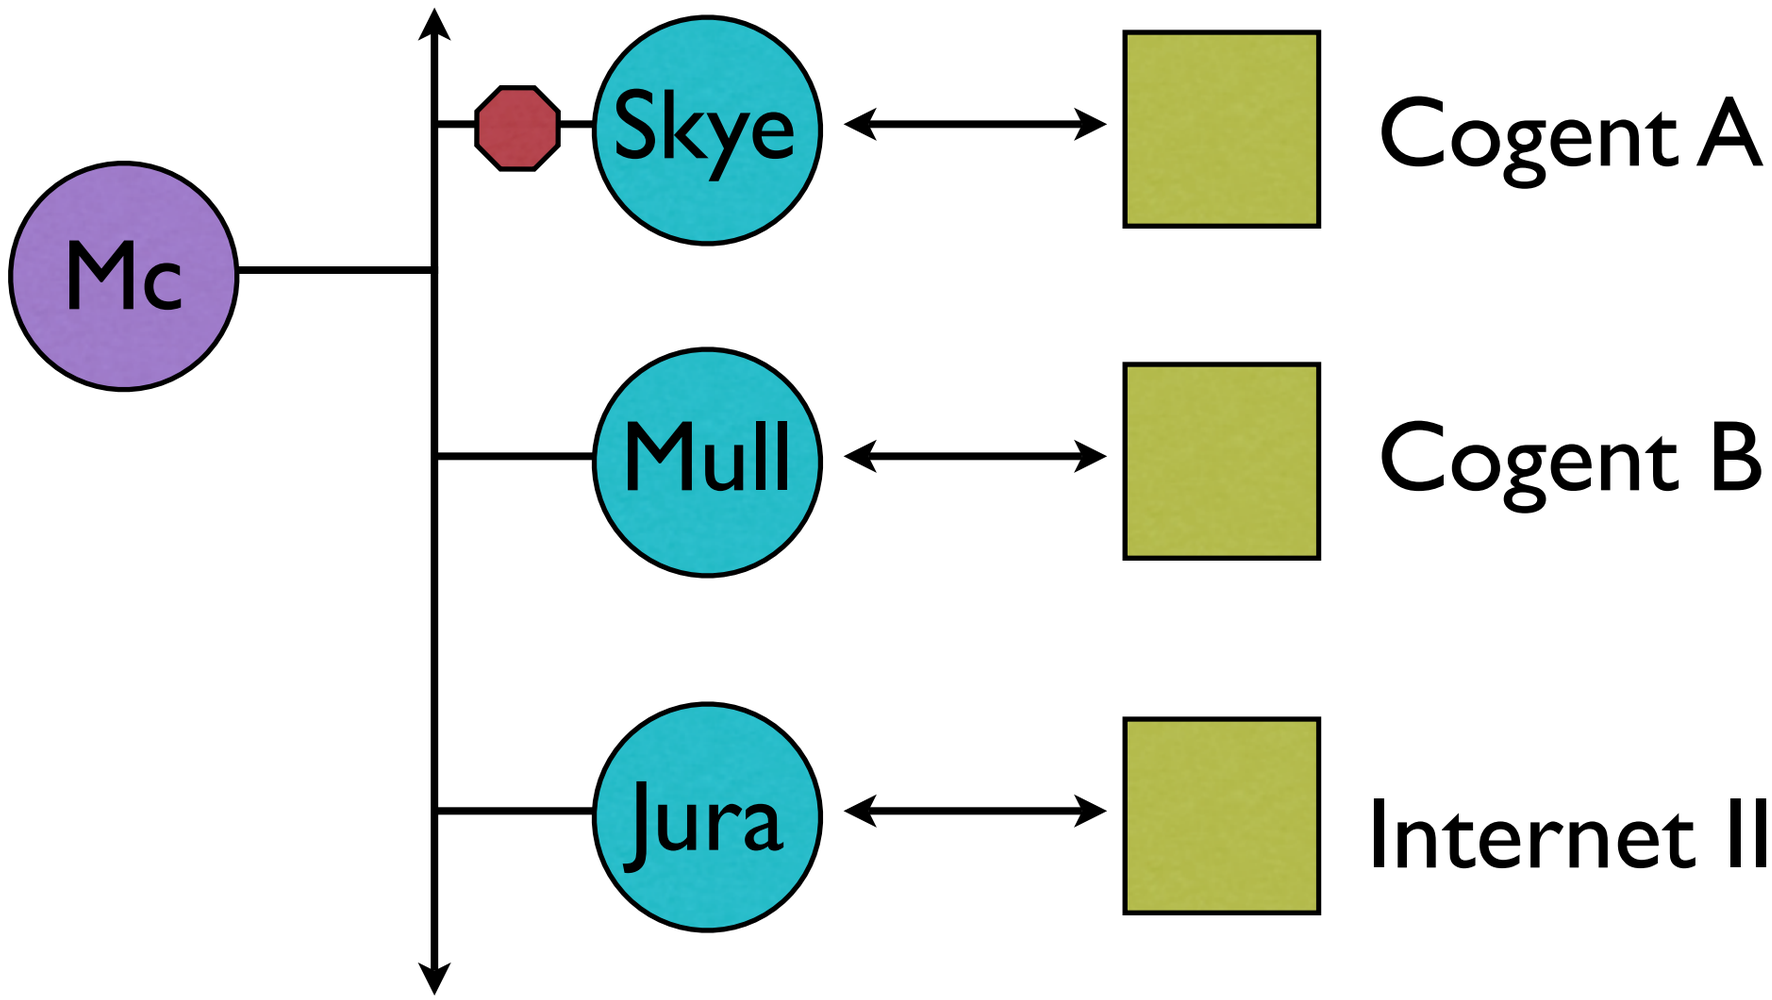

Router Configuration

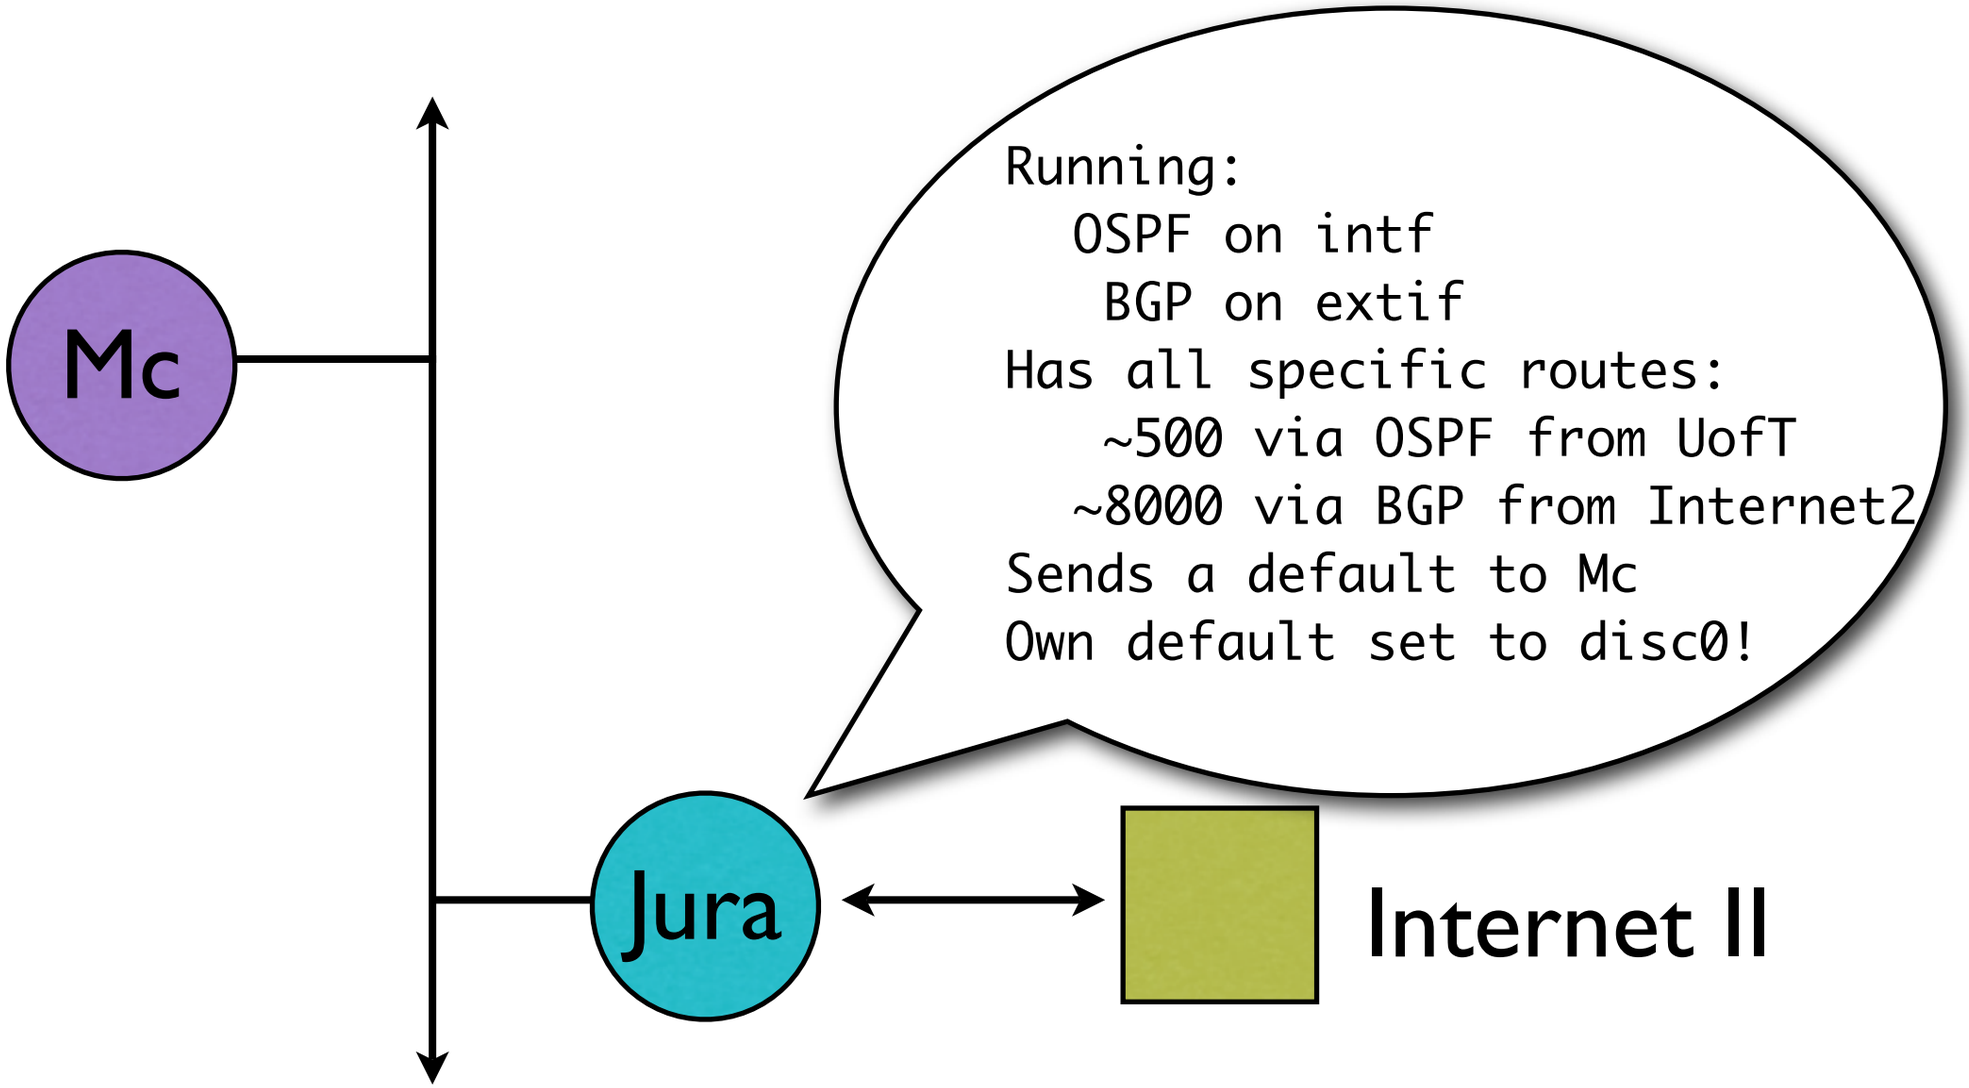

Router Configuration

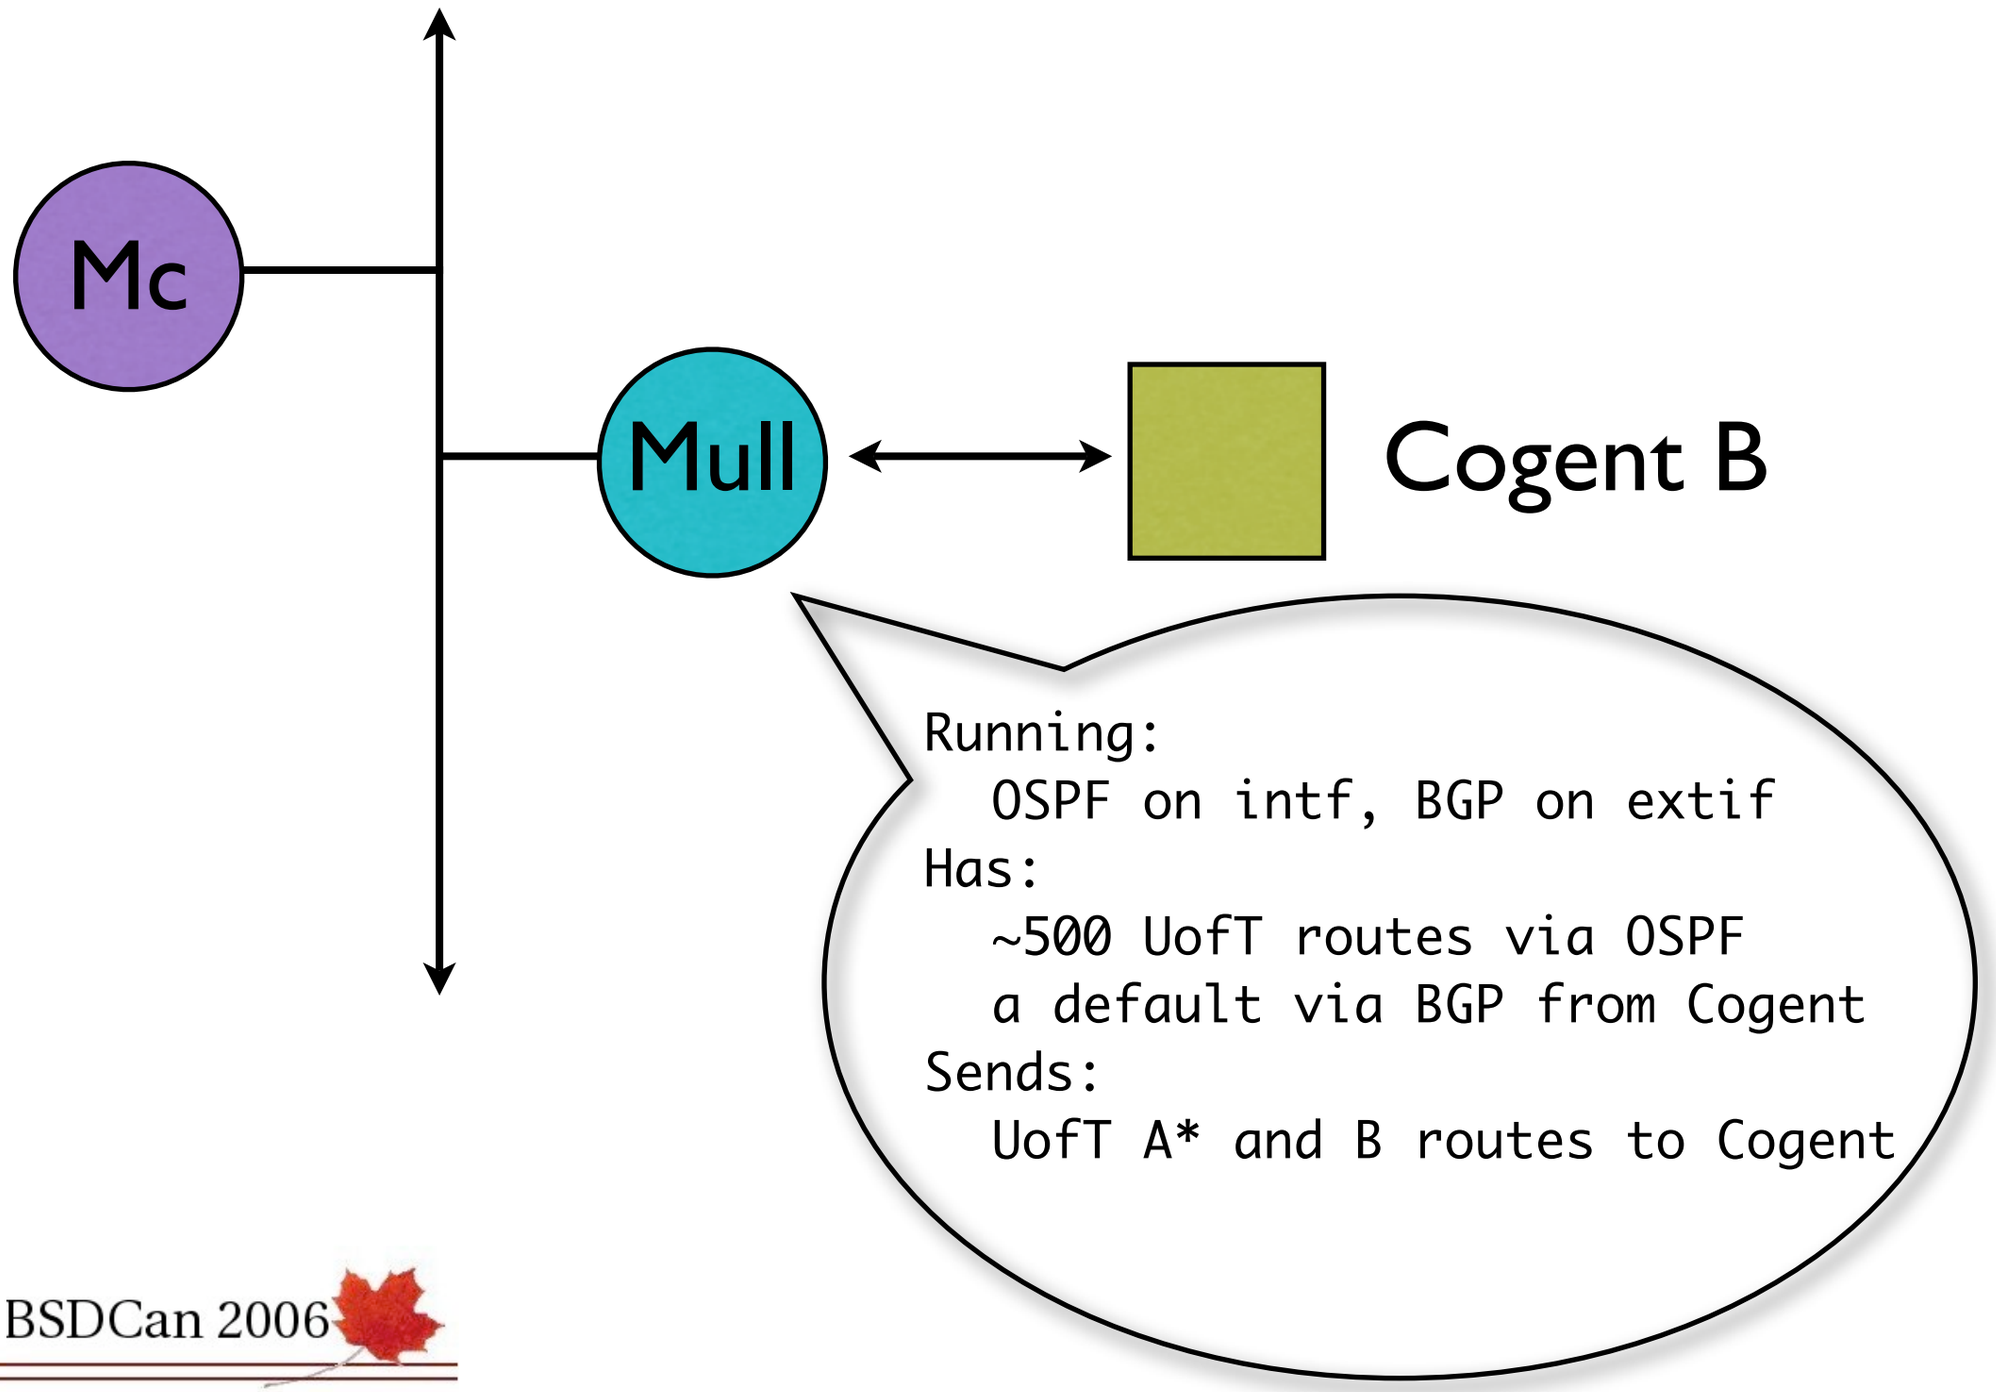

Router Configuration

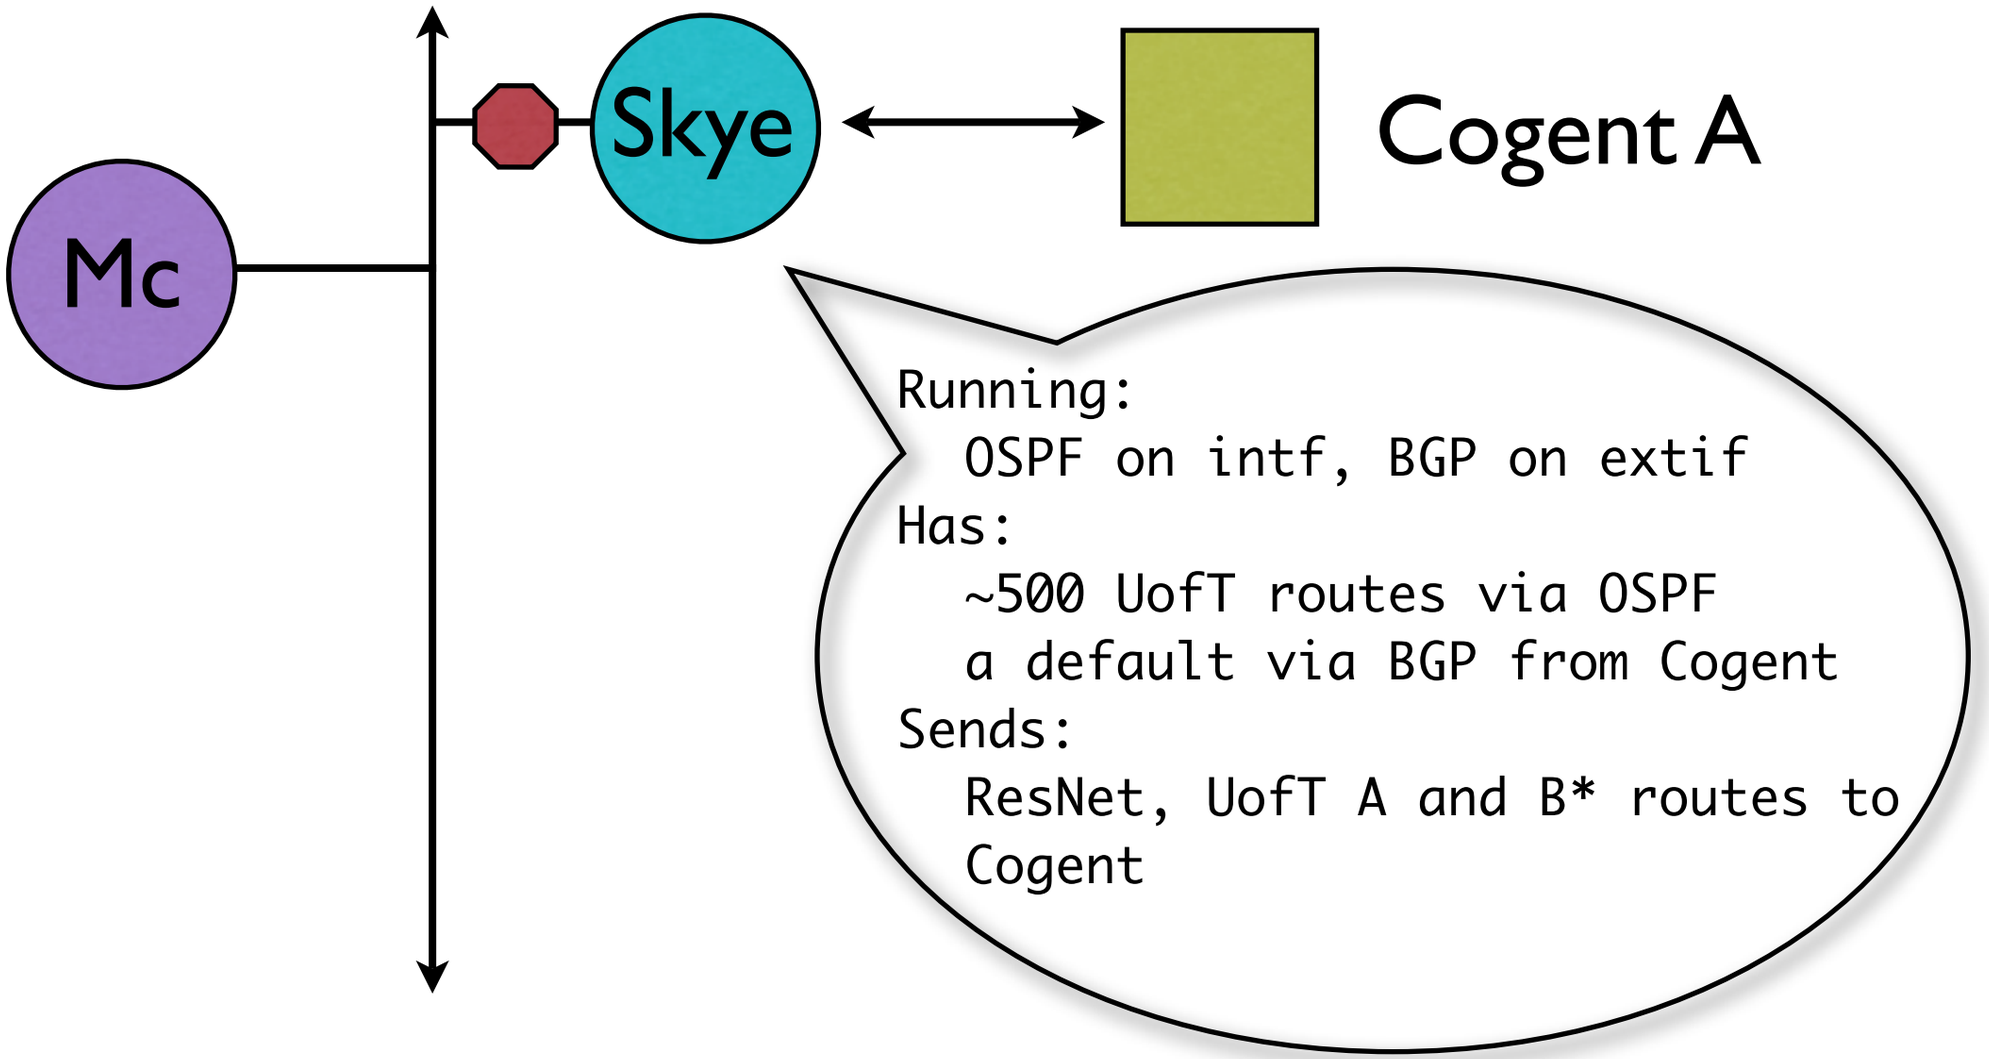

Router Configuration

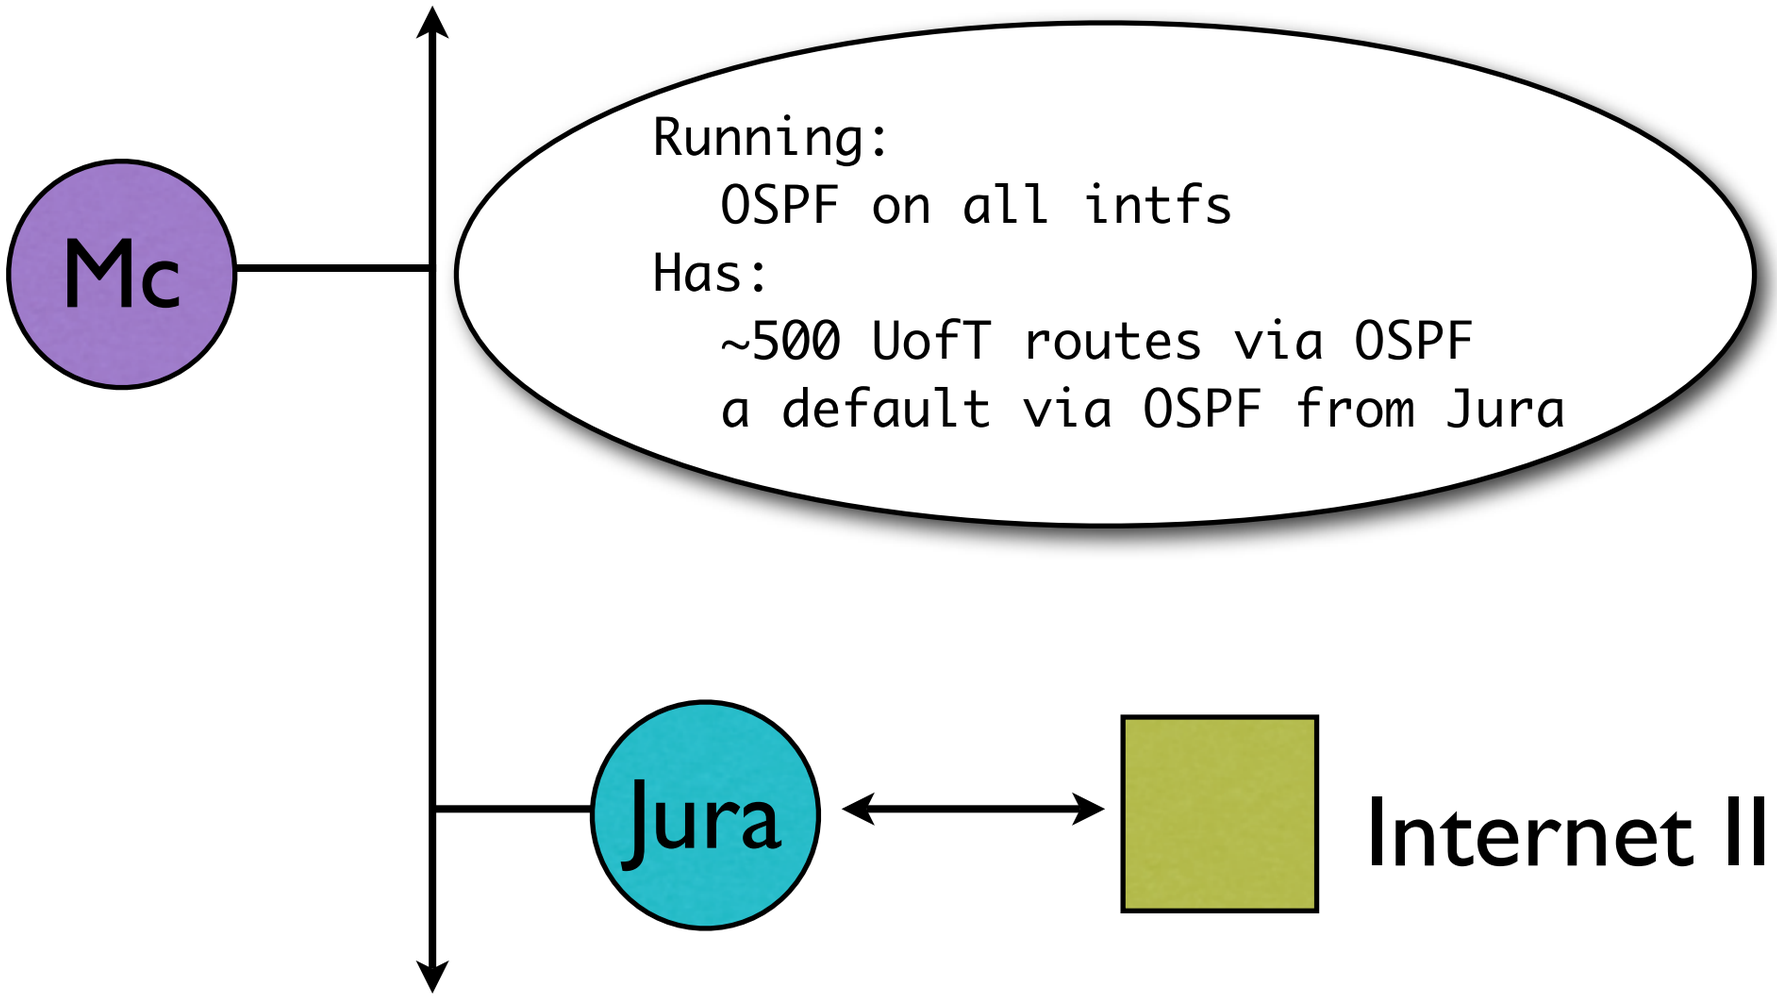

Quagga Configuration

- Very Cisco IOS like in its syntax. E.g.

```
# conf t
(config)# interface em0
(config-if)# description connection to ISP
(config-if)# ip address 11.12.13.14/28
(config-if)# exit
(config)# exit
# conf t
(config)# router bgp 777
(config-router)# bgp router-id 12.34.56.78
(config-router)# network 123.1.2.0/24
(config-router)# redistribute static
(config-router)# neighbor 210.1.2.3 remote as 999
(config-router)# exit
(config)# exit
#
```

Quagga Router Status

- Very Cisco IOS like in its operation. E.g.

```
# show ip route
Codes: K - kernel route, C - connected, S - static
  B - BGP, > - selected route, * FIB route
S>* 0.0.0.0/0 [10/0] via 120.100.90.191, disc0
B>* 6.1.0.0/16 [20/0] via 212.211.12.11, yk0, 2w3d22h
B>* 6.4.0.0/18 [20/0] via 212.211.12.11, yk0, 2w3d22h
B>* 6.9.0.0/18 [20/0] via 212.211.12.11, yk0, 2w3d22h
....

# show bgp neighbors
BGP neighbor is 212.211.12.11 remote AS 999, local AS 123, external link
BGP version 4, remote router ID 212.211.1.2
BGP state = Established, up for 2w3d22h
```

Firewall Configuration

- Standard ipfw rules generated by a shell script

```
#!/bin/sh

