Participatory Networking

Rodrigo Fonseca
Join work with Andrew Ferguson, Arjun Guha, Jordan Place

BROWN

Networking in the Cloud

Networking in the Cloud

Can the cloud provider get **help** in confirguring the network?

A few motivating examples

Large-scale processing

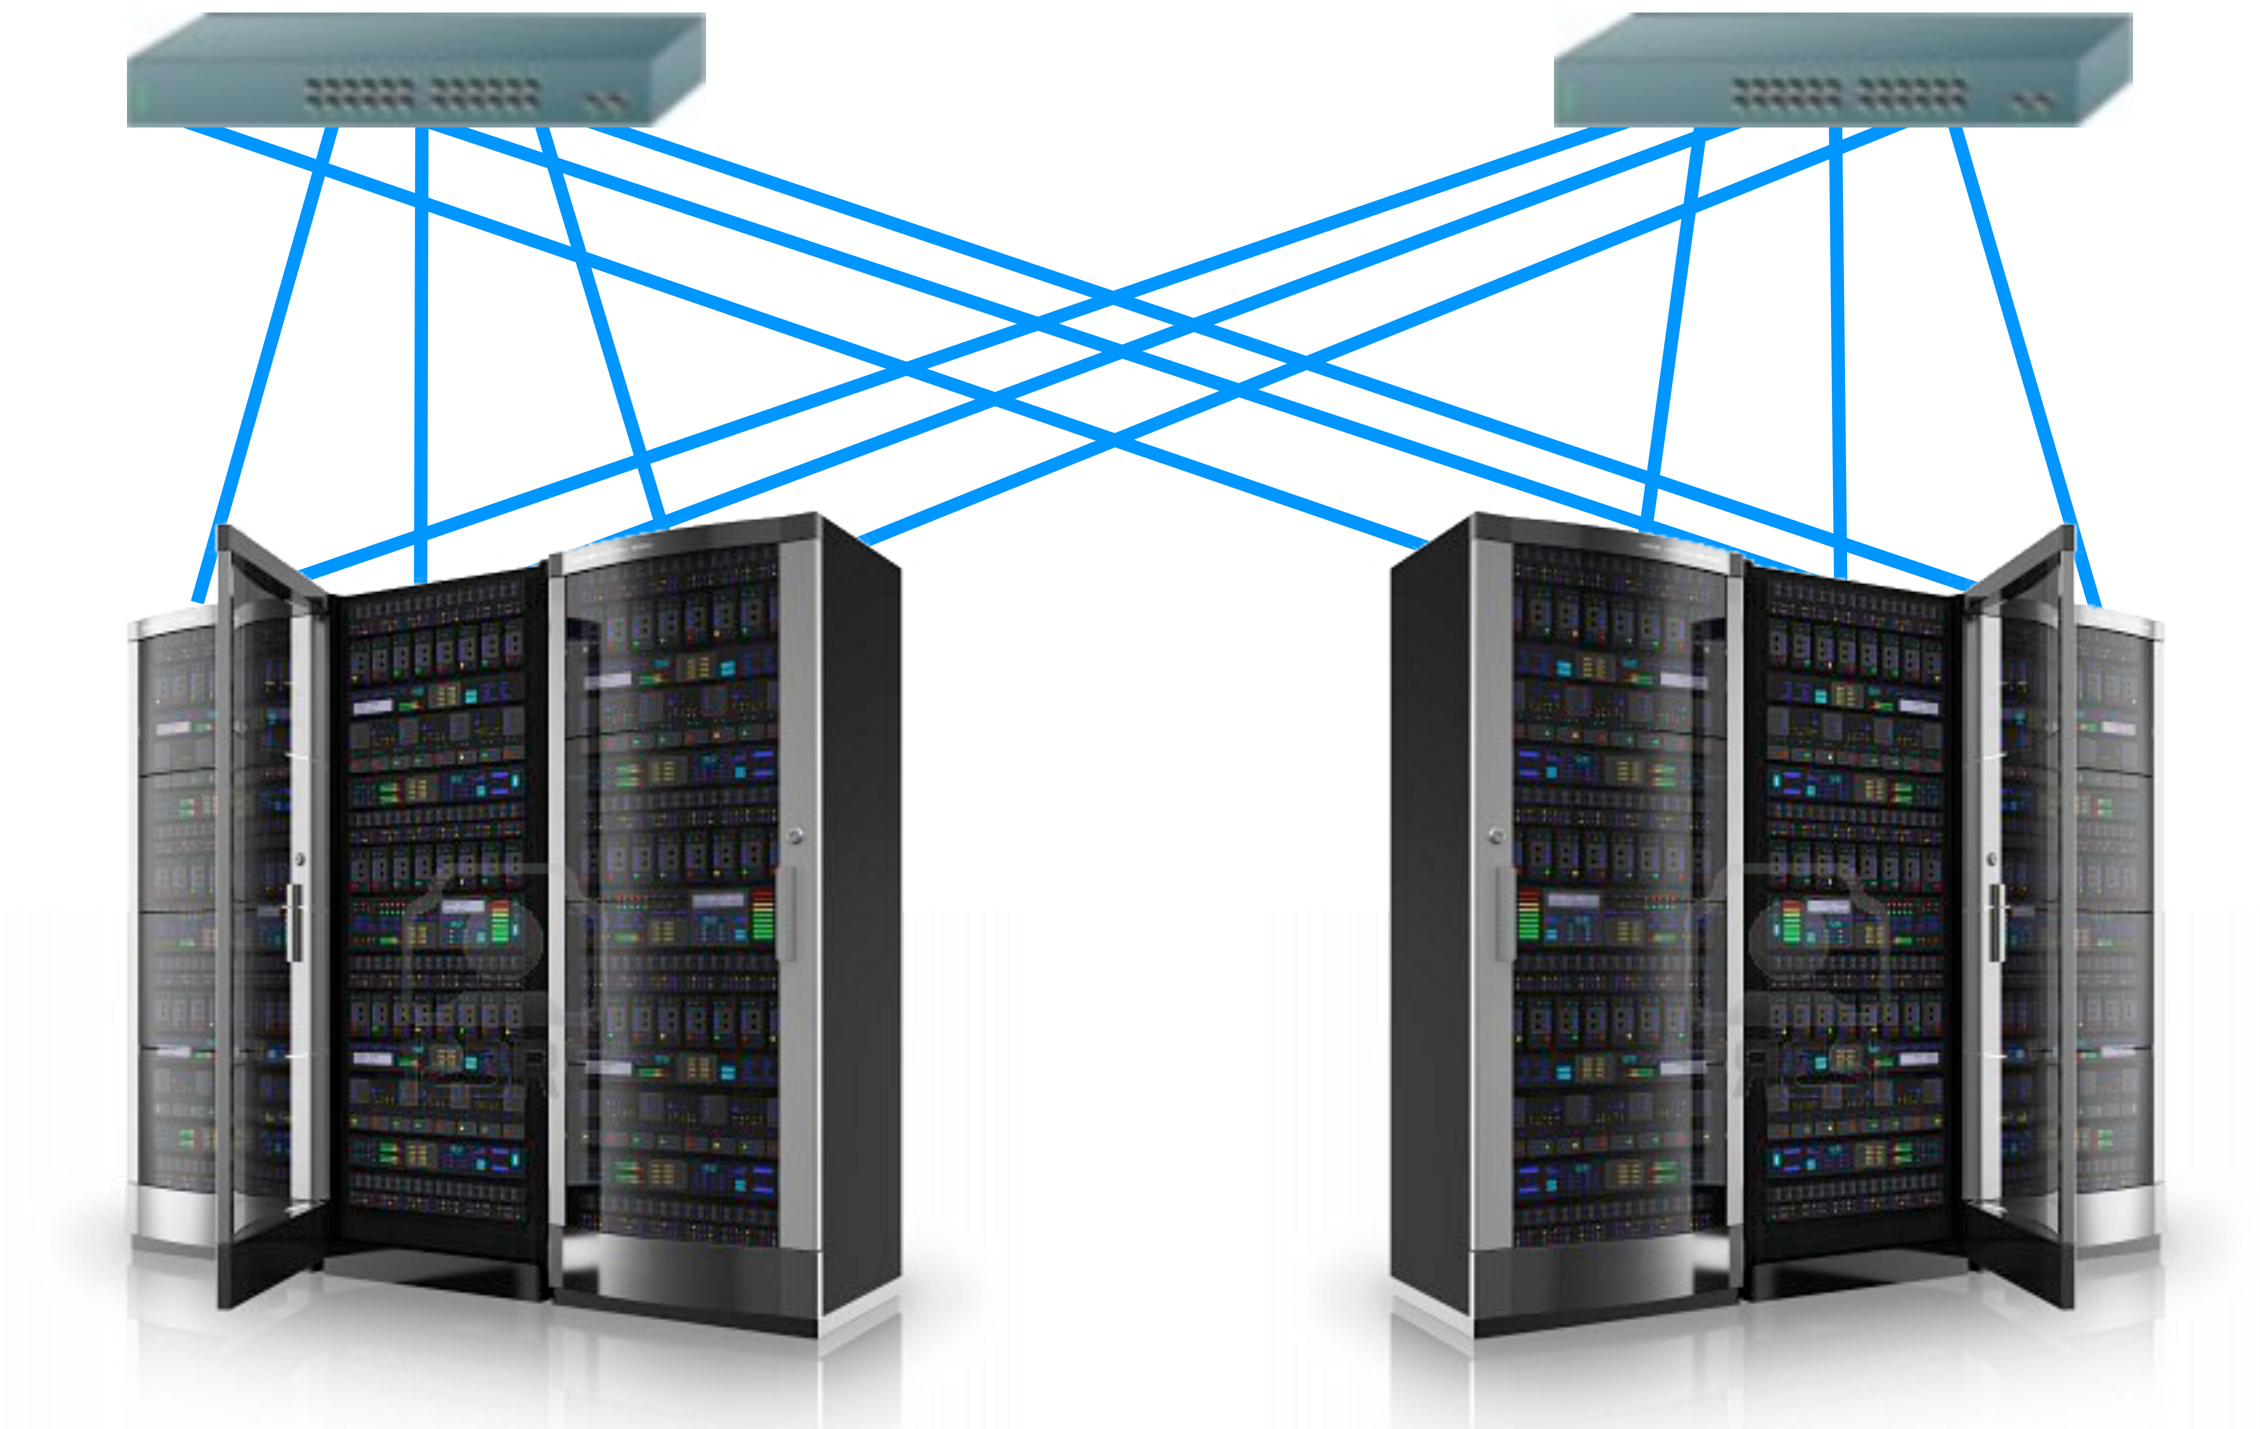

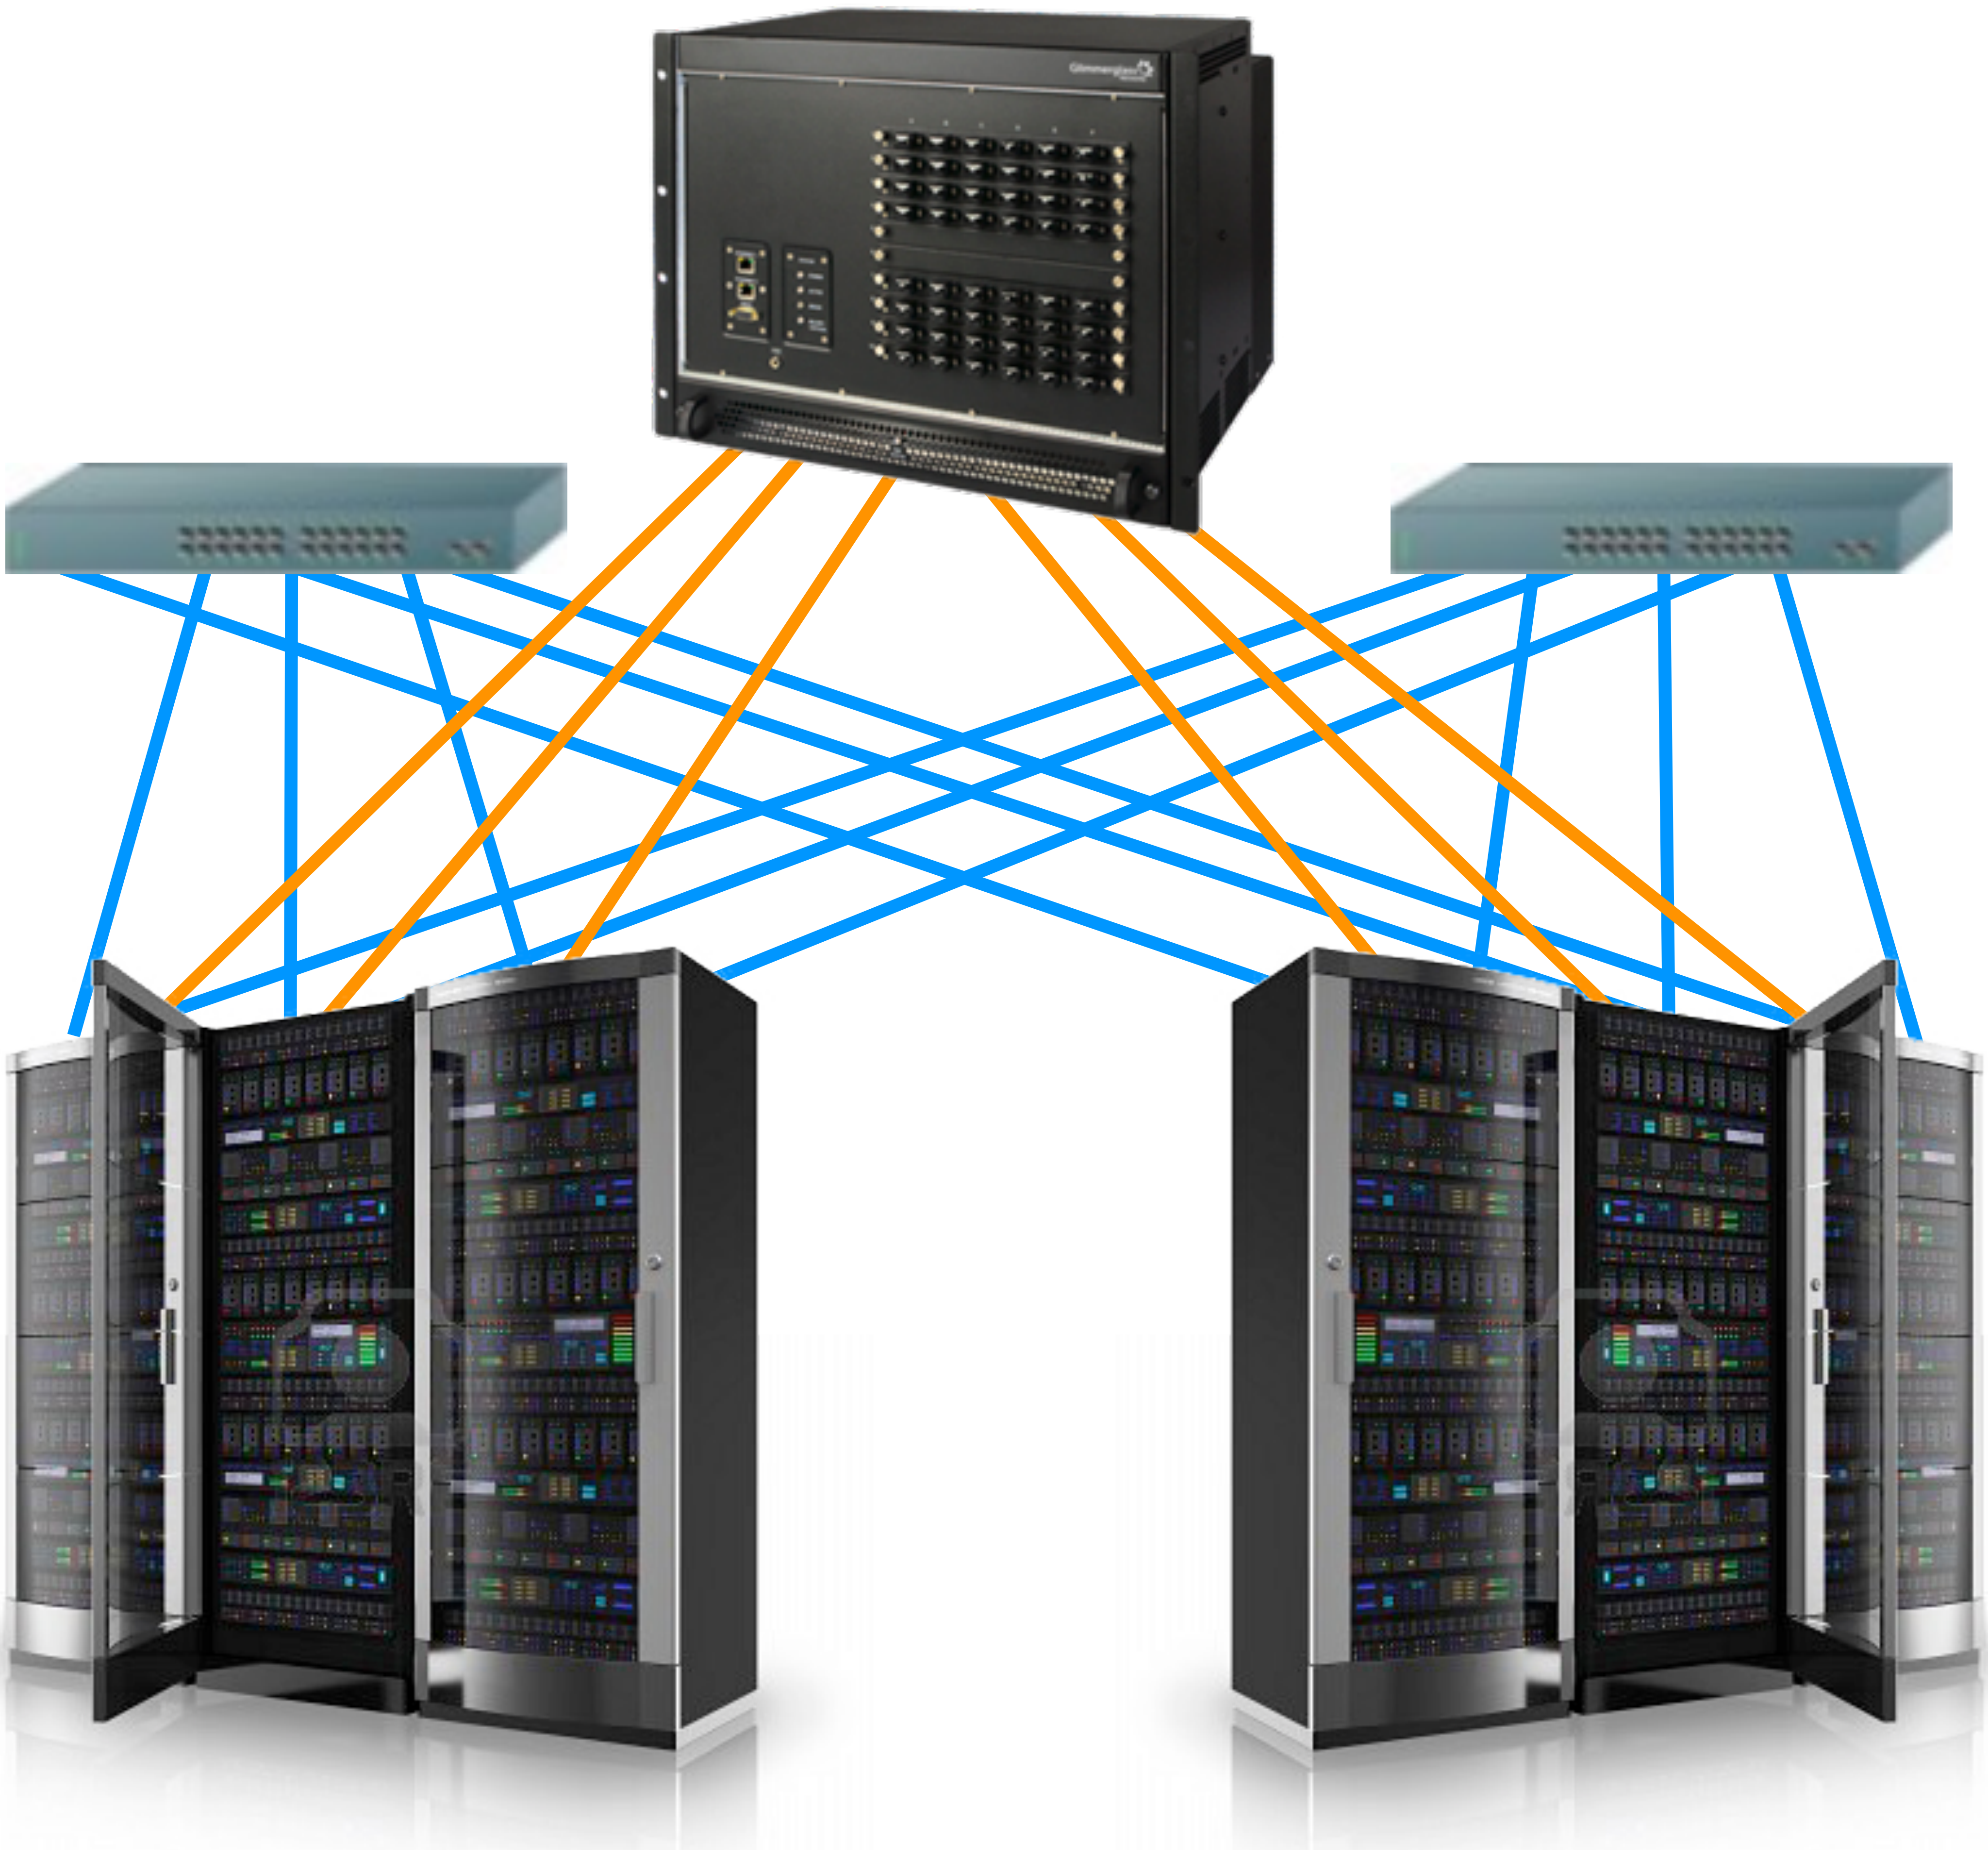

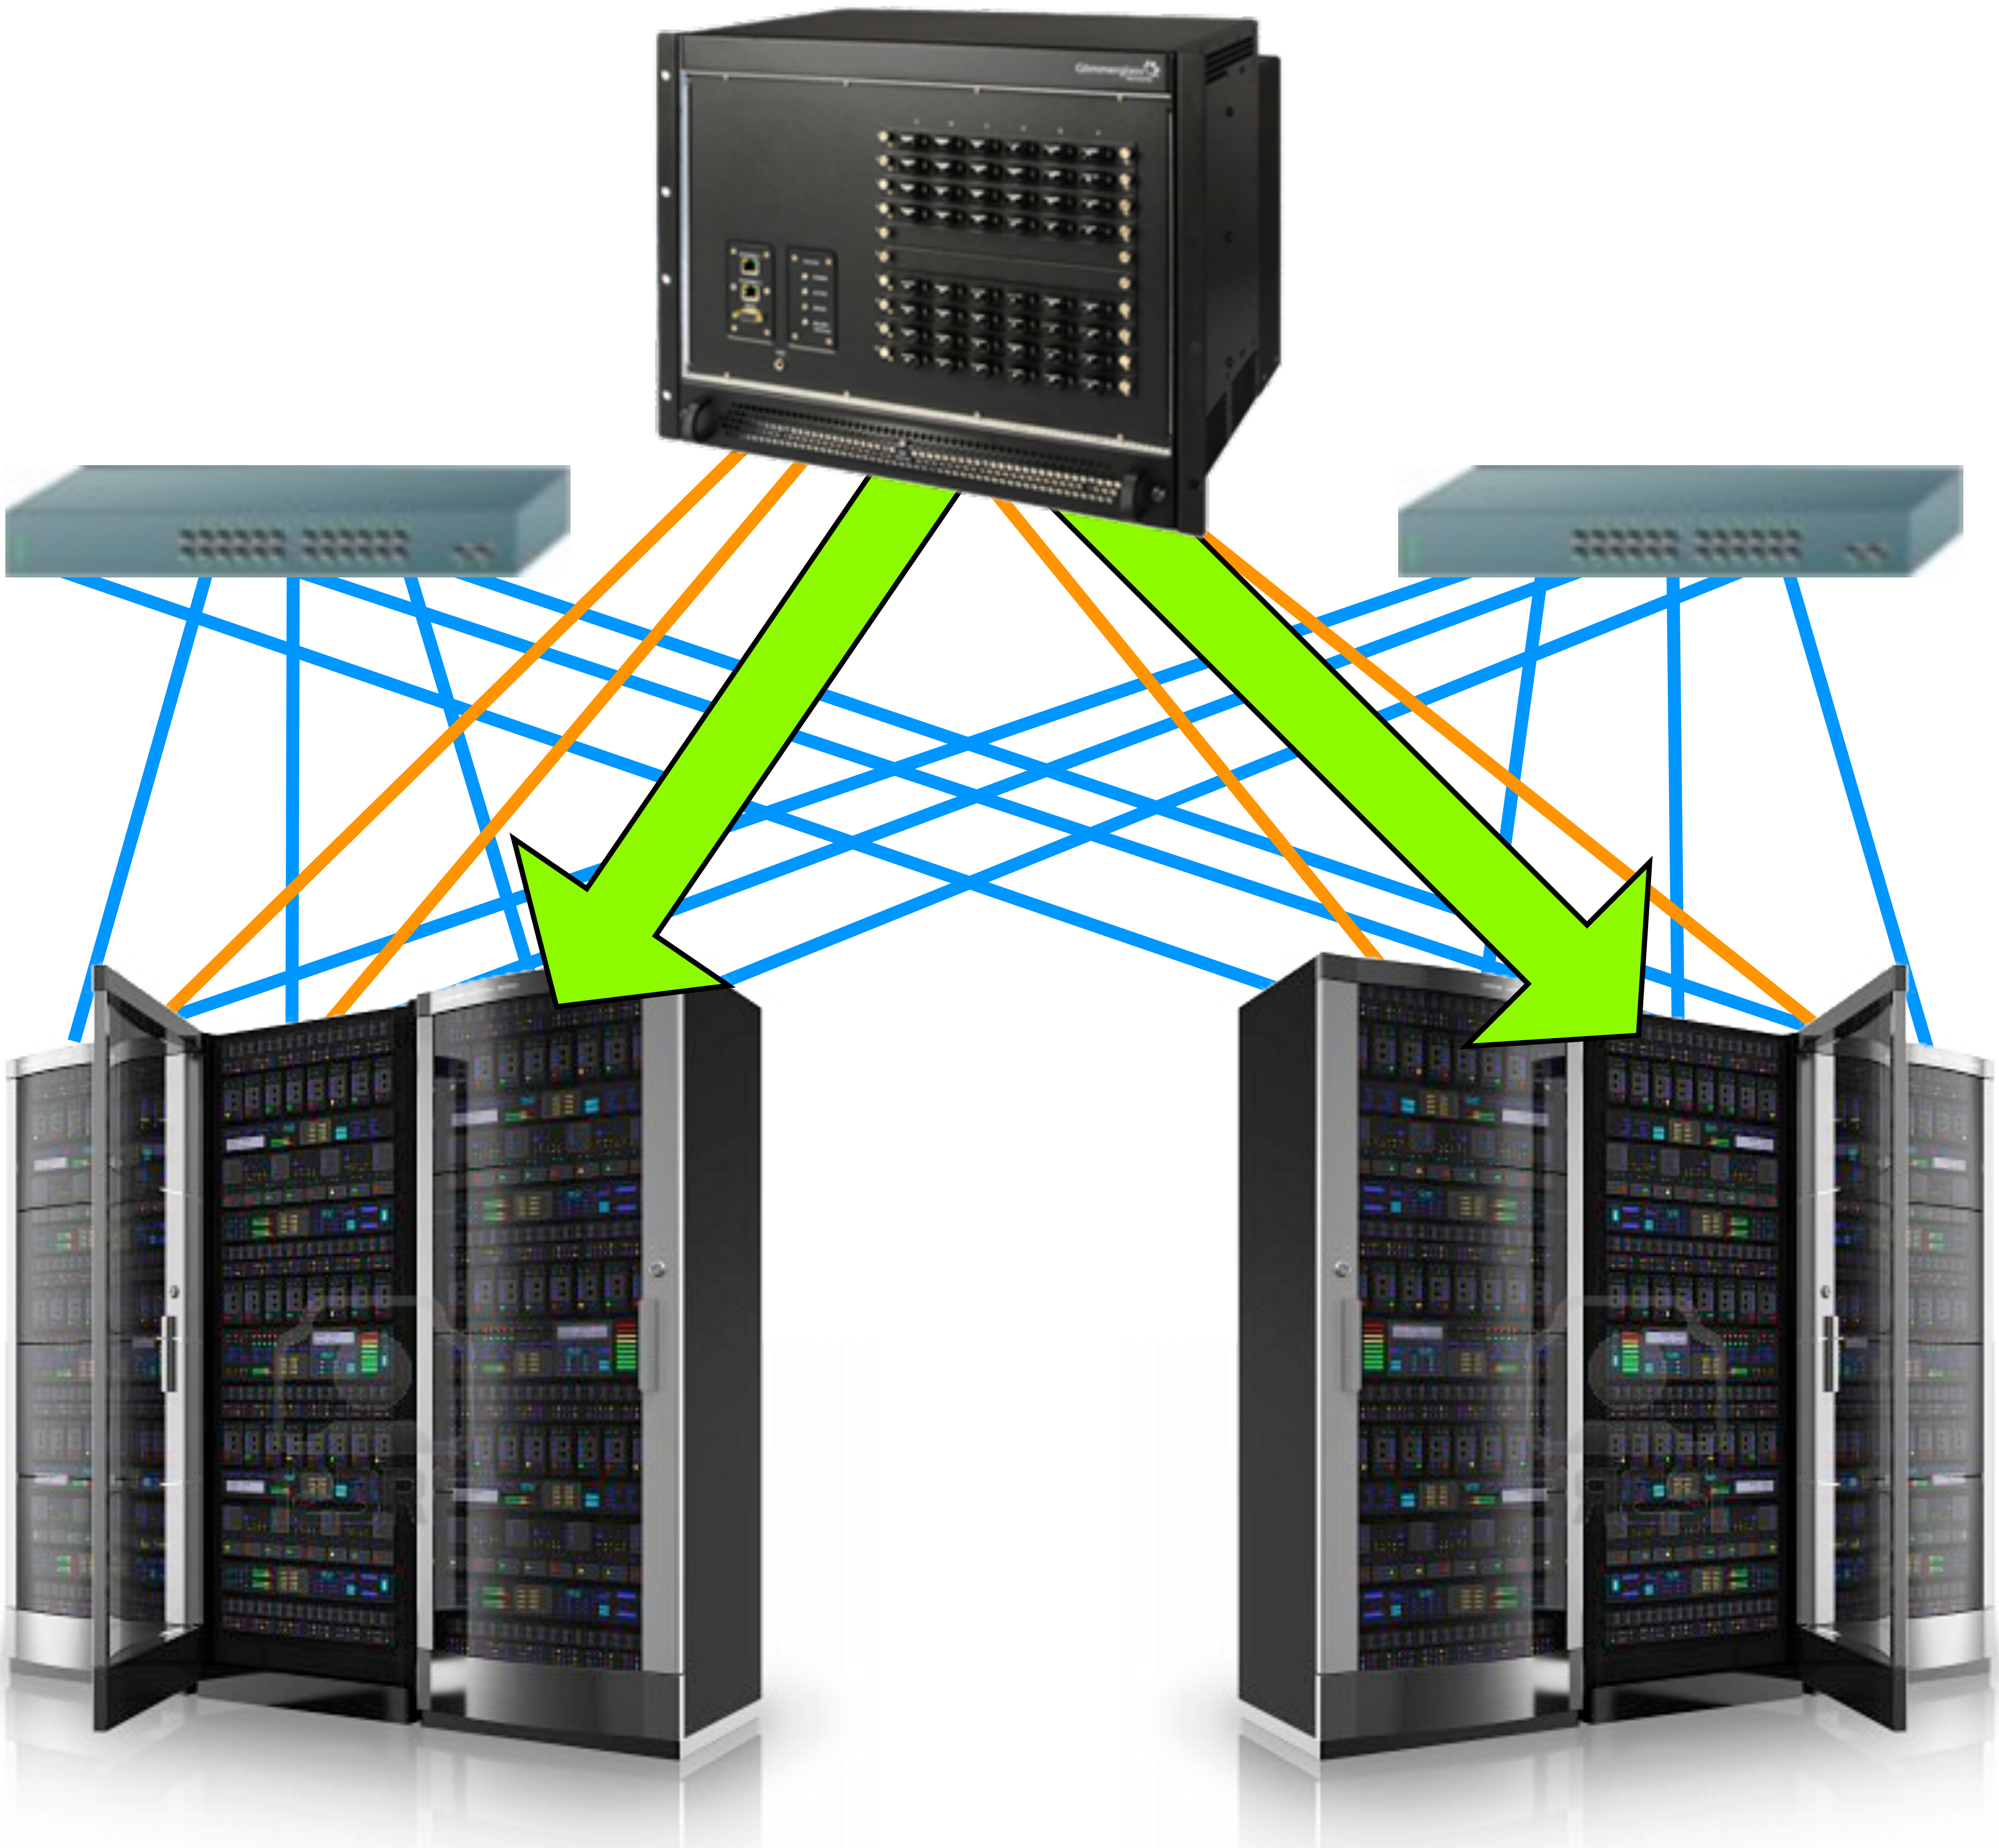

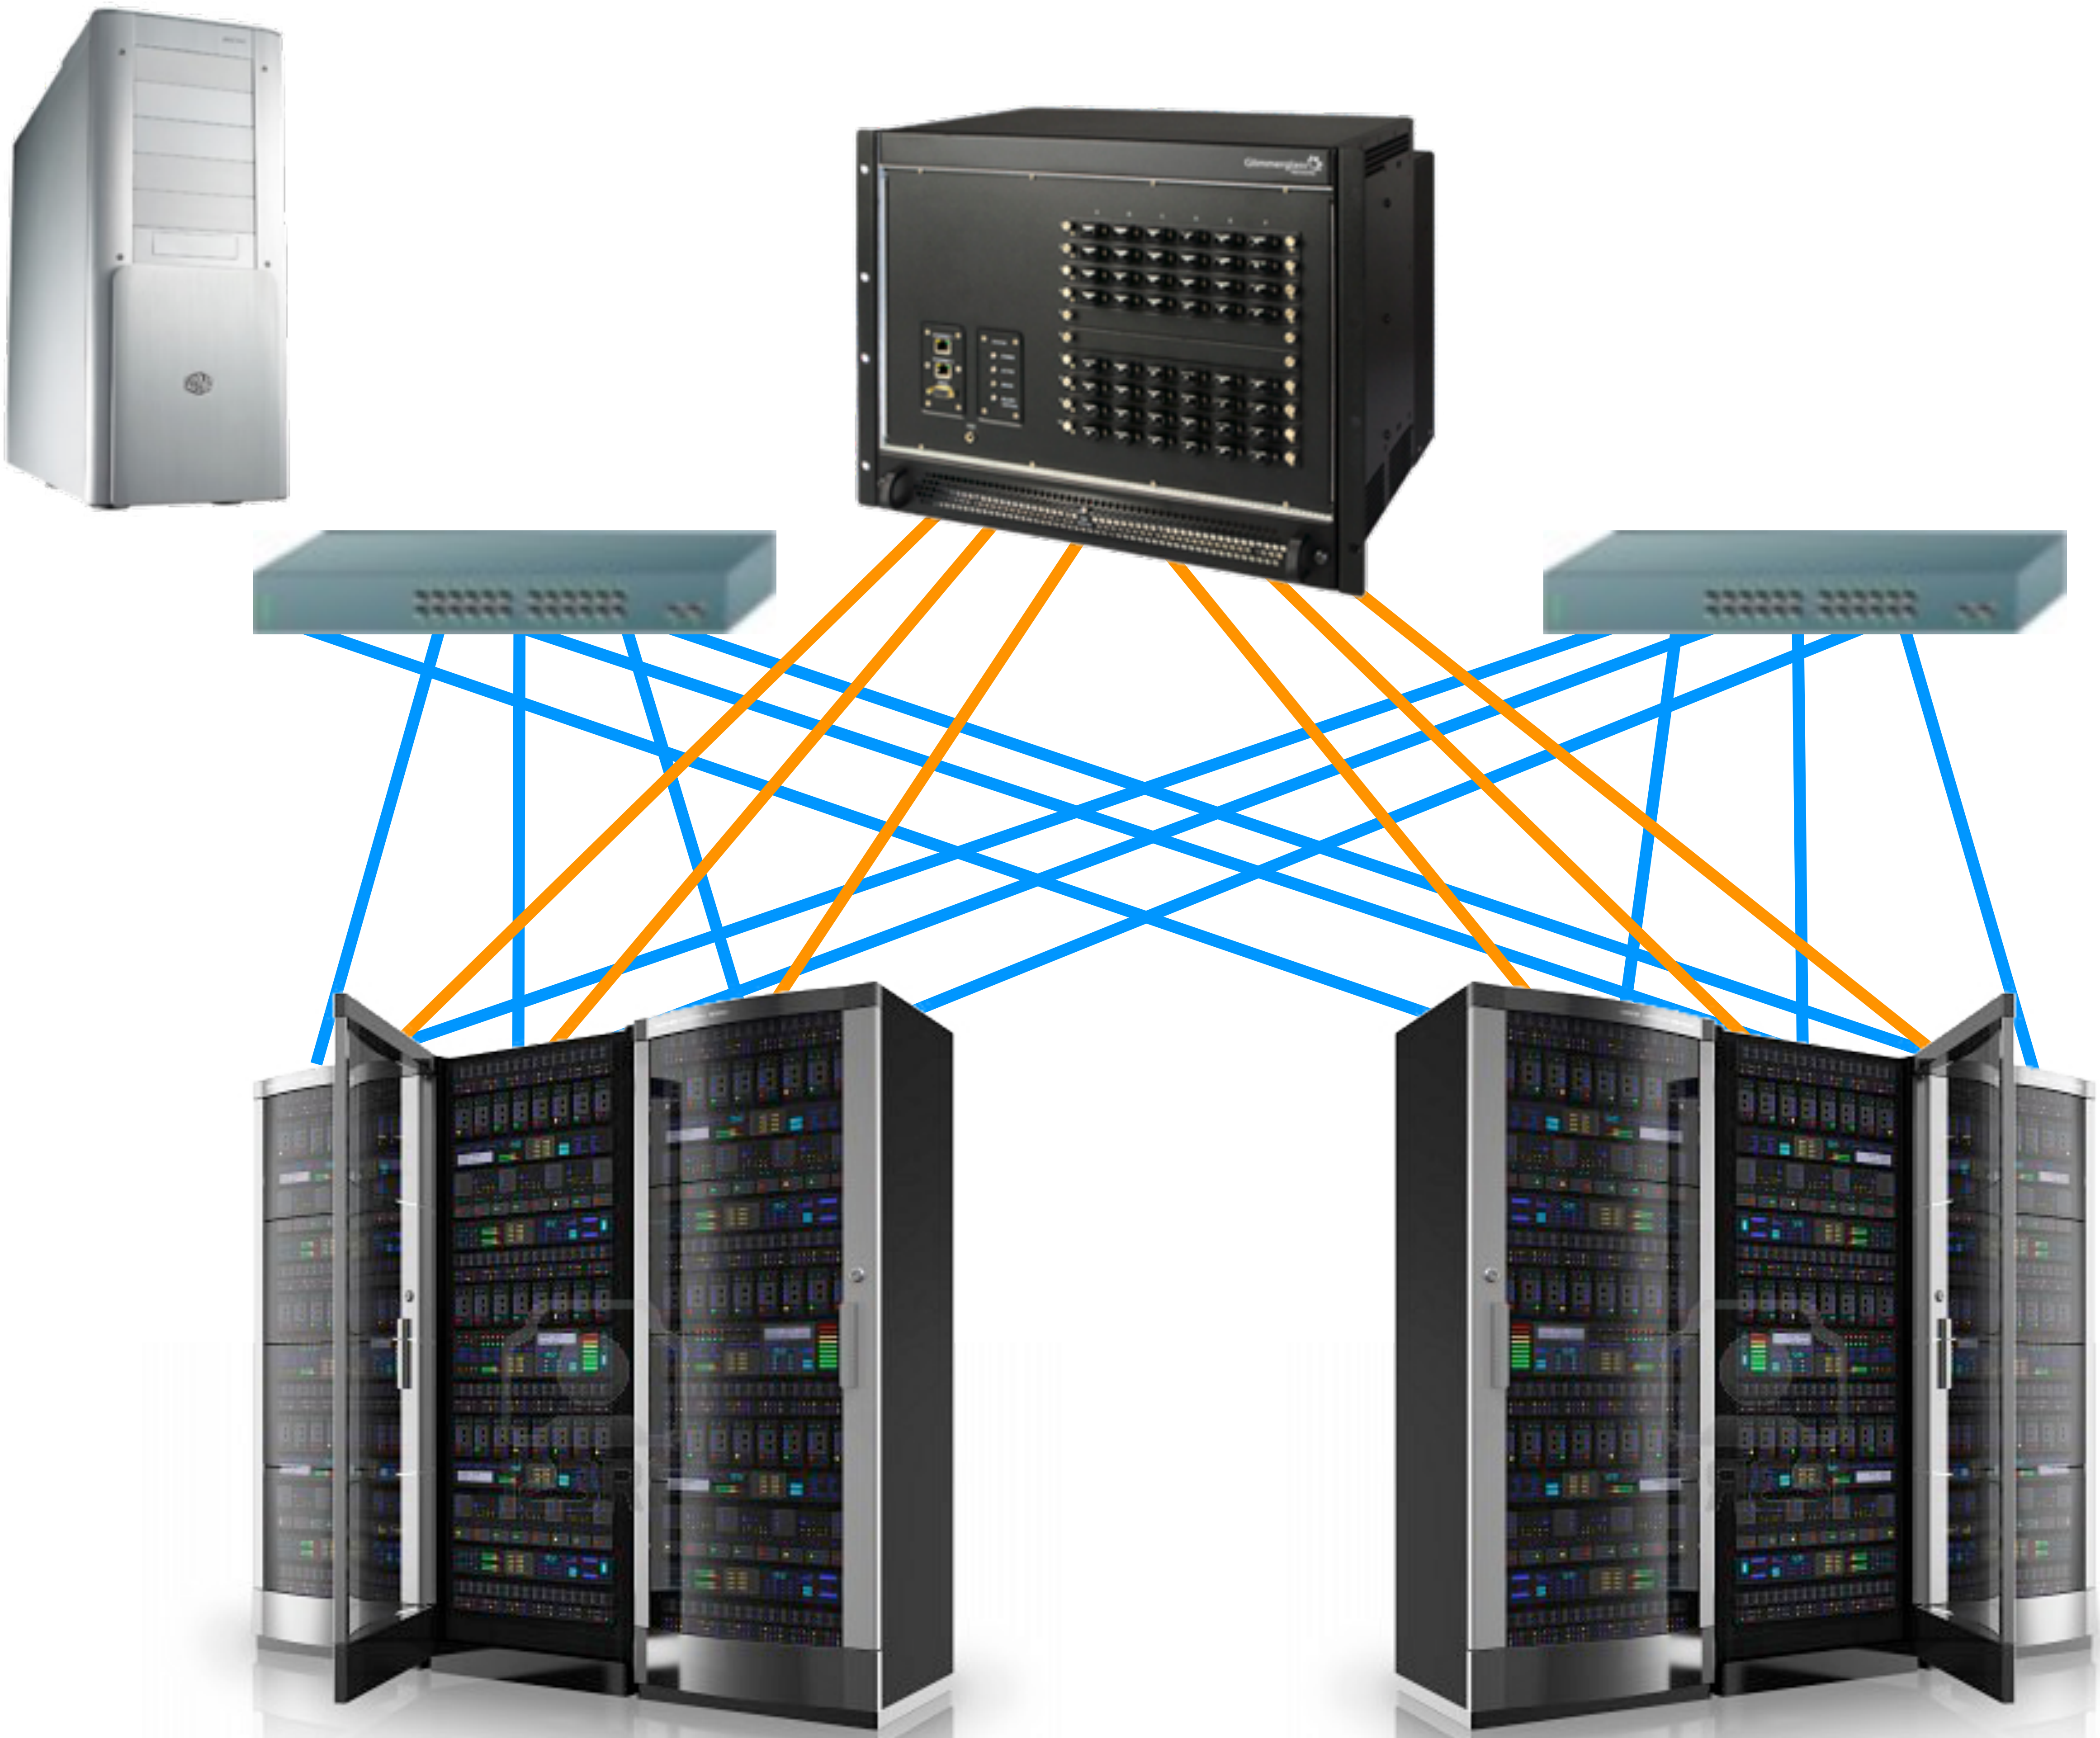

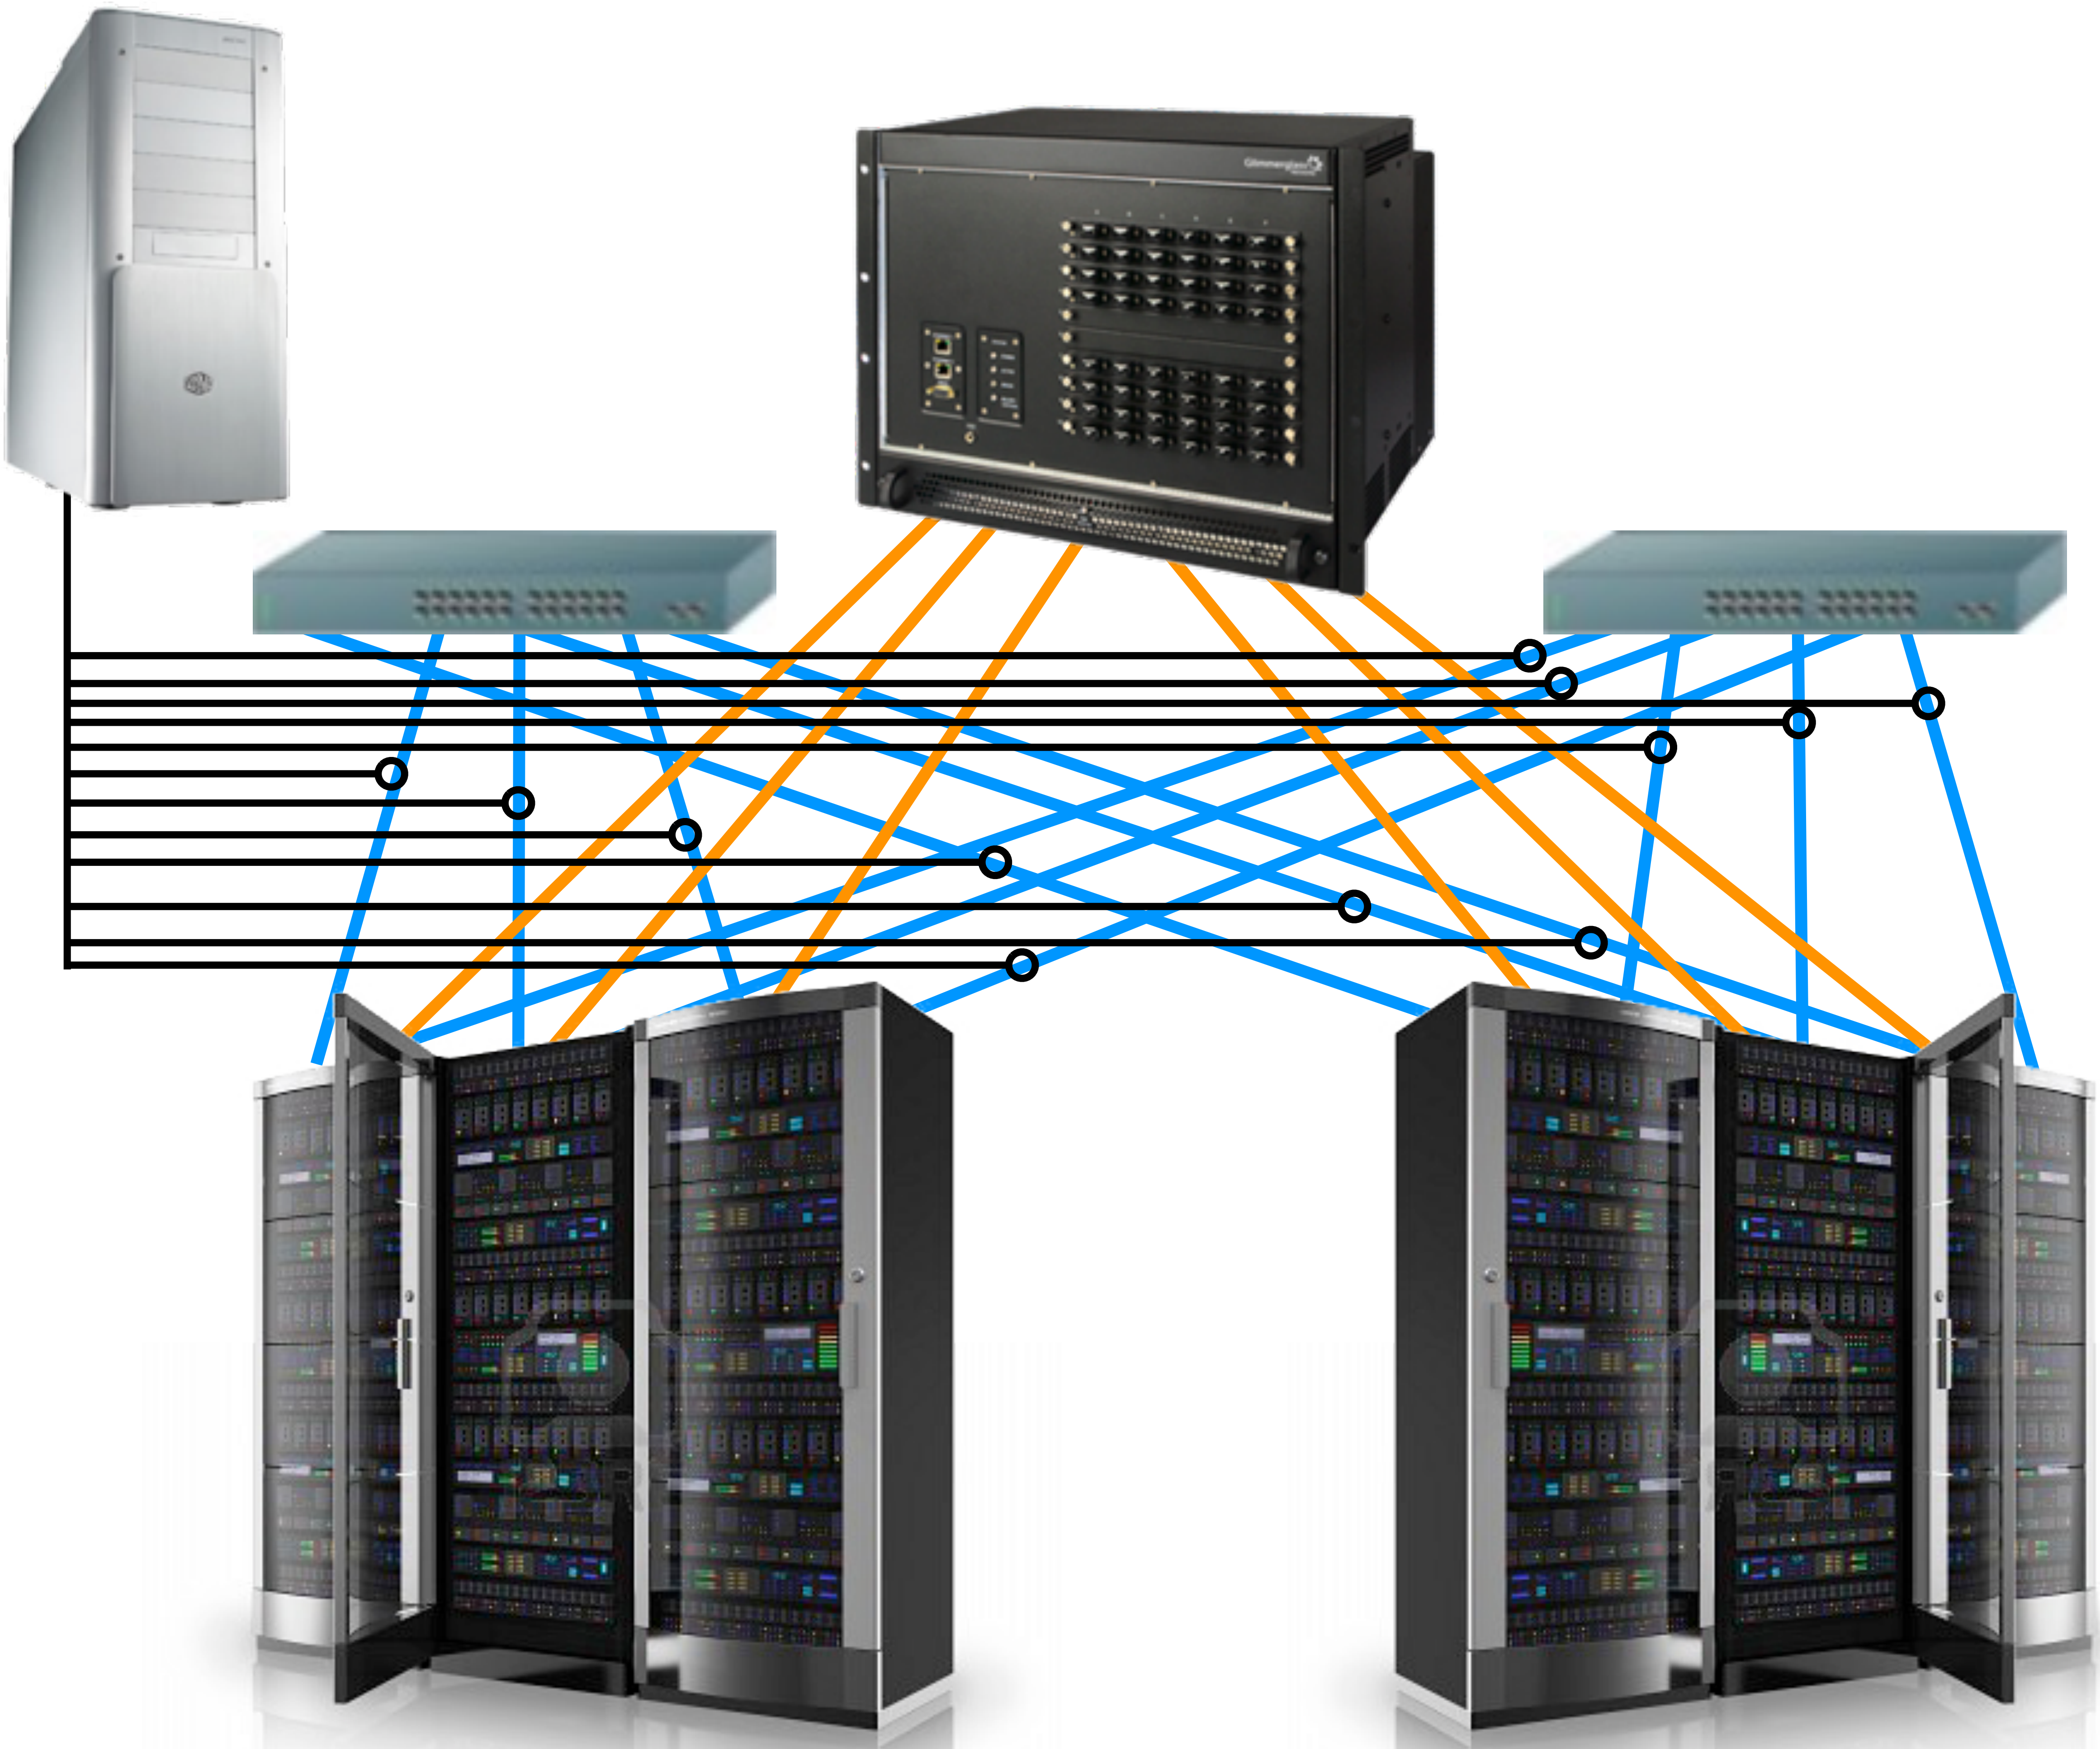

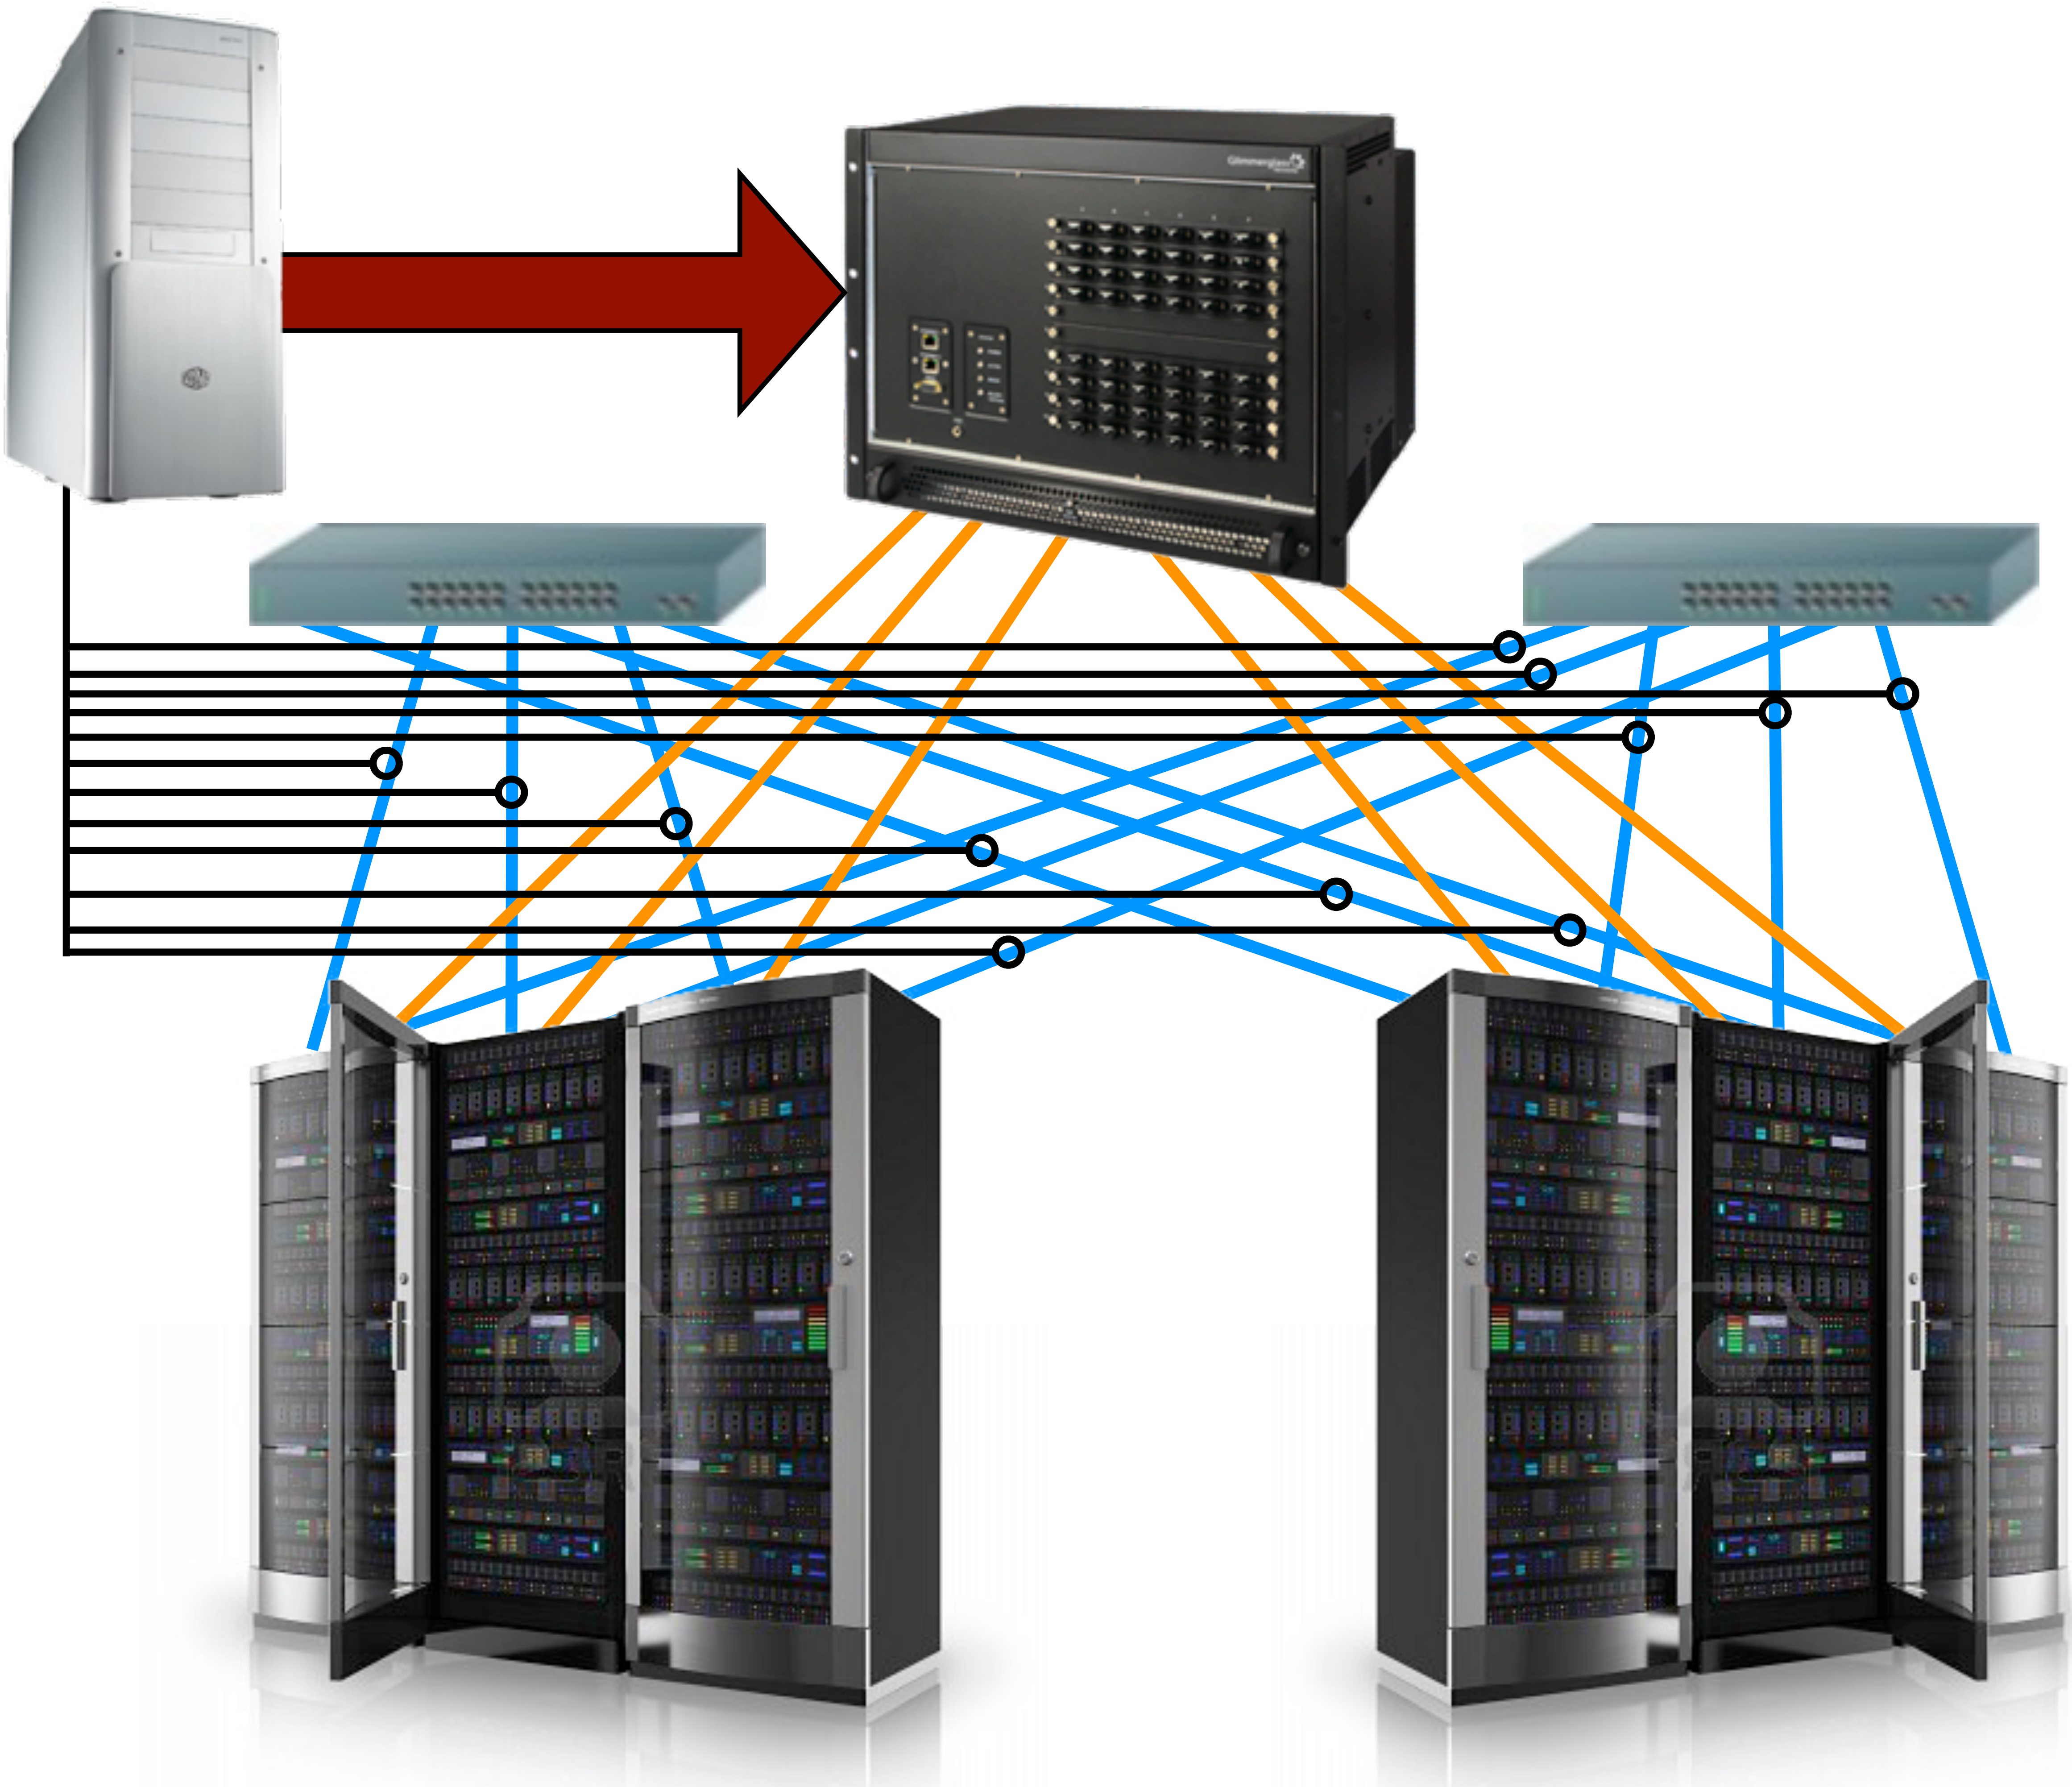

Defending from attacks

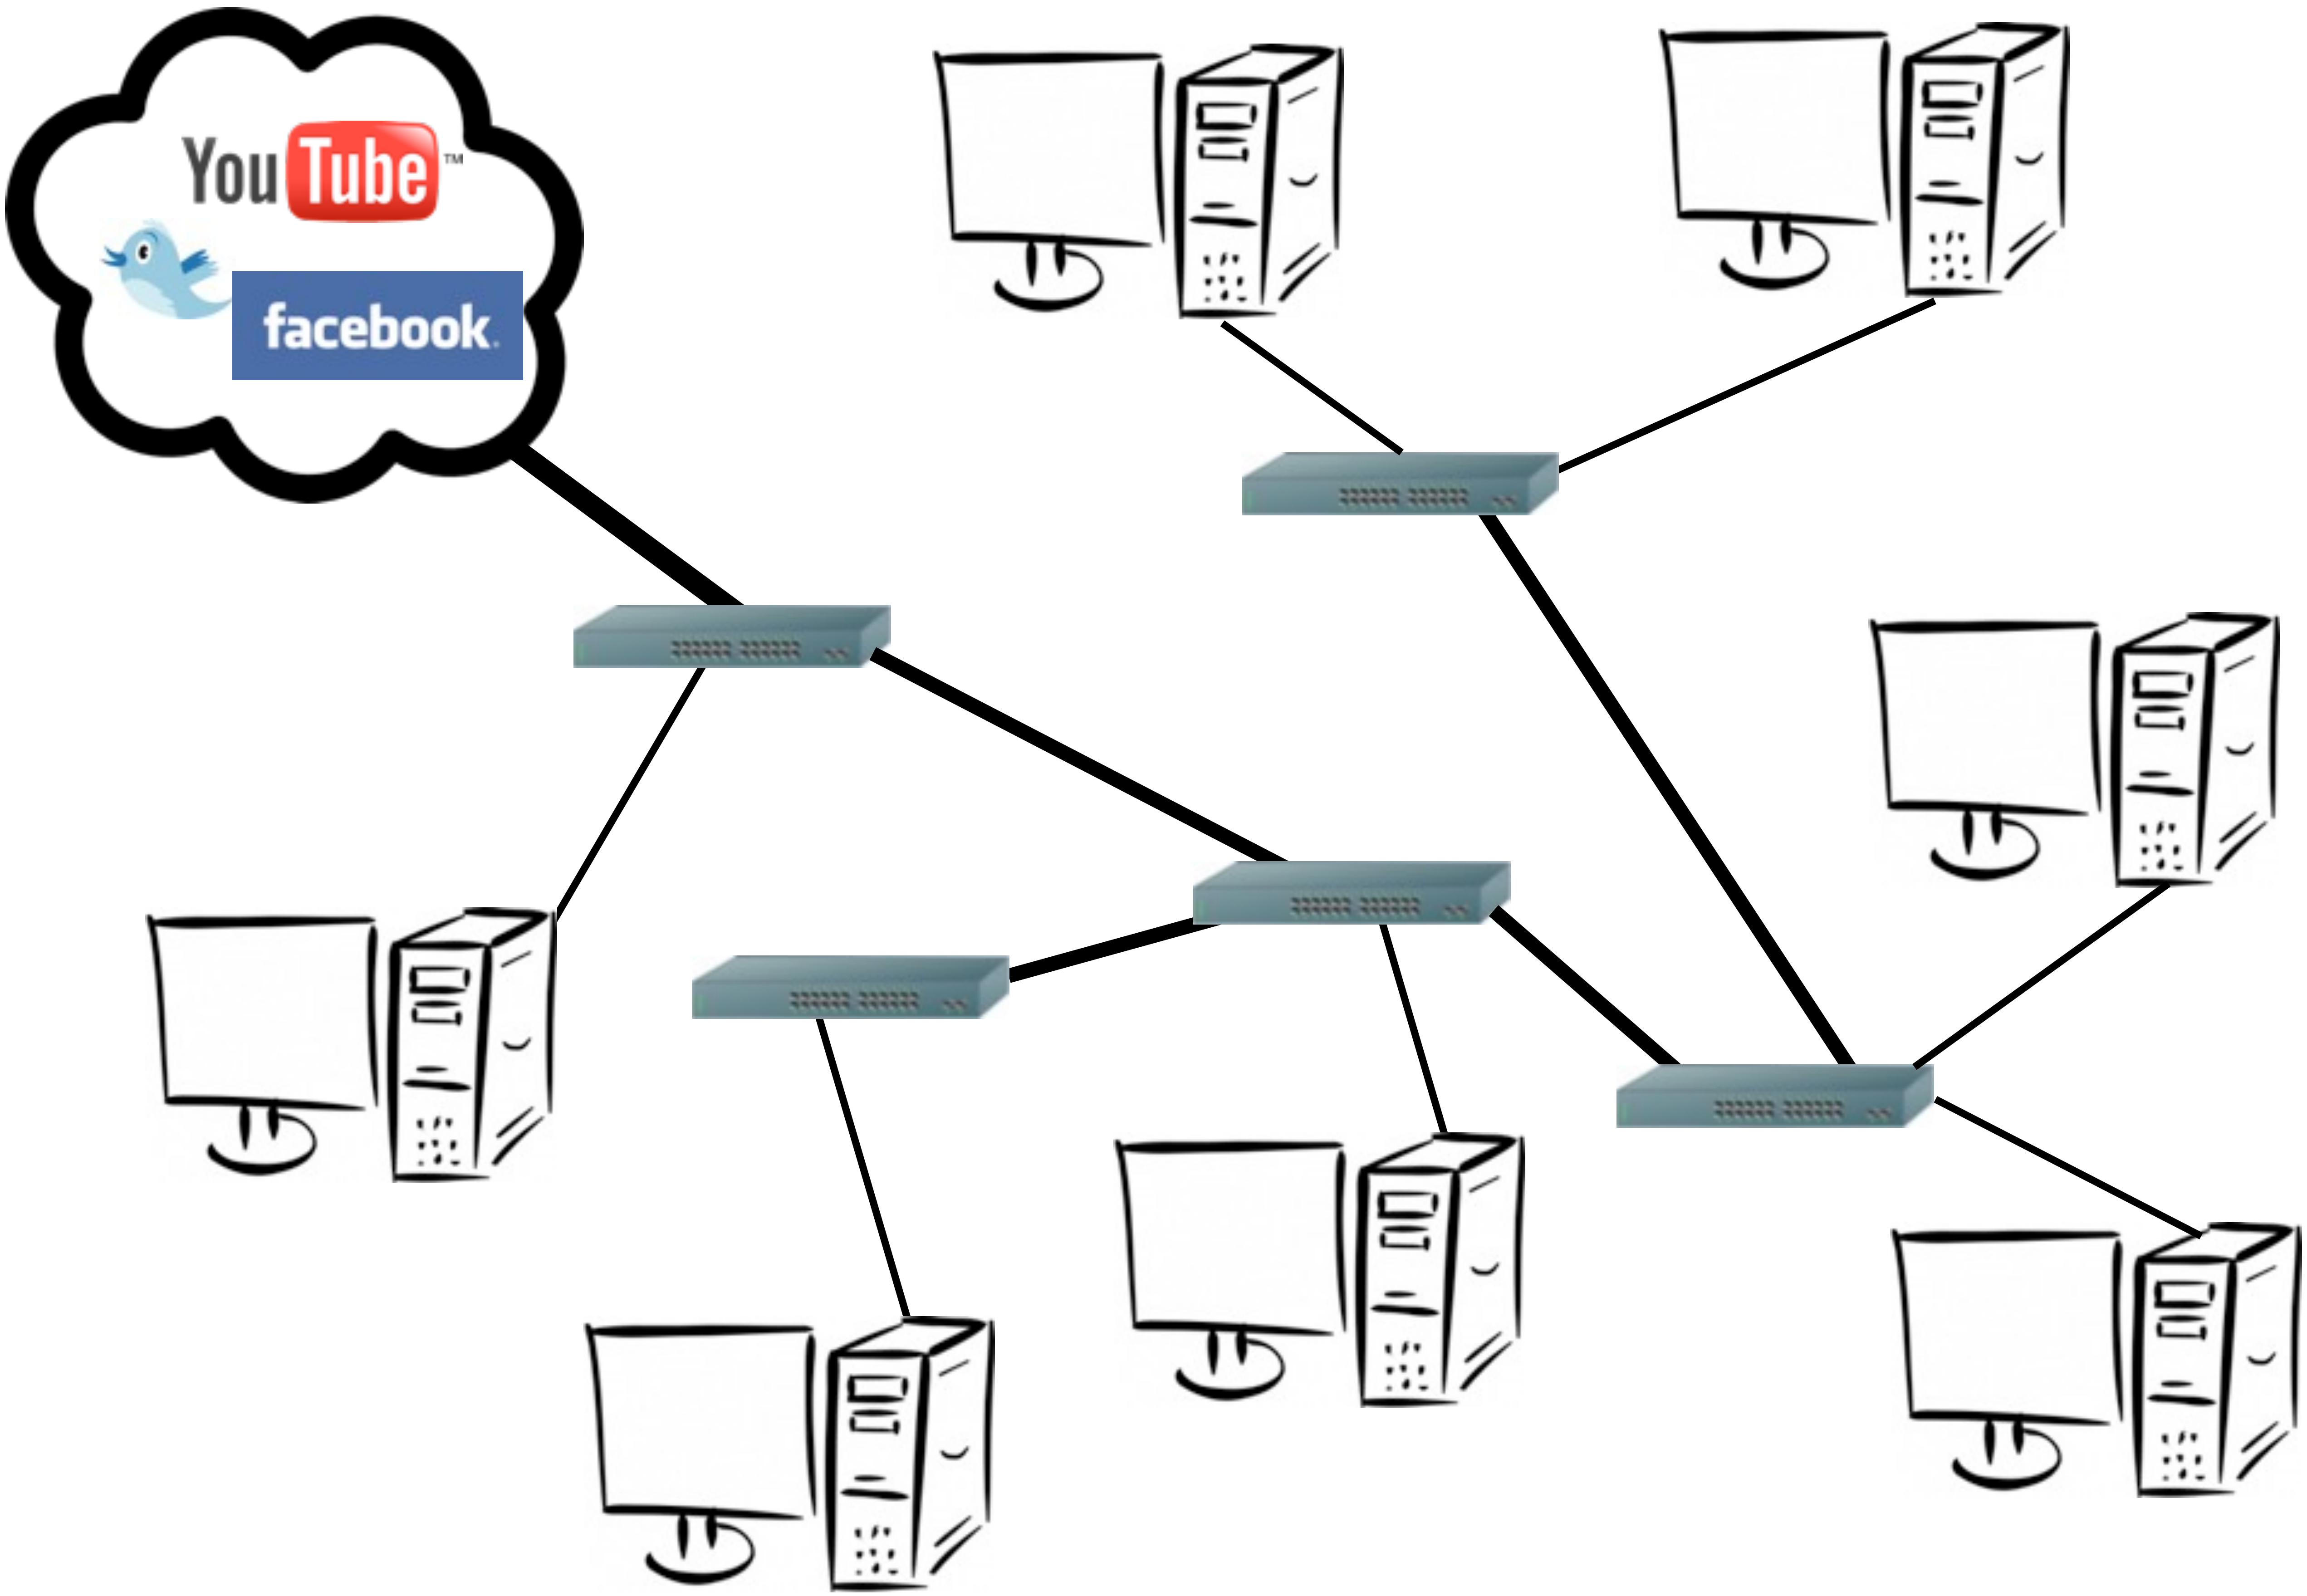

YouTube
facebook

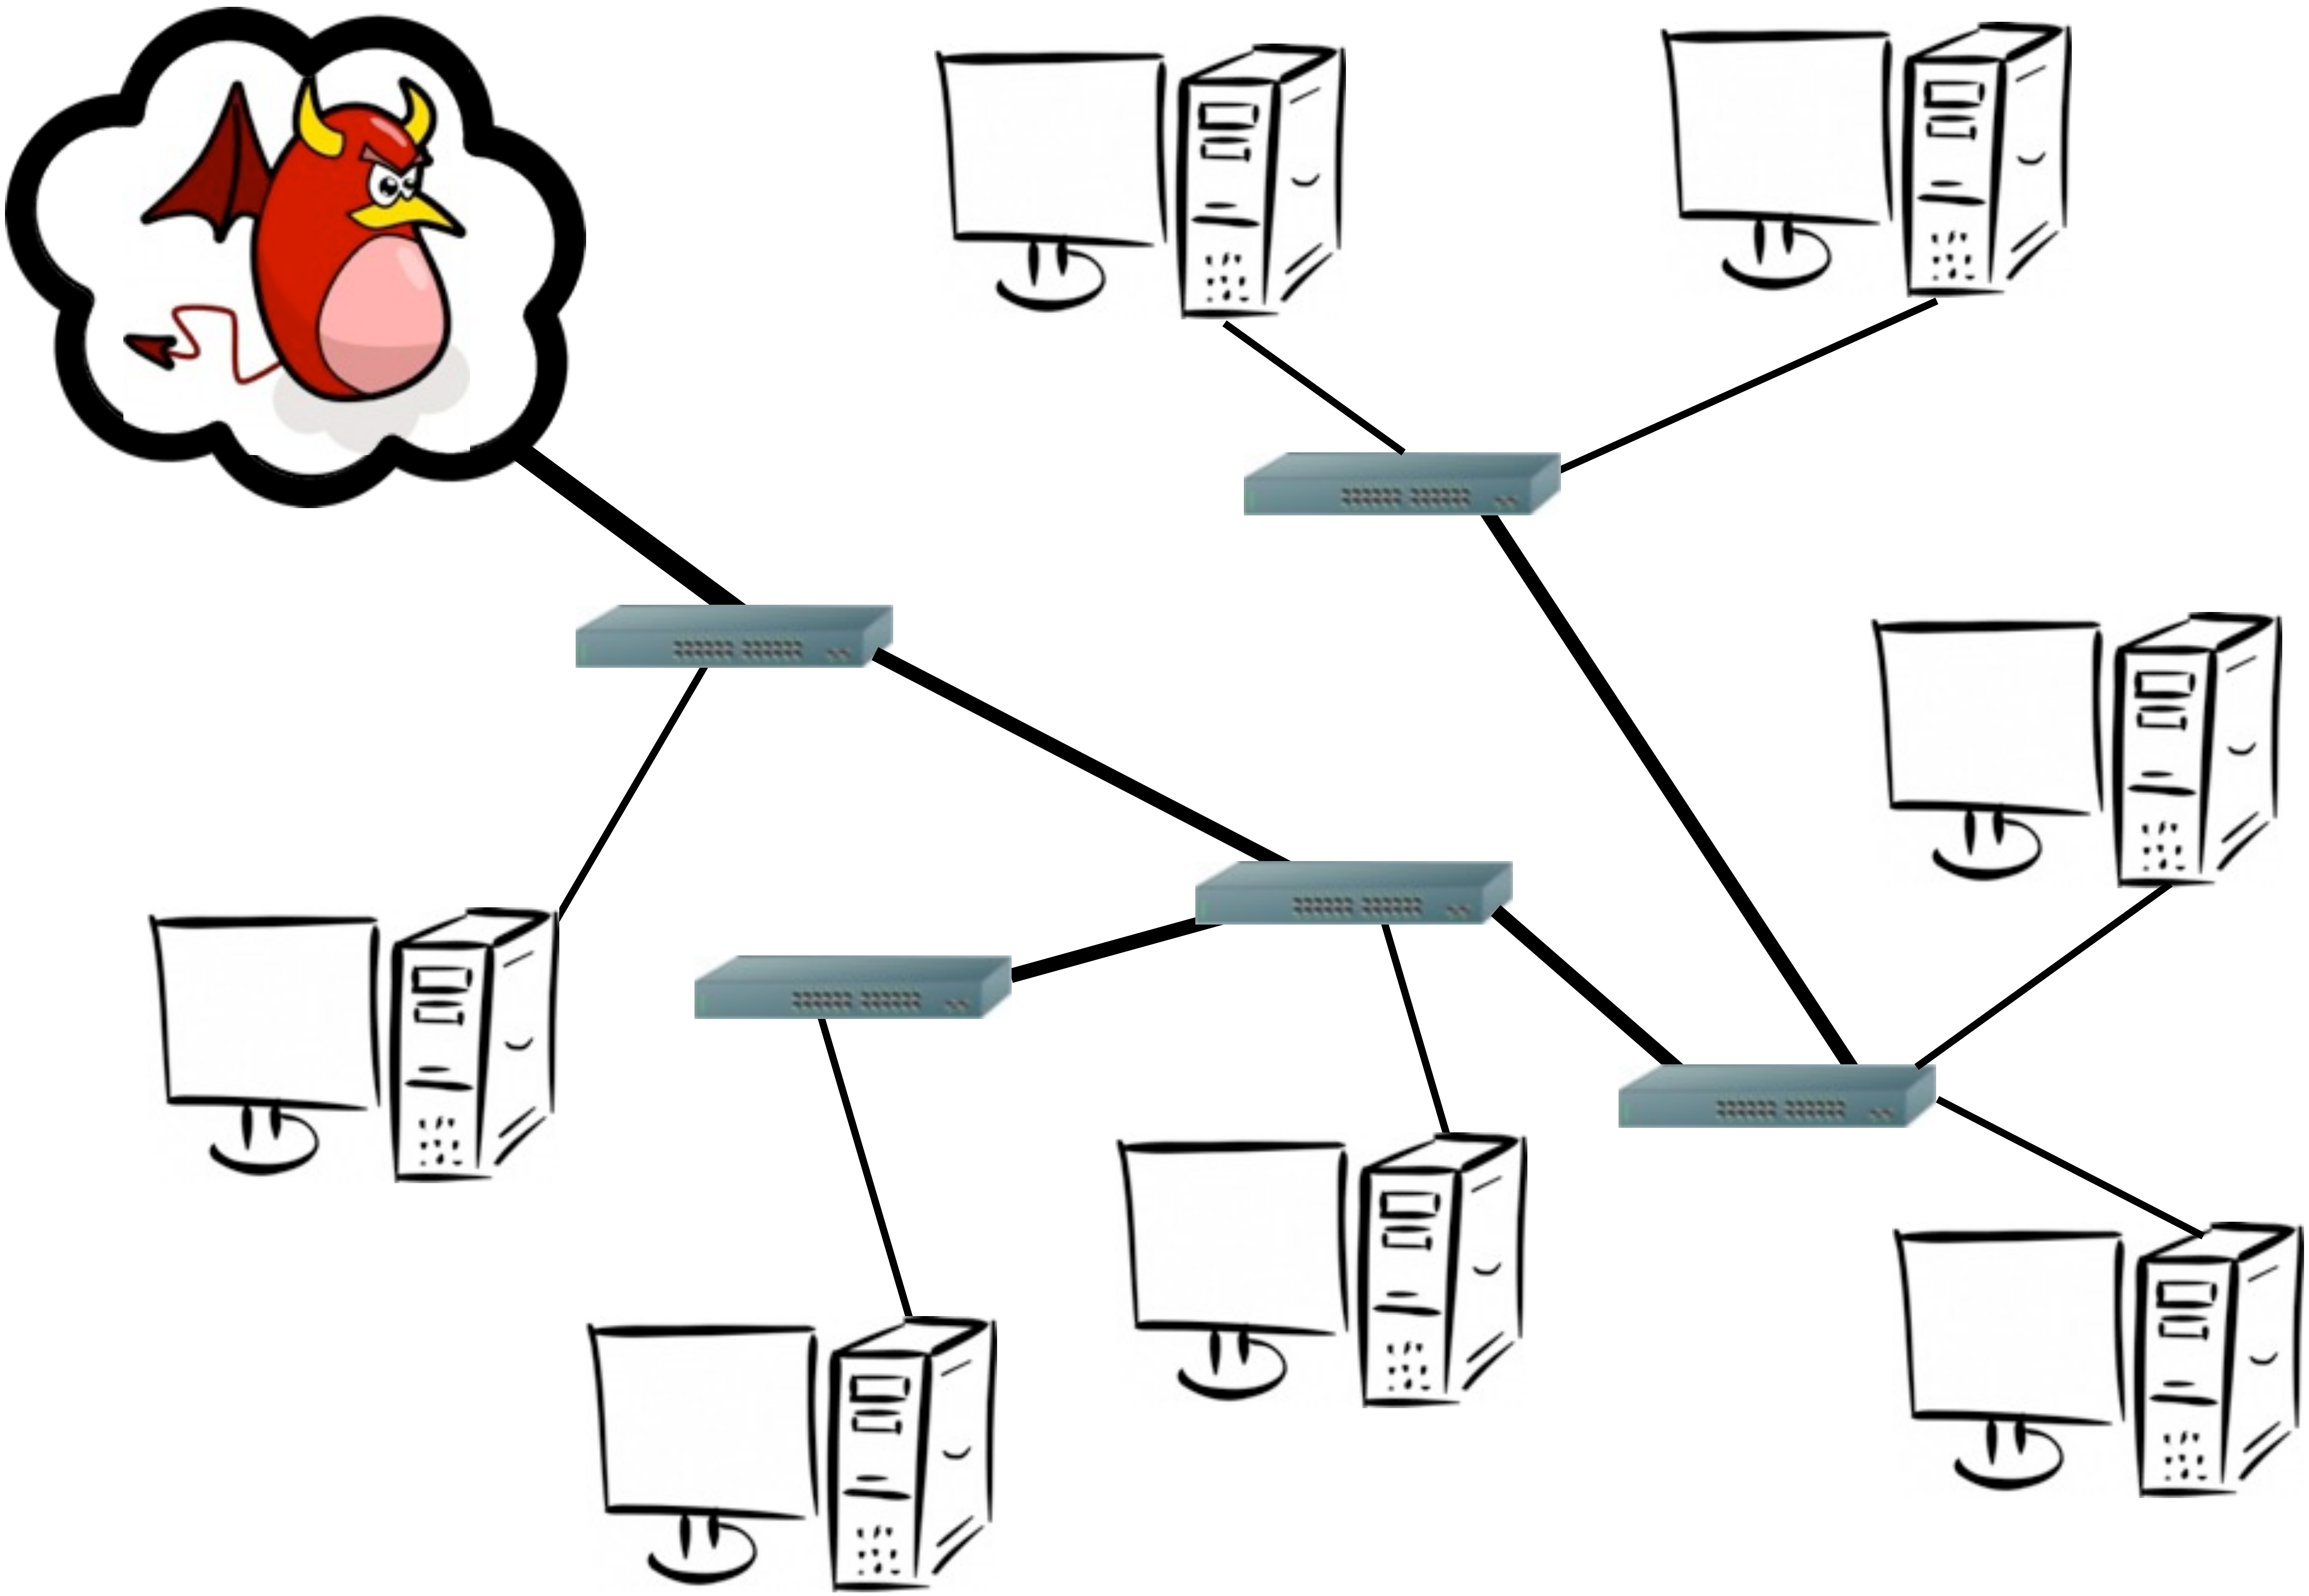

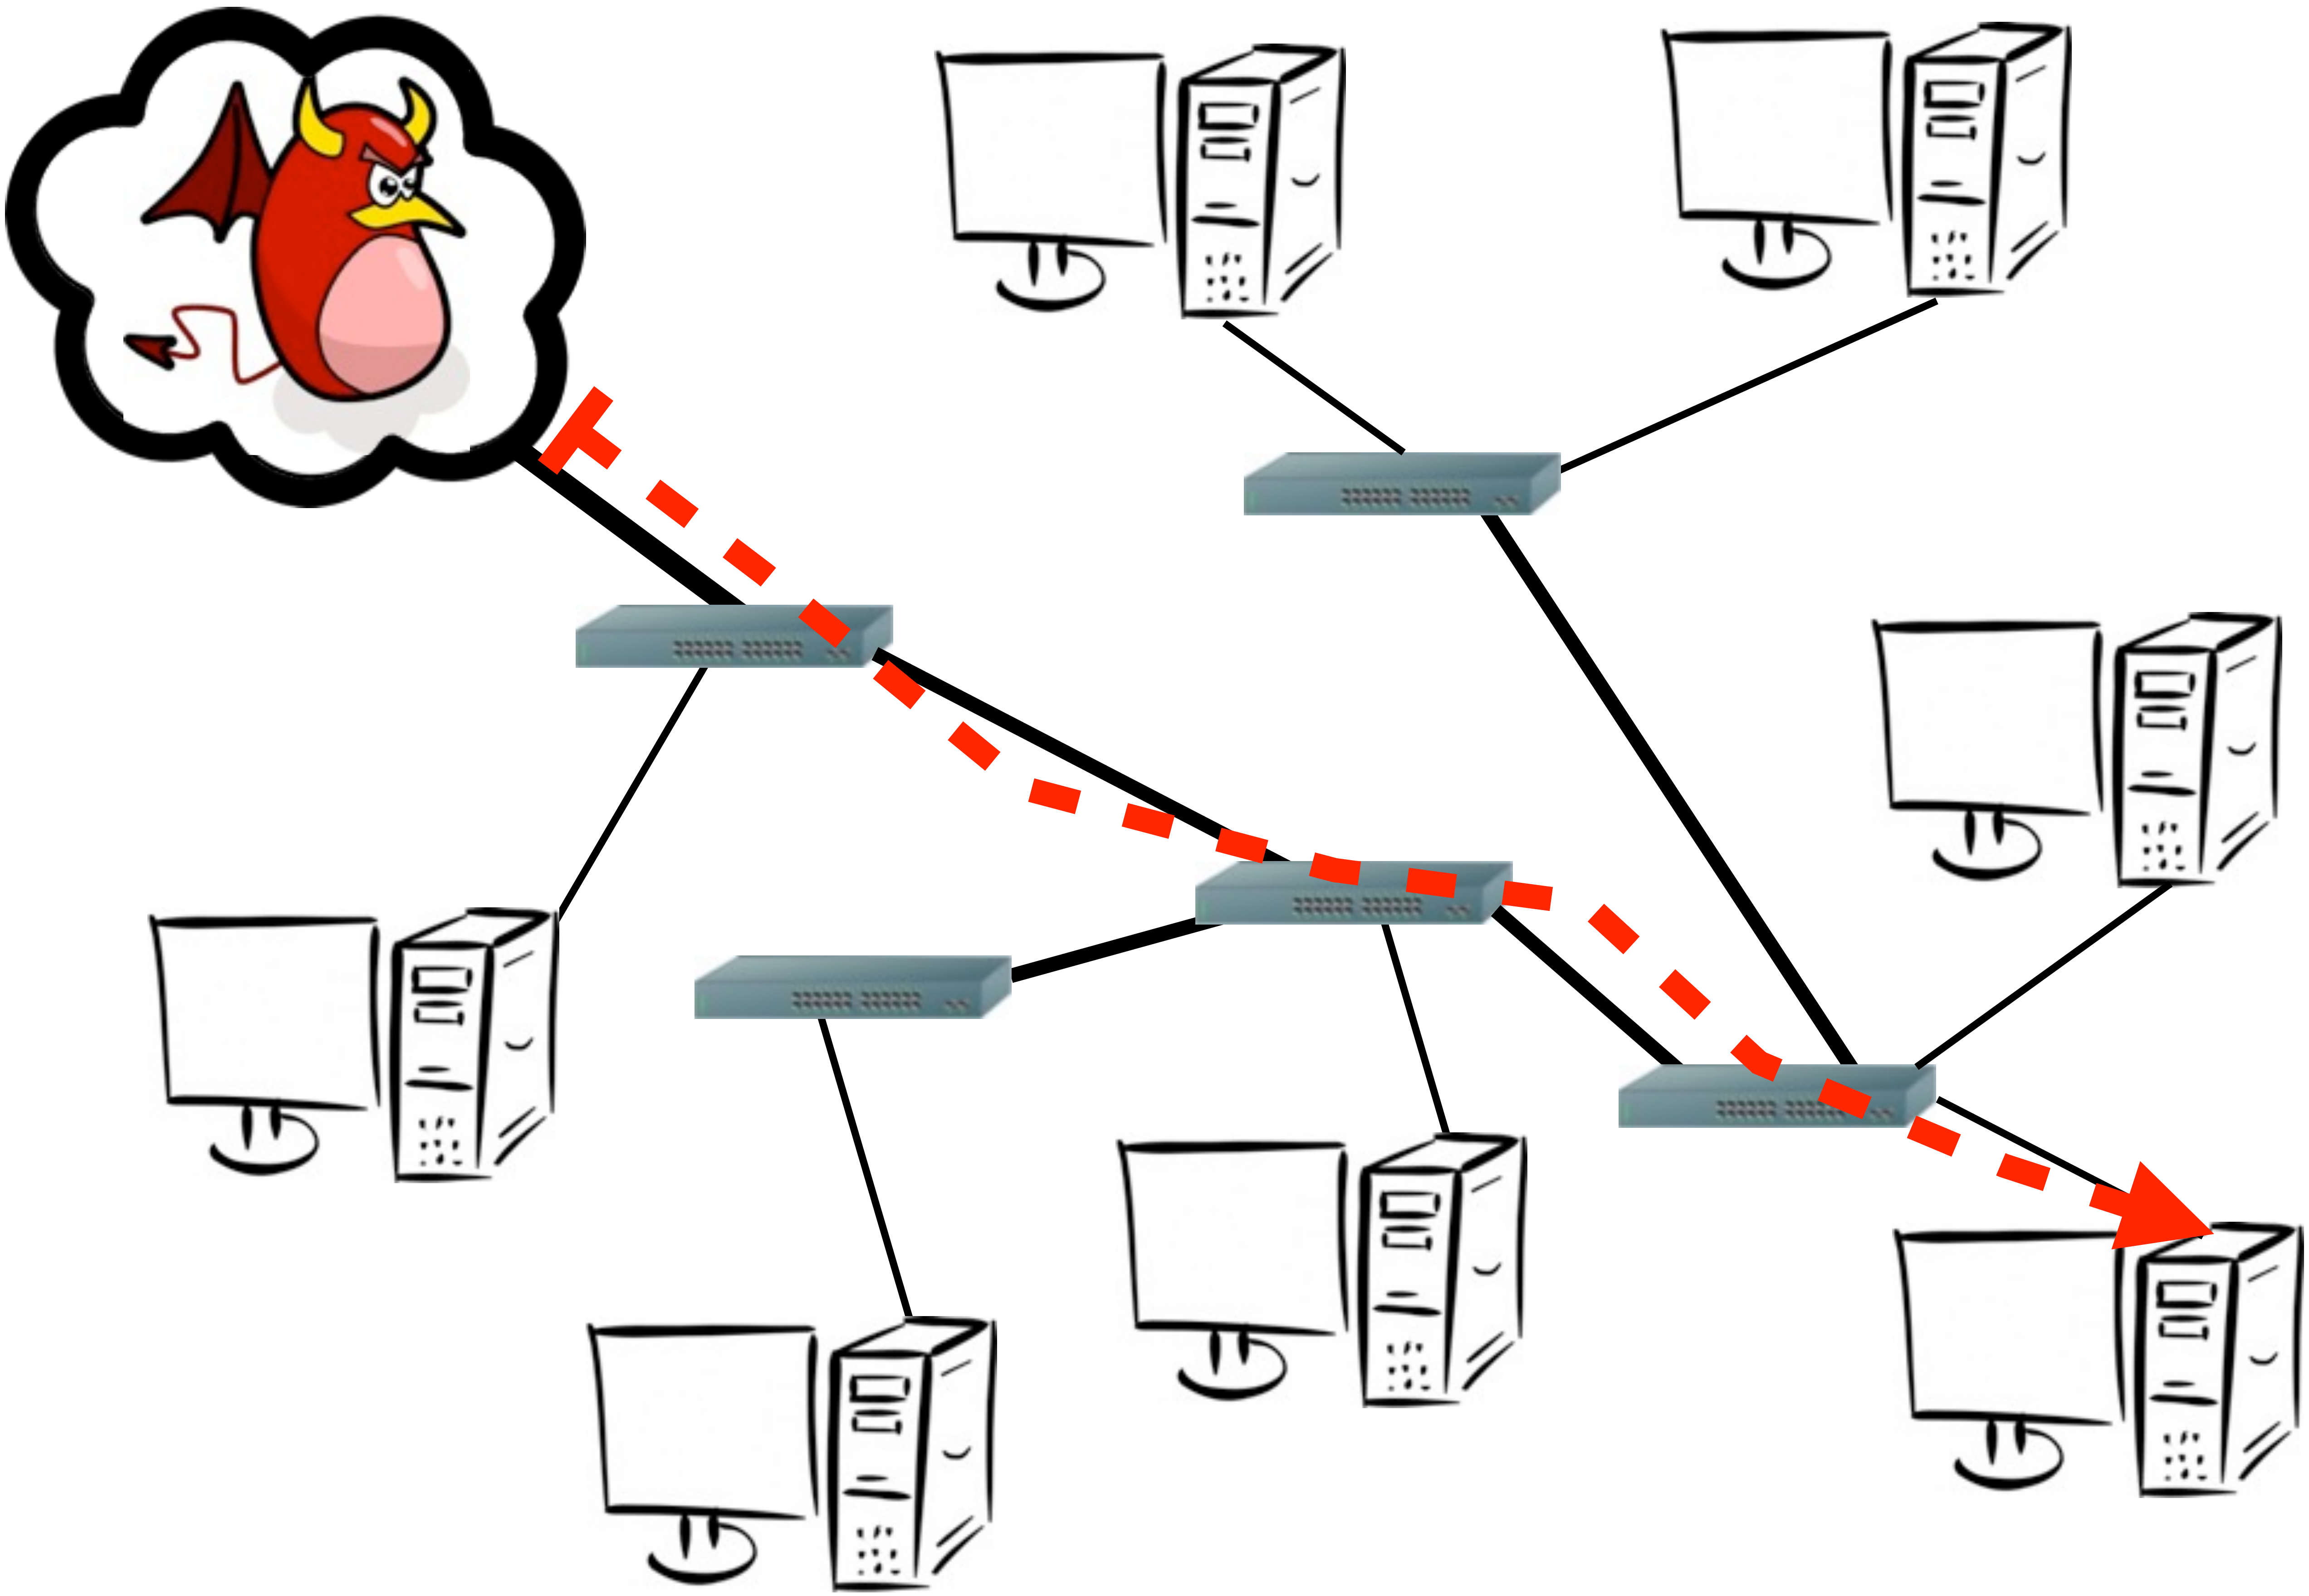

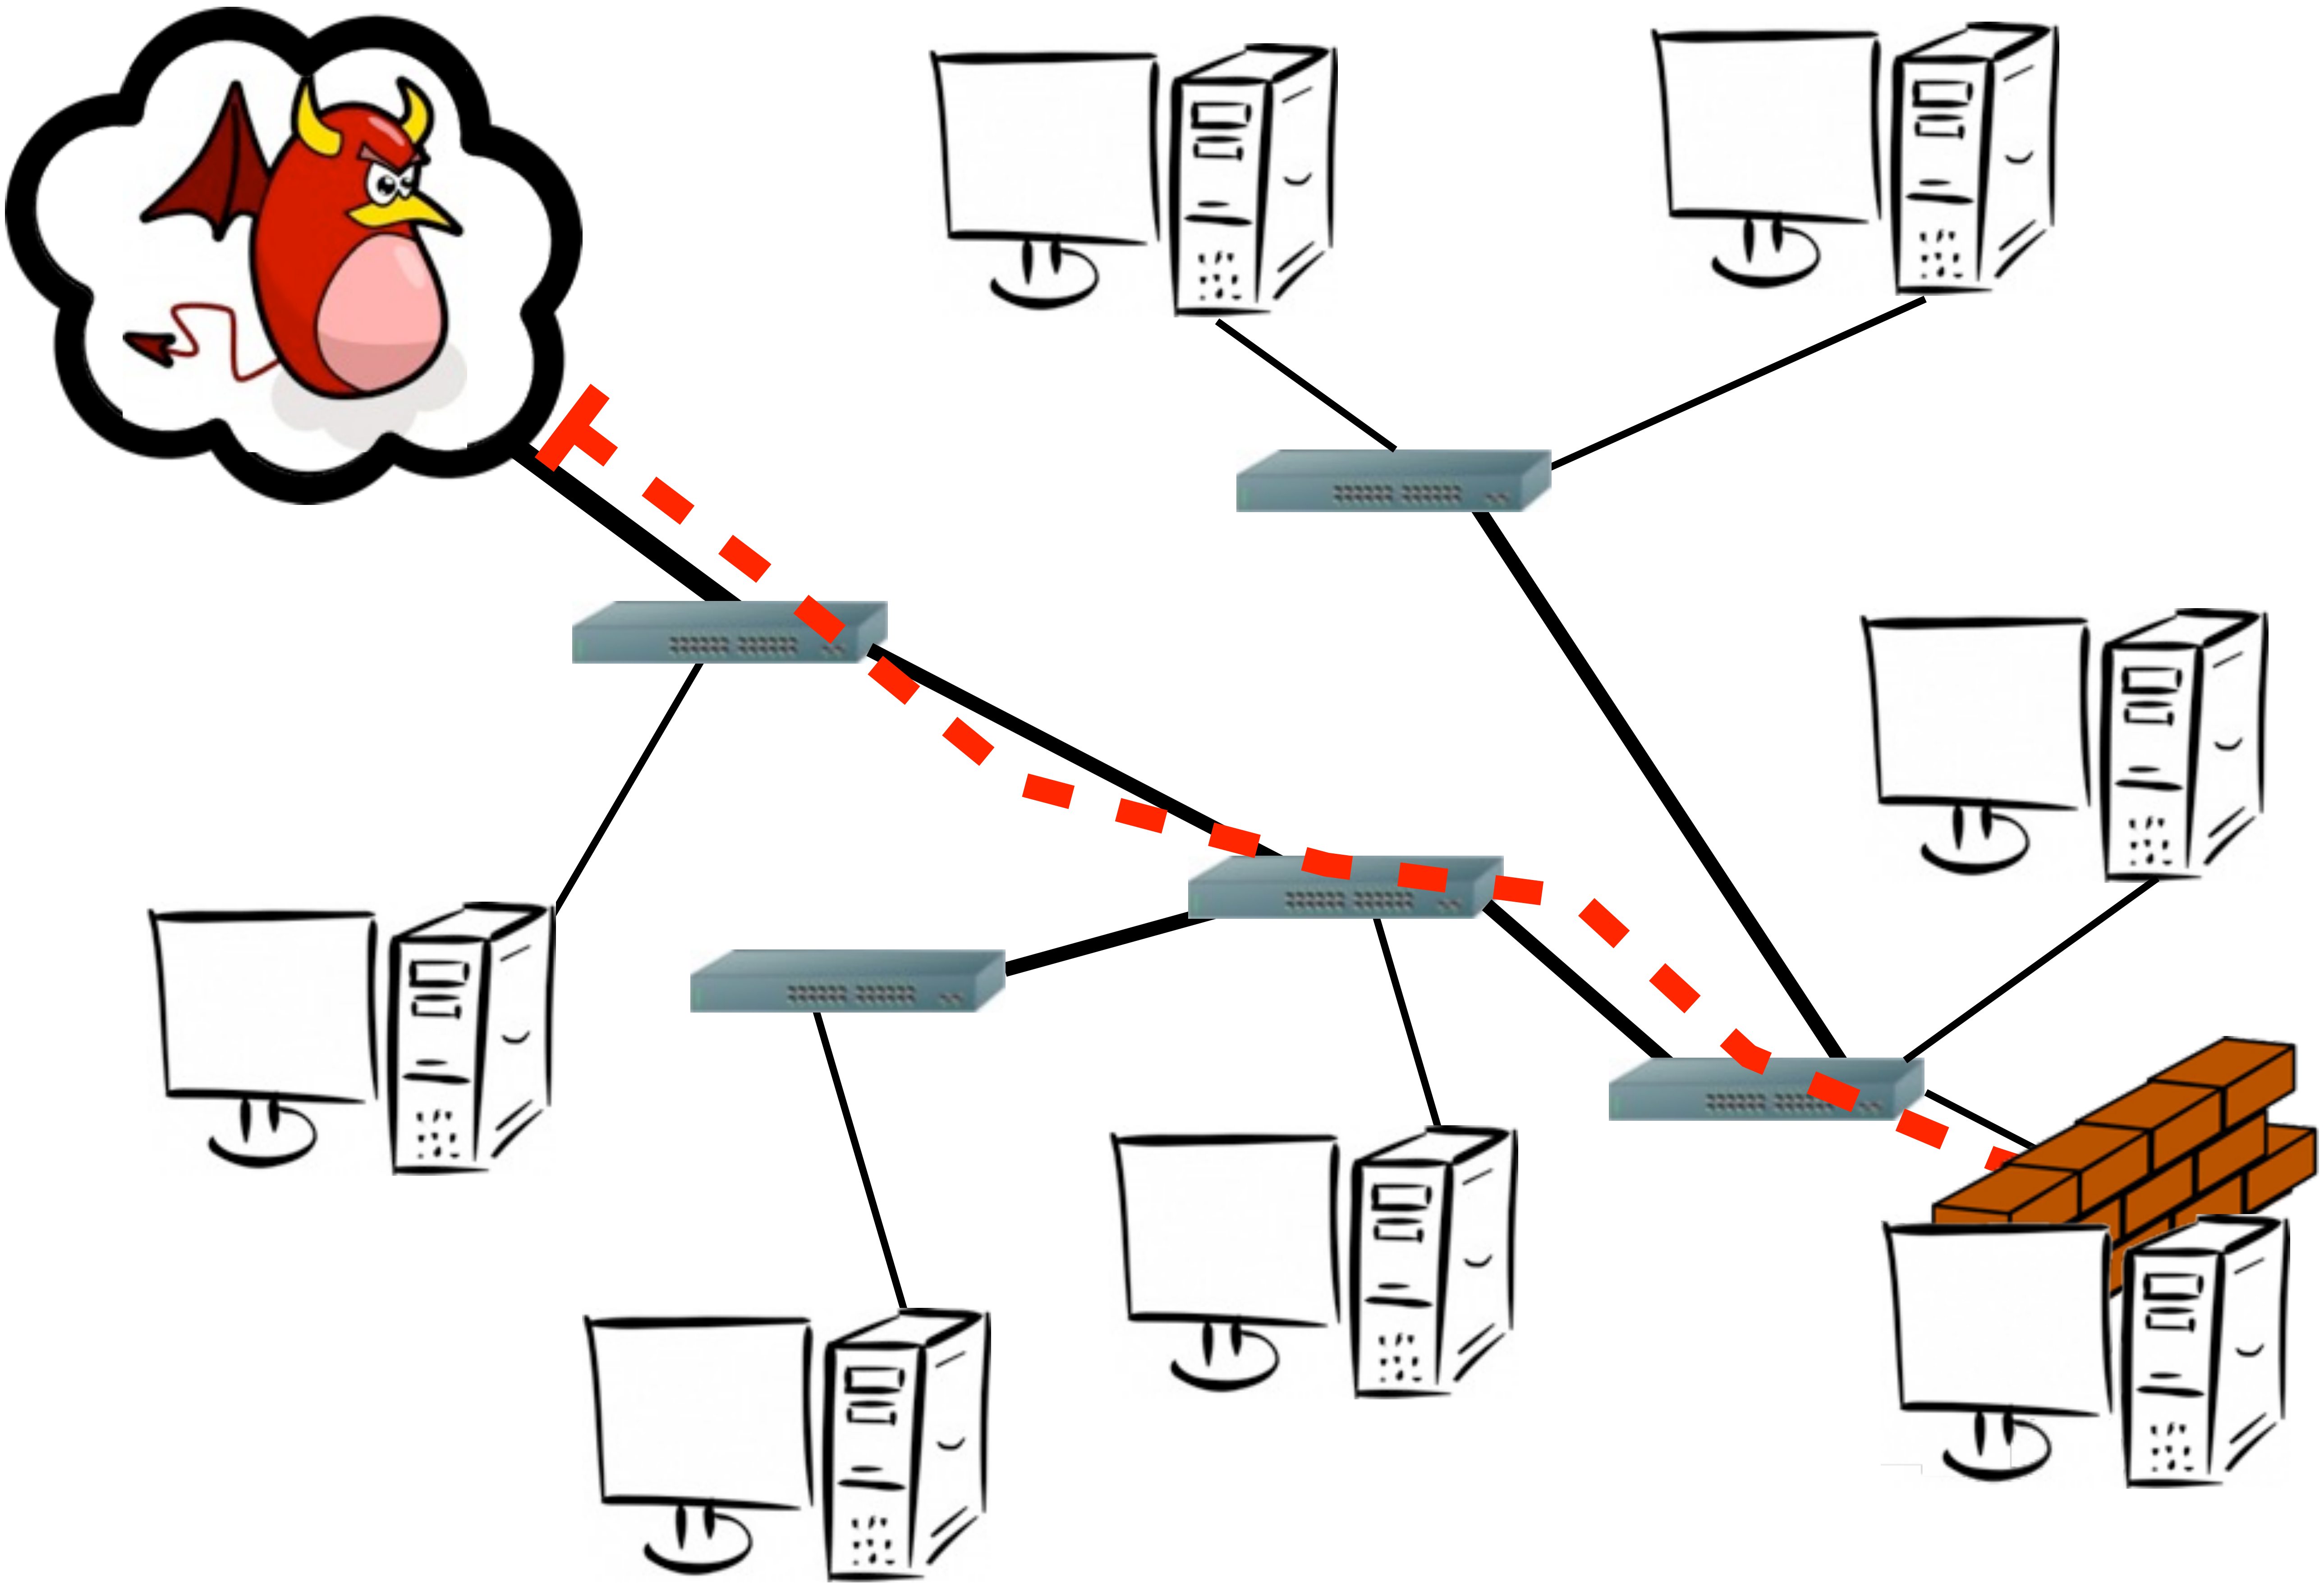

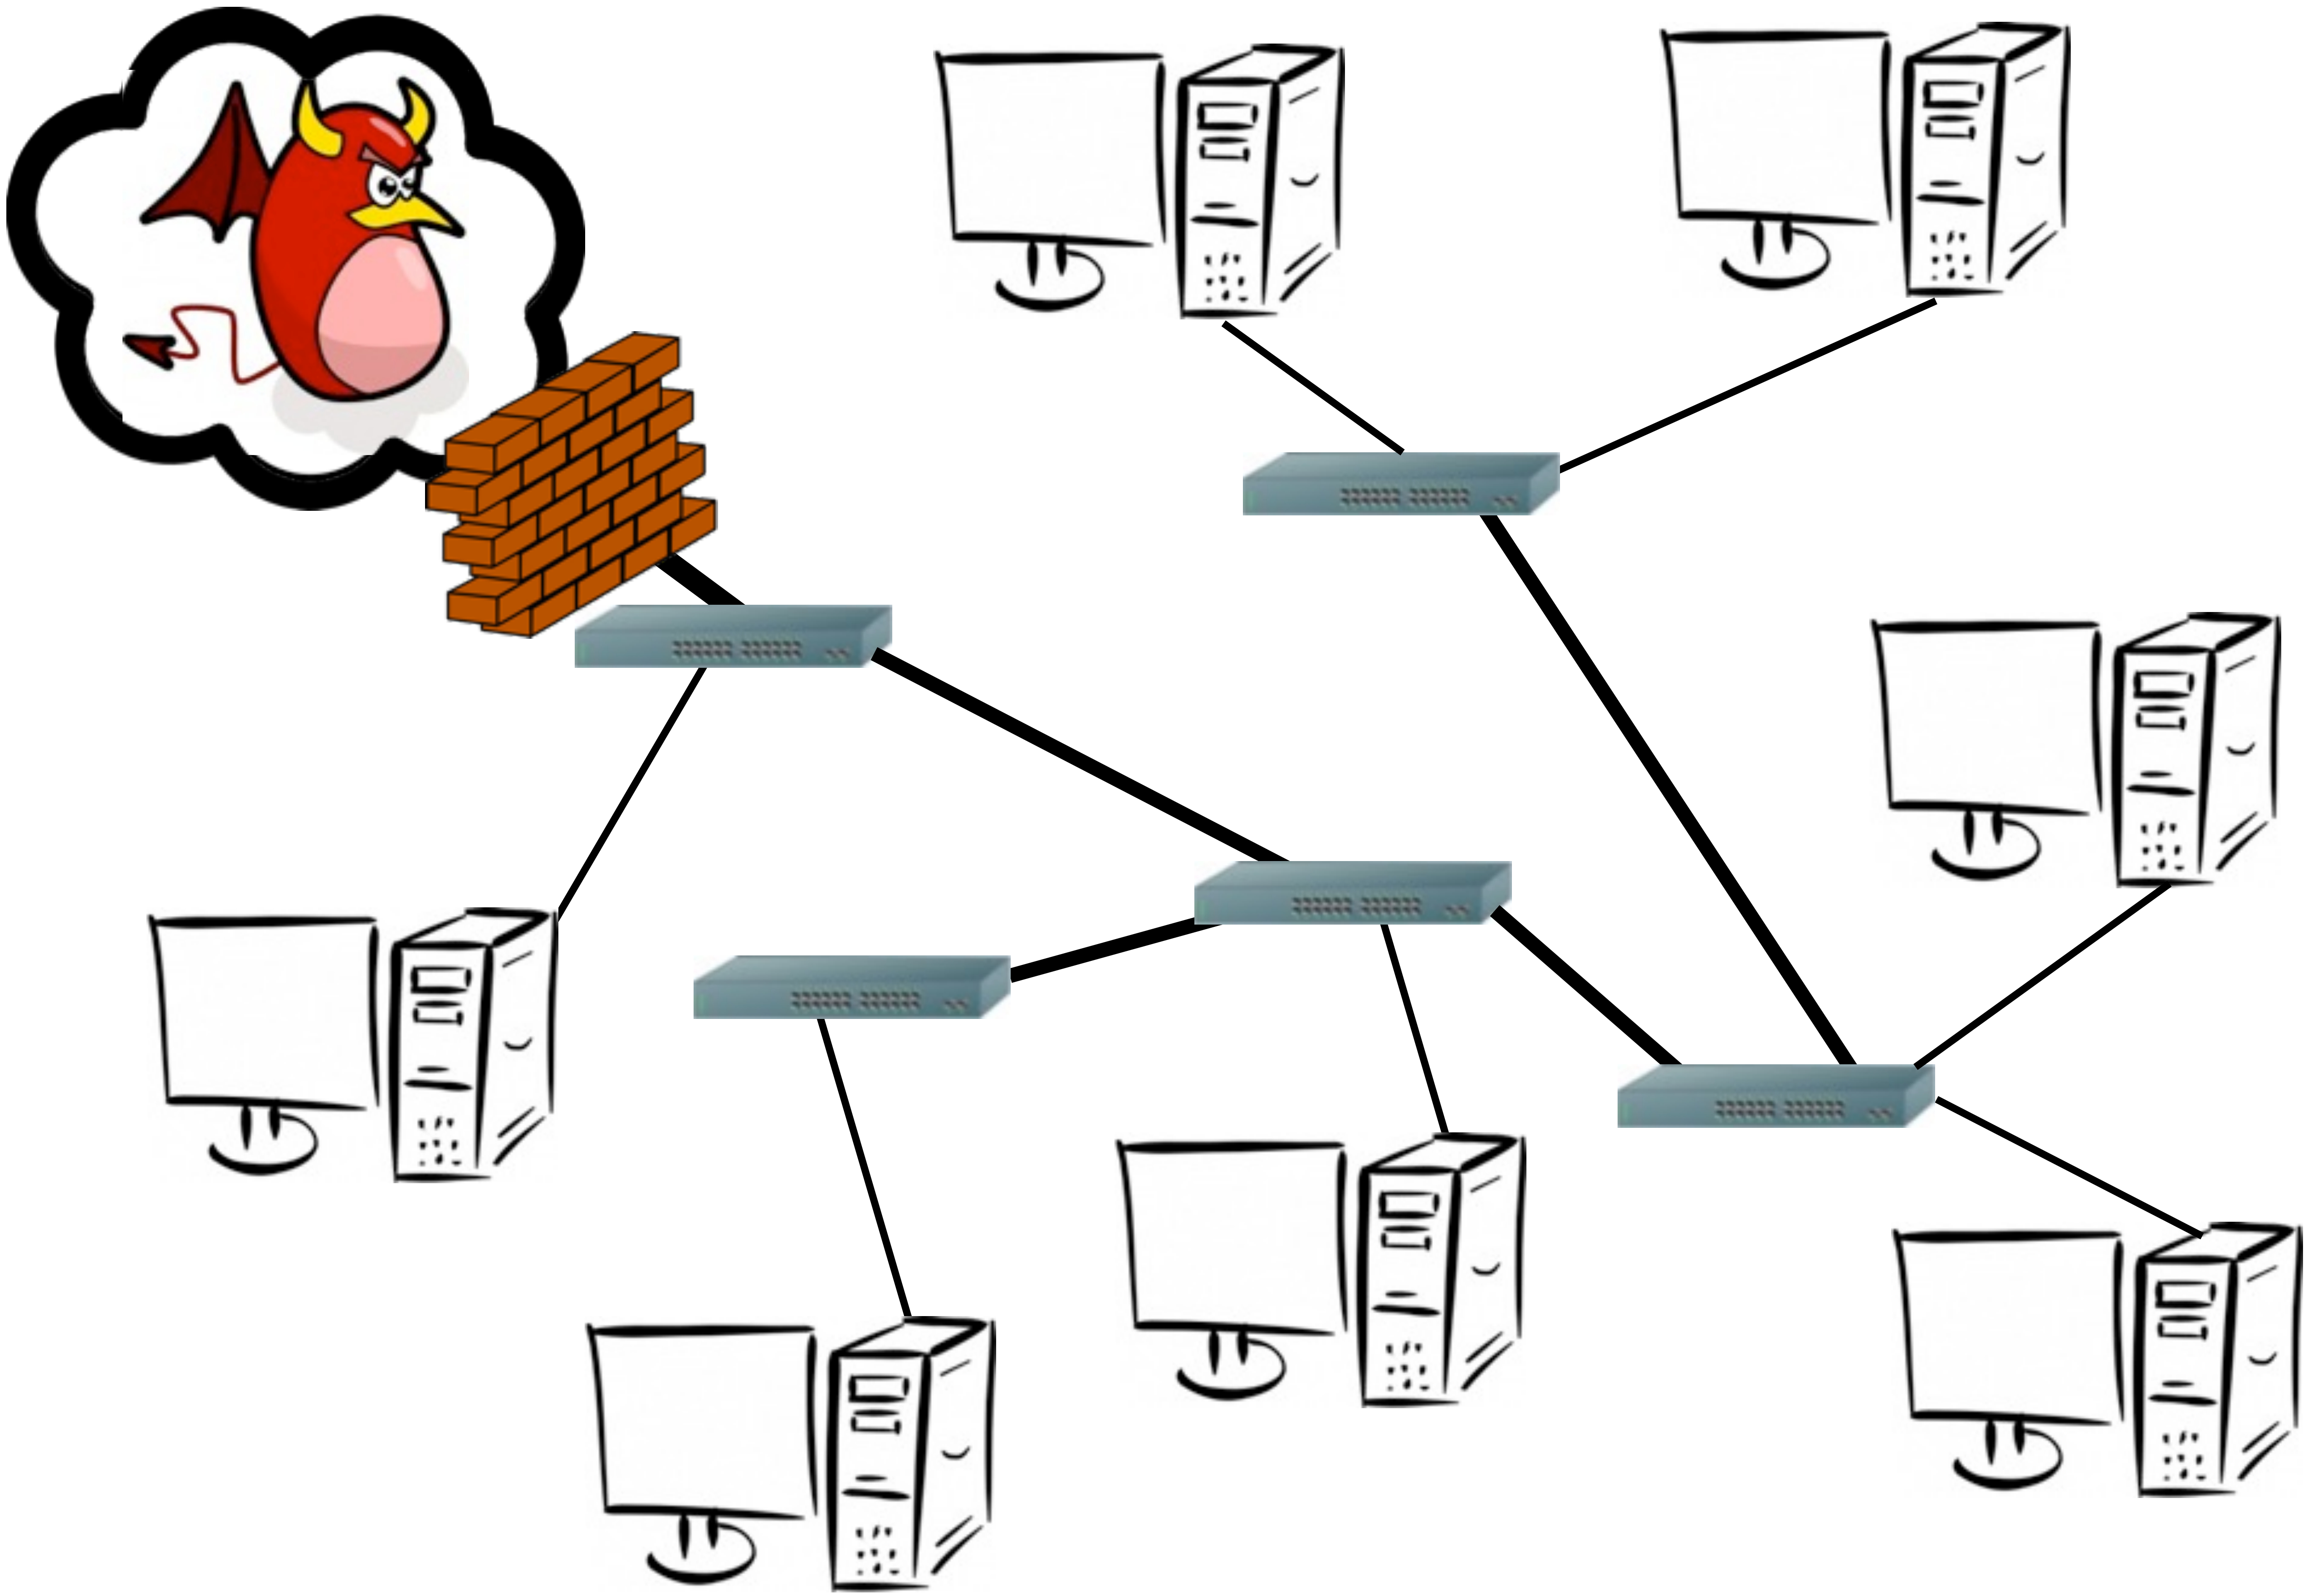

Untrusted VMs

Production Platform

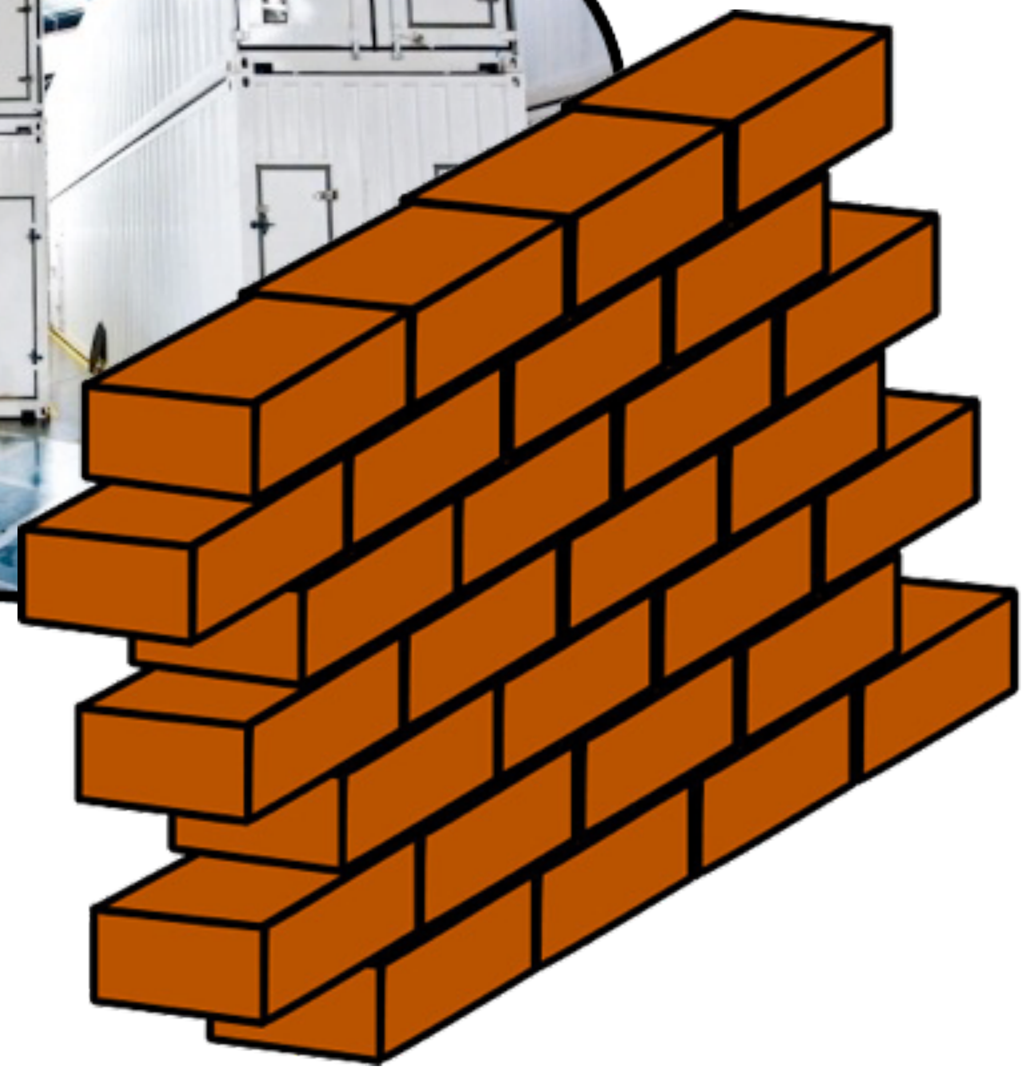

Production Platform

Based on "Delusional Boot: Securing Cloud Hypervisors without Massive Re-Engineering" (EuroSys 2012)

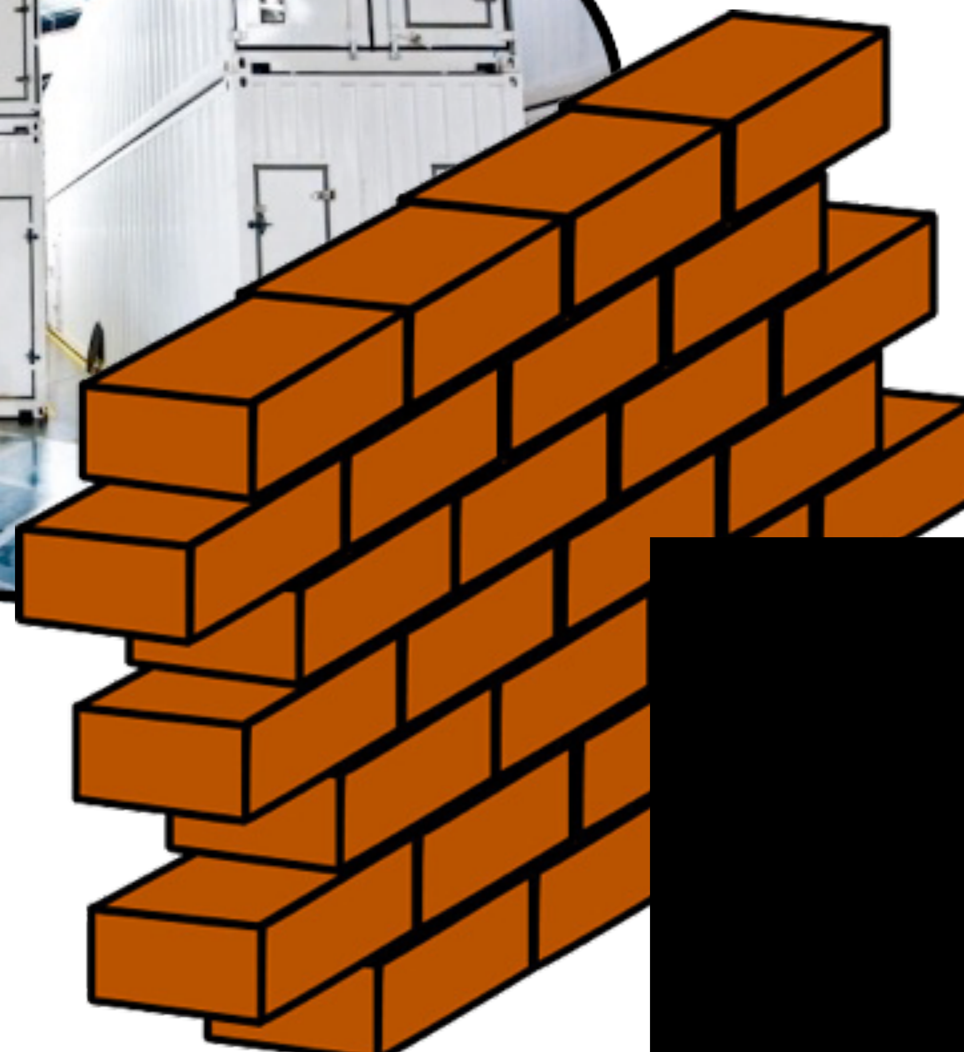

Boot
Service

Production
Platform