# Shell variable definitions

#fw_cmd="/sbin/ipfw -q"
#fw_cmd="echo"
fw_cmd="/sbin/ipfw"
...

block_specific_ports () {
...
    for i in 42 '67-69' 135 137 138 139 445 593 707 4444
    do
        $fw_cmd add deny udp from any to any $i
        $fw_cmd add deny tcp from any to any $i
    done
...
```

Firewall Configuration

- Generates the following rules

```
....
04400 deny udp from any to any dst-port 42
04500 deny tcp from any to any dst-port 42
04600 deny udp from any to any dst-port 67-69
04700 deny tcp from any to any dst-port 67-69
04800 deny udp from any to any dst-port 135
04900 deny tcp from any to any dst-port 135
05000 deny udp from any to any dst-port 137
05100 deny tcp from any to any dst-port 137
05200 deny udp from any to any dst-port 138
05300 deny tcp from any to any dst-port 138
05400 deny udp from any to any dst-port 139
05500 deny tcp from any to any dst-port 139
....
```

- With some significant hit rates

```
04400      993982      76998318 deny udp from any to any dst-port 42
04500       50674       2268196 deny tcp from any to any dst-port 42
04600     1123622      89220640 deny udp from any to any dst-port 67-69
04700       29041       1311372 deny tcp from any to any dst-port 67-69
04800     1224973      95173144 deny udp from any to any dst-port 135
04900     3997326     191659369 deny tcp from any to any dst-port 135
05000    17172122    1444445154 deny udp from any to any dst-port 137