Based on "Delusional Boot: Securing Cloud Hypervisors without Massive Re-Engineering" (EuroSys 2012)

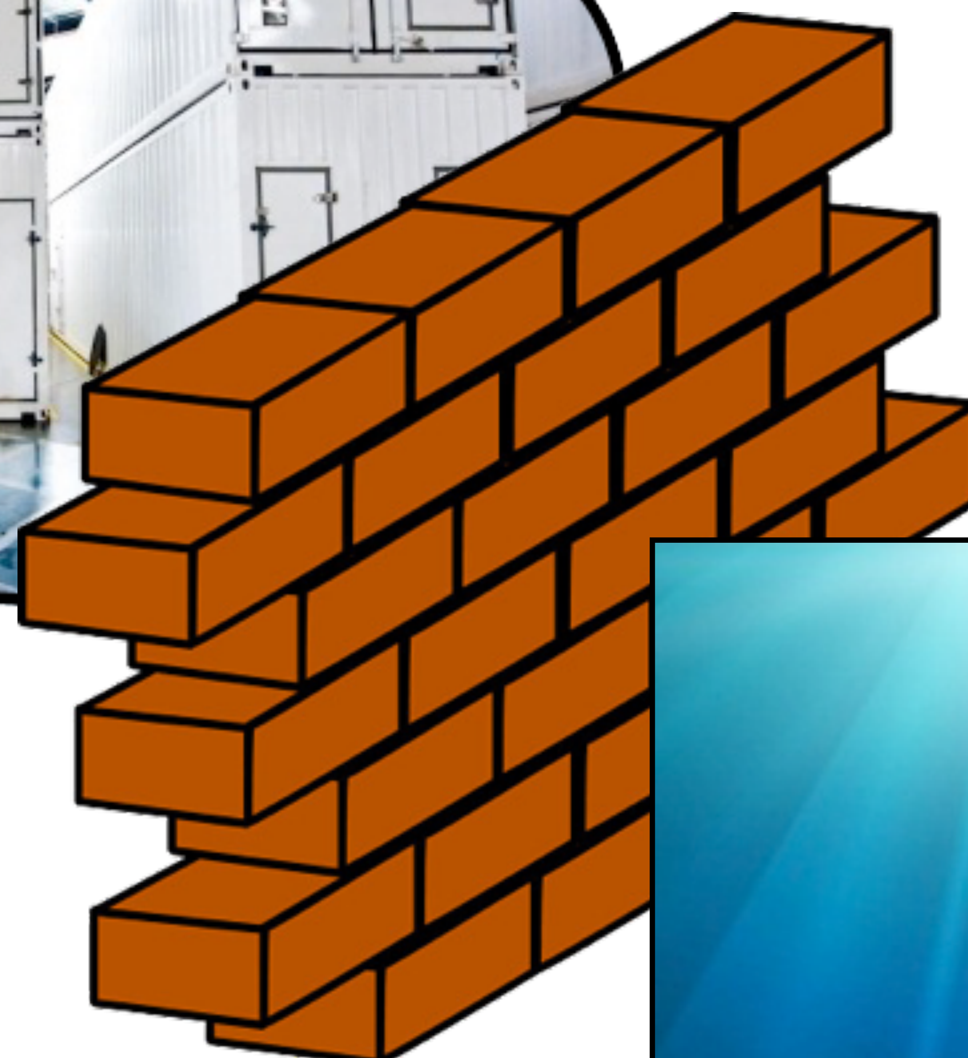

Boot
Service

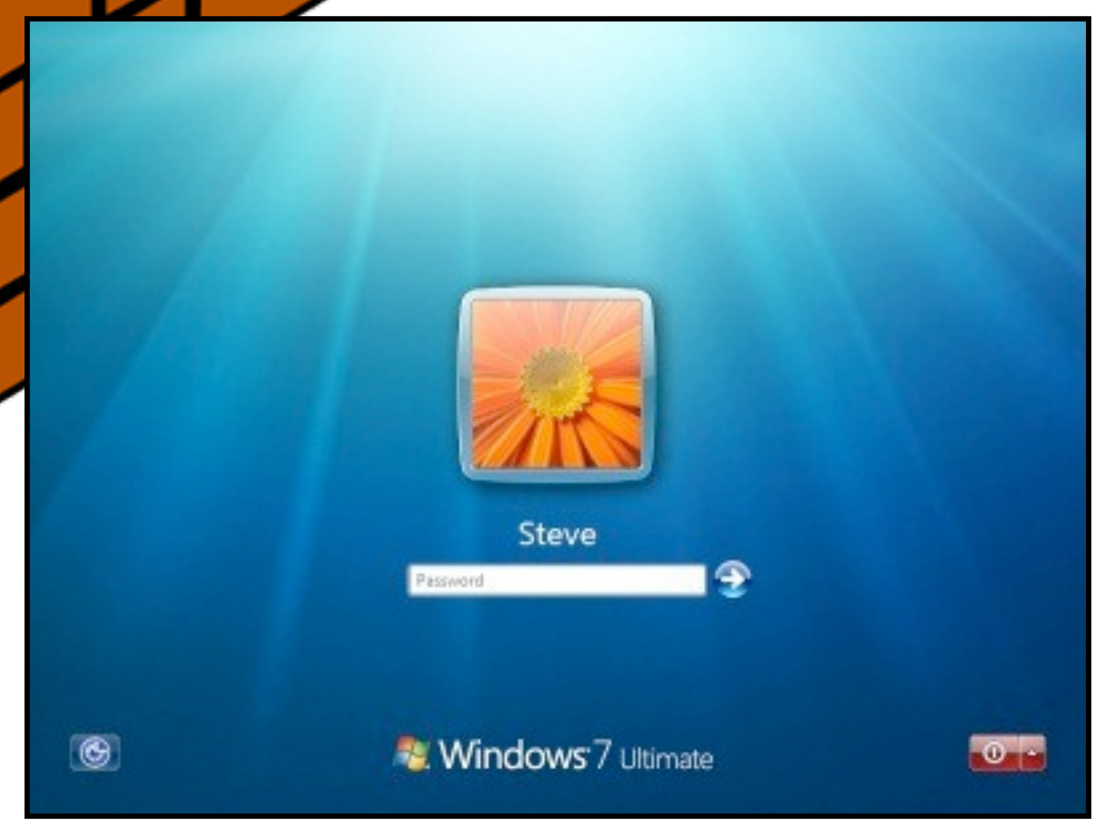

Production
Platform

Boot
Service

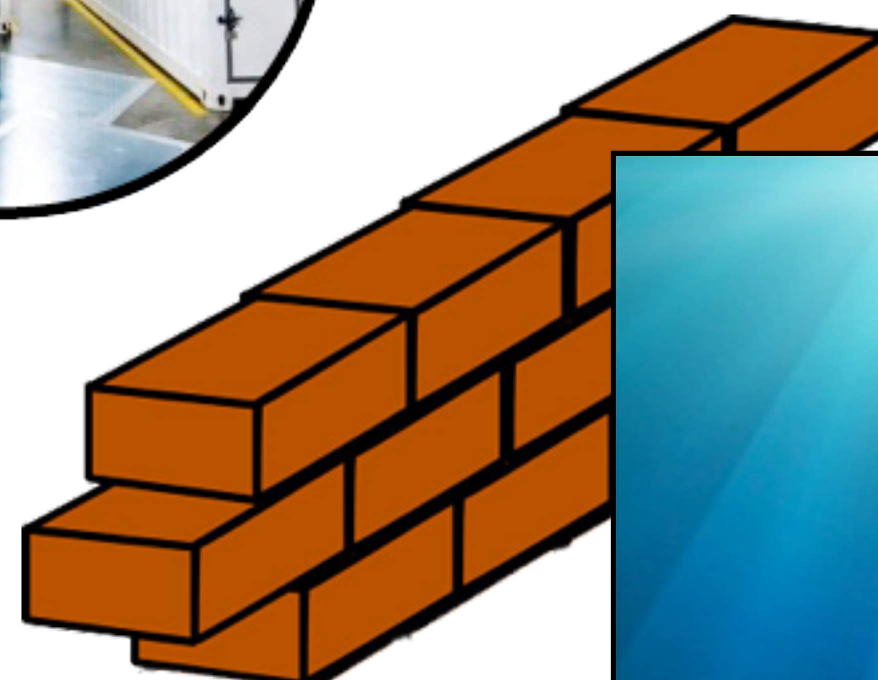

Production
Platform

Boot Service

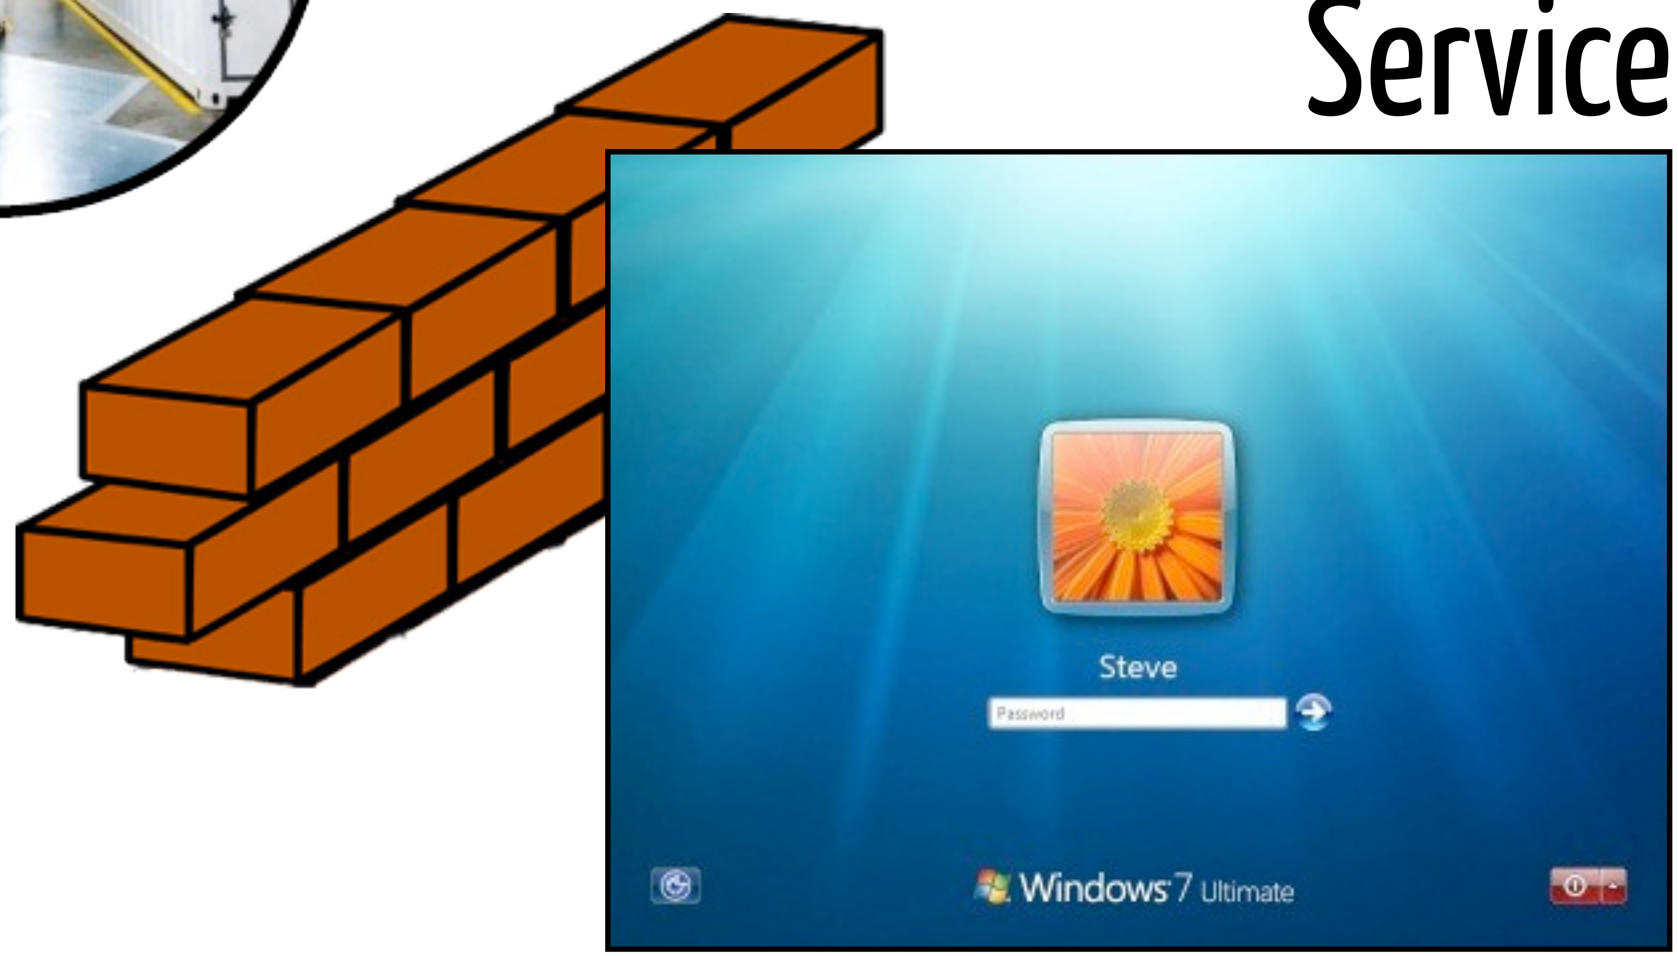

Production Platform

Production
Platform

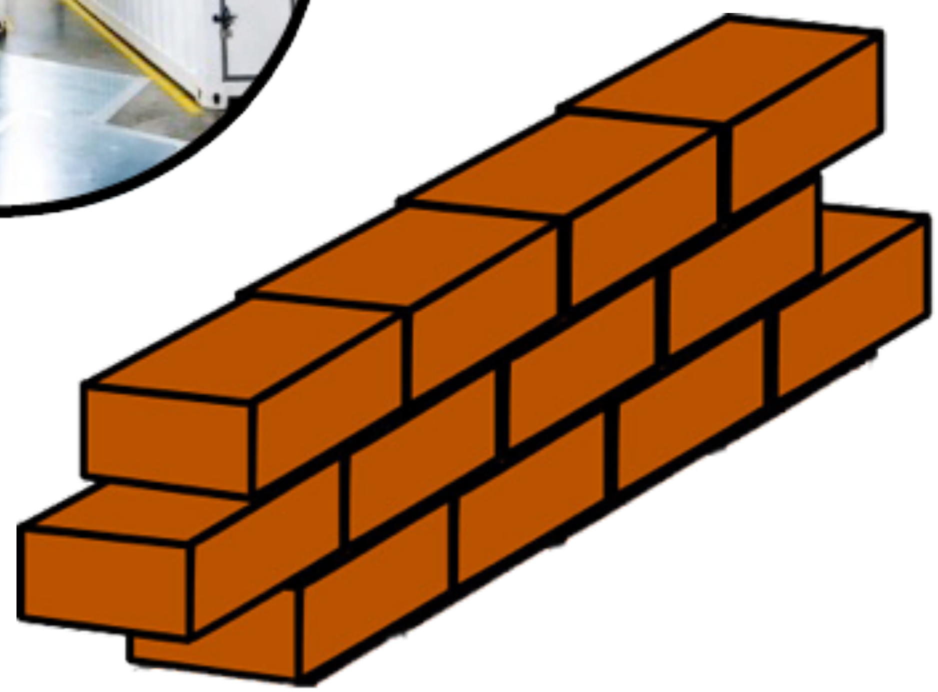

Boot
Service

Based on "Delusional Boot: Securing Cloud Hypervisors without Massive Re-Engineering" (EuroSys 2012)