05100        4820        199364 deny tcp from any to any dst-port 137
05200      648353      80342588 deny udp from any to any dst-port 138
05300        4365        178676 deny tcp from any to any dst-port 138
05400     1360813     103639114 deny udp from any to any dst-port 139
05500     3295785     157996089 deny tcp from any to any dst-port 139
05600     1670810     127109315 deny udp from any to any dst-port 445
05700     9001654     432332437 deny tcp from any to any dst-port 445
05800         736         53546 deny udp from any to any dst-port 593
05900        4292        175240 deny tcp from any to any dst-port 593
06000      356367      27086091 deny udp from any to any dst-port 707
06100        4014        164292 deny tcp from any to any dst-port 707
06200      653978      70351571 deny udp from any to any dst-port 4444
06300     6258924     278074726 deny tcp from any to any dst-port 4444
```

Traffic Shaping Configuration

- This fellow was running a public RH mirror...

```
#!/bin/sh

#fw_cmd="echo"
fw_cmd="/sbin/ipfw"
...
rh_hosts="{ $eyetap or $horus or $mann5 or $commando }"
...
initialize () {
        $fw_cmd -f flush
        $fw_cmd -f pipe flush
        $fw_cmd -f queue flush
}
steve_mann_traffic_shaping () {
#       $fw_cmd add pipe 3 ip from $eyetap to any in recv $uoft_if
        $fw_cmd add pipe 3 ip from $rh_hosts to any in recv $uoft_if
        $fw_cmd pipe 3 config bw 50Kbit/s
}
```