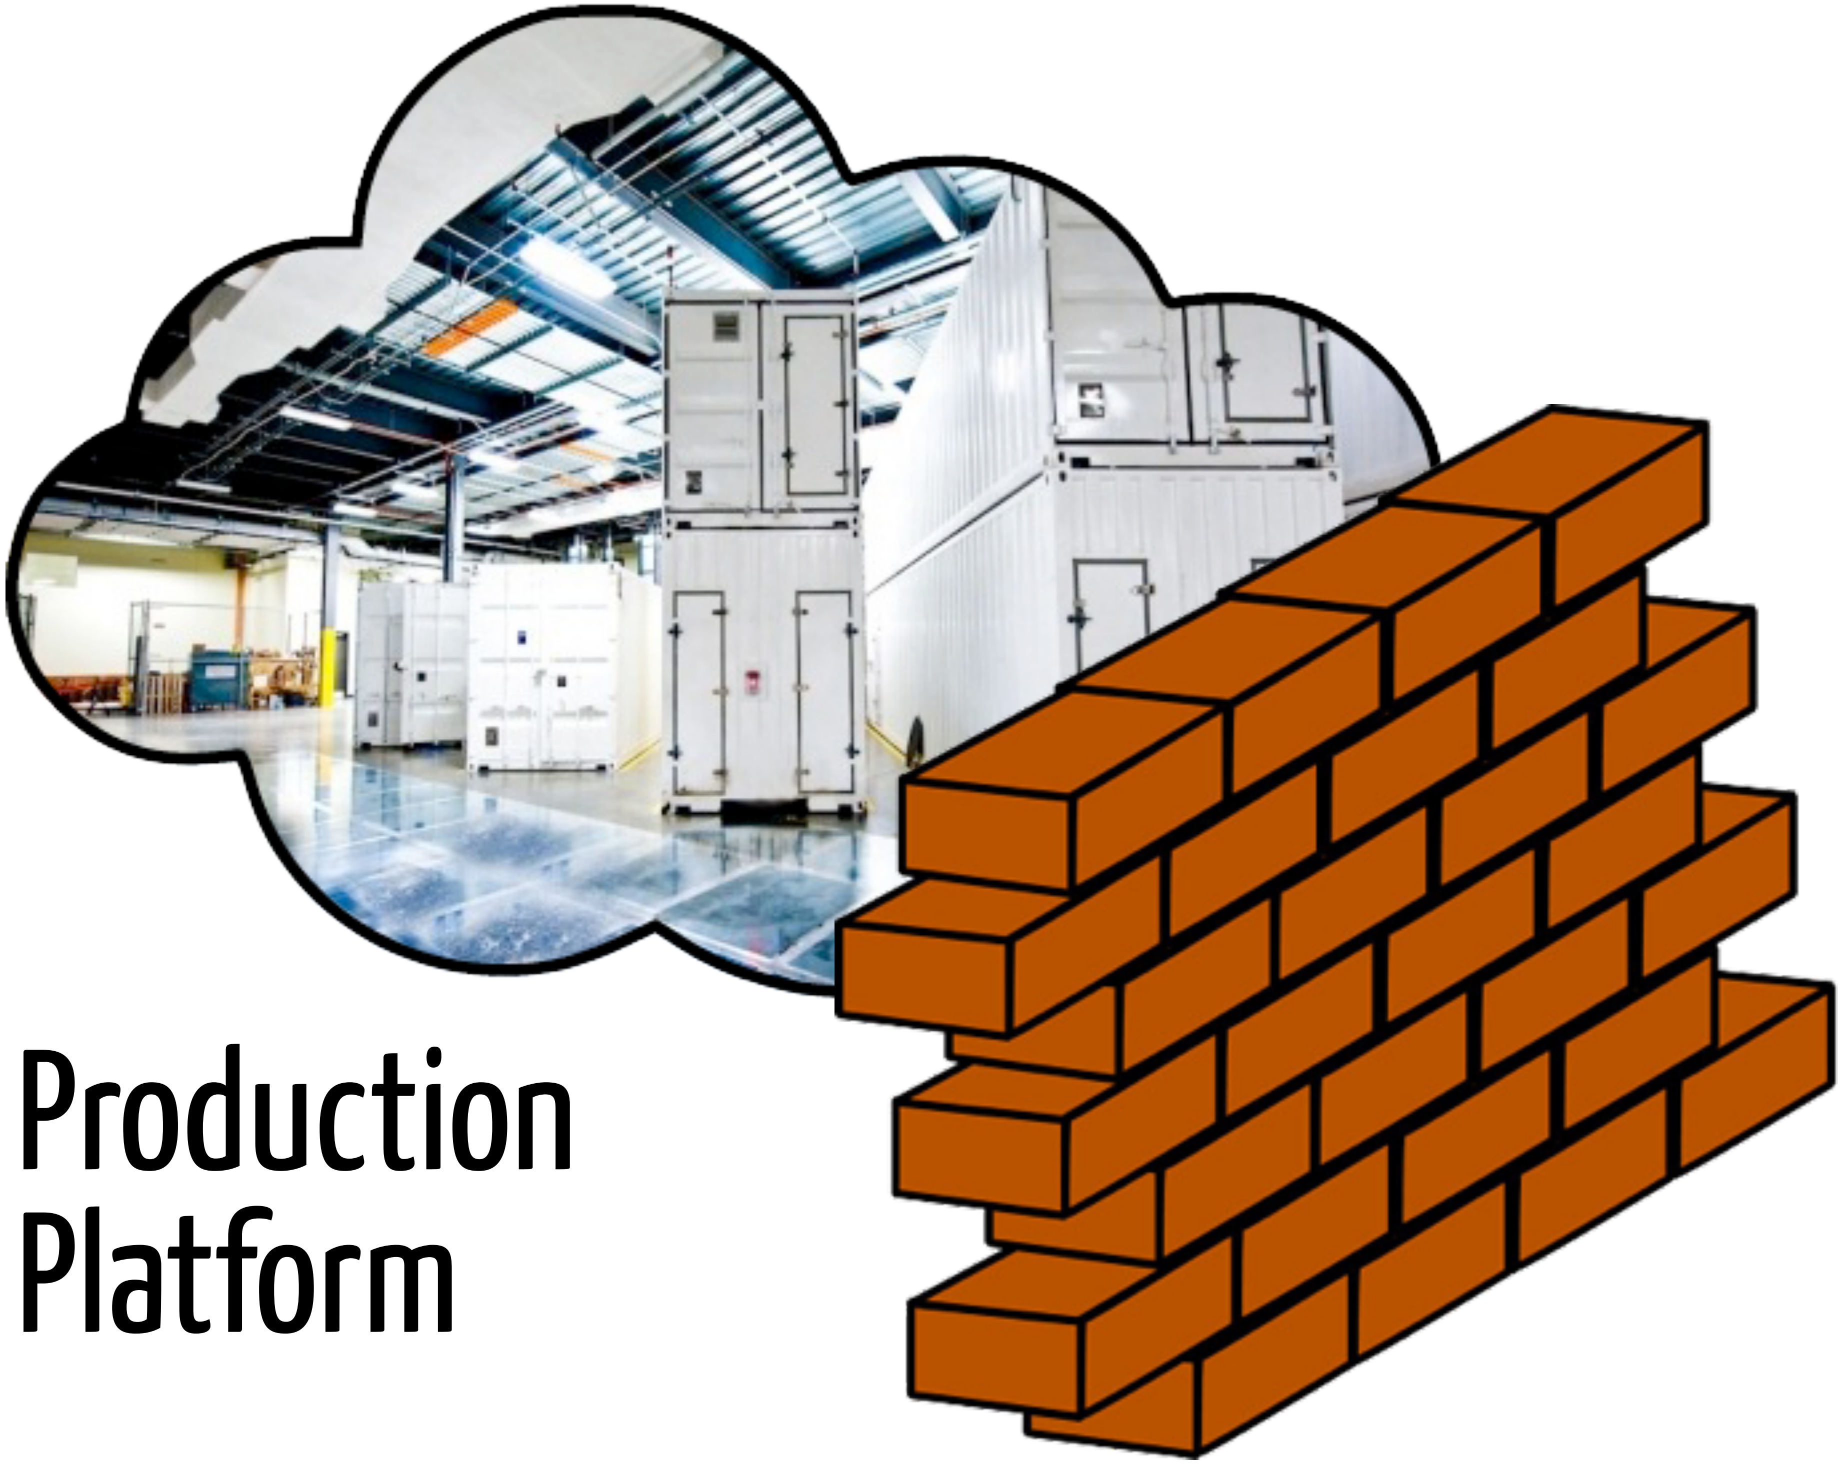

Boot
Service

Production
Platform

Based on “Delusional Boot: Securing Cloud Hypervisors without Massive Re-Engineering” (EuroSys 2012)

Proposal

Participatory Networking

Participatory Networking

Participatory Networking

PANE

Participatory Networking

1. Requests

PANE

Participatory Networking

1. Requests
2. Hints

PANE

Participatory Networking

1. Requests
2. Hints
3. Queries

PANE

Participatory Networking

Participatory Networking

Safe?

Participatory Networking

Safe?

Secure?

Participatory Networking

Safe? Secure? Fair?

Participatory Networking

Safe?

Secure?

Fair?

Practical?

Participatory Networking

Safe?

Secure?

Fair?

Practical?

Efficient?

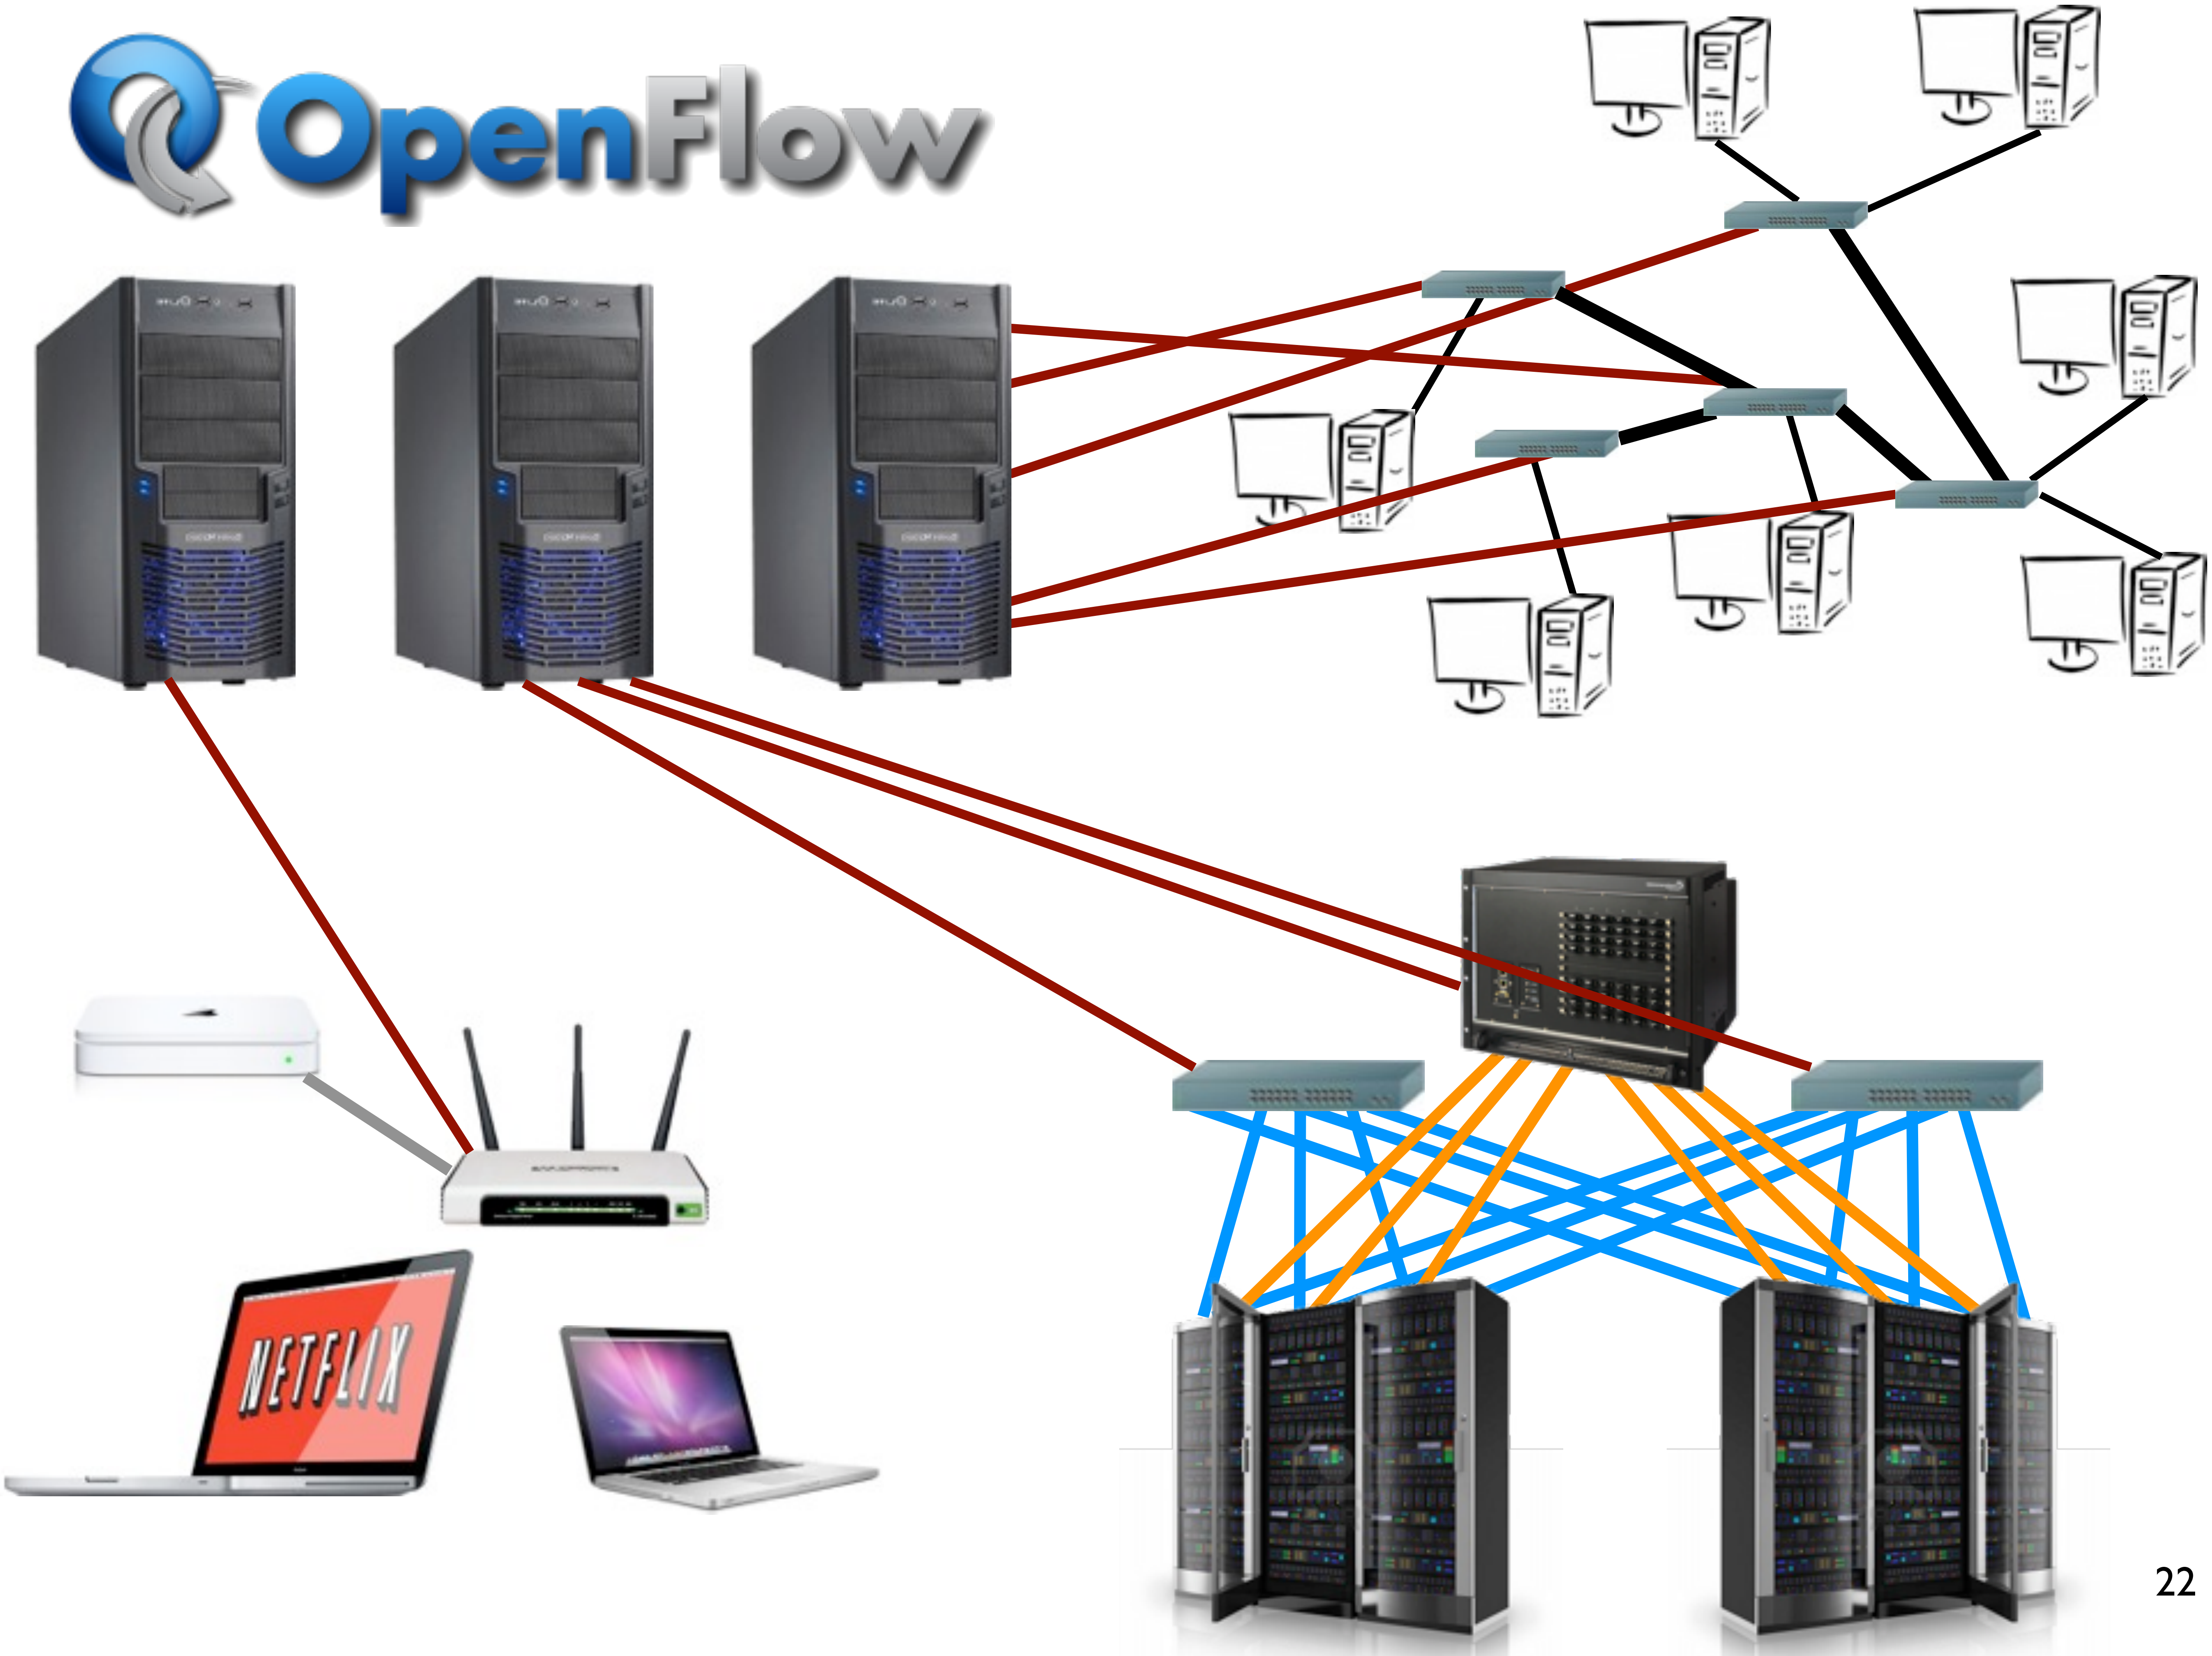
OpenFlow
NETFLIX

Participatory Networking

Participatory Networking

- End-user API for SDNs

Participatory Networking

- End-user API for SDNs
- Exposes existing mechanisms

Participatory Networking

- End-user API for SDNs
- Exposes existing mechanisms
- No effect on unmodified applications

Outline of PANE

1. Privilege Delegation Semantics

Outline of PANE

1. Privilege Delegation Semantics
2. Protocol Sketch

Outline of PANE

1. Privilege Delegation Semantics
2. Protocol Sketch
3. Dynamic Flow Processing

Outline of PANE

1. Privilege Delegation Semantics
2. Protocol Sketch
3. Dynamic Flow Processing
4. Current Status

Outline of PANE

Privilege Delegation Semantics

Shares

Shares

Flowgroup

Shares

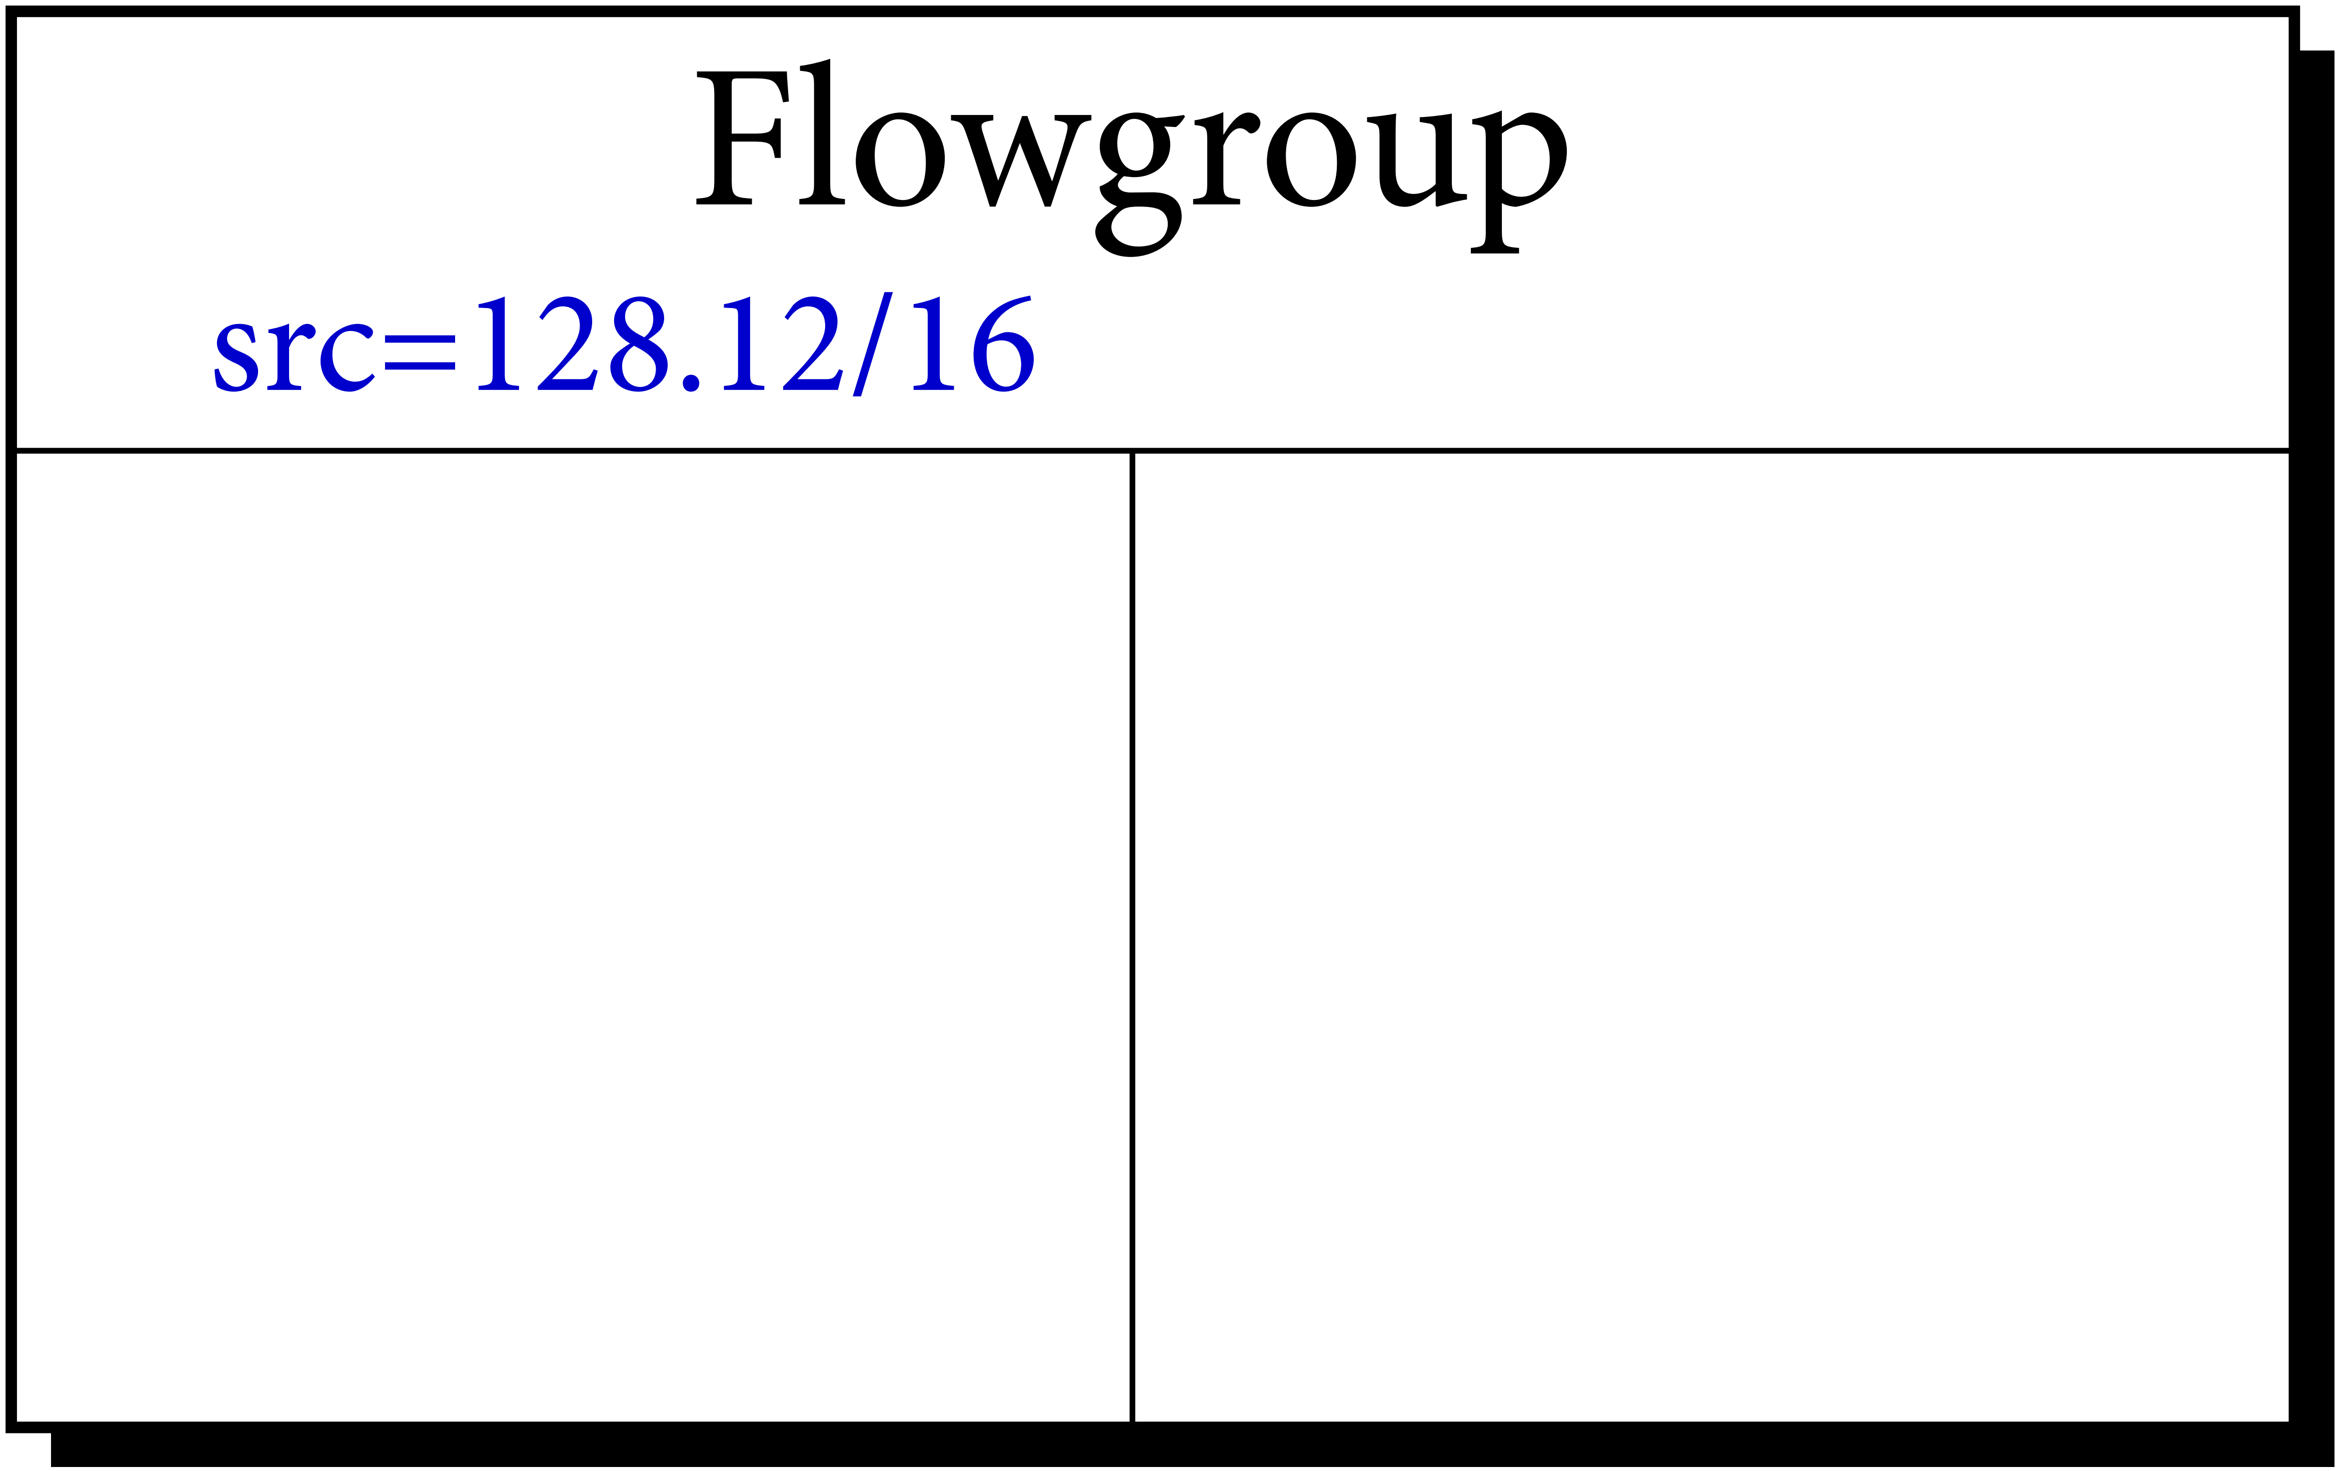

Shares

Flowgroup

src=128.12/16 ∧ dst.port ≤1024

Shares

Flowgroup

src=128.12/16 $\wedge$ dst.port $\leq$1024

Privileges

Shares

Flowgroup	
$src=128.12/16 \wedge dst.port \leq 1024$	
	Privileges deny, allow

Shares

Flowgroup

src=128.12/16 $\wedge$ dst.port $\leq$1024

Privileges

deny, allow

bandwidth: 5Mb/s

limit: 10Mb/s

Shares

Flowgroup src=128.12/16 $\wedge$ dst.port $\leq$1024	
	Privileges deny, allow bandwidth: 5Mb/s limit: 10Mb/s *hint* *query*

Shares

Flowgroup src=128.12/16 ∧ dst.port ≤1024	
Speakers	Privileges deny, allow bandwidth: 5Mb/s limit: 10Mb/s *hint* *query*

Shares

Flowgroup src=128.12/16 ∧ dst.port ≤1024	
Speakers	Privileges
Alice Bob	deny, allow bandwidth: 5Mb/s limit: 10Mb/s *hint* *query*