ResNet Routing Revisited

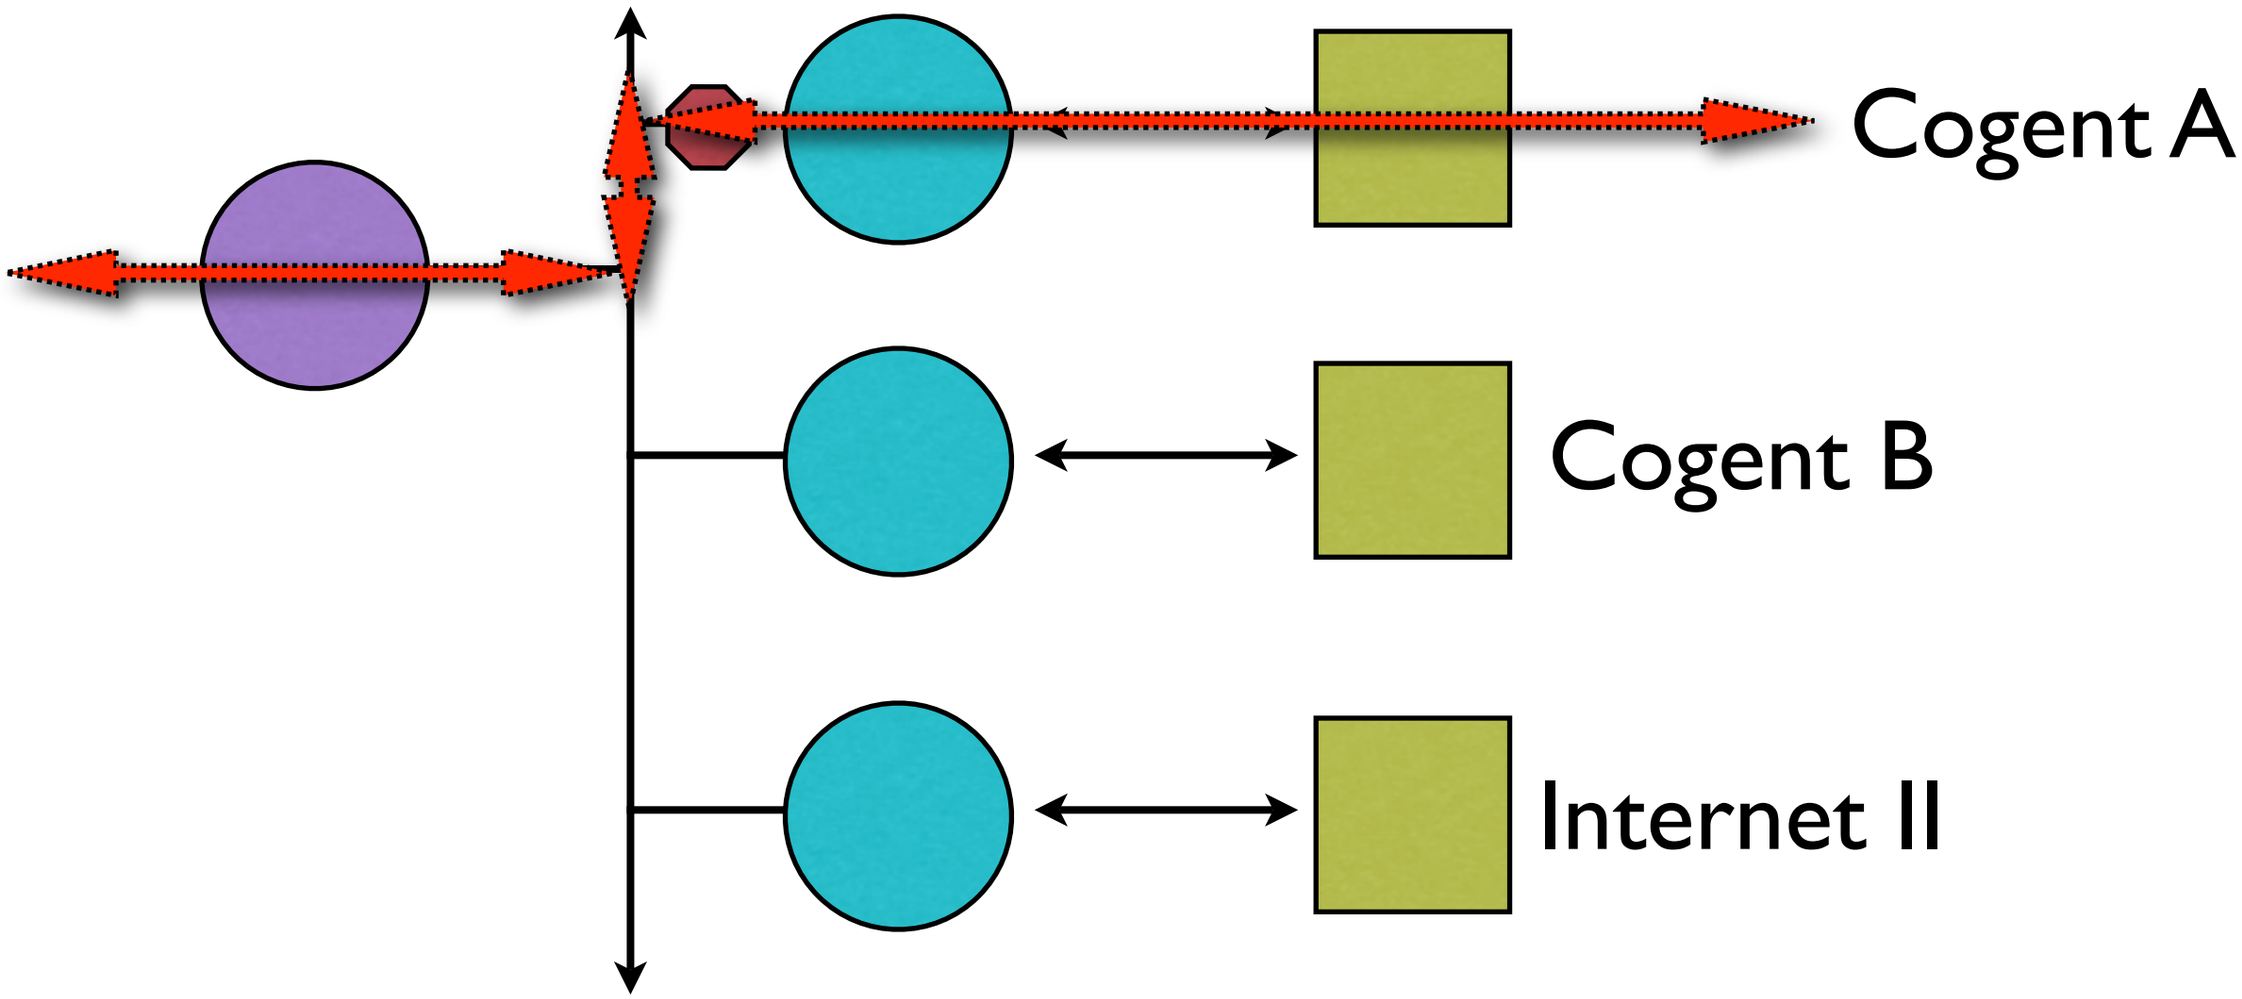

ResNet Policy Using fwd

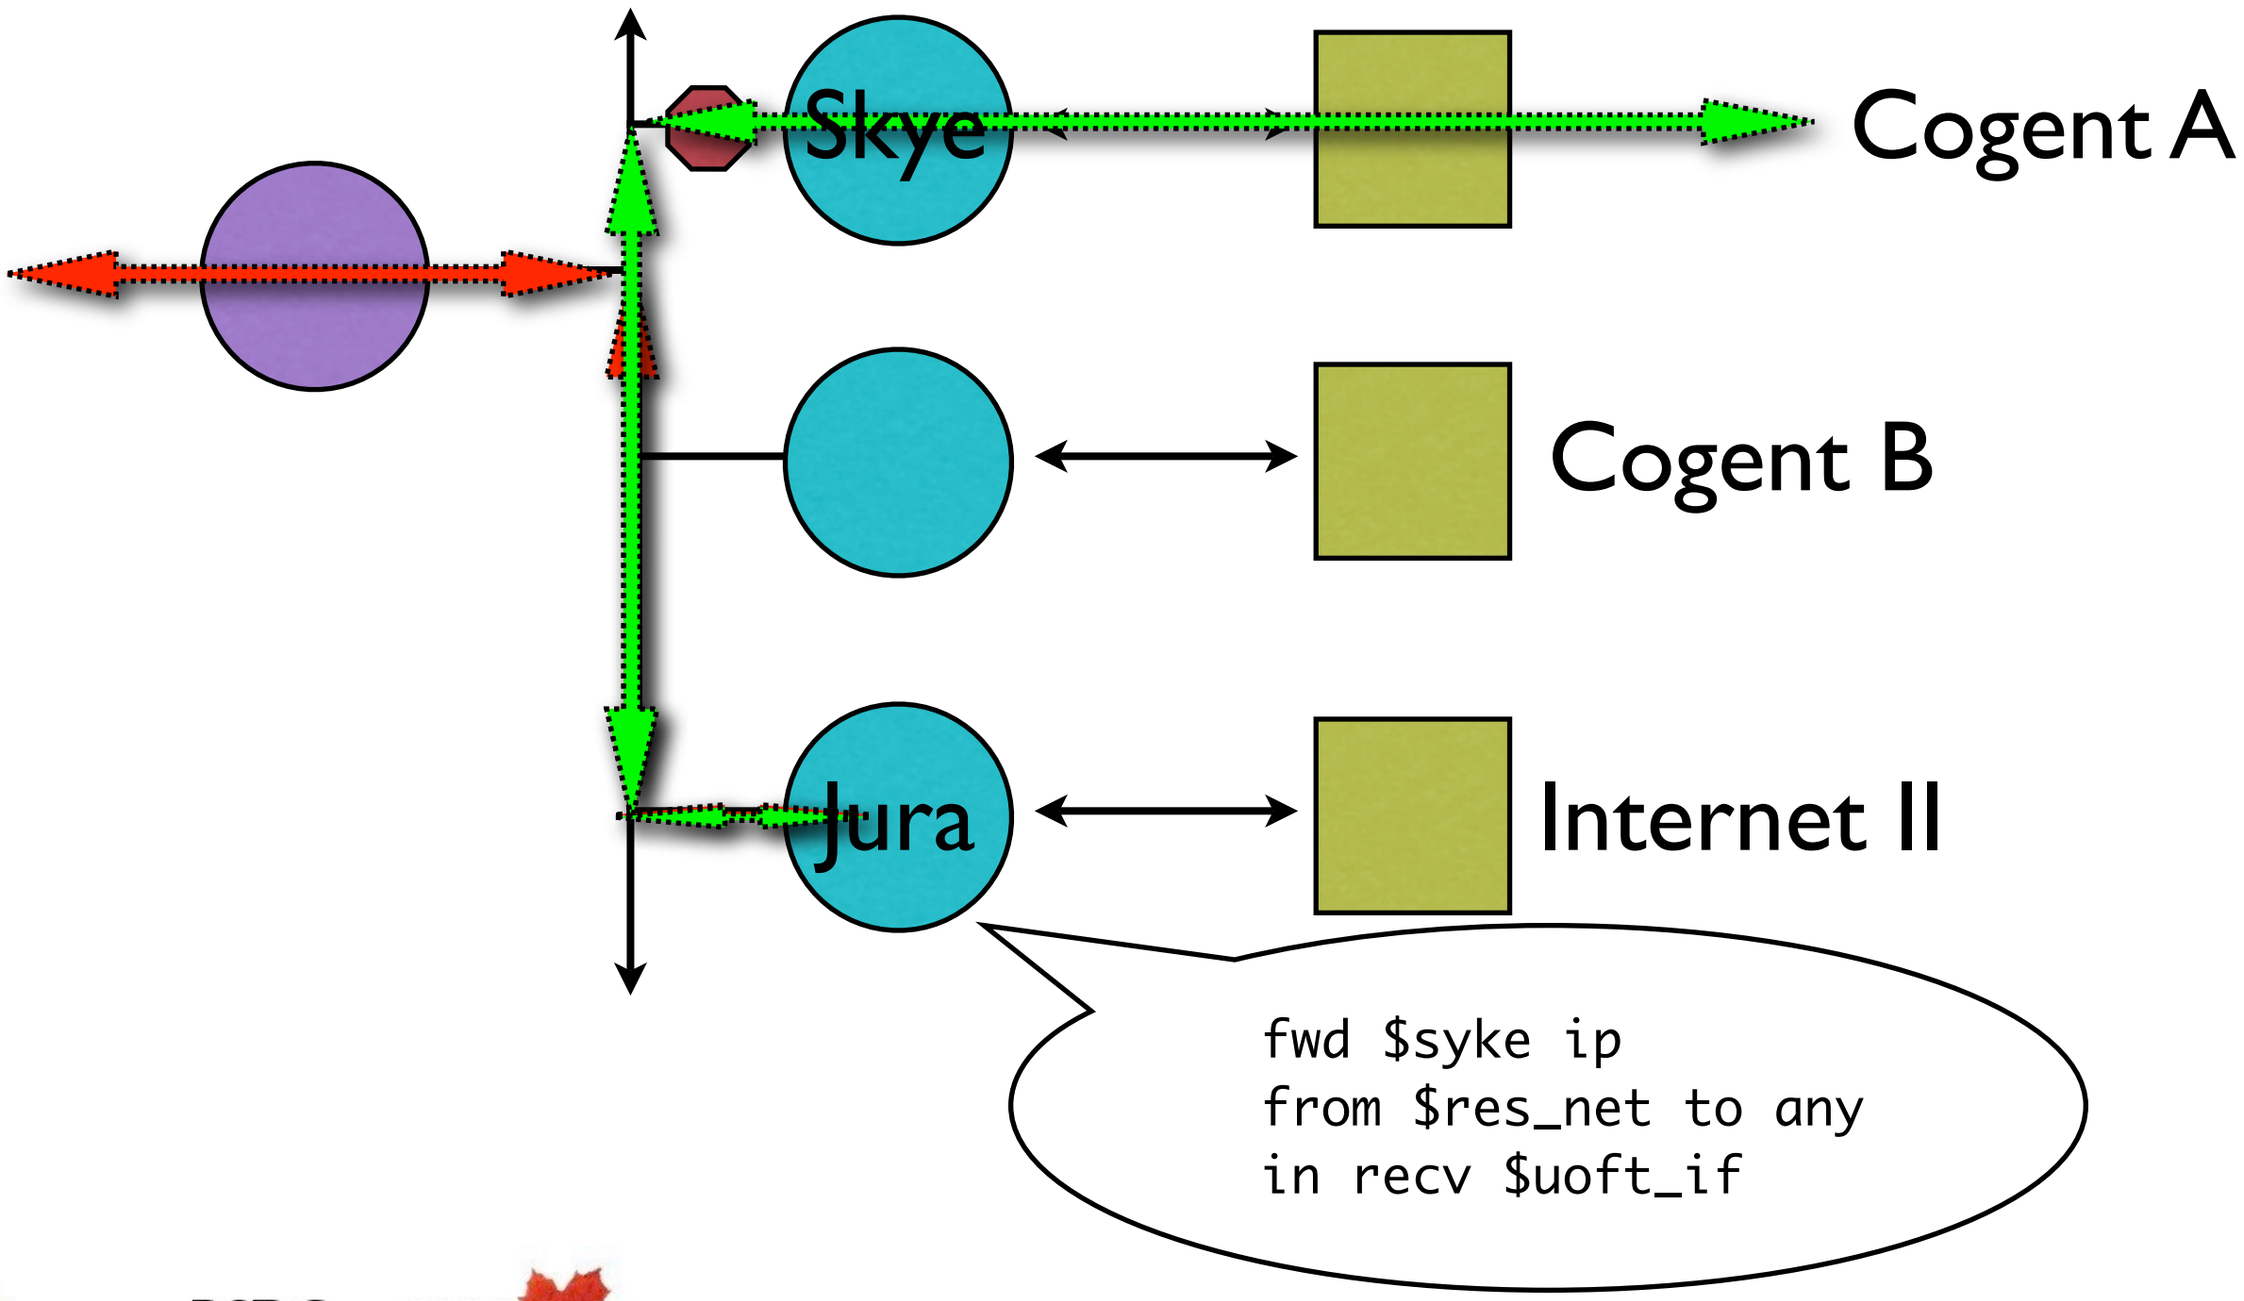

UofT A Routing Revisited