Shares

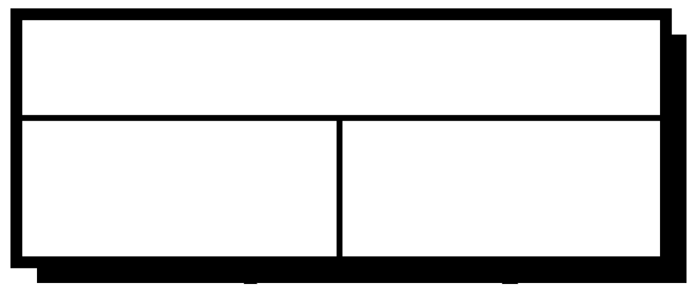

Delegation

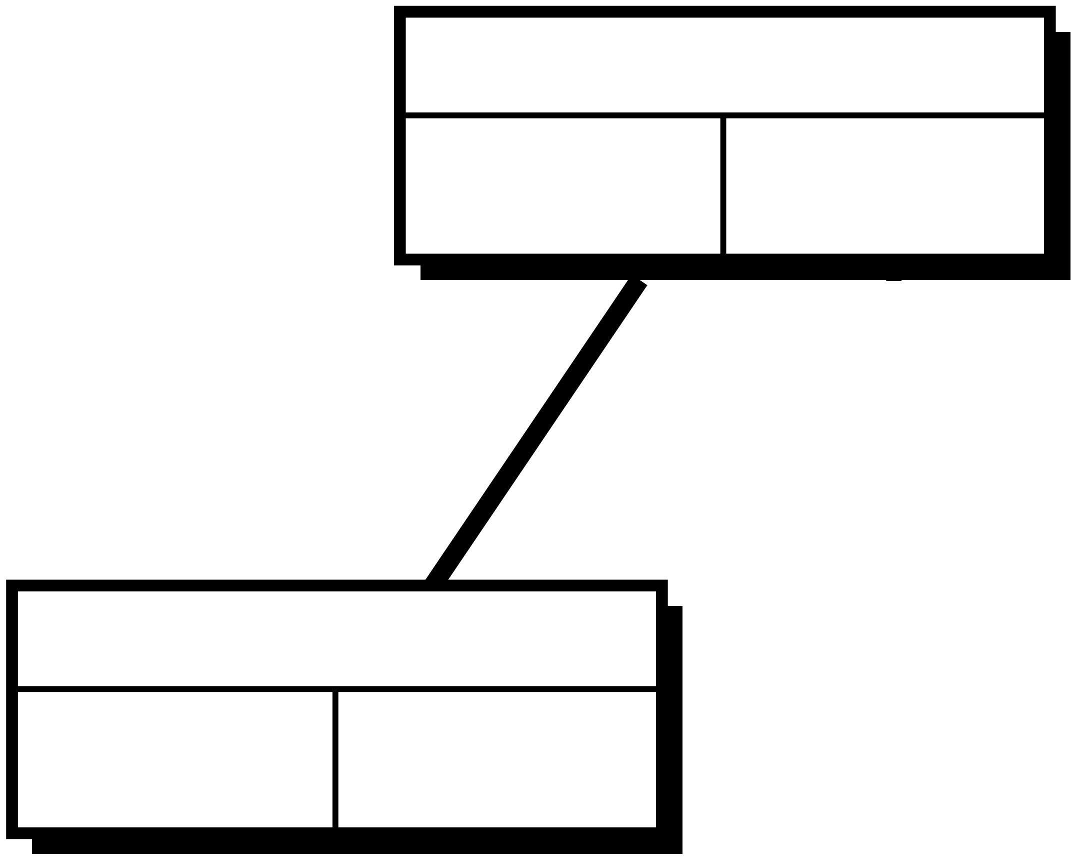

Delegation

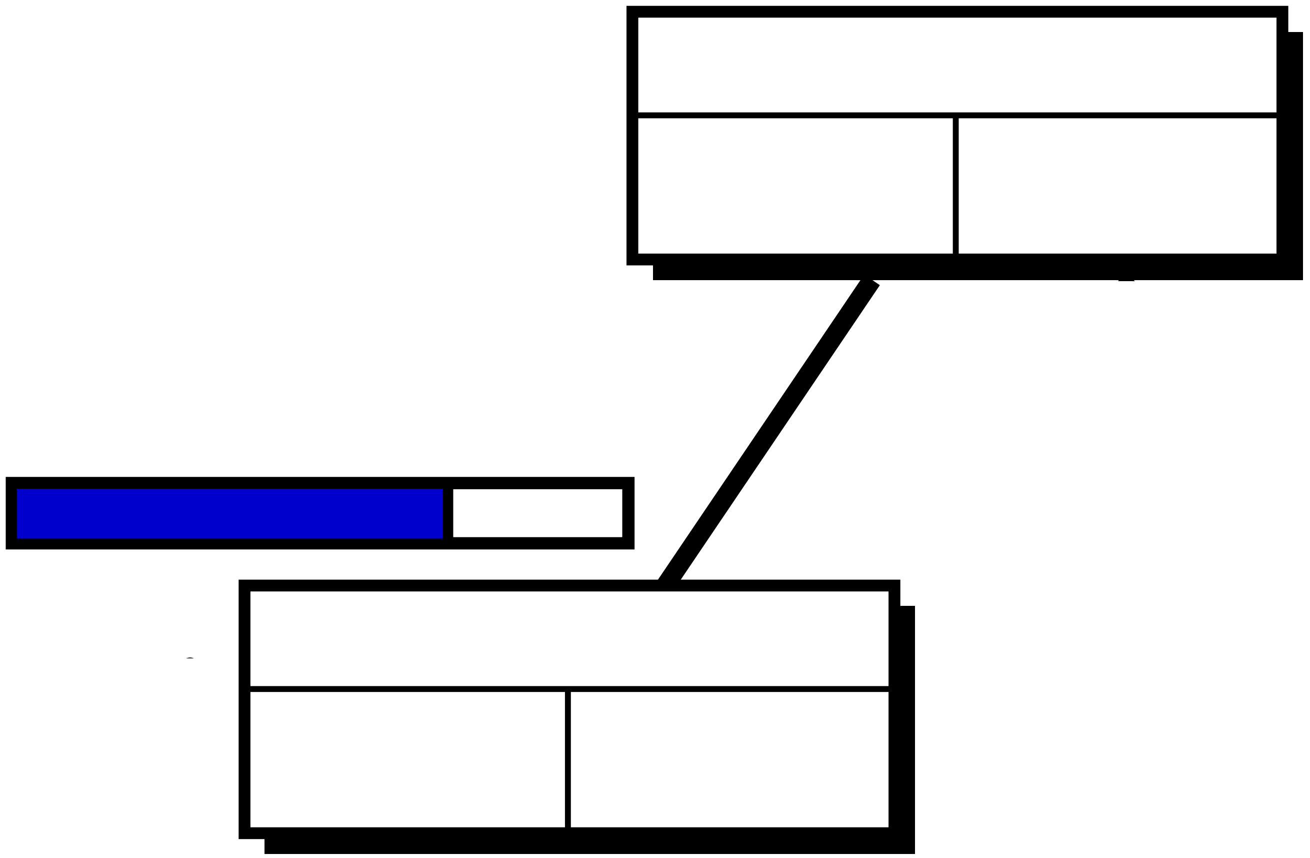

Delegation

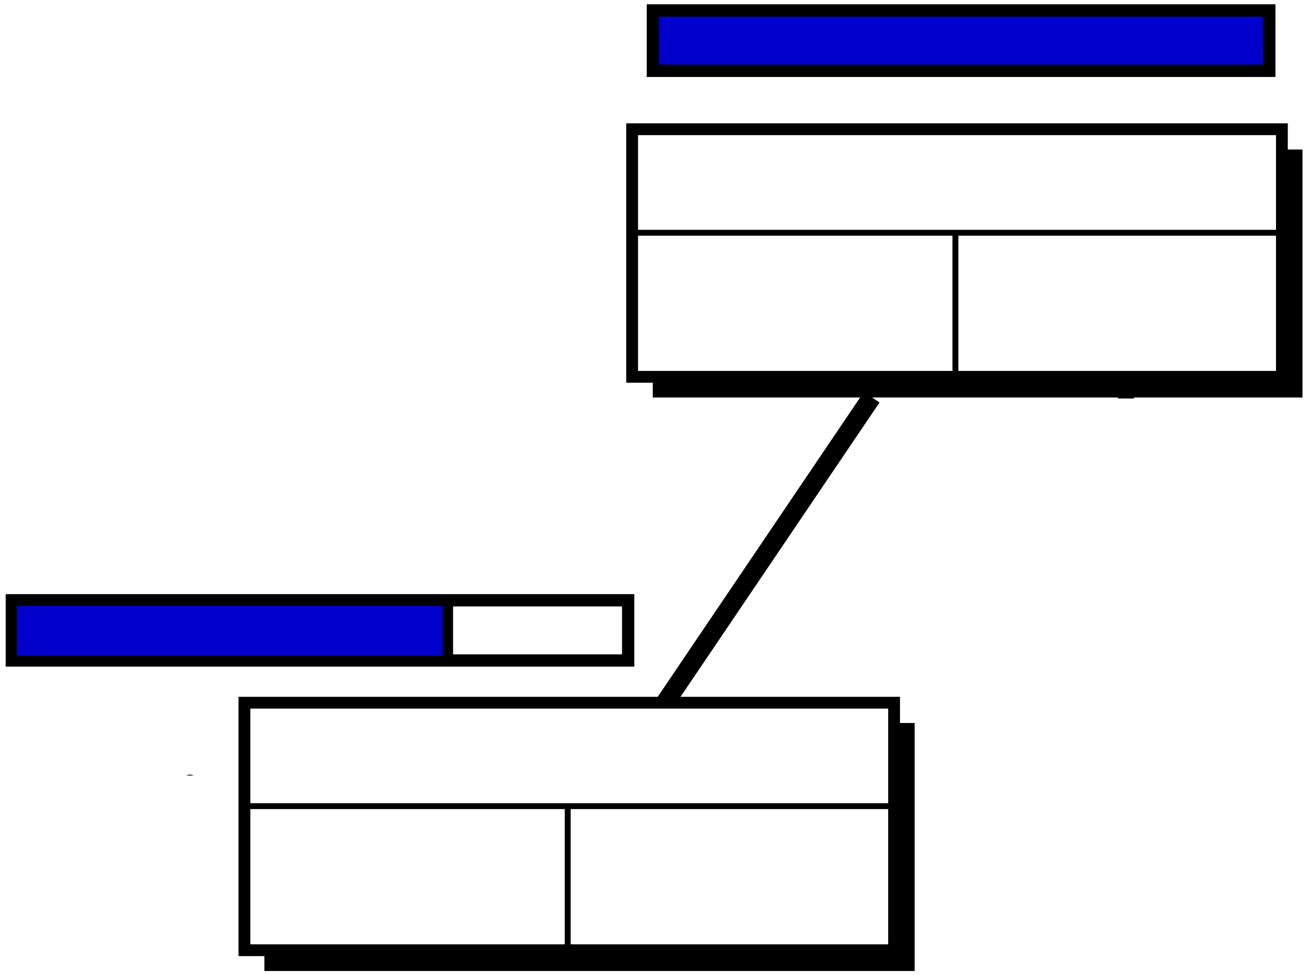

Delegation

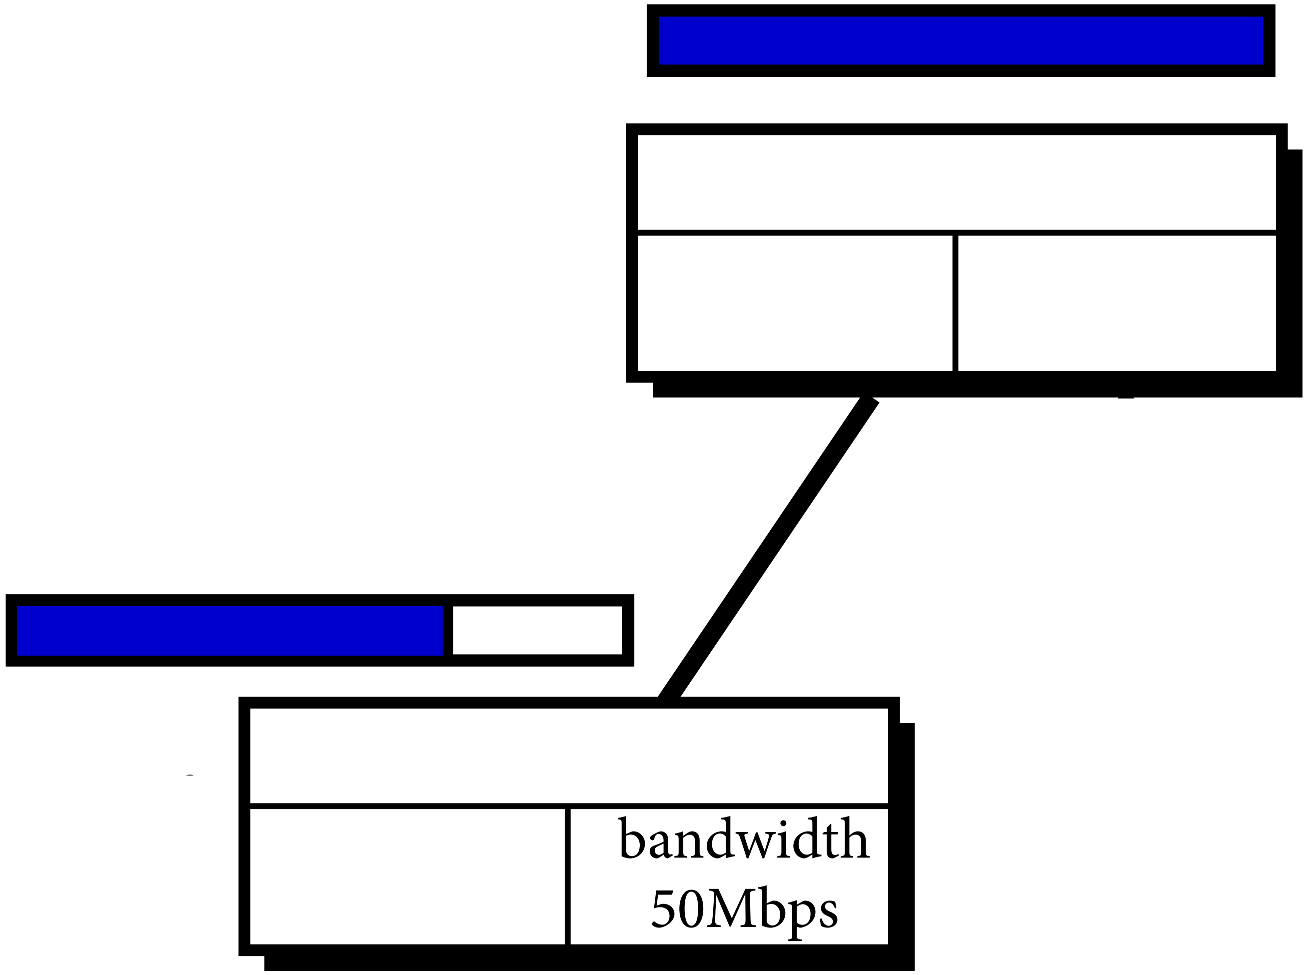

Delegation

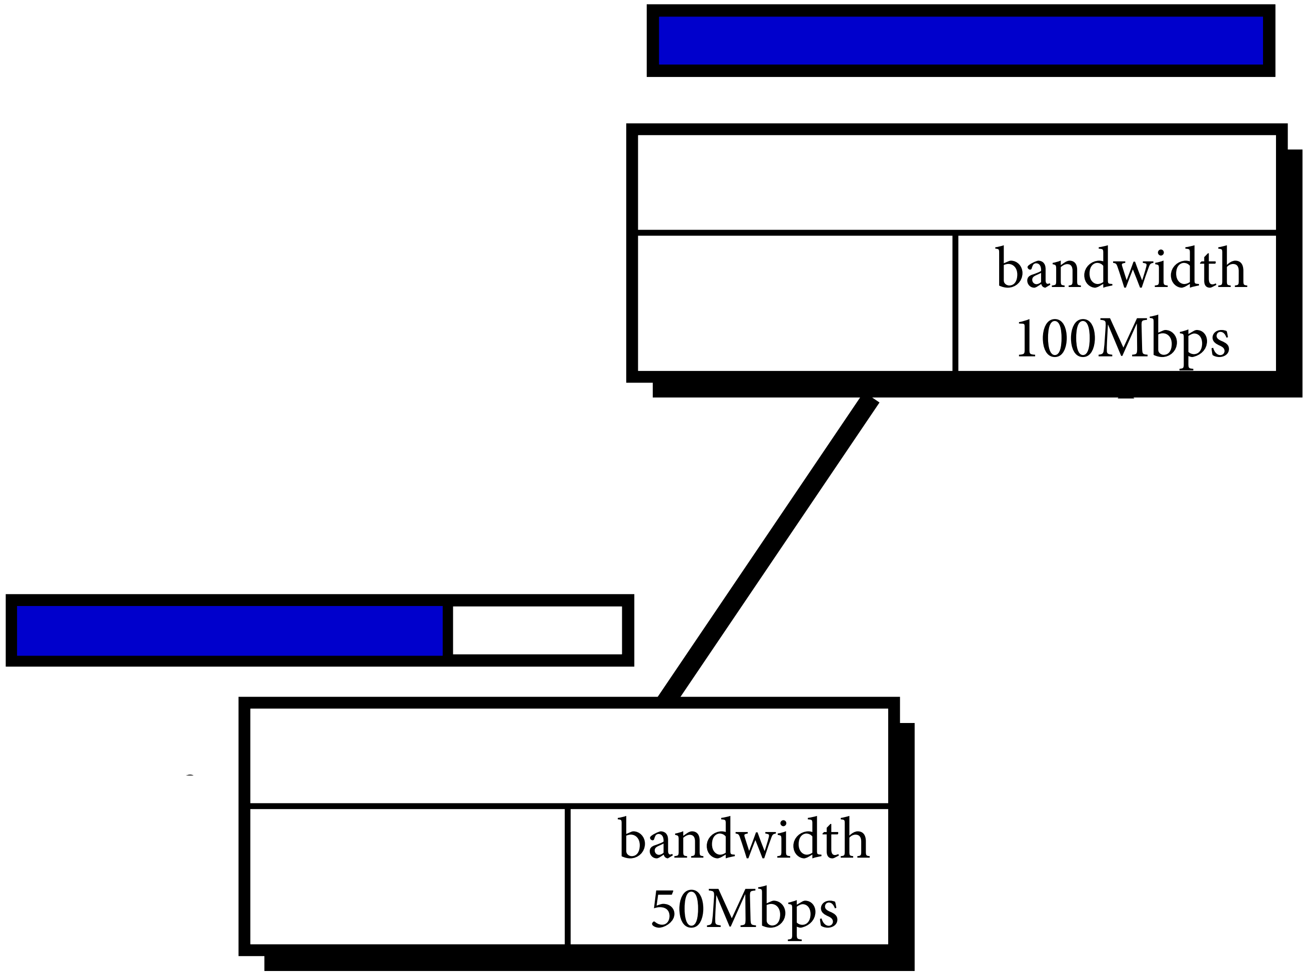

Delegation

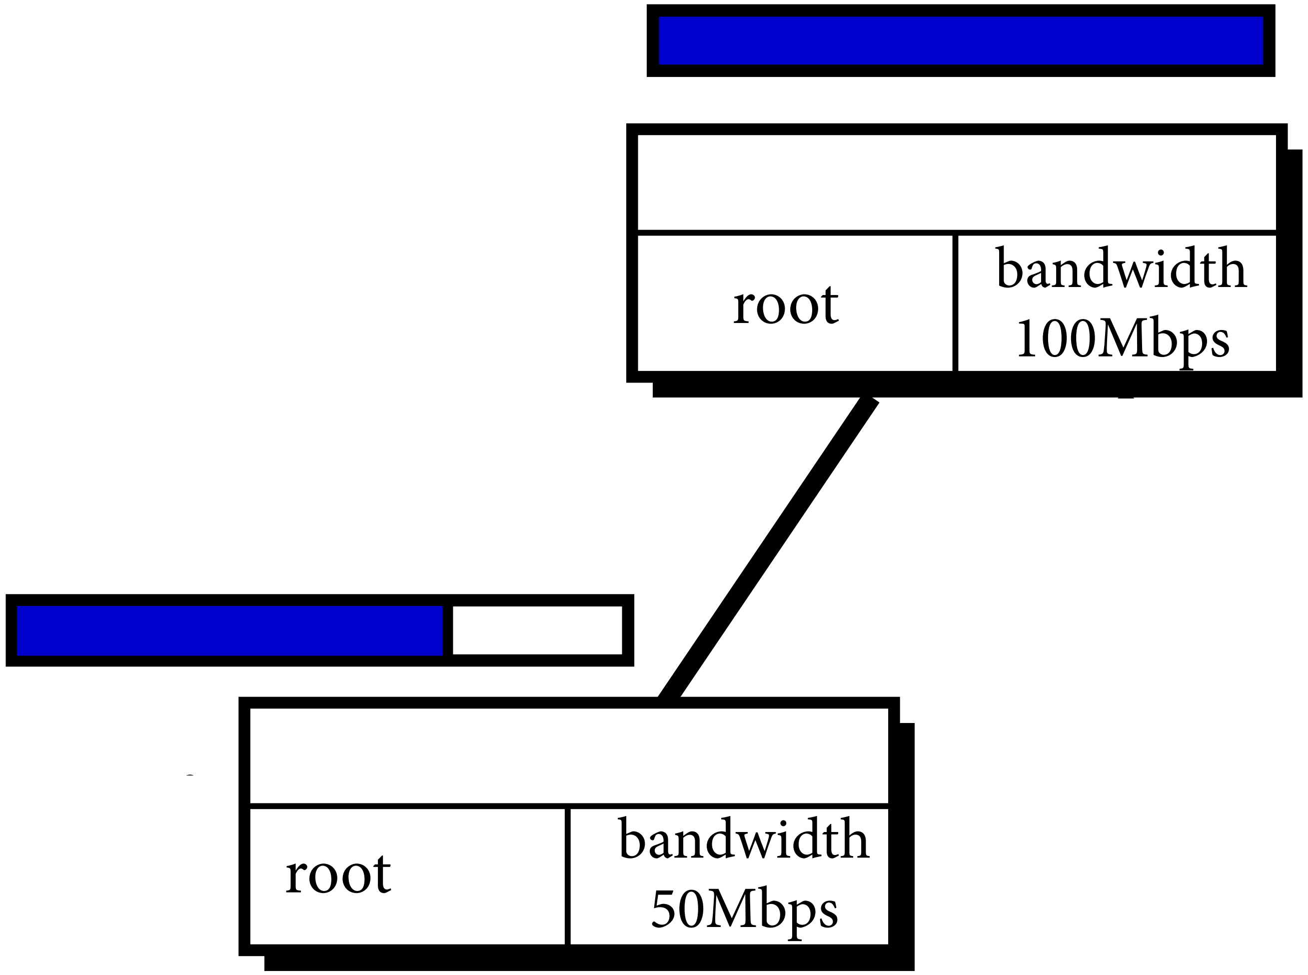

Delegation

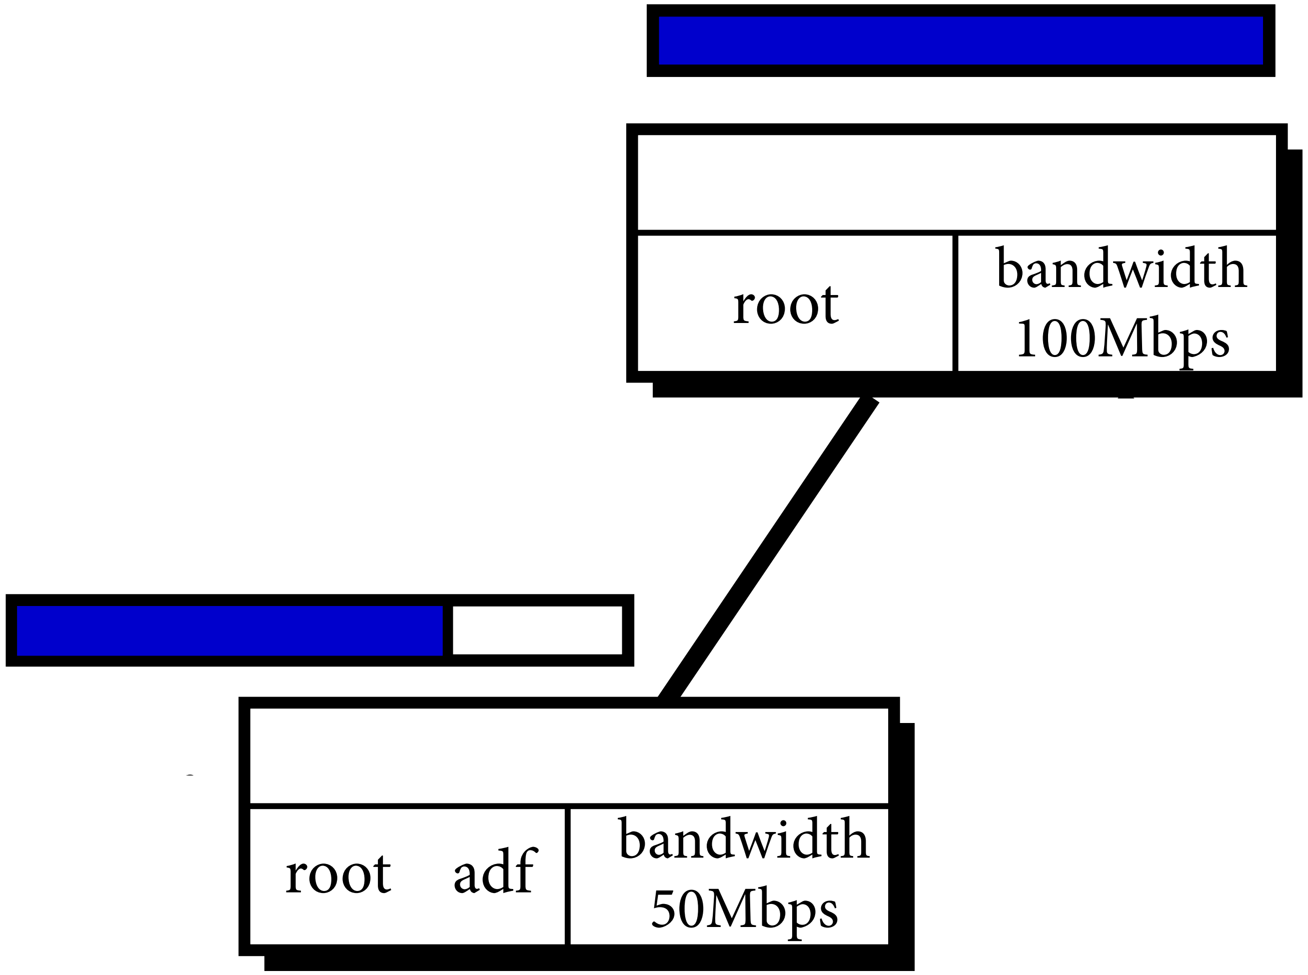

Delegation

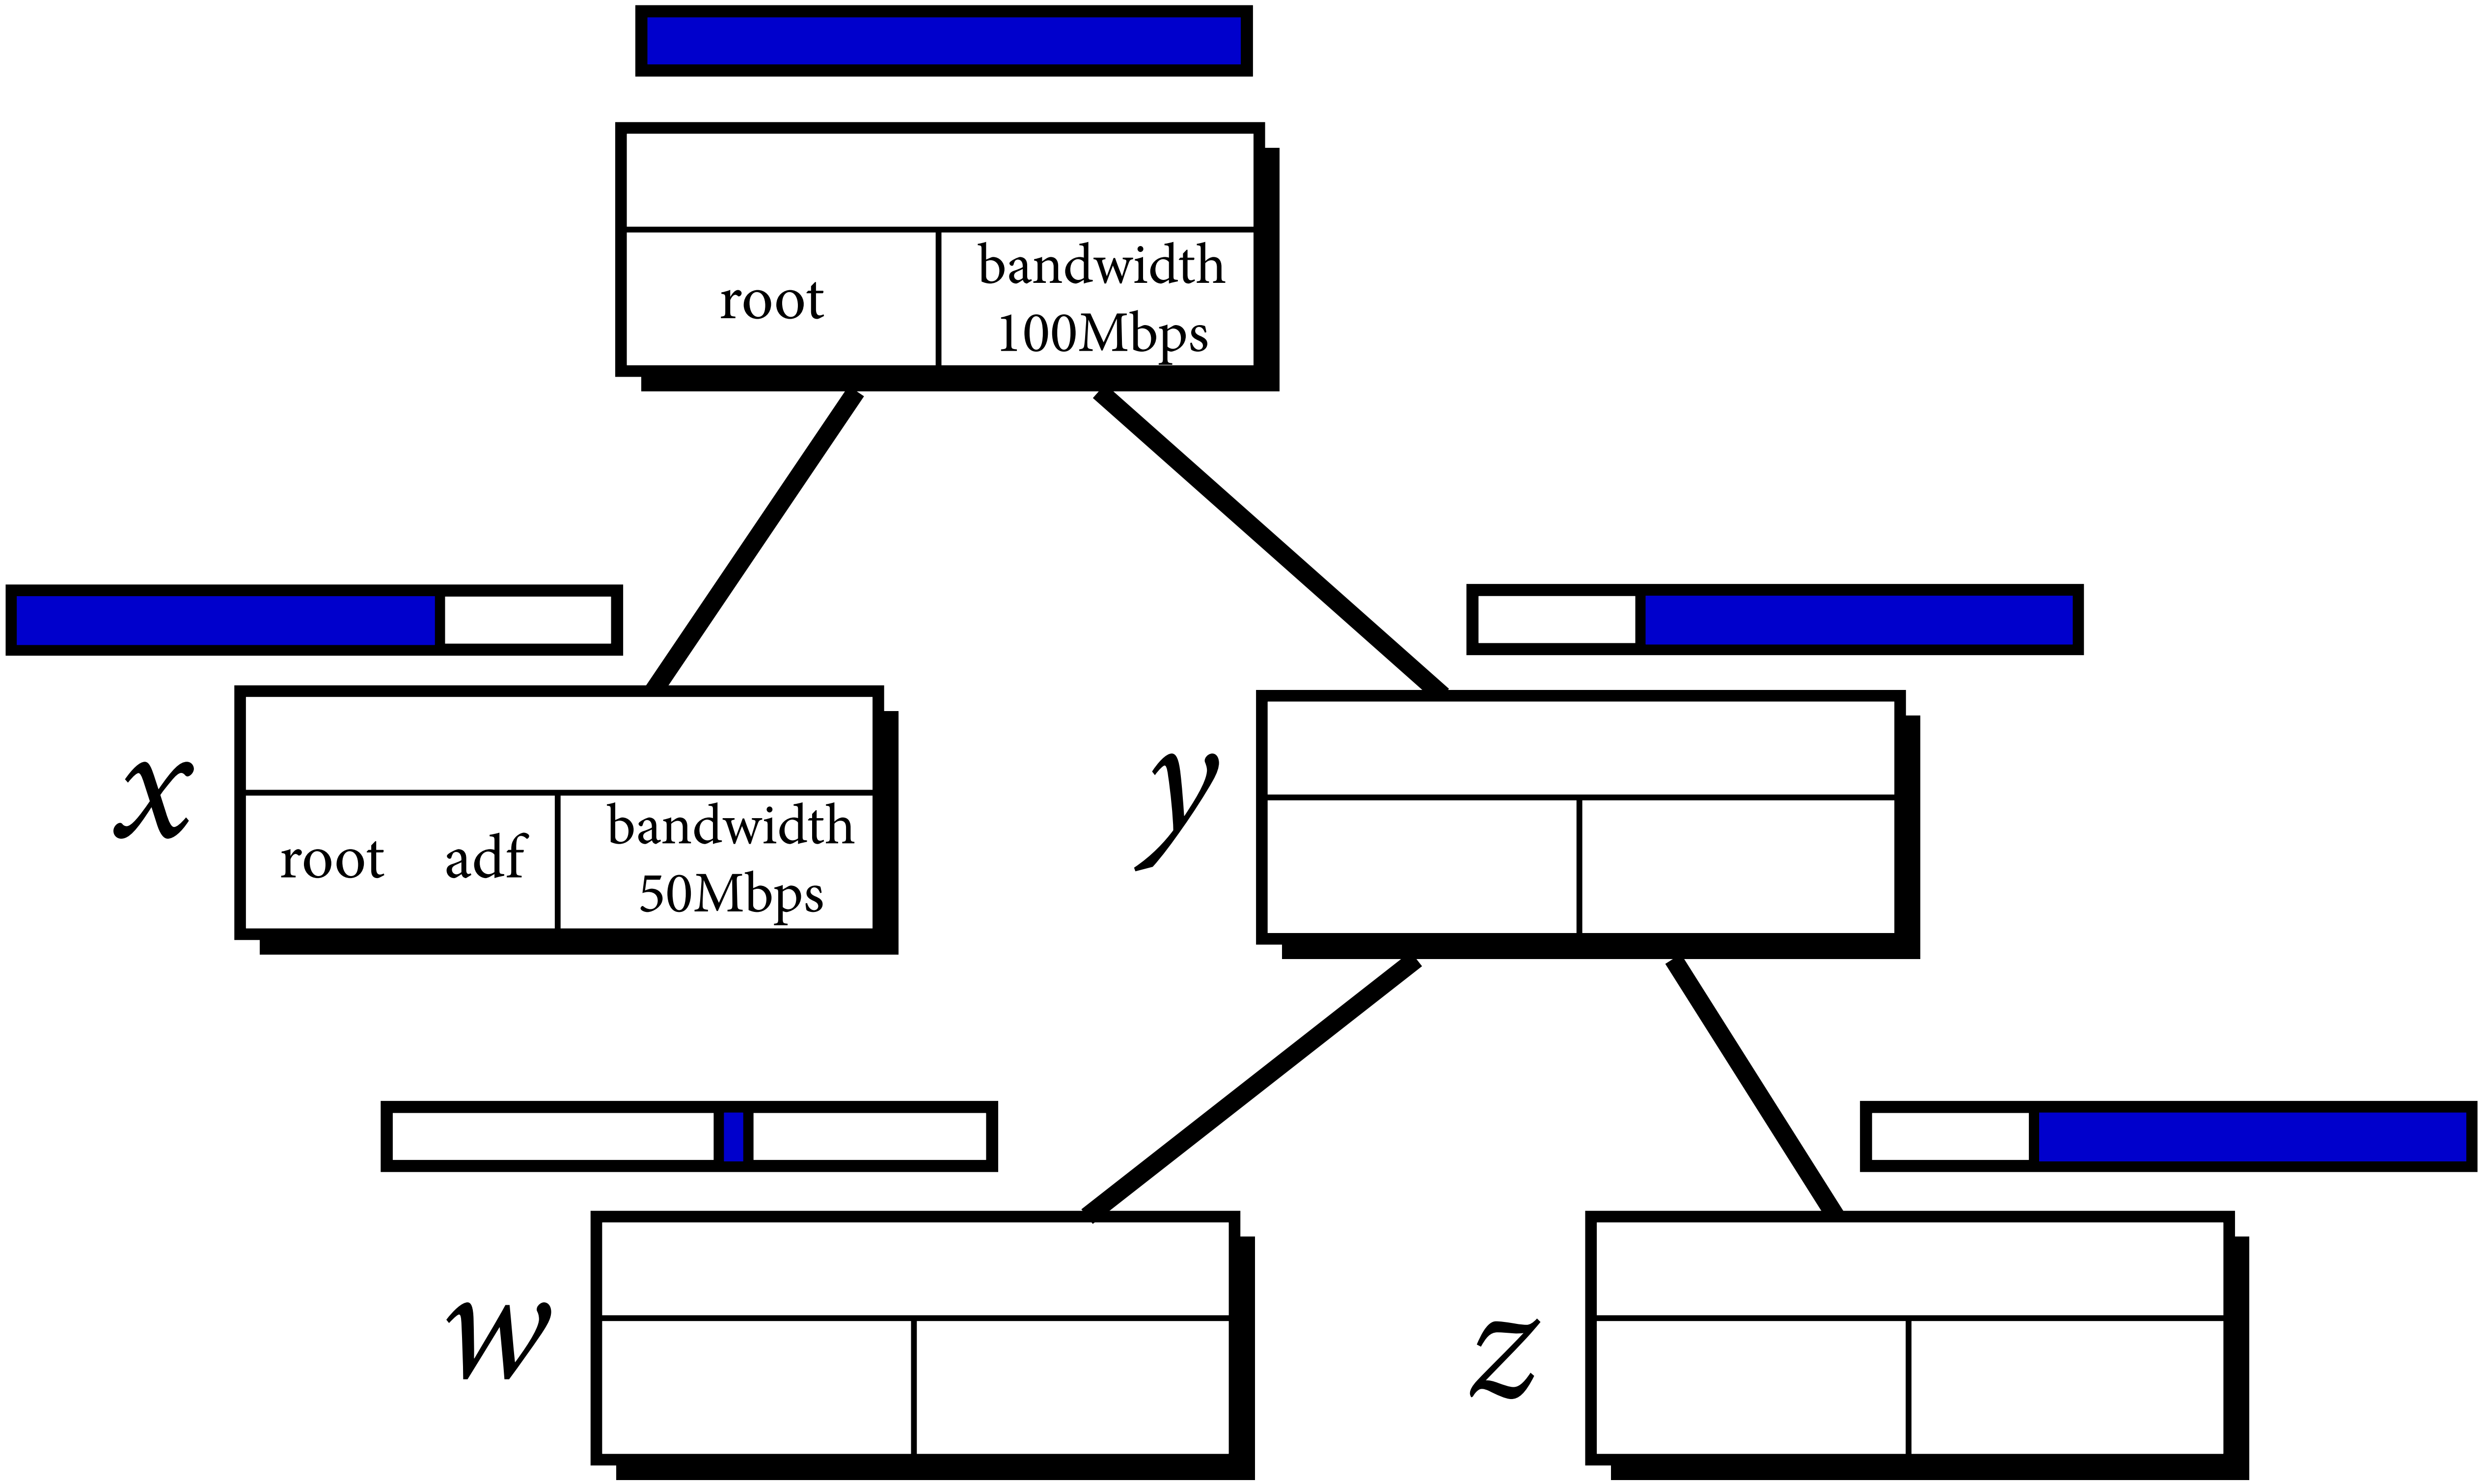

Delegation

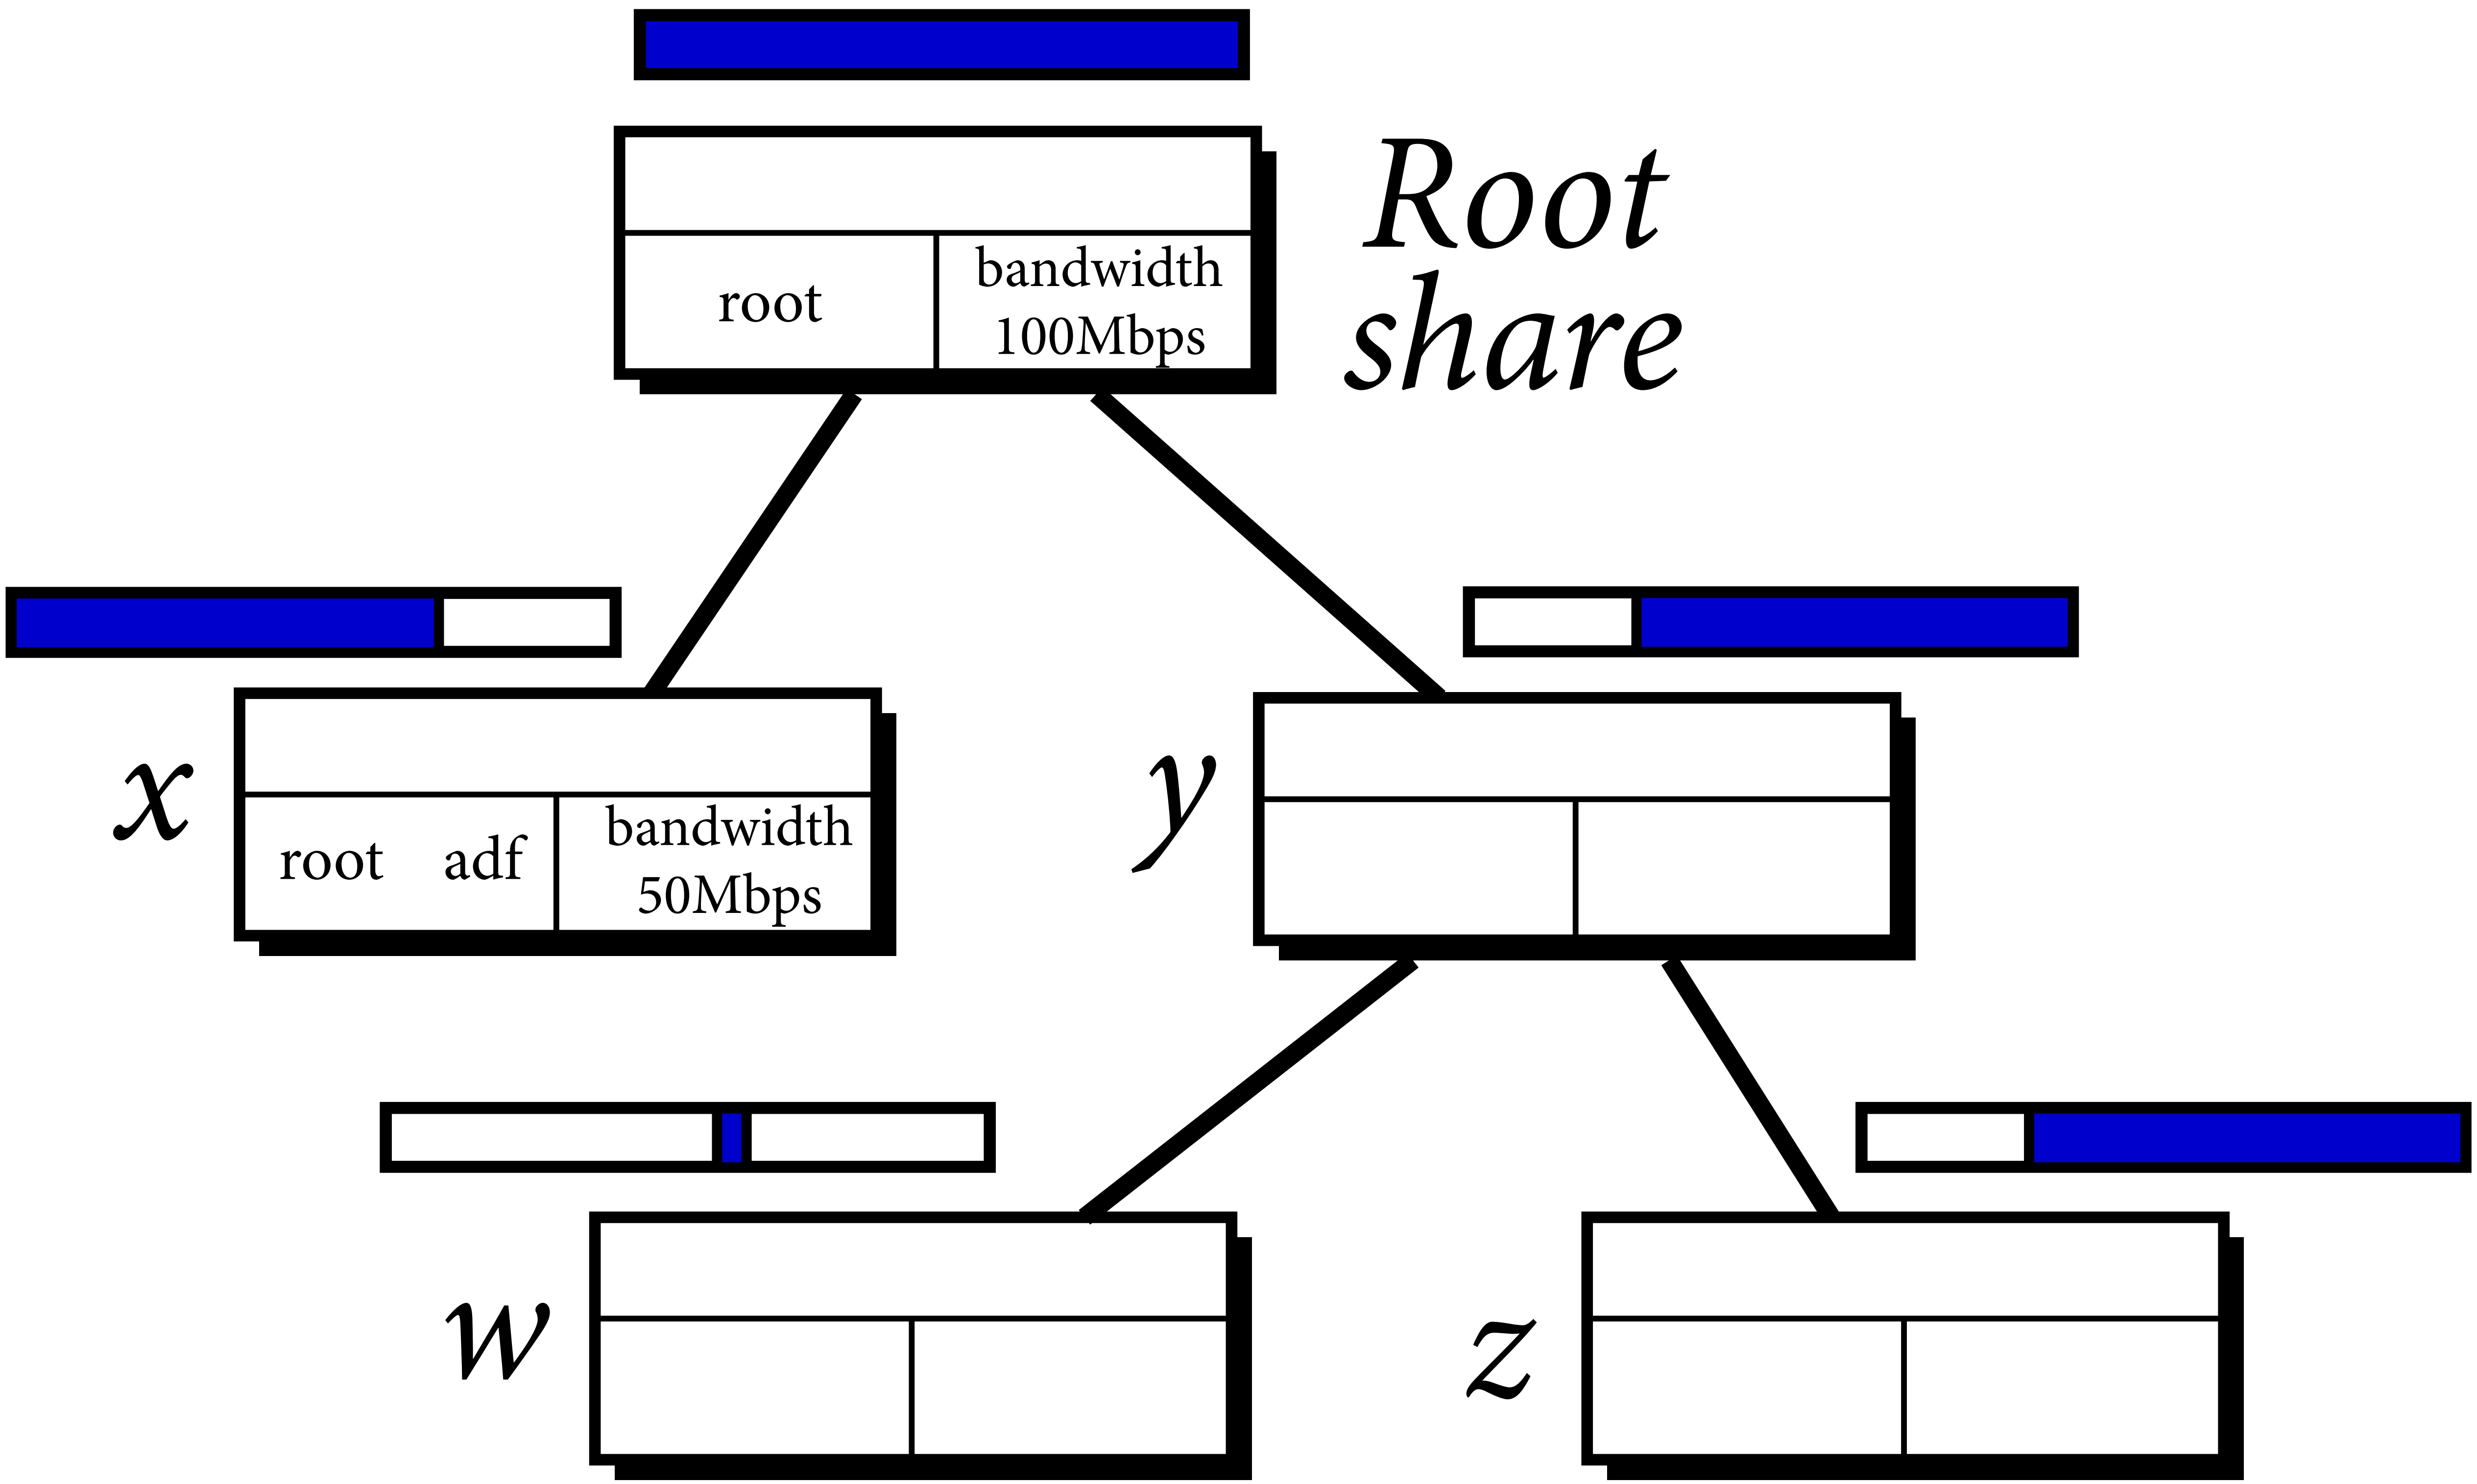
Root share
root
bandwidth 100Mbps
x
root adf
bandwidth 50Mbps
y
w
z

Delegation

Flowgroup

src=128.12/16 ∧ dst.port ≤1024

Speakers	Privileges
Alice Bob	deny, allow bandwidth: 5Mb/s limit: 10Mb/s *hint* *query*

Dynamic Context

Flowgroup	
src=128.12/16 ∧ dst.port ≤1024	
Speakers	Privileges
Alice Bob	deny, allow bandwidth: 5Mb/s limit: 10Mb/s *hint* *query*

PANE

Dynamic Context

Flowgroup	
src=128.12/16 ∧ dst.port ≤1024	
Speakers	Privileges
Alice Bob	deny, allow bandwidth: 5Mb/s limit: 10Mb/s *hint* *query*

Reserve 2 Mbps
from now to +5min?

PANE

Dynamic Context

Flowgroup

src=128.12/16 ∧ dst.port ≤1024

Speakers	Privileges
Alice	deny, allow
Bob	bandwidth: 5Mb/s
	limit: 10Mb/s
	hint
	query

Yes

PANE

Dynamic Context

Flowgroup

src=128.12/16 ∧ dst.port ≤1024

Speakers	Privileges
Alice Bob	deny, allow bandwidth: 5Mb/s limit: 10Mb/s *hint* *query*

This traffic will be short and bursty

PANE

Dynamic Context

Flowgroup	
src=128.12/16 ∧ dst.port ≤1024	
Speakers	**Privileges**
Alice Bob	deny, allow bandwidth: 5Mb/s limit: 10Mb/s *hint* *query*

OK

PANE

Dynamic Context

Flowgroup	
src=128.12/16 ∧ dst.port ≤1024	
Speakers	Privileges
Alice Bob	deny, allow bandwidth: 5Mb/s limit: 10Mb/s *hint* *query*

How much web traffic in the last hour?

PANE

Dynamic Context

Flowgroup	
src=128.12/16 ∧ dst.port ≤1024	
Speakers	**Privileges**
Alice Bob	deny, allow bandwidth: 5Mb/s limit: 10Mb/s *hint* *query*

67,560 bytes

PANE

Dynamic Context

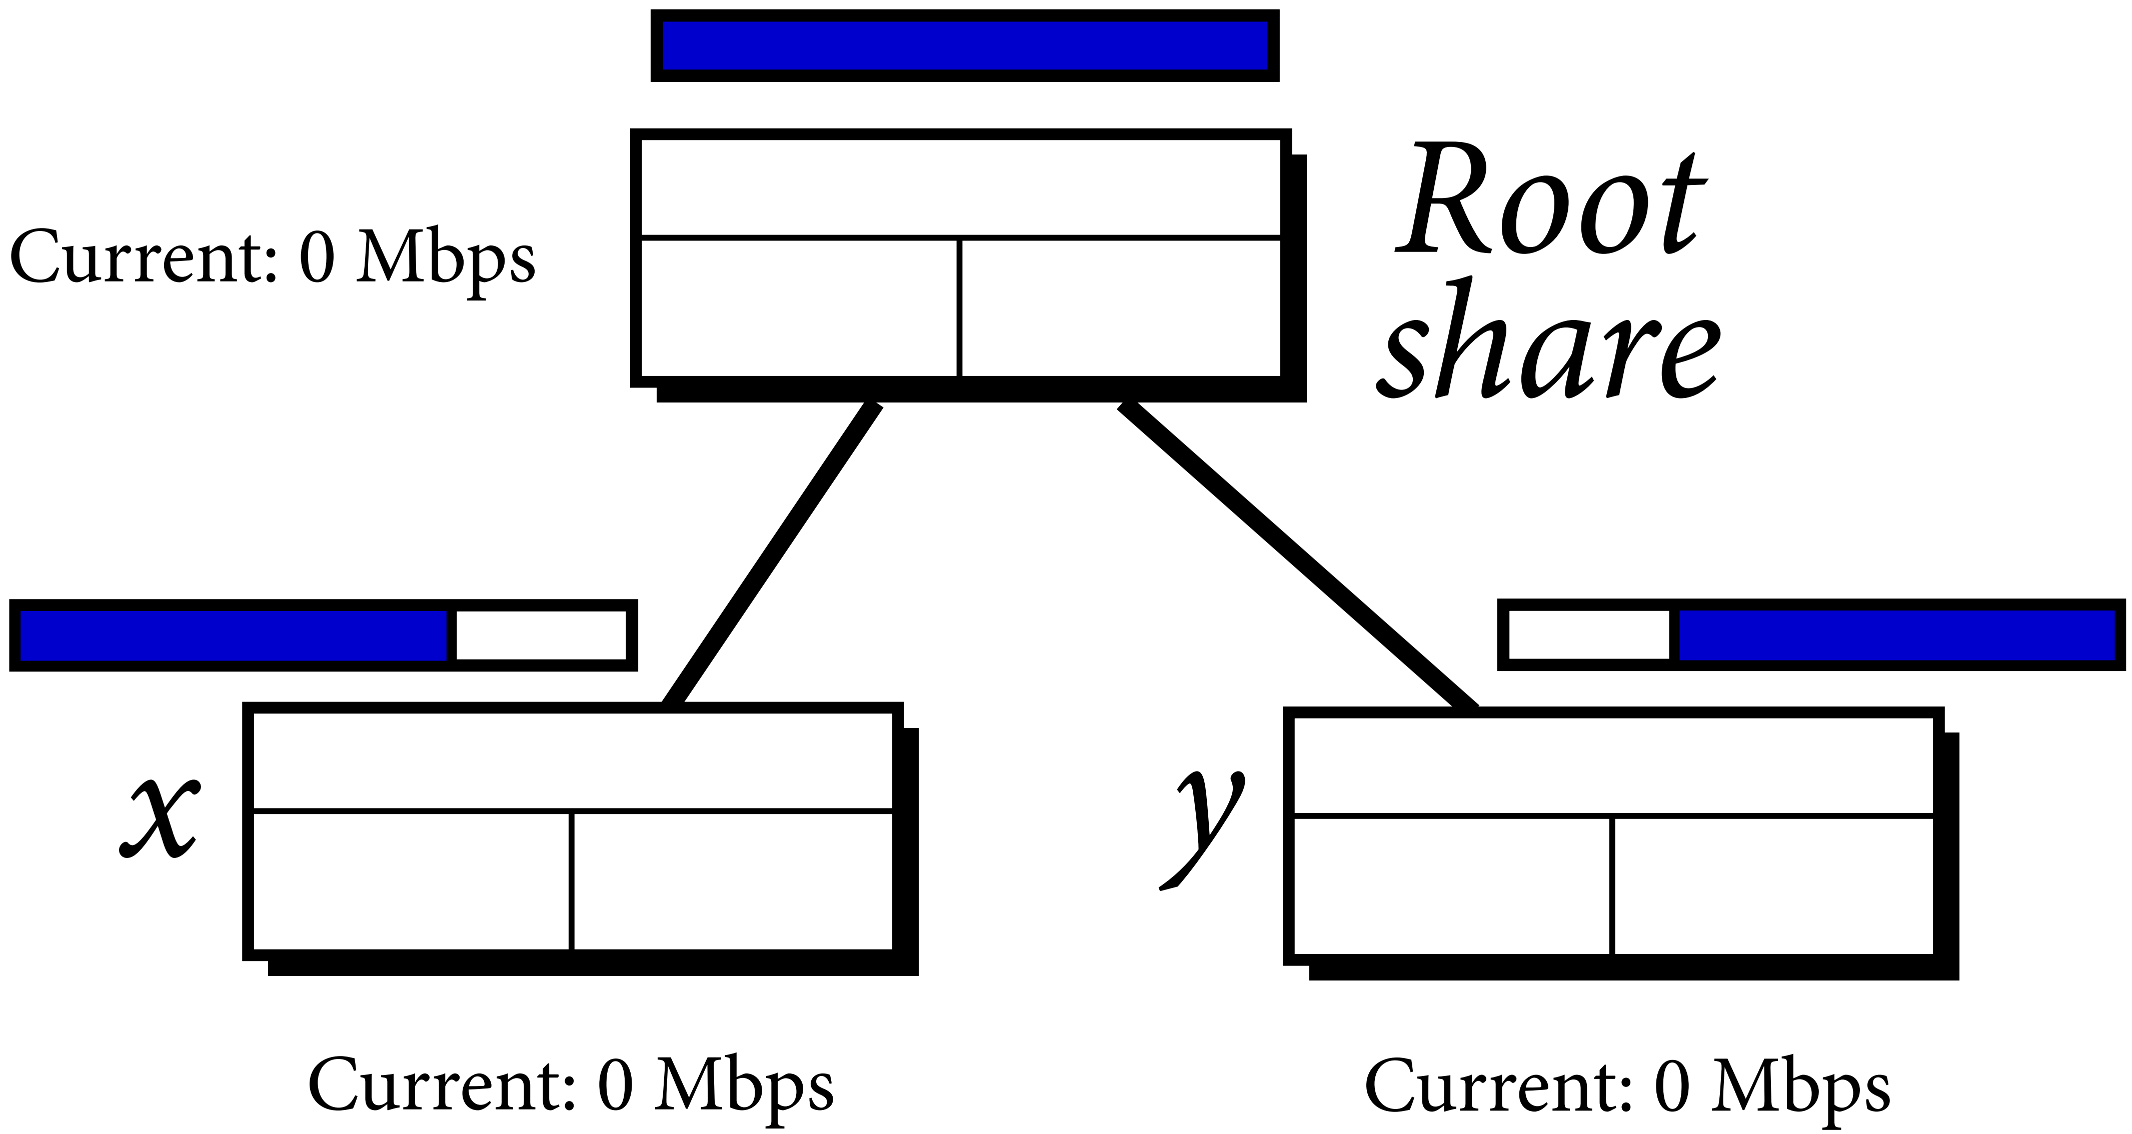

PANE

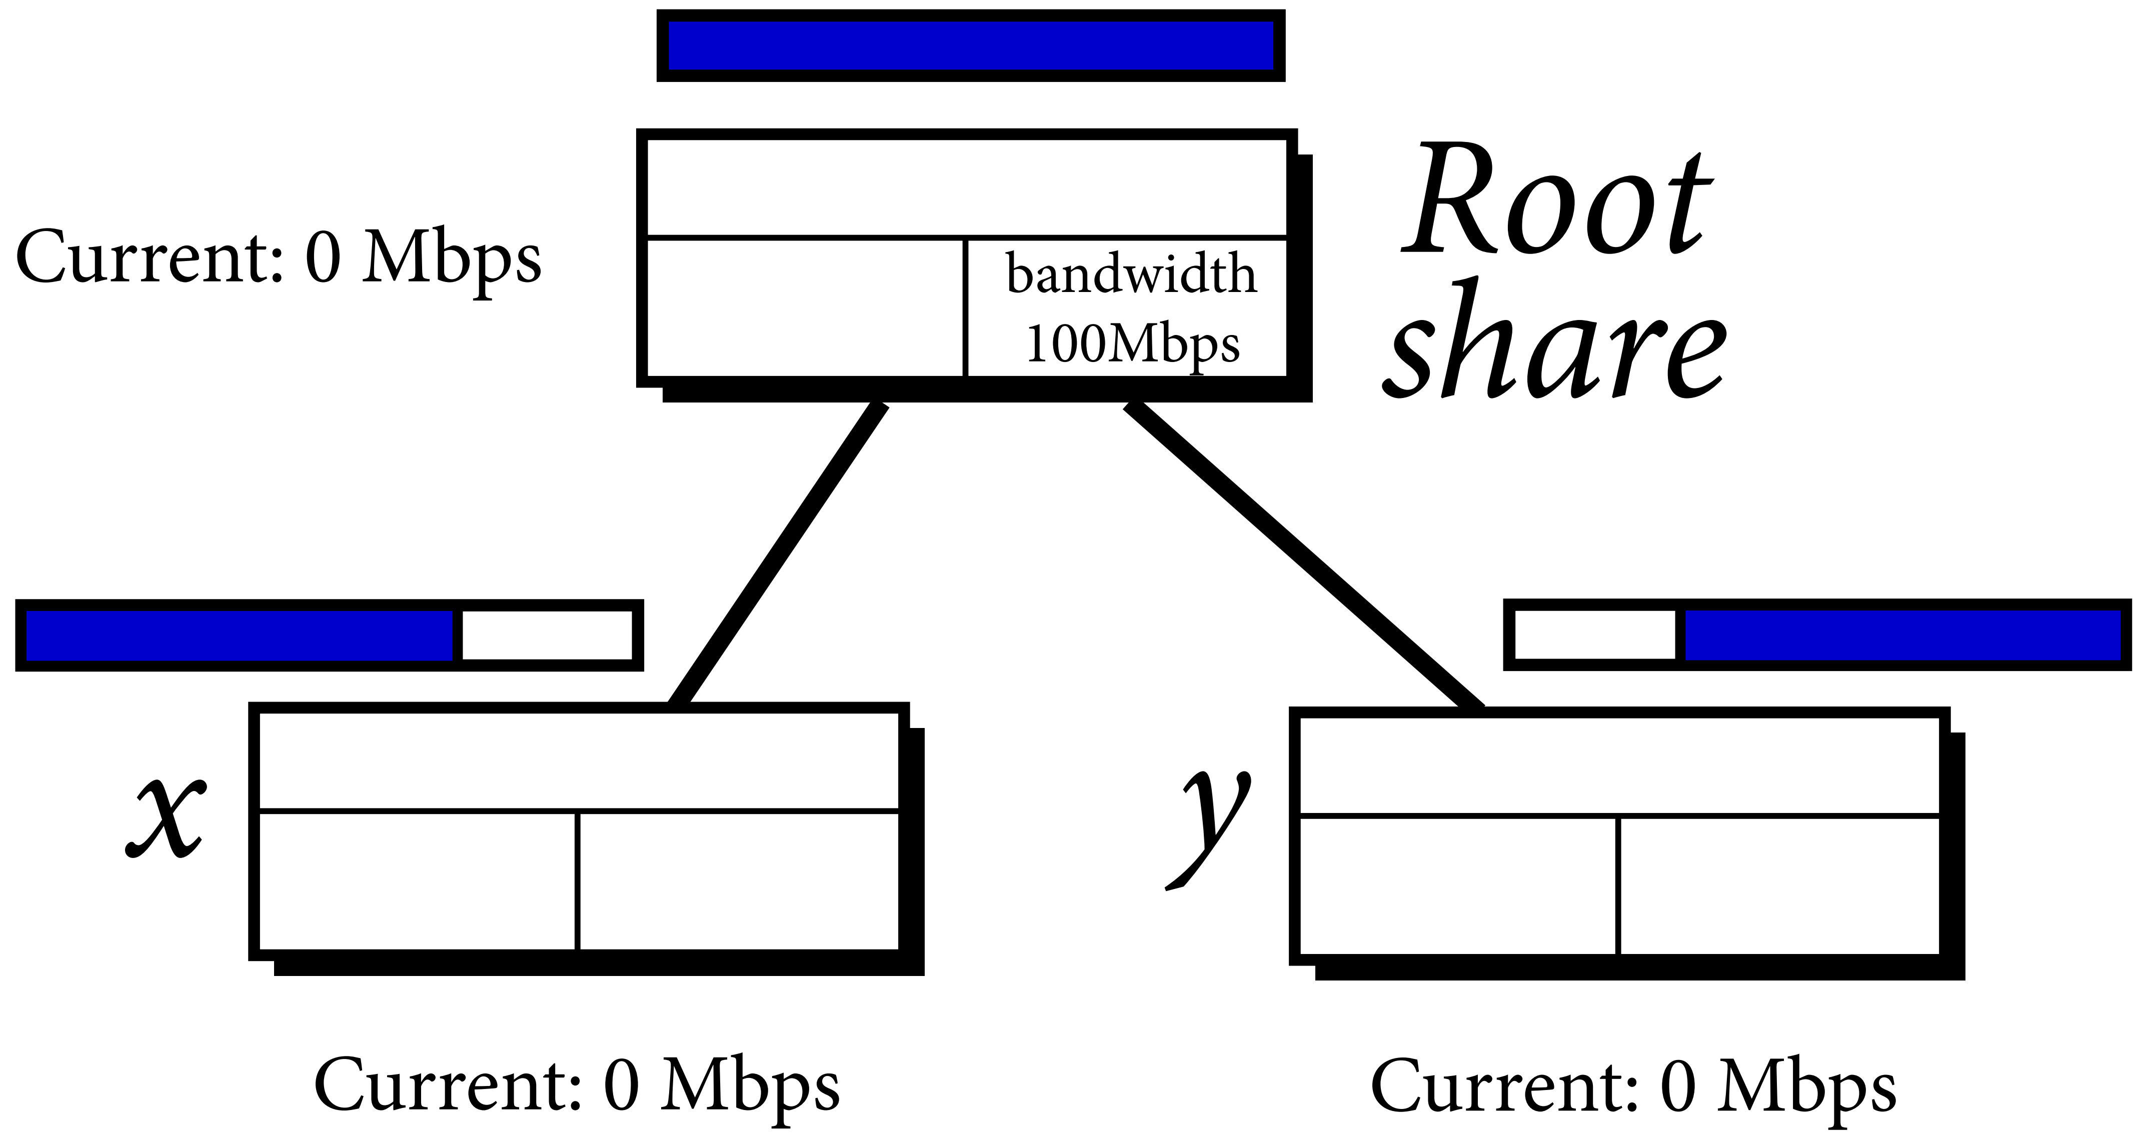
Current: 0 Mbps
bandwidth
100Mbps
Root
share
x
y
Current: 0 Mbps
Current: 0 Mbps

PANE

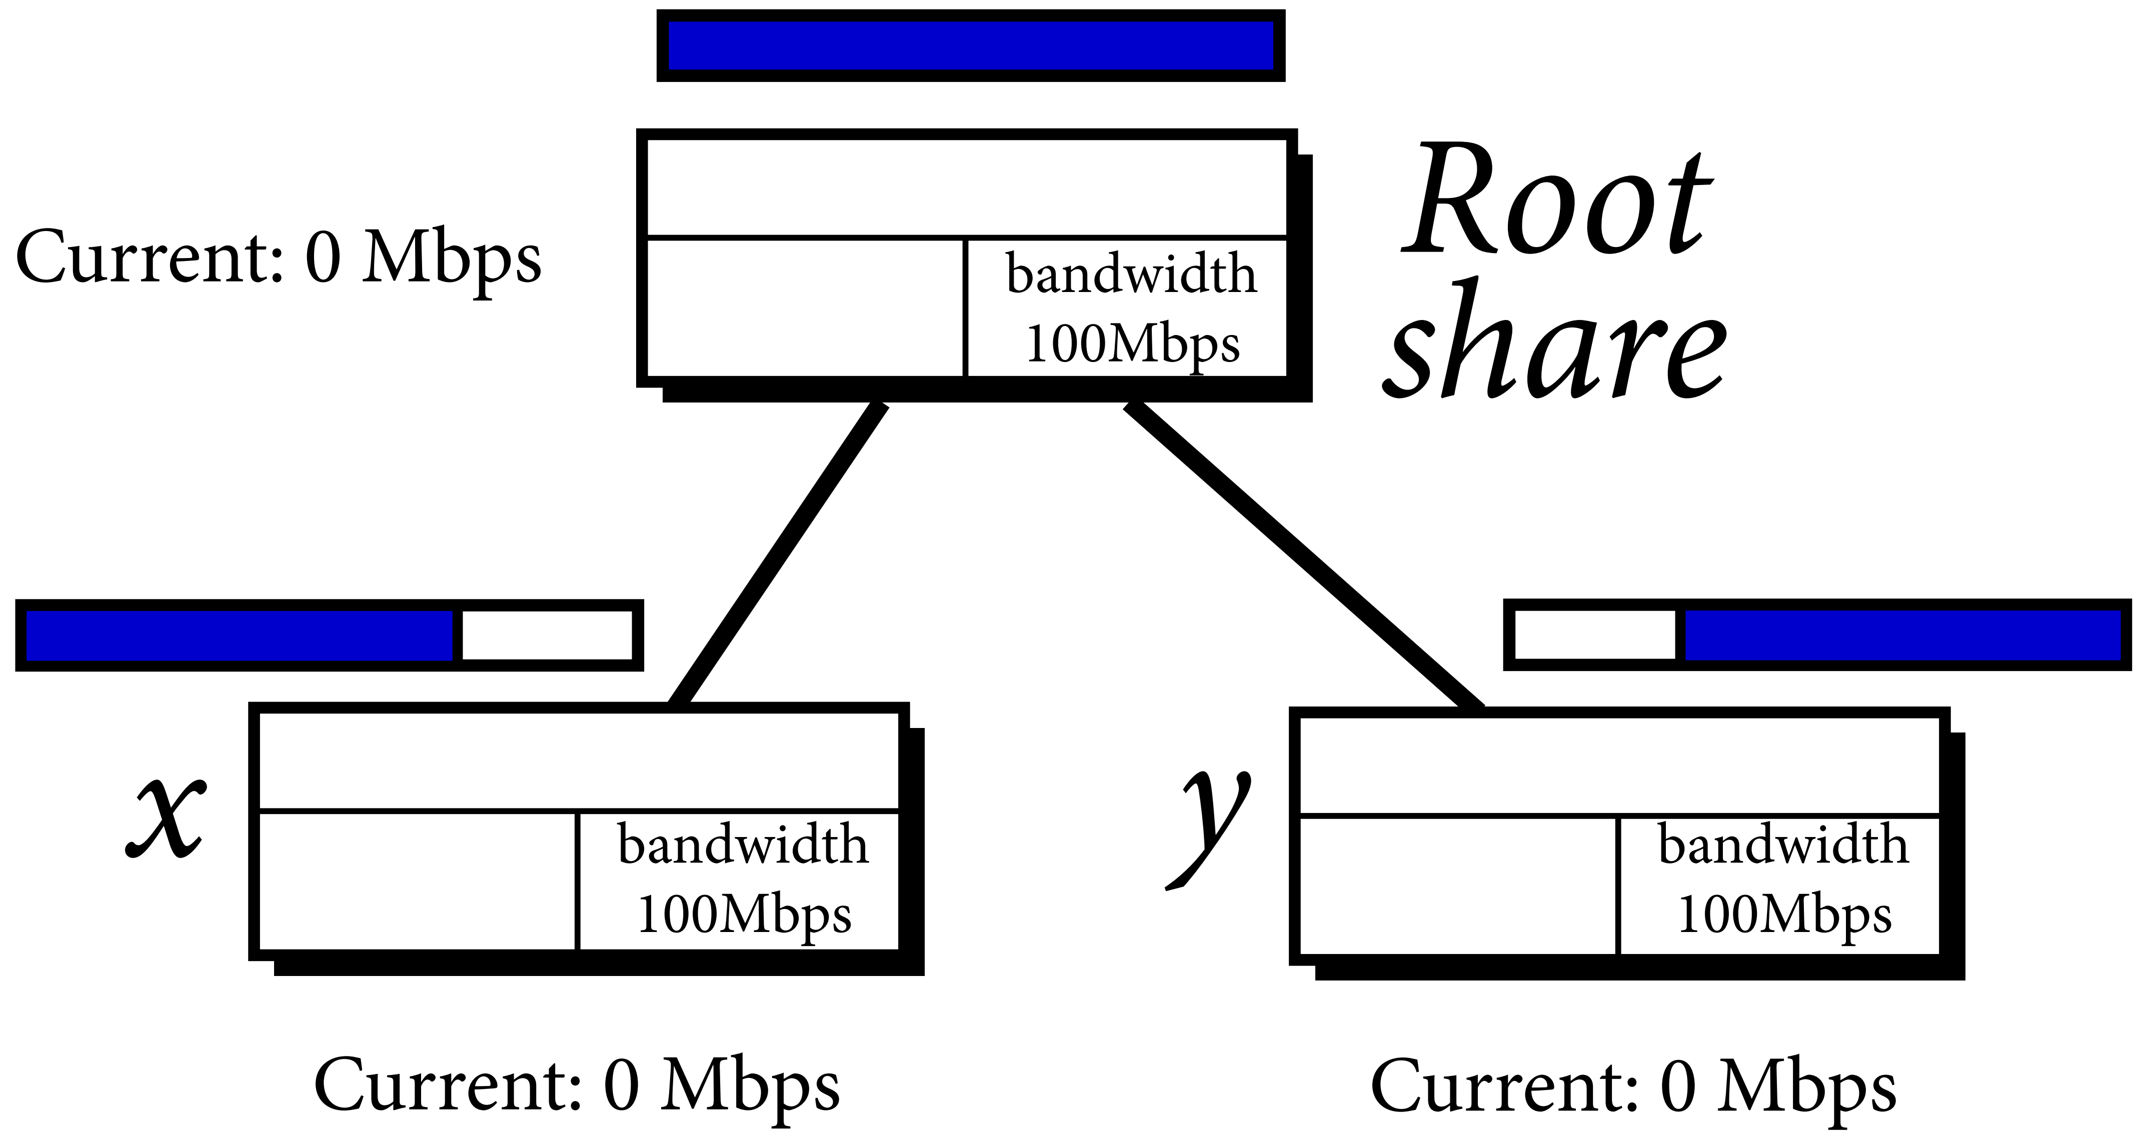
Root share
Current: 0 Mbps
bandwidth 100Mbps
x
bandwidth 100Mbps
Current: 0 Mbps
y
bandwidth 100Mbps
Current: 0 Mbps

PANE

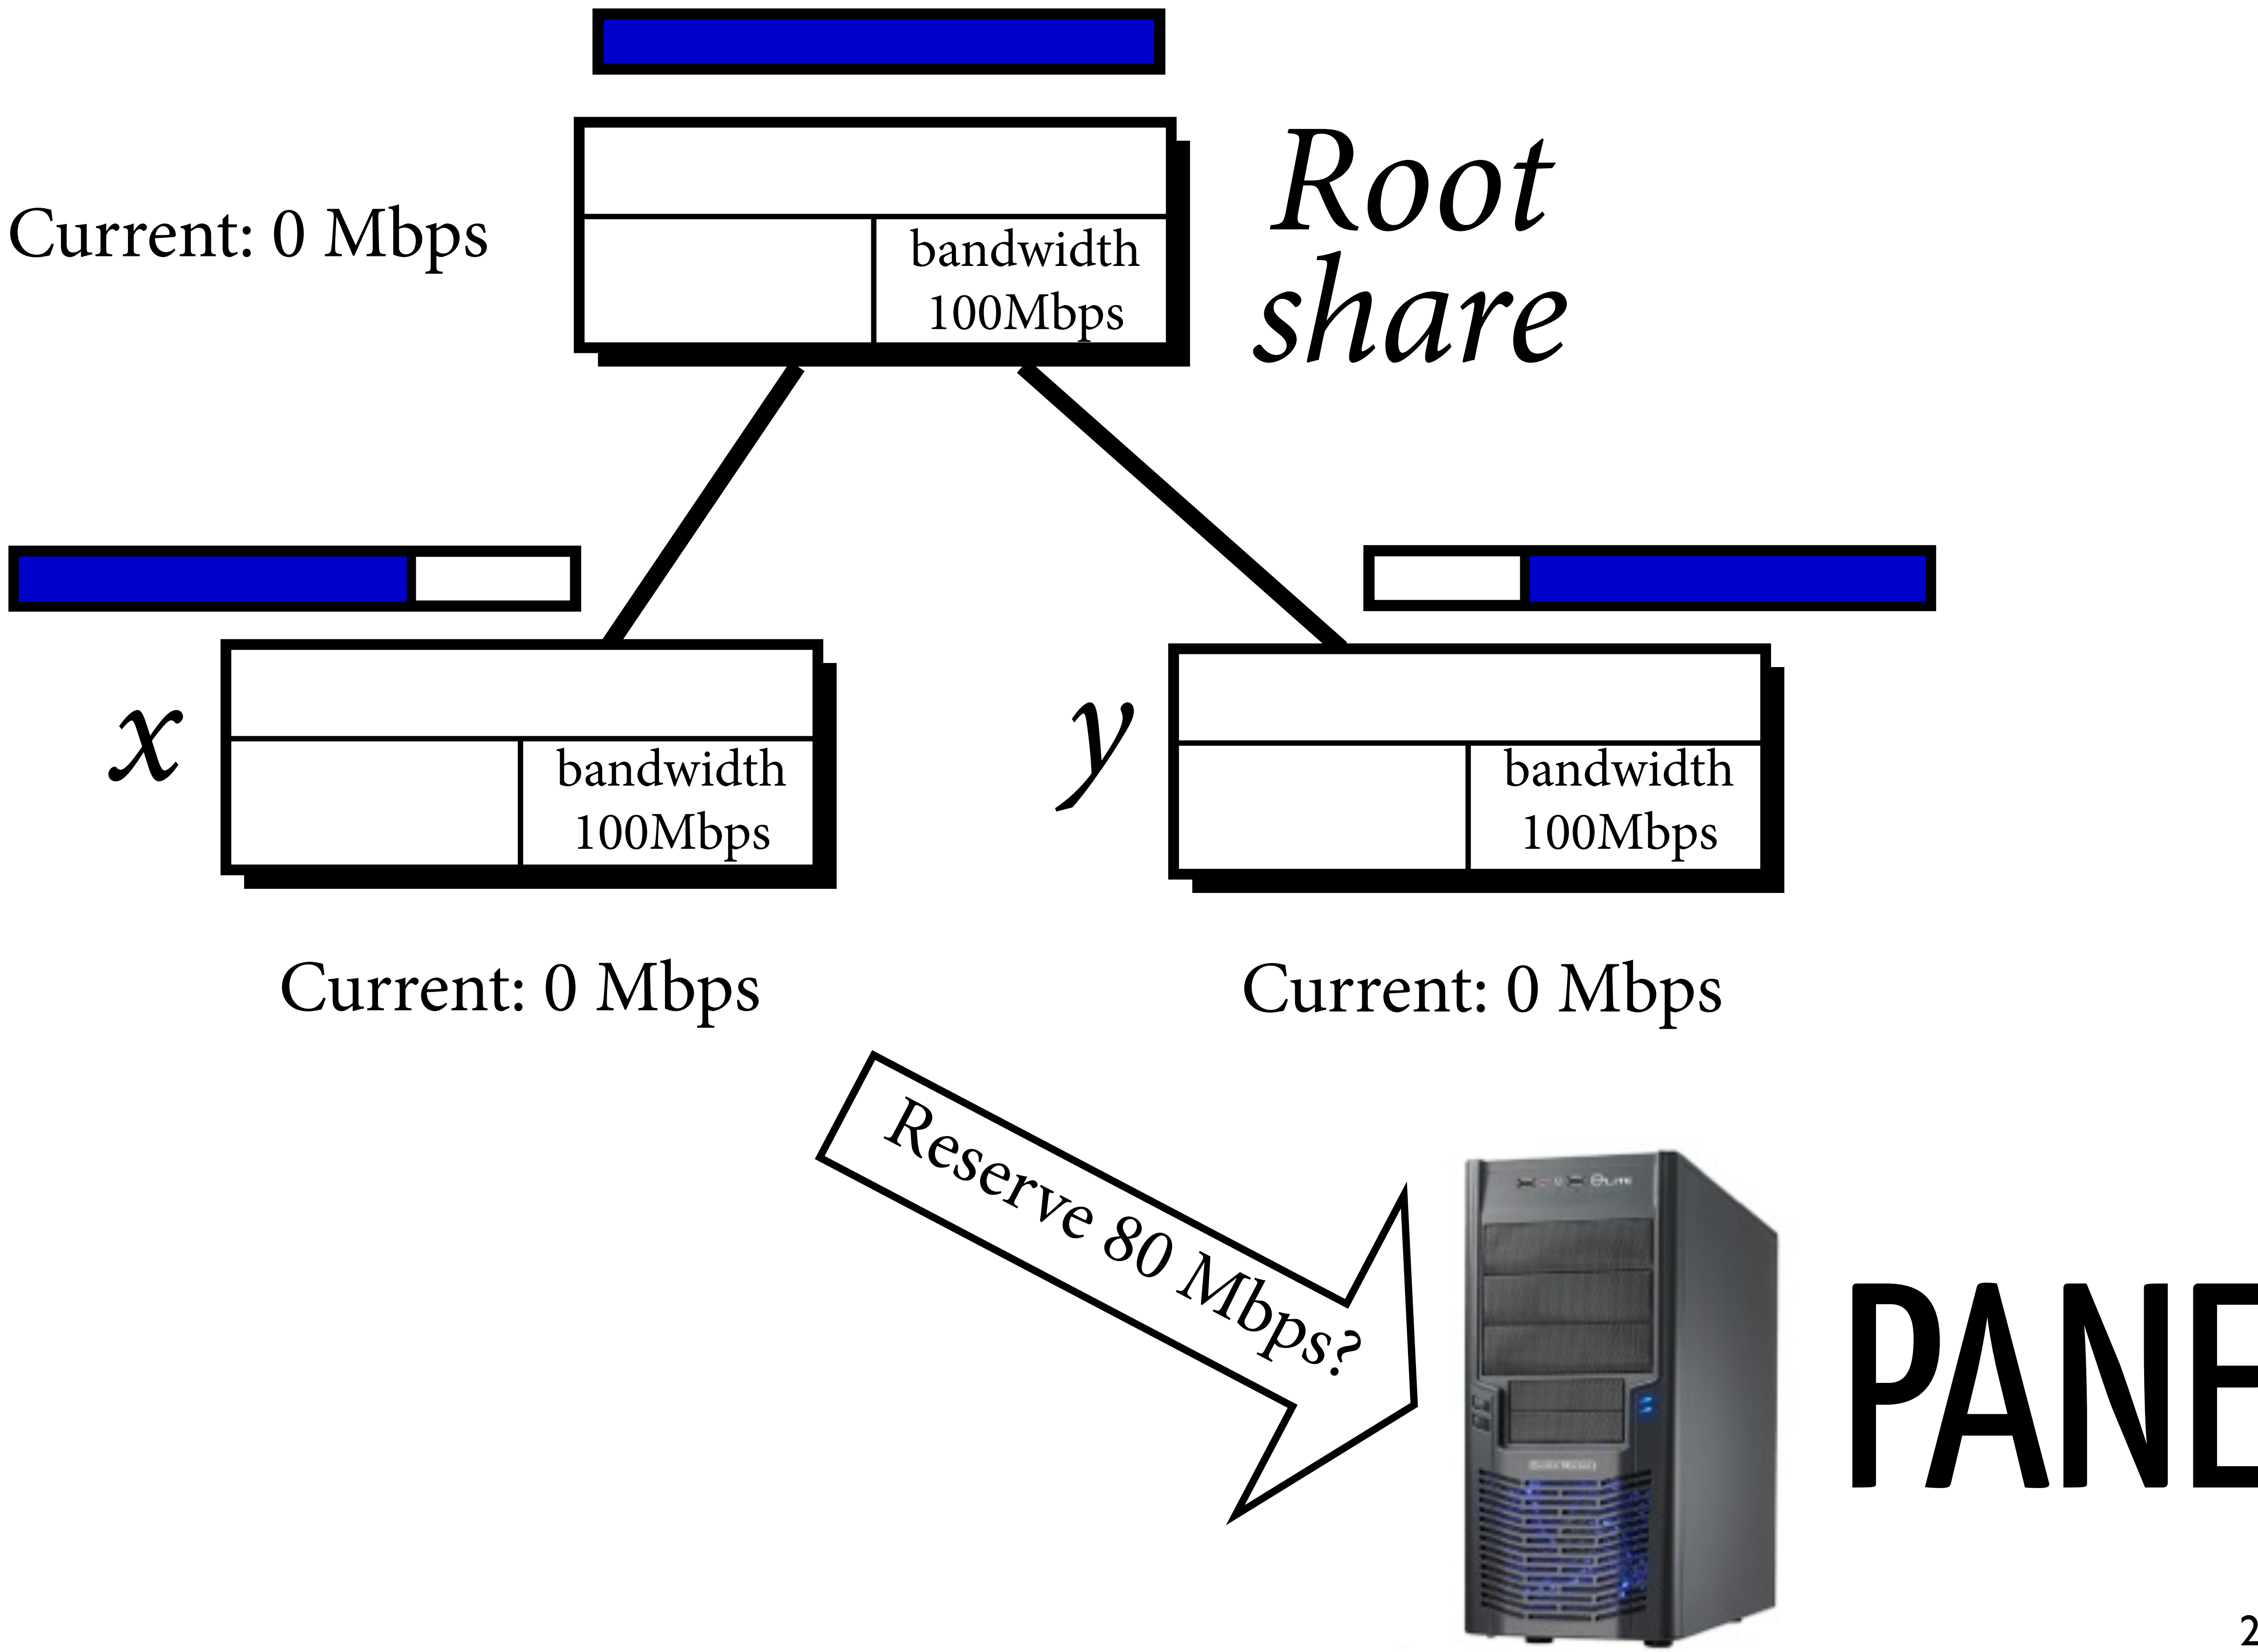
Root share
Current: 0 Mbps
bandwidth 100Mbps
x
bandwidth 100Mbps
Current: 0 Mbps
y
bandwidth 100Mbps
Current: 0 Mbps
Reserve 80 Mbps?
PANE