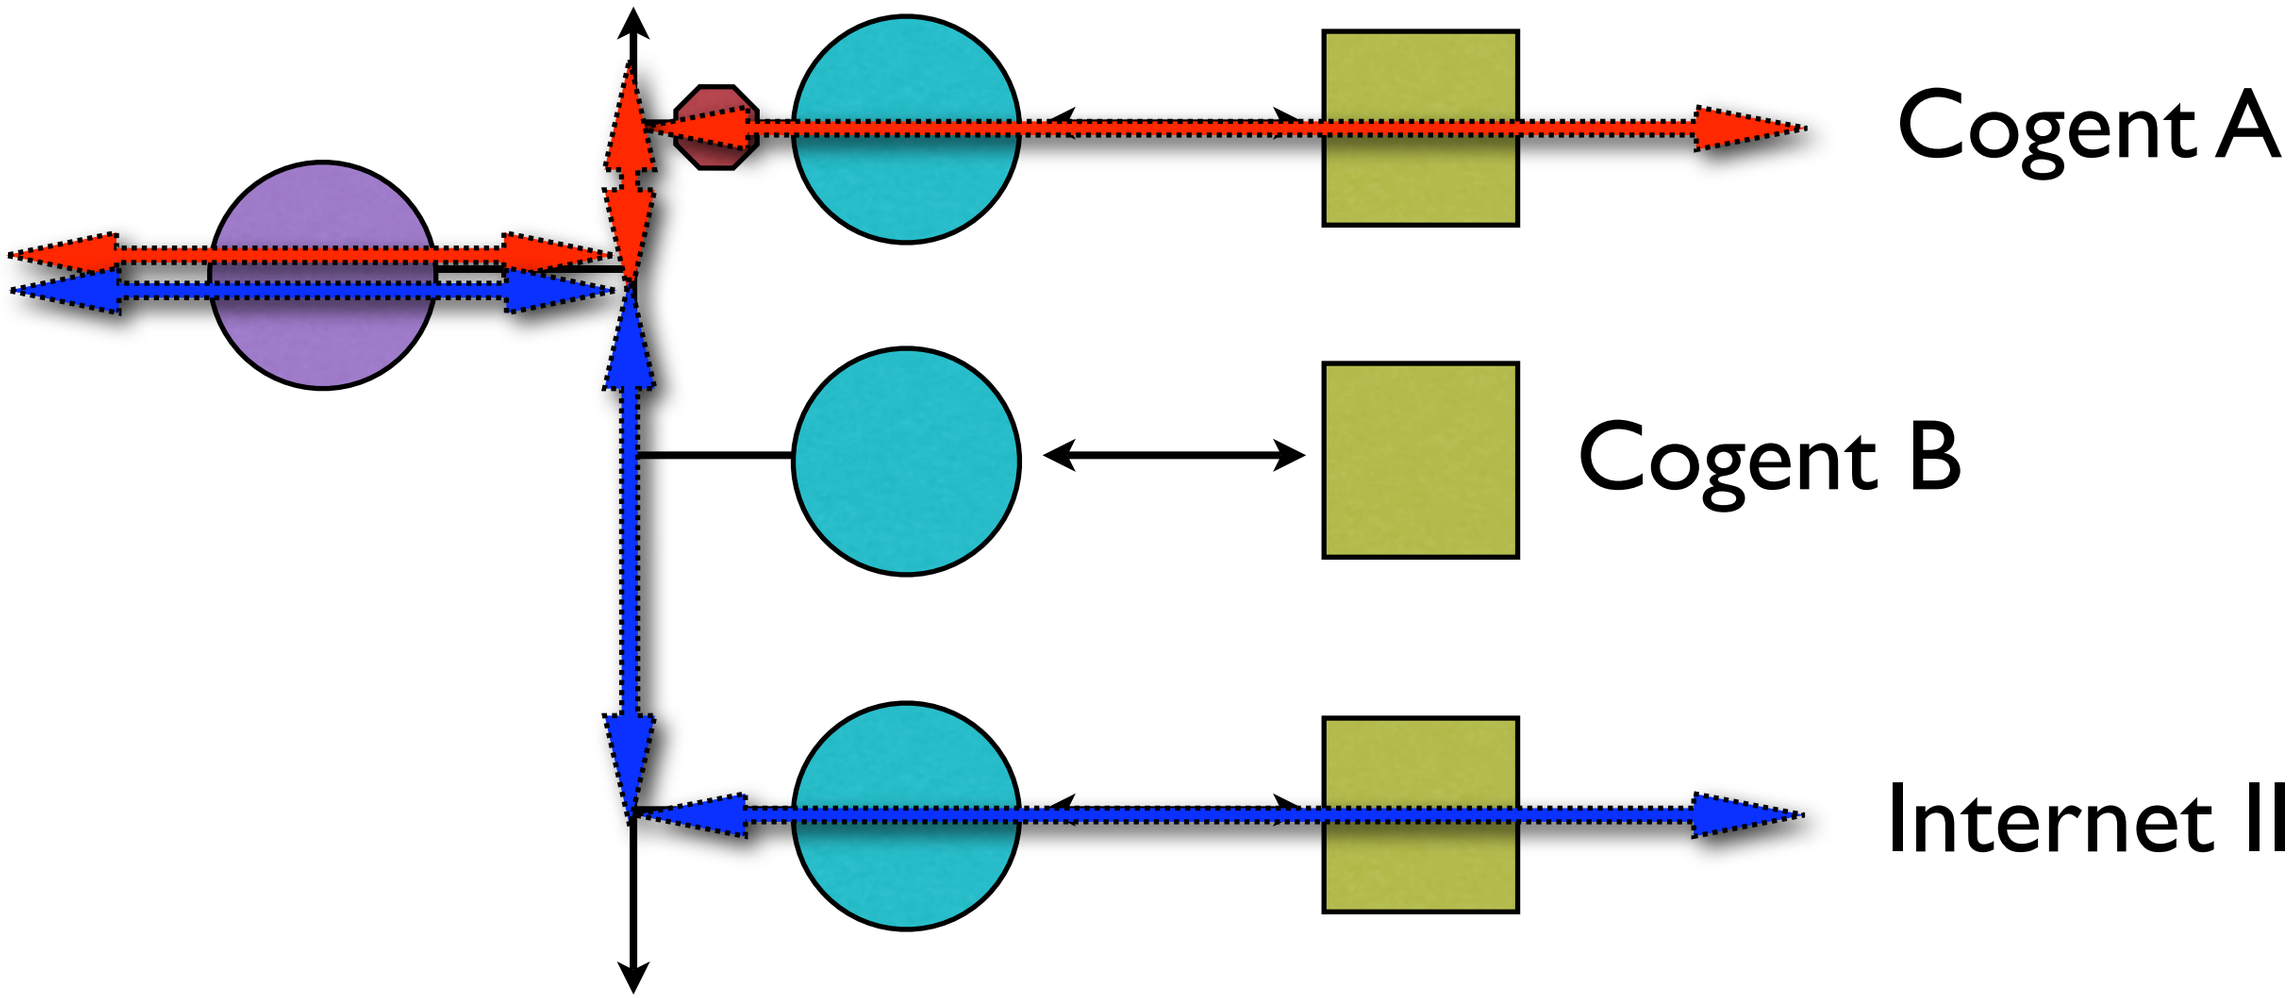

UofT A Policy Using fwd

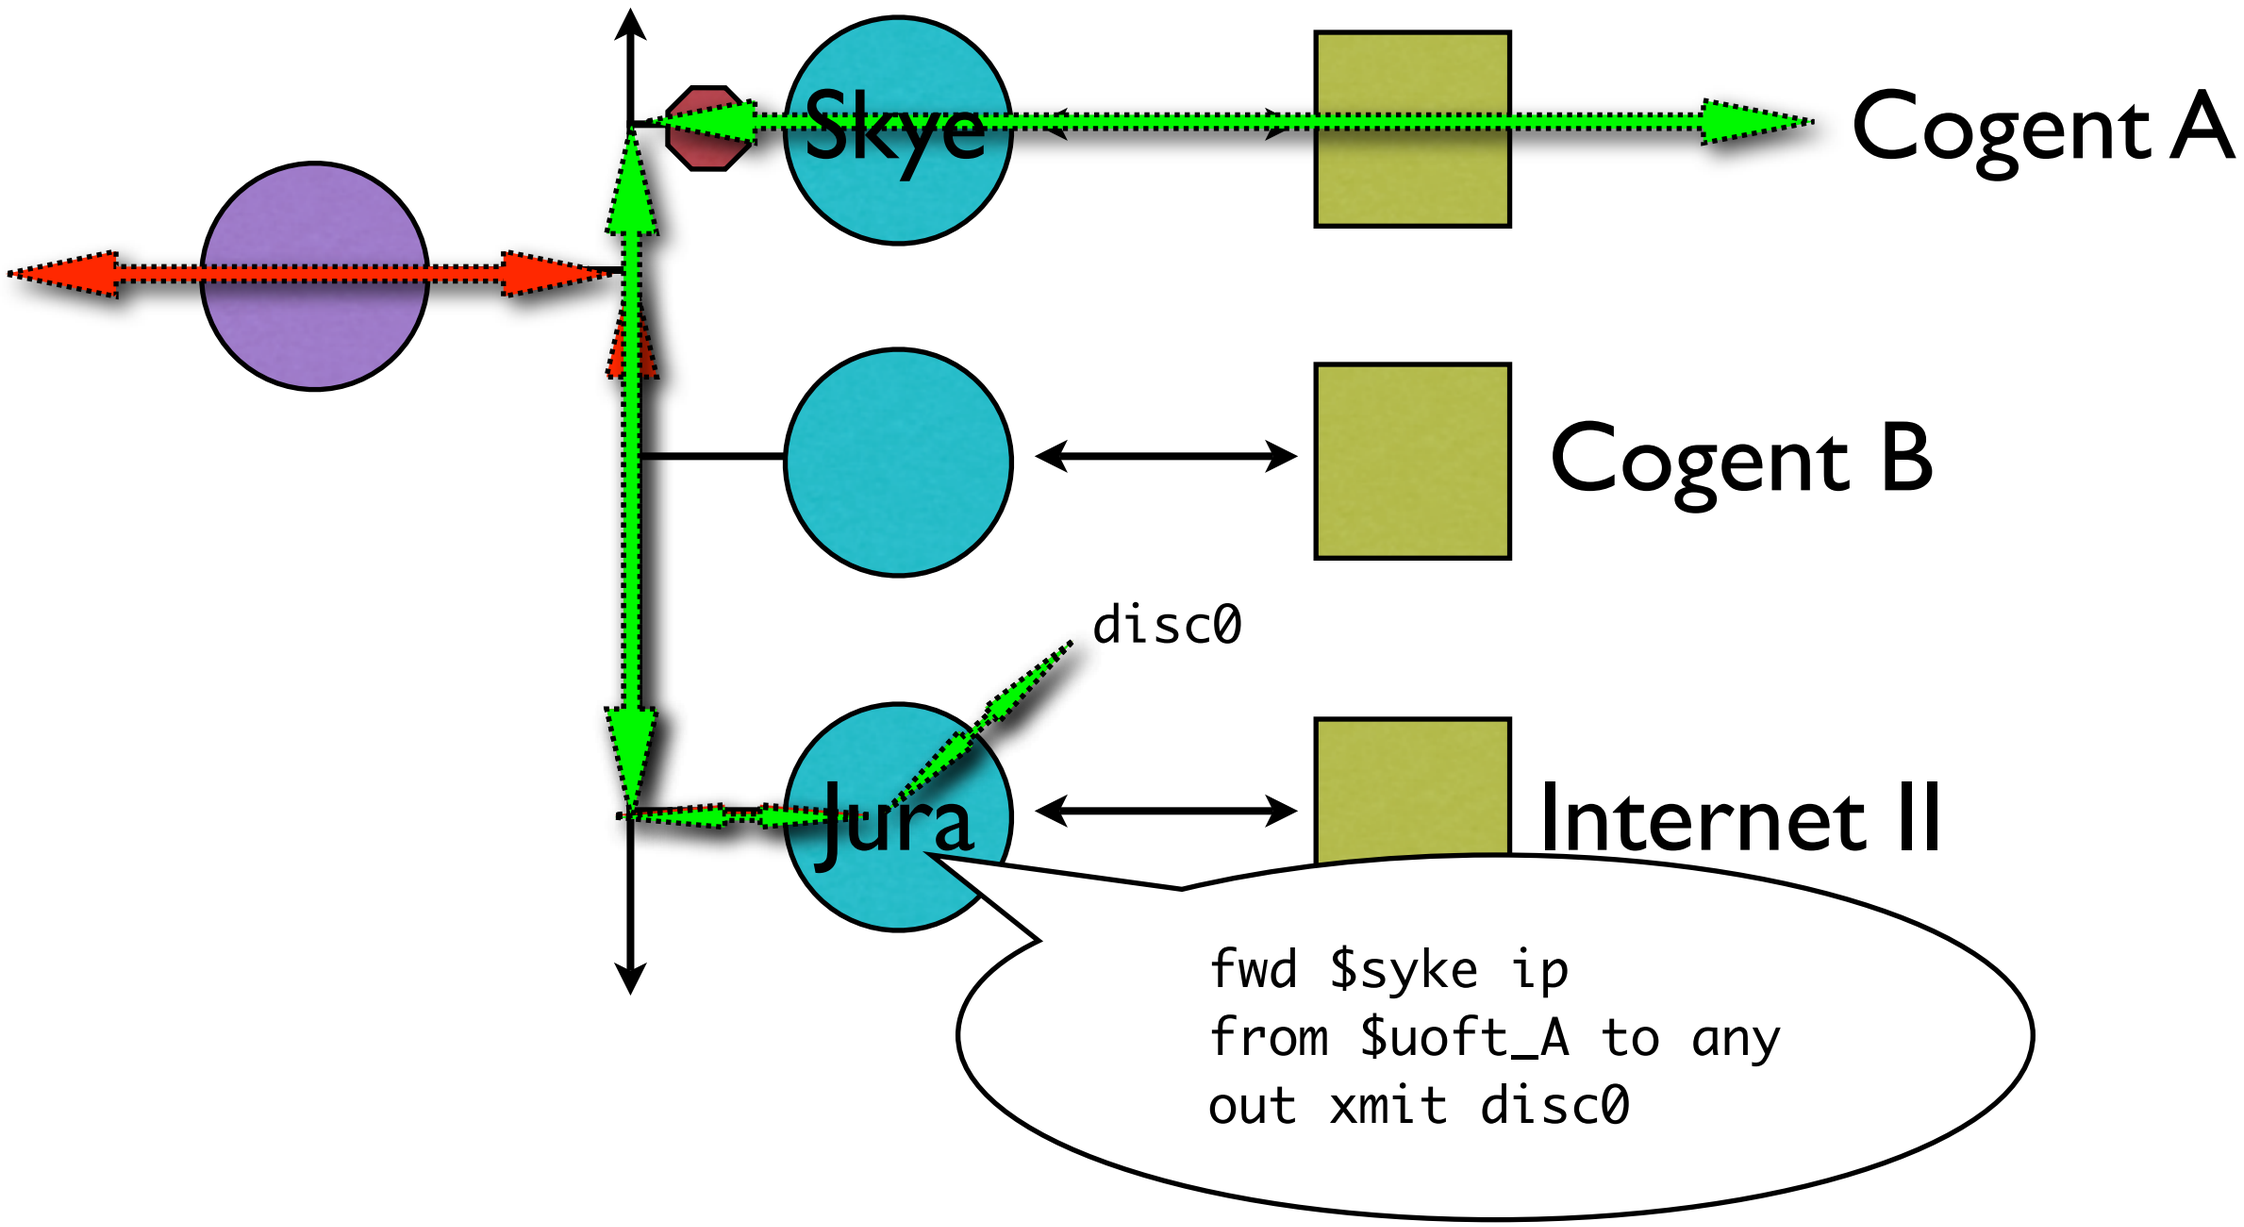

Cacti Statistics

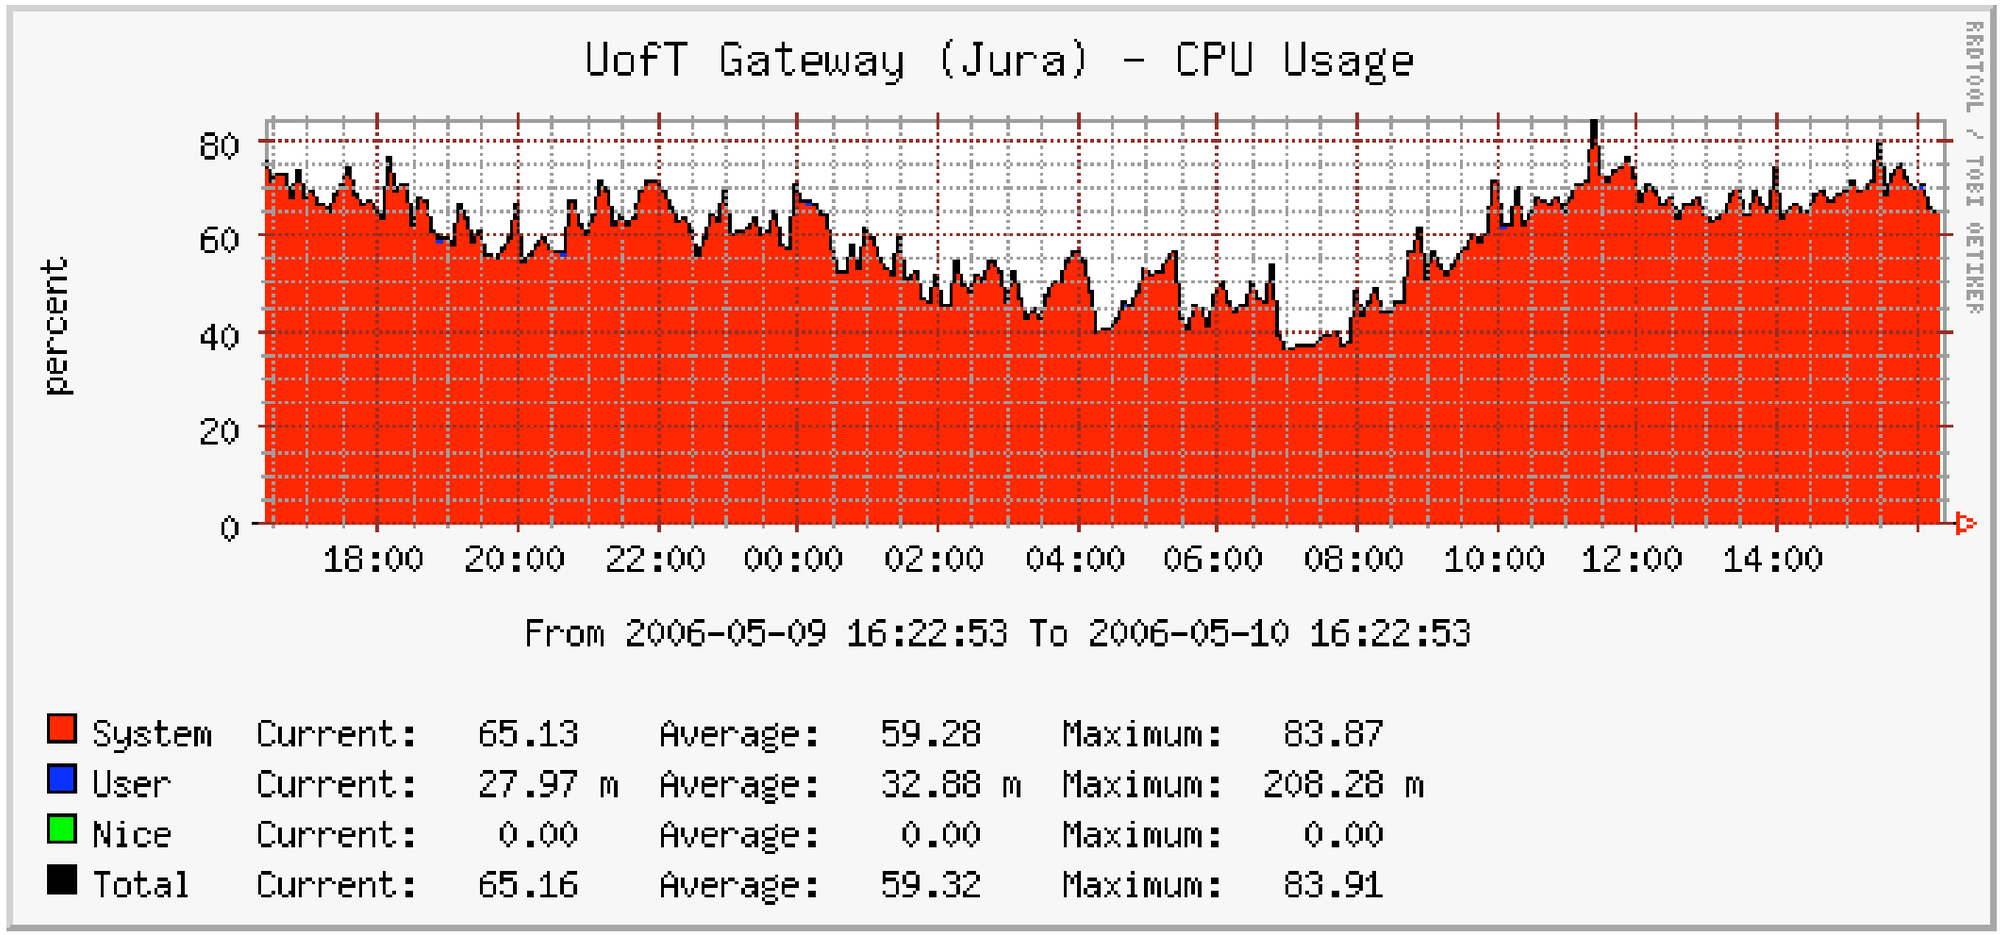

Cacti Statistics

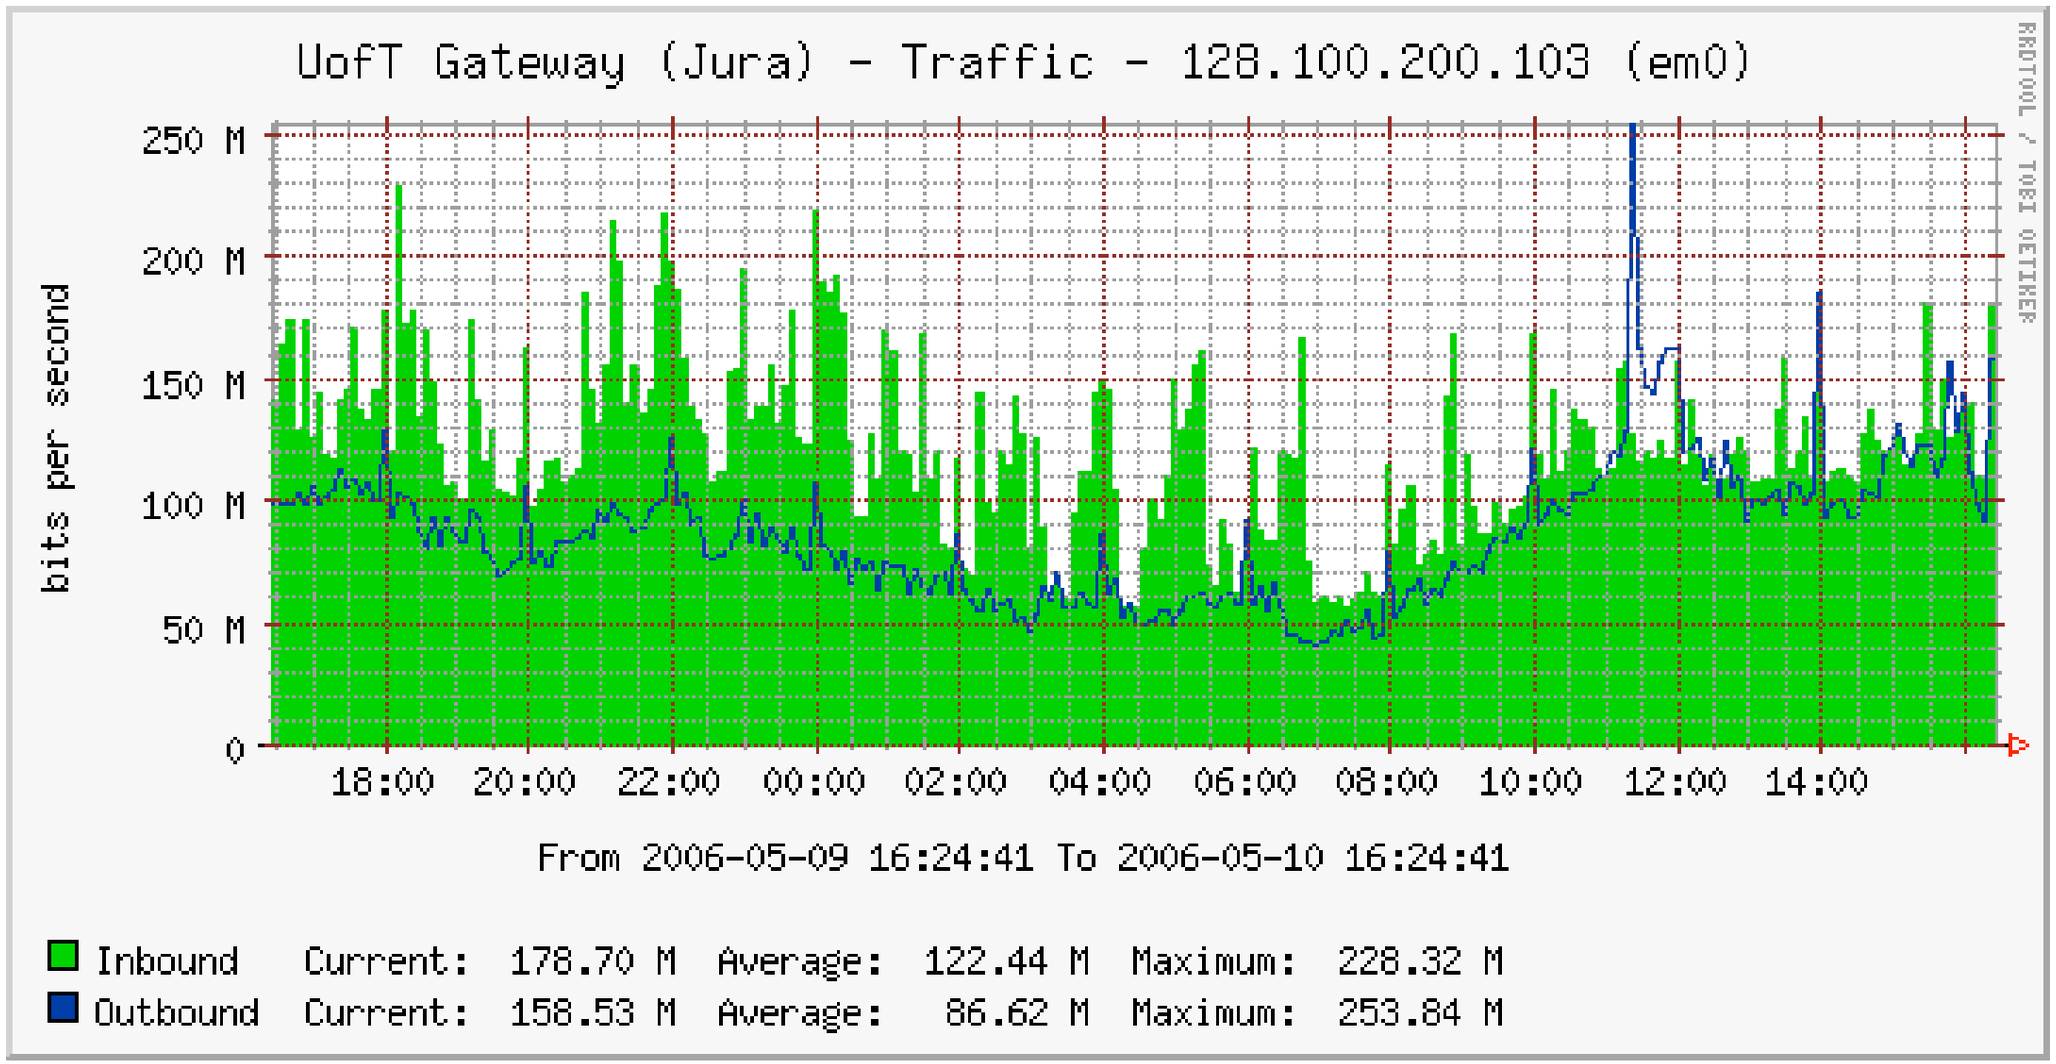

Cacti Statistics

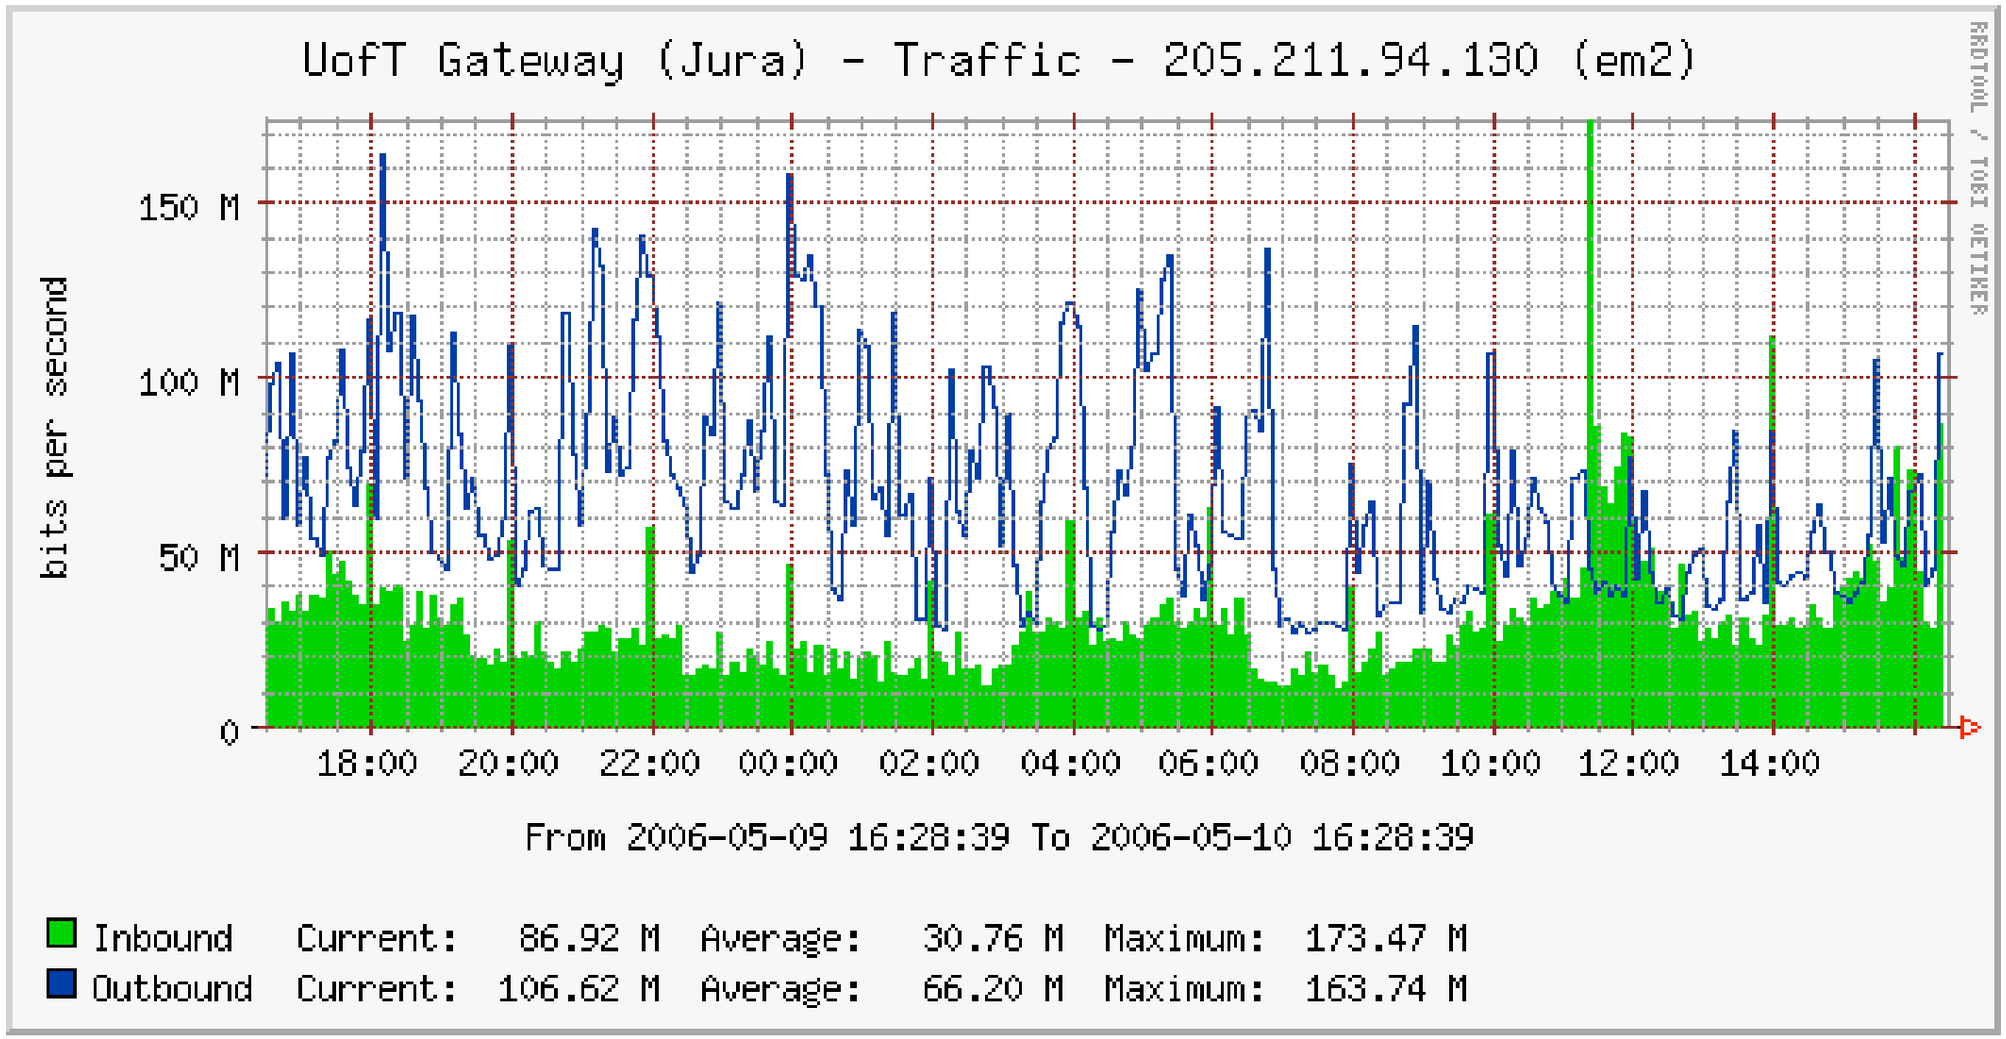

Cacti Statistics

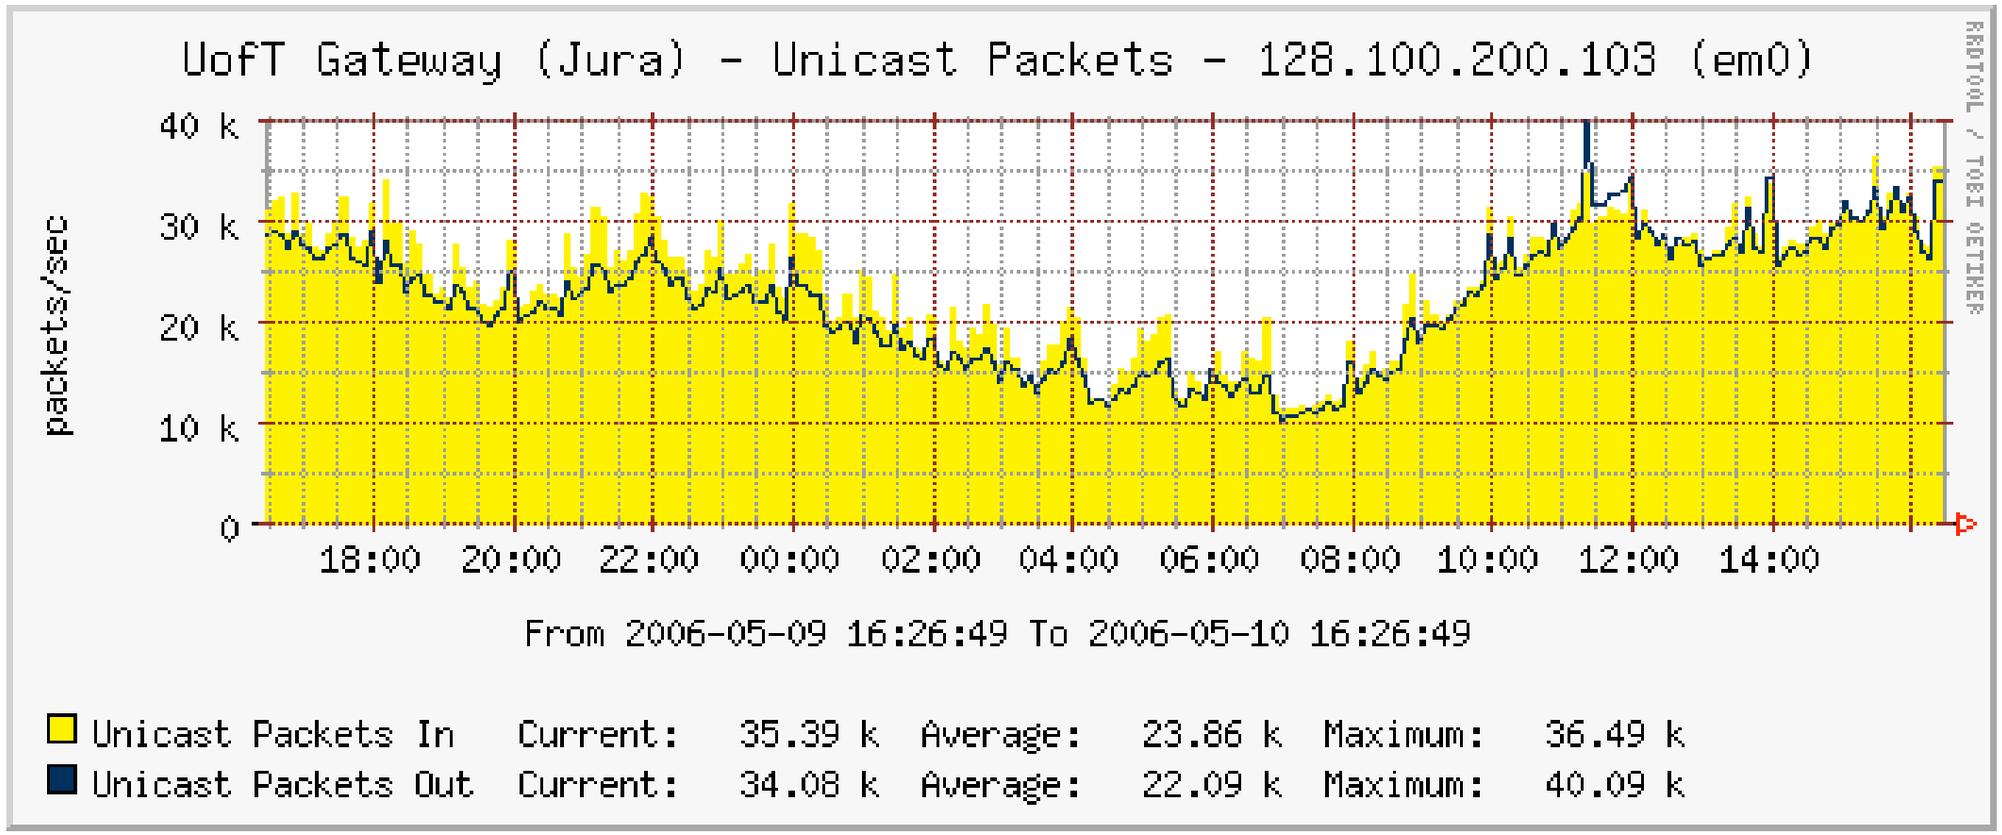

Cacti Statistics

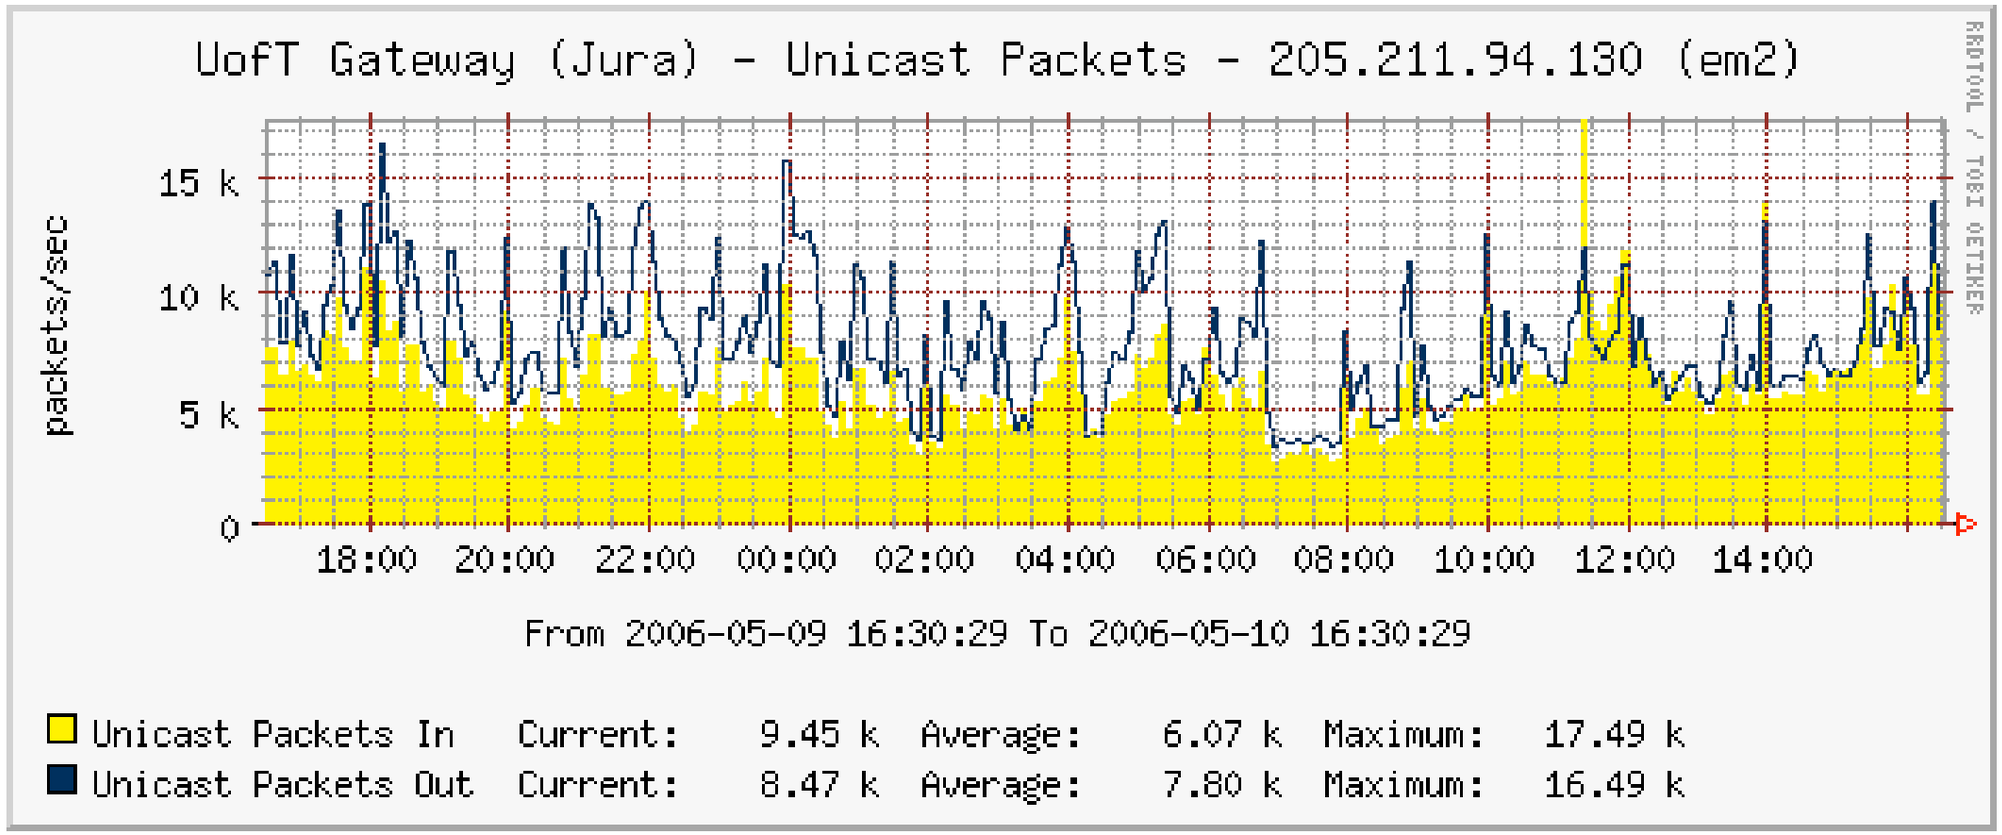

Conclusions

- Routers have been in service for over 3 years
- Not for the feint of heart
- In-house UNIX expertise is needed
- Testing configurations is difficult
- Performance uncertainties at current bit rates