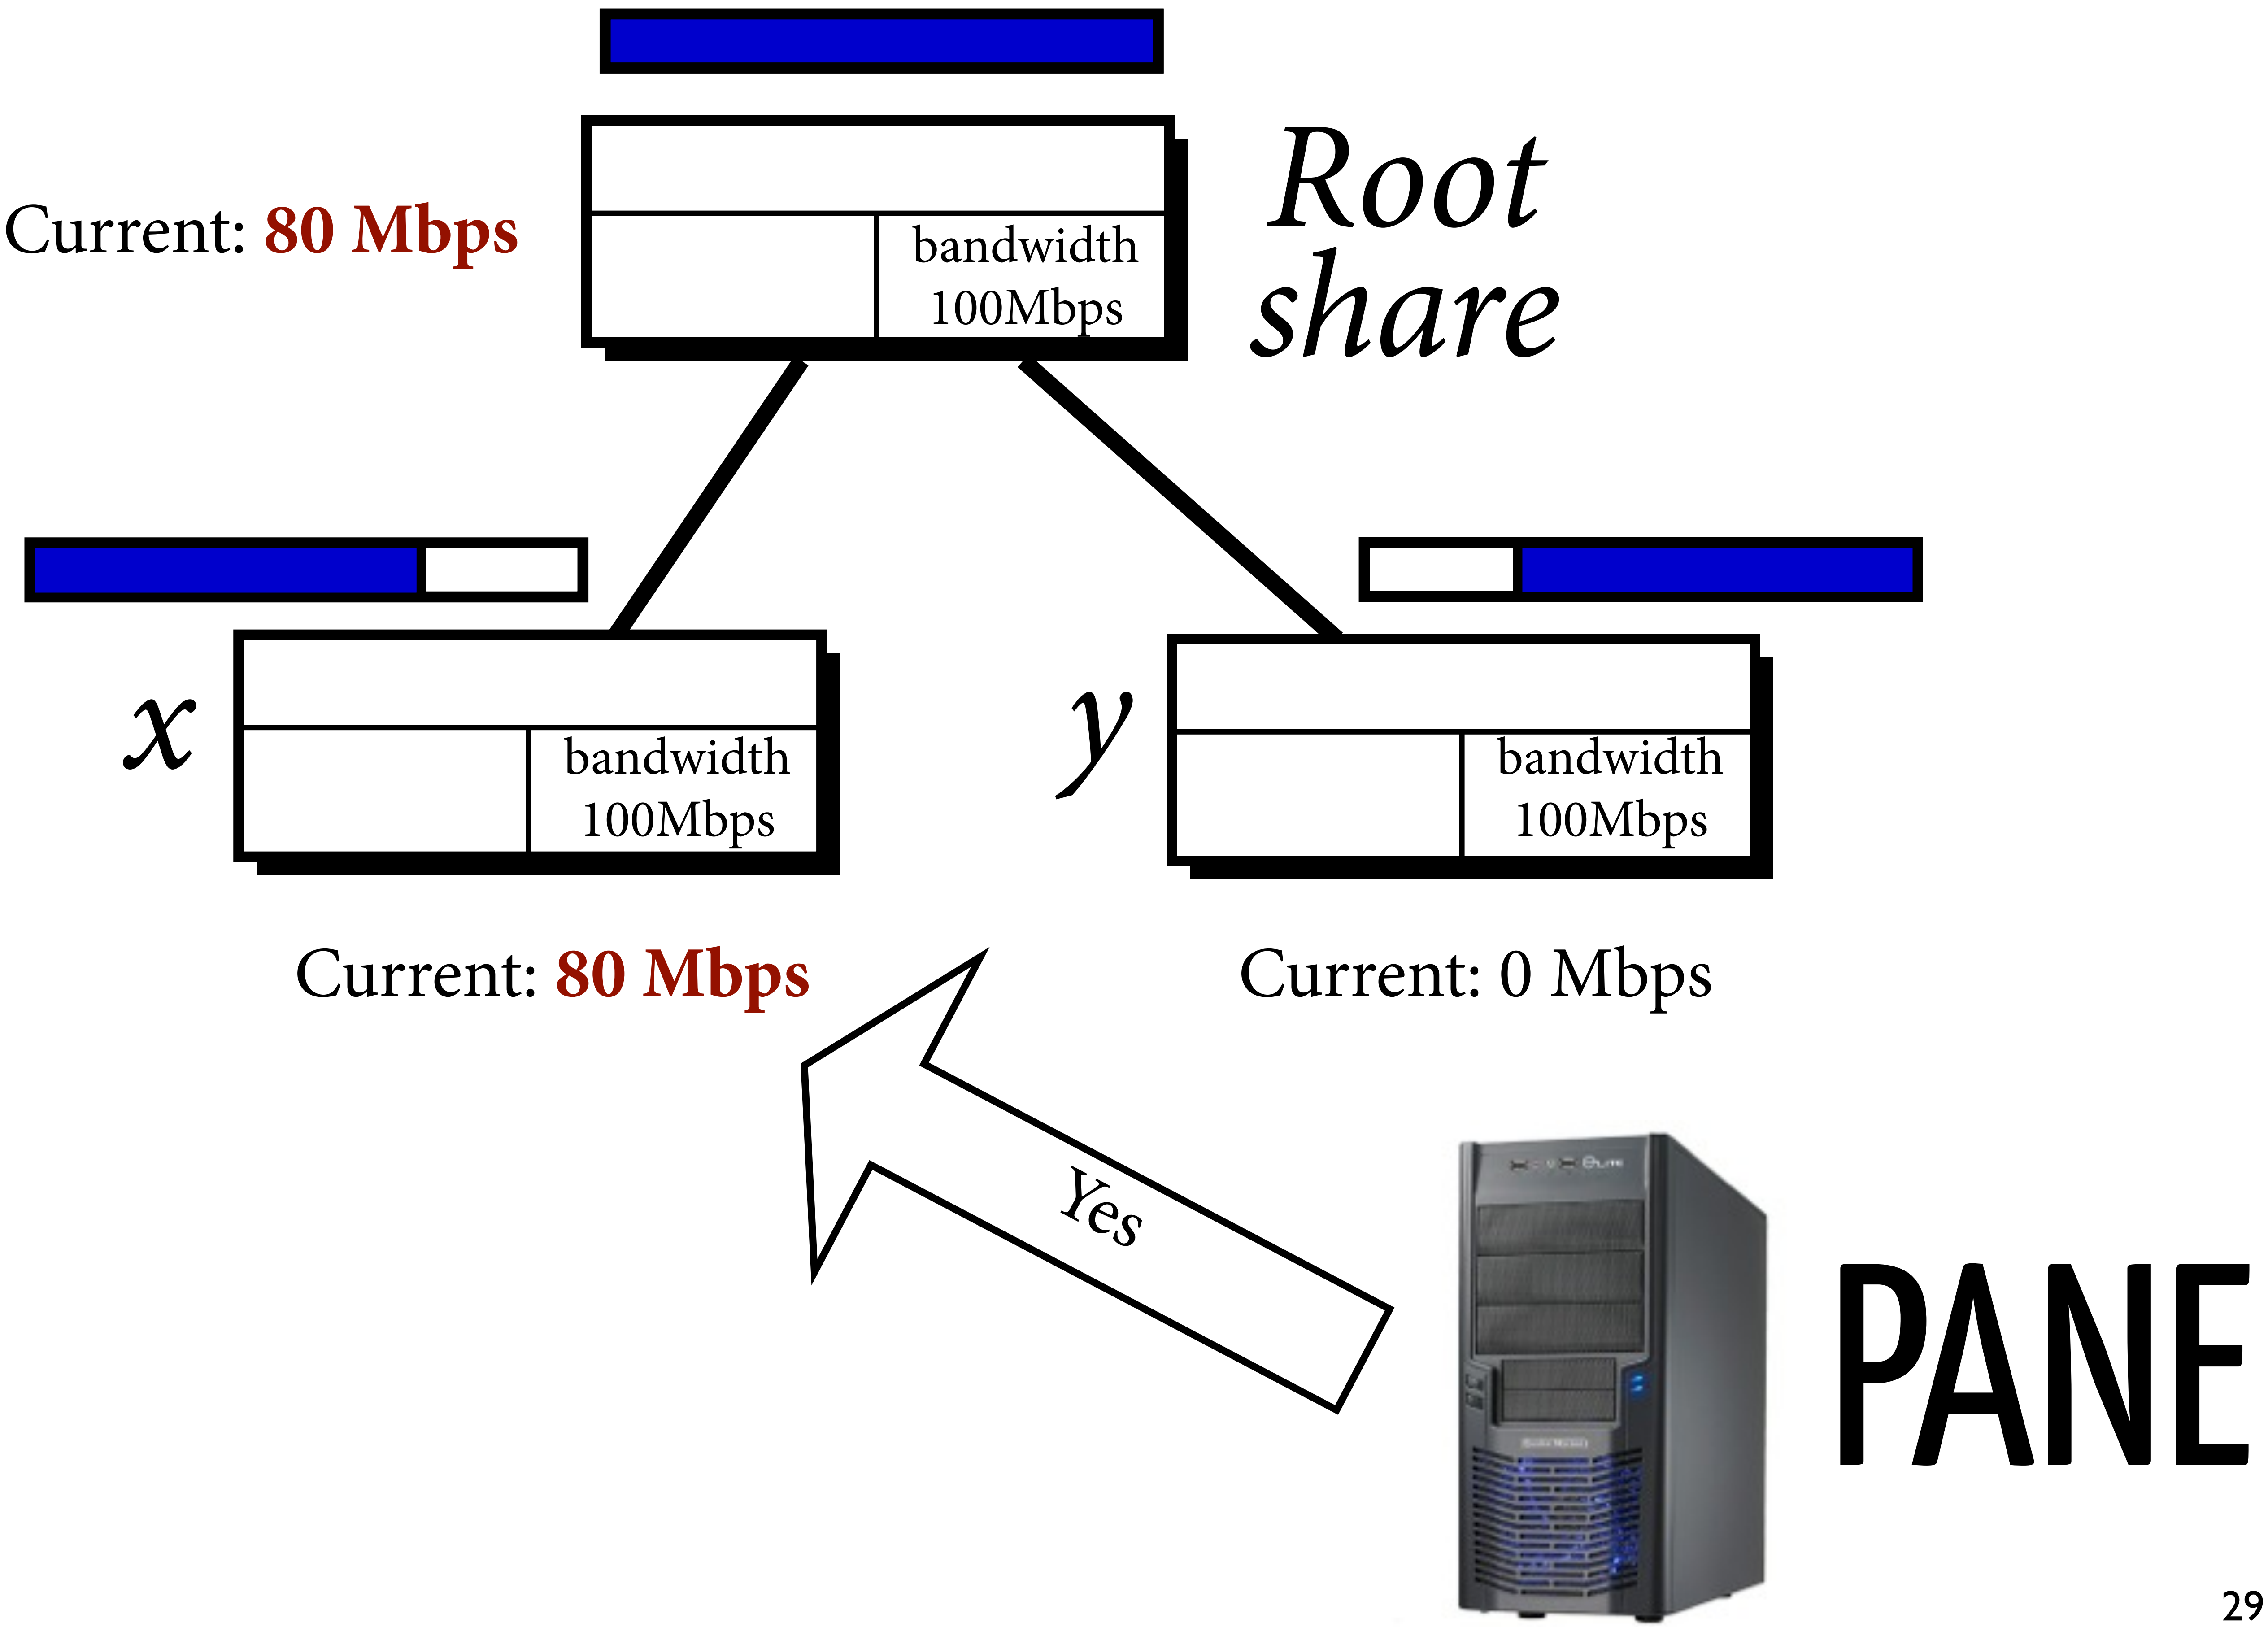
Root share
Current: 80 Mbps
bandwidth
100Mbps
x
bandwidth
100Mbps
Current: 80 Mbps
y
bandwidth
100Mbps
Current: 0 Mbps
Yes
PANE

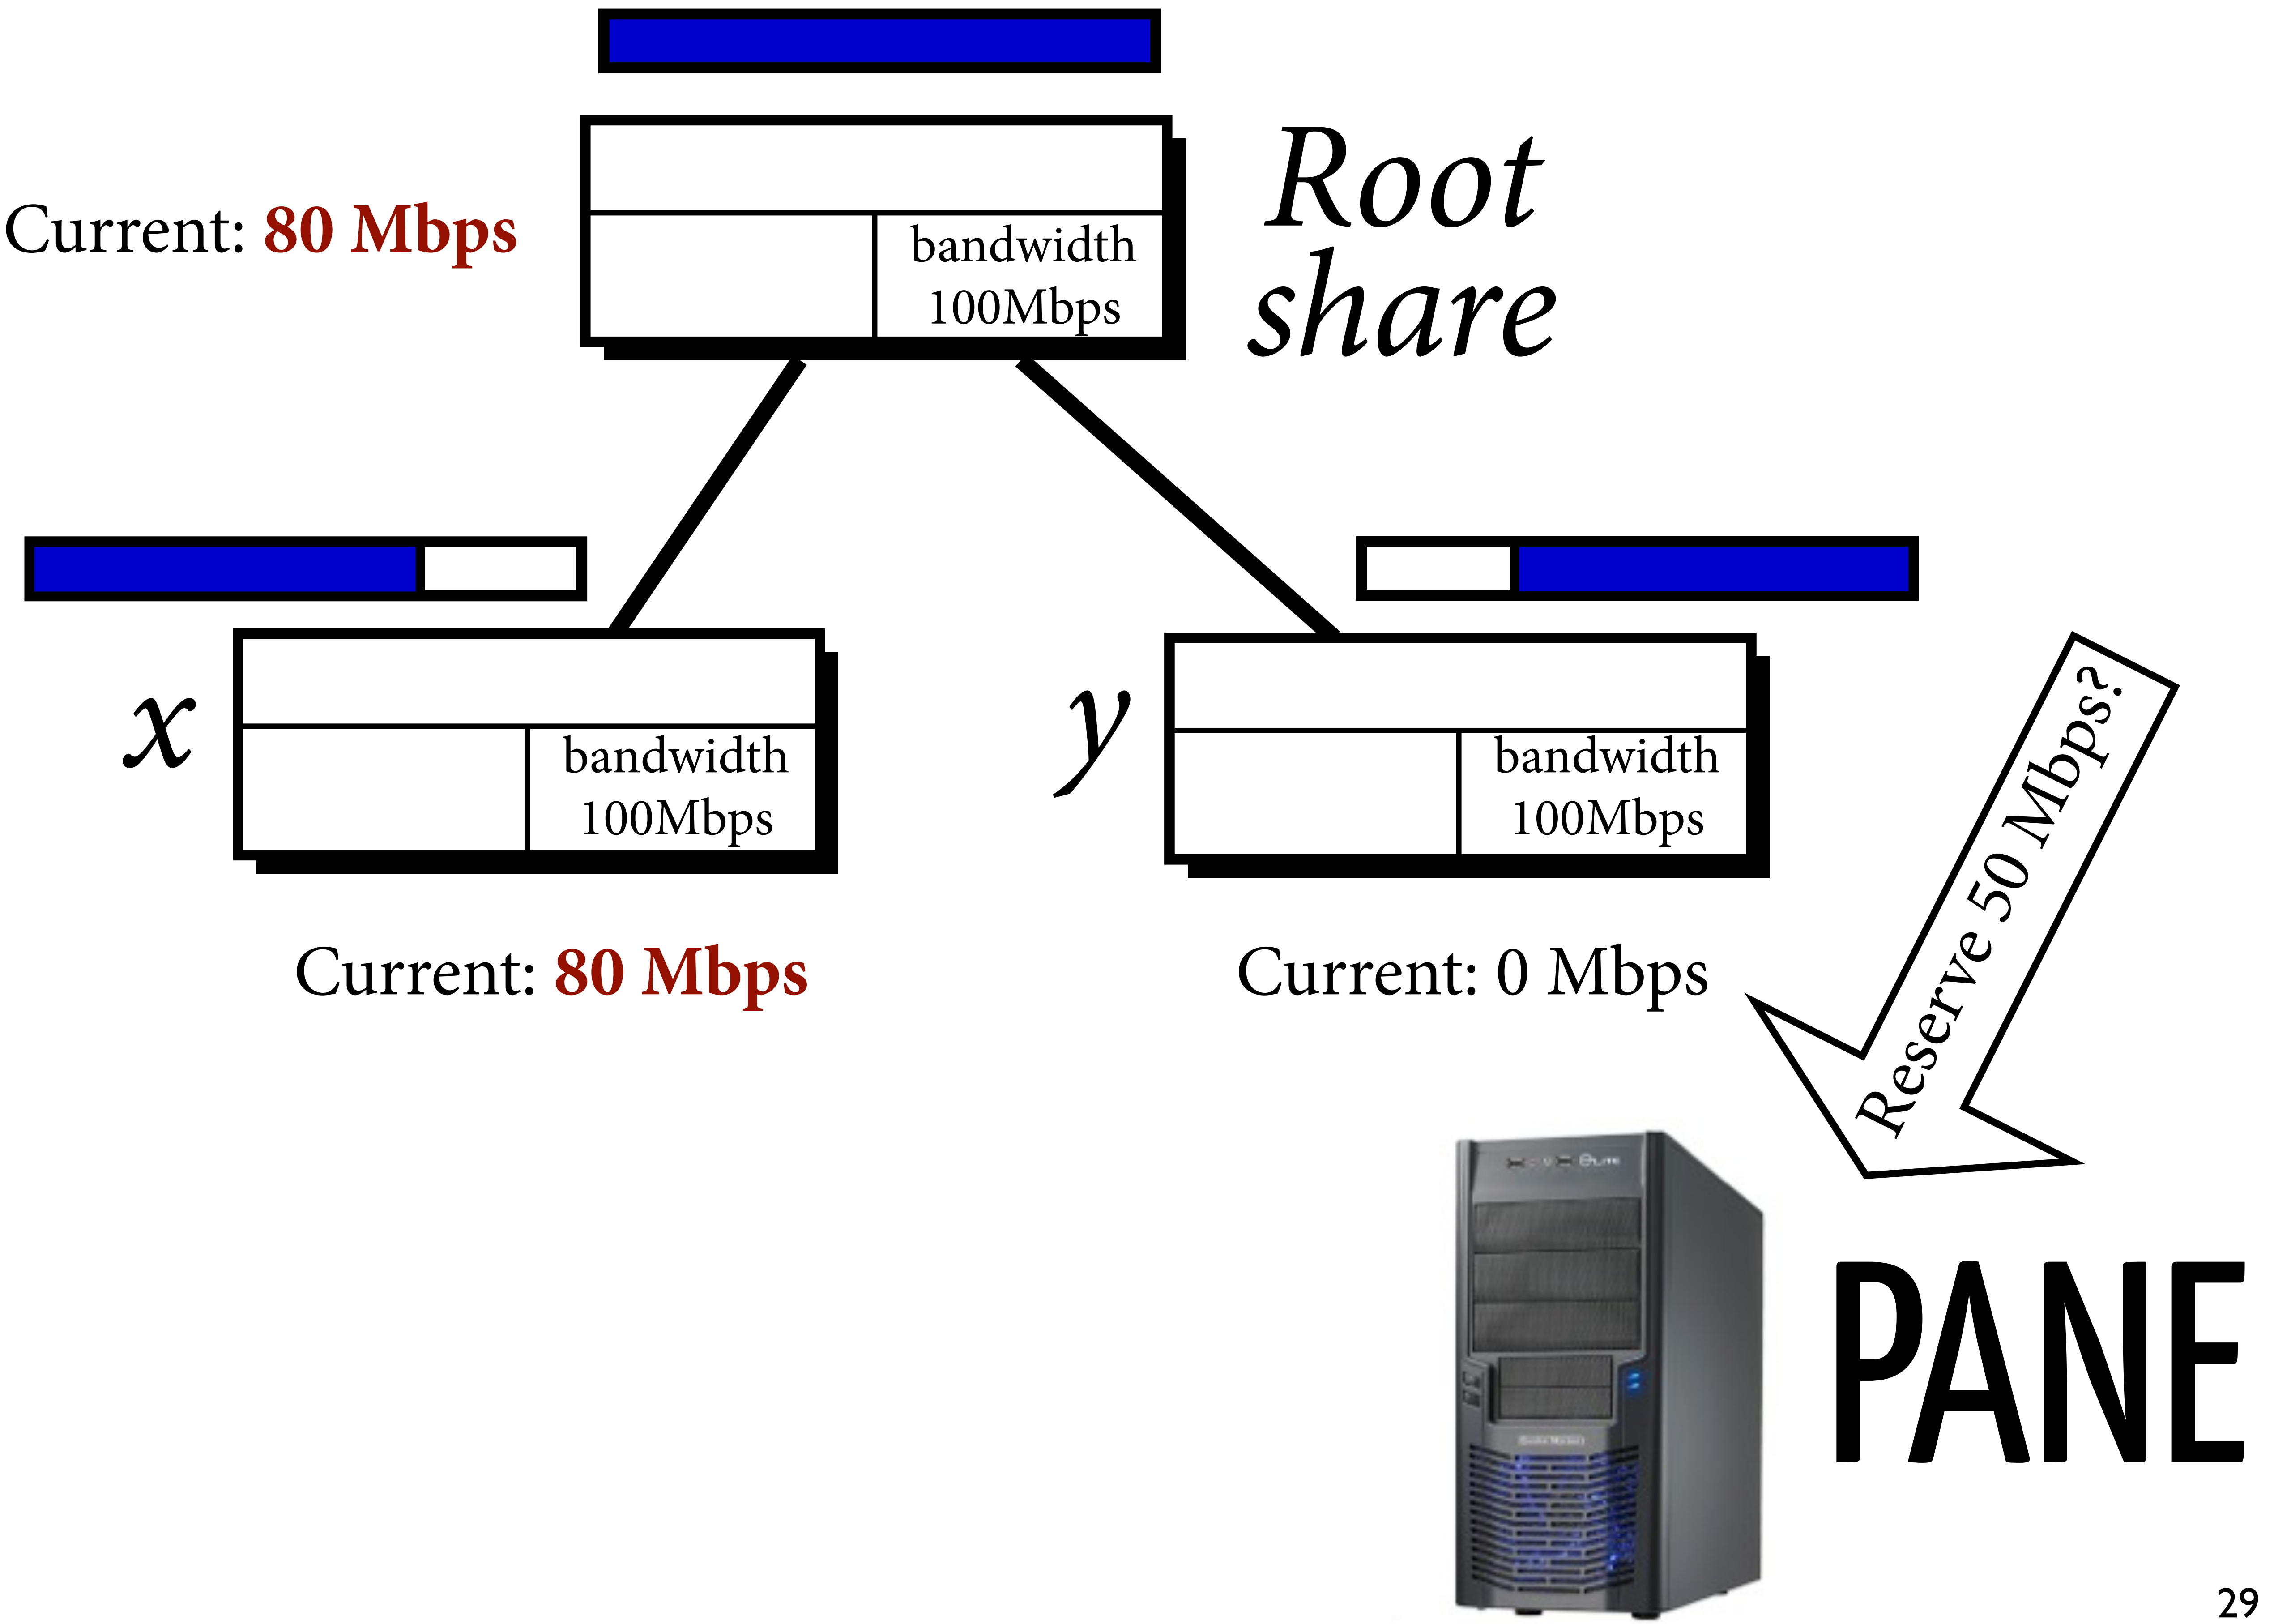
Current: 80 Mbps
bandwidth
100Mbps
Root
share
x
bandwidth
100Mbps
y
bandwidth
100Mbps
Current: 80 Mbps
Current: 0 Mbps
Reserve 50 Mbps?
PANE

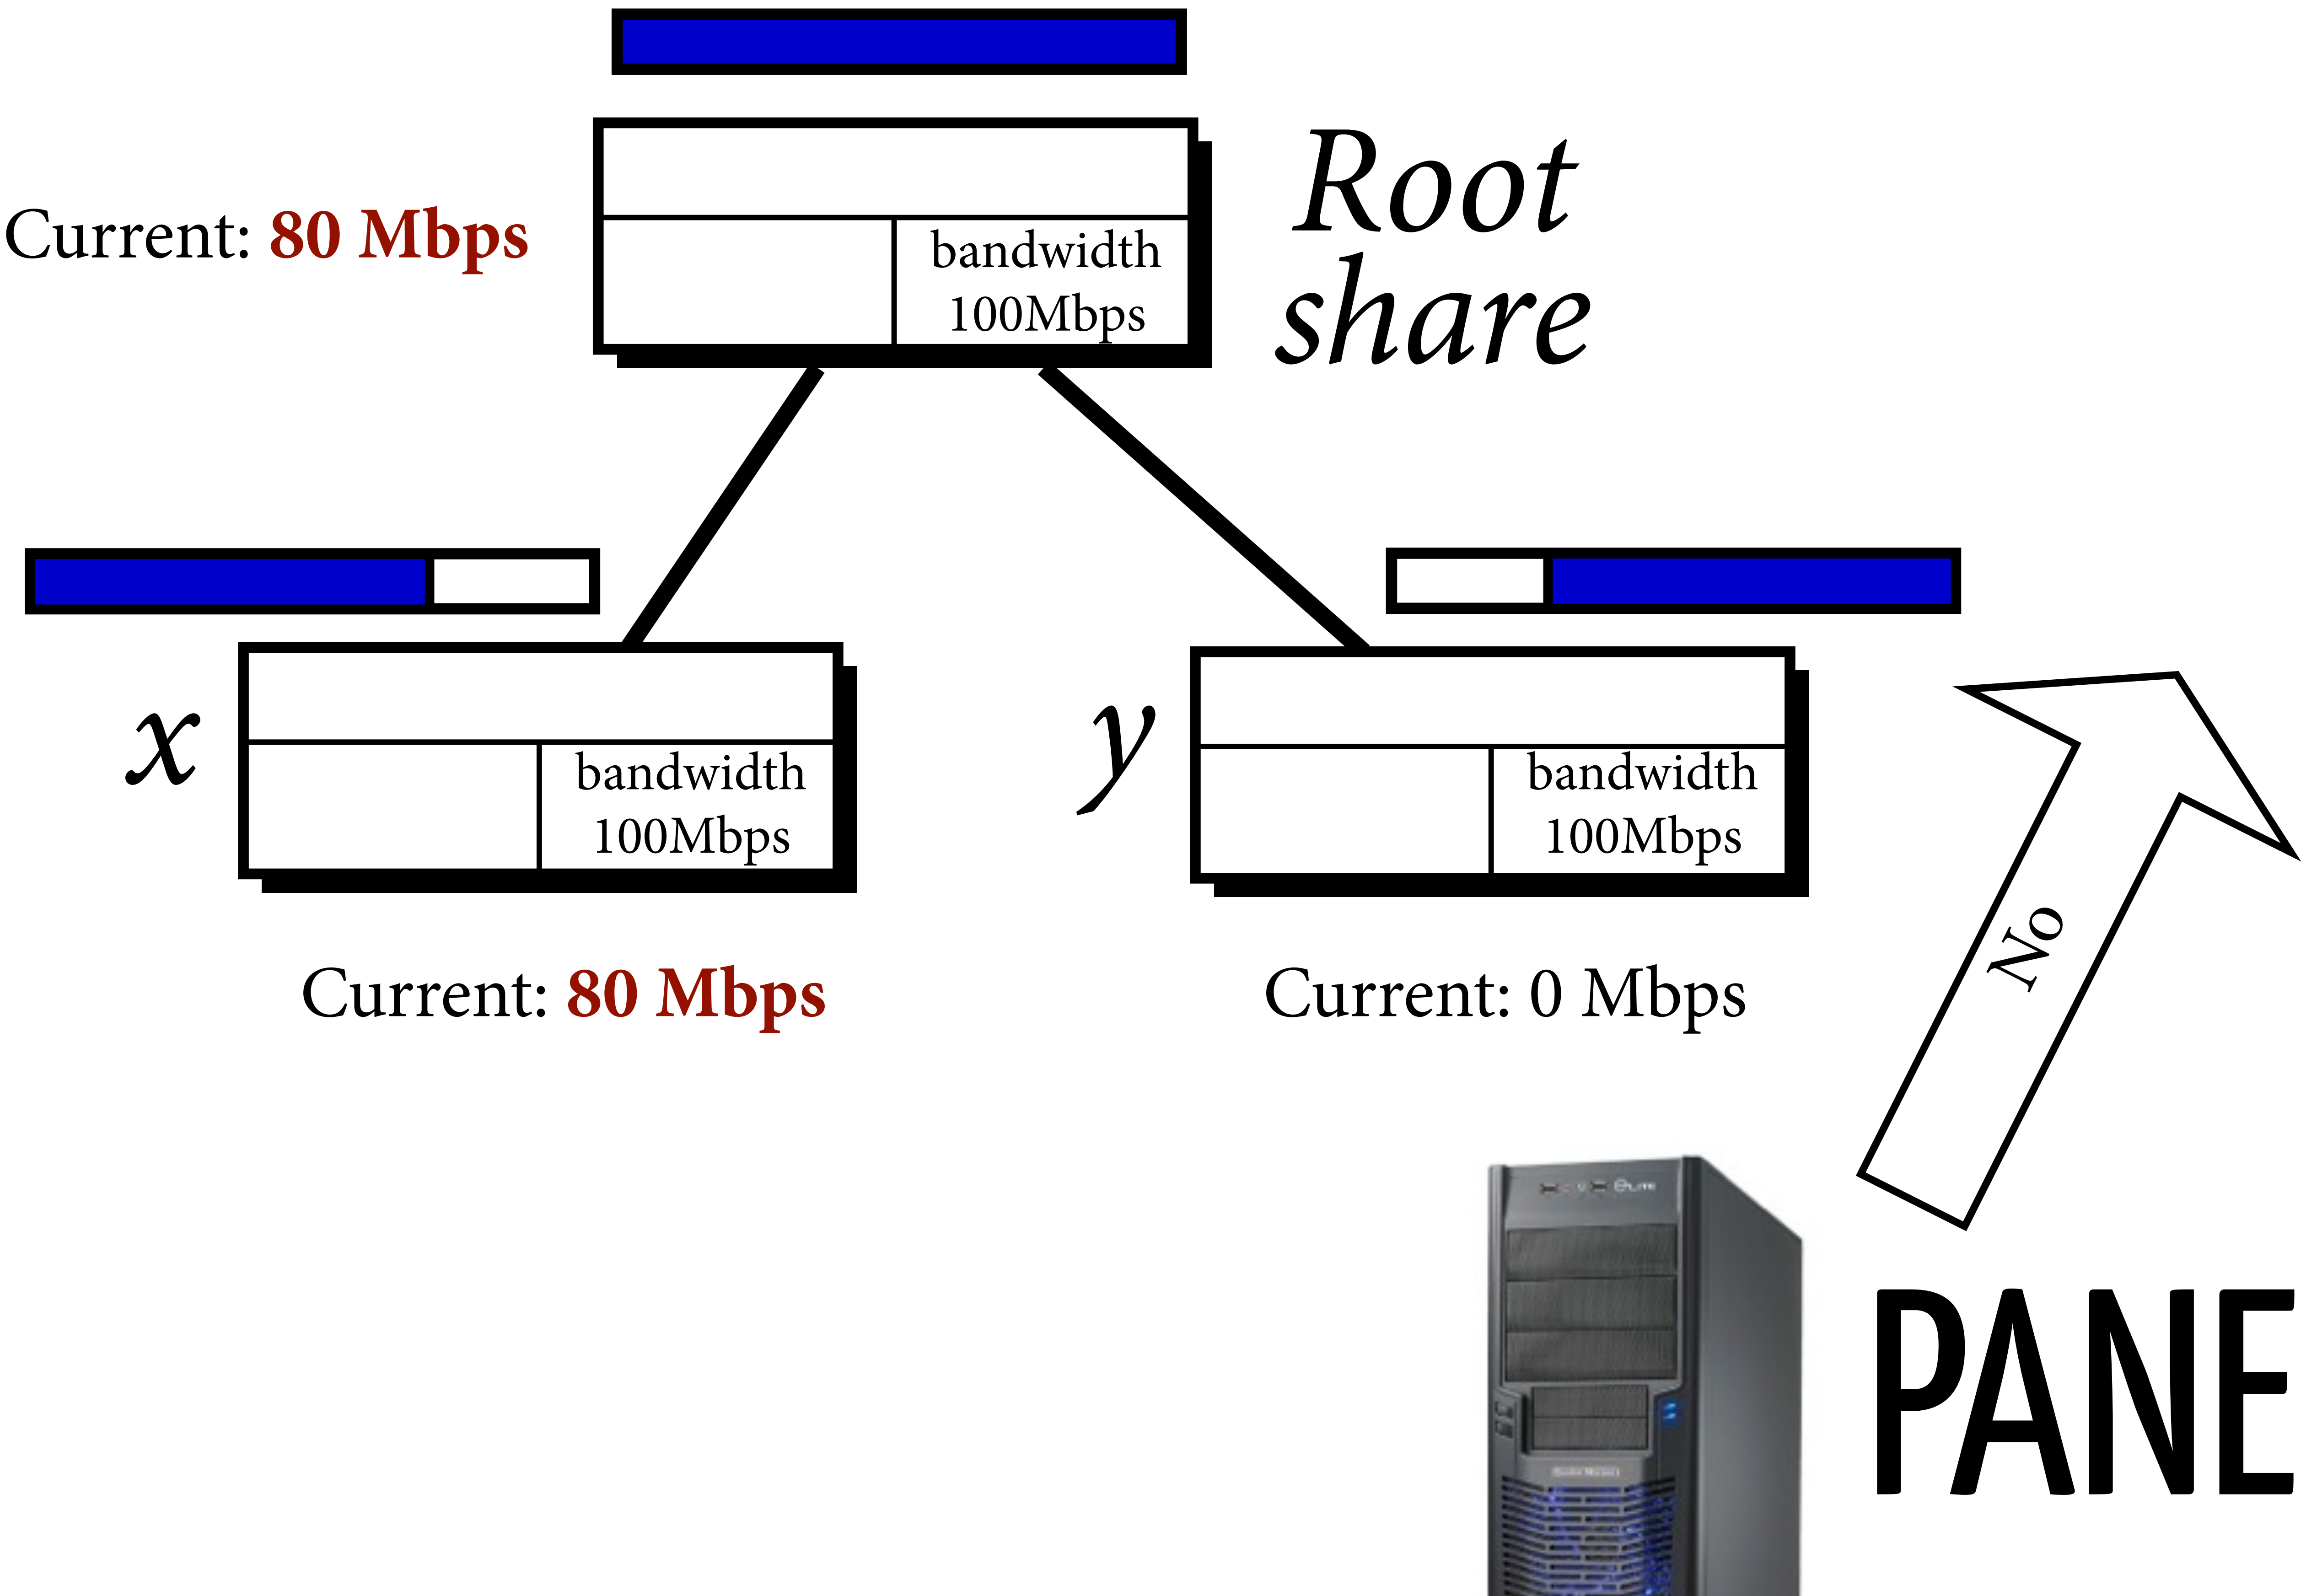
Root share
bandwidth 100Mbps
Current: 80 Mbps
x
bandwidth 100Mbps
Current: 80 Mbps
y
bandwidth 100Mbps
Current: 0 Mbps
No
PANE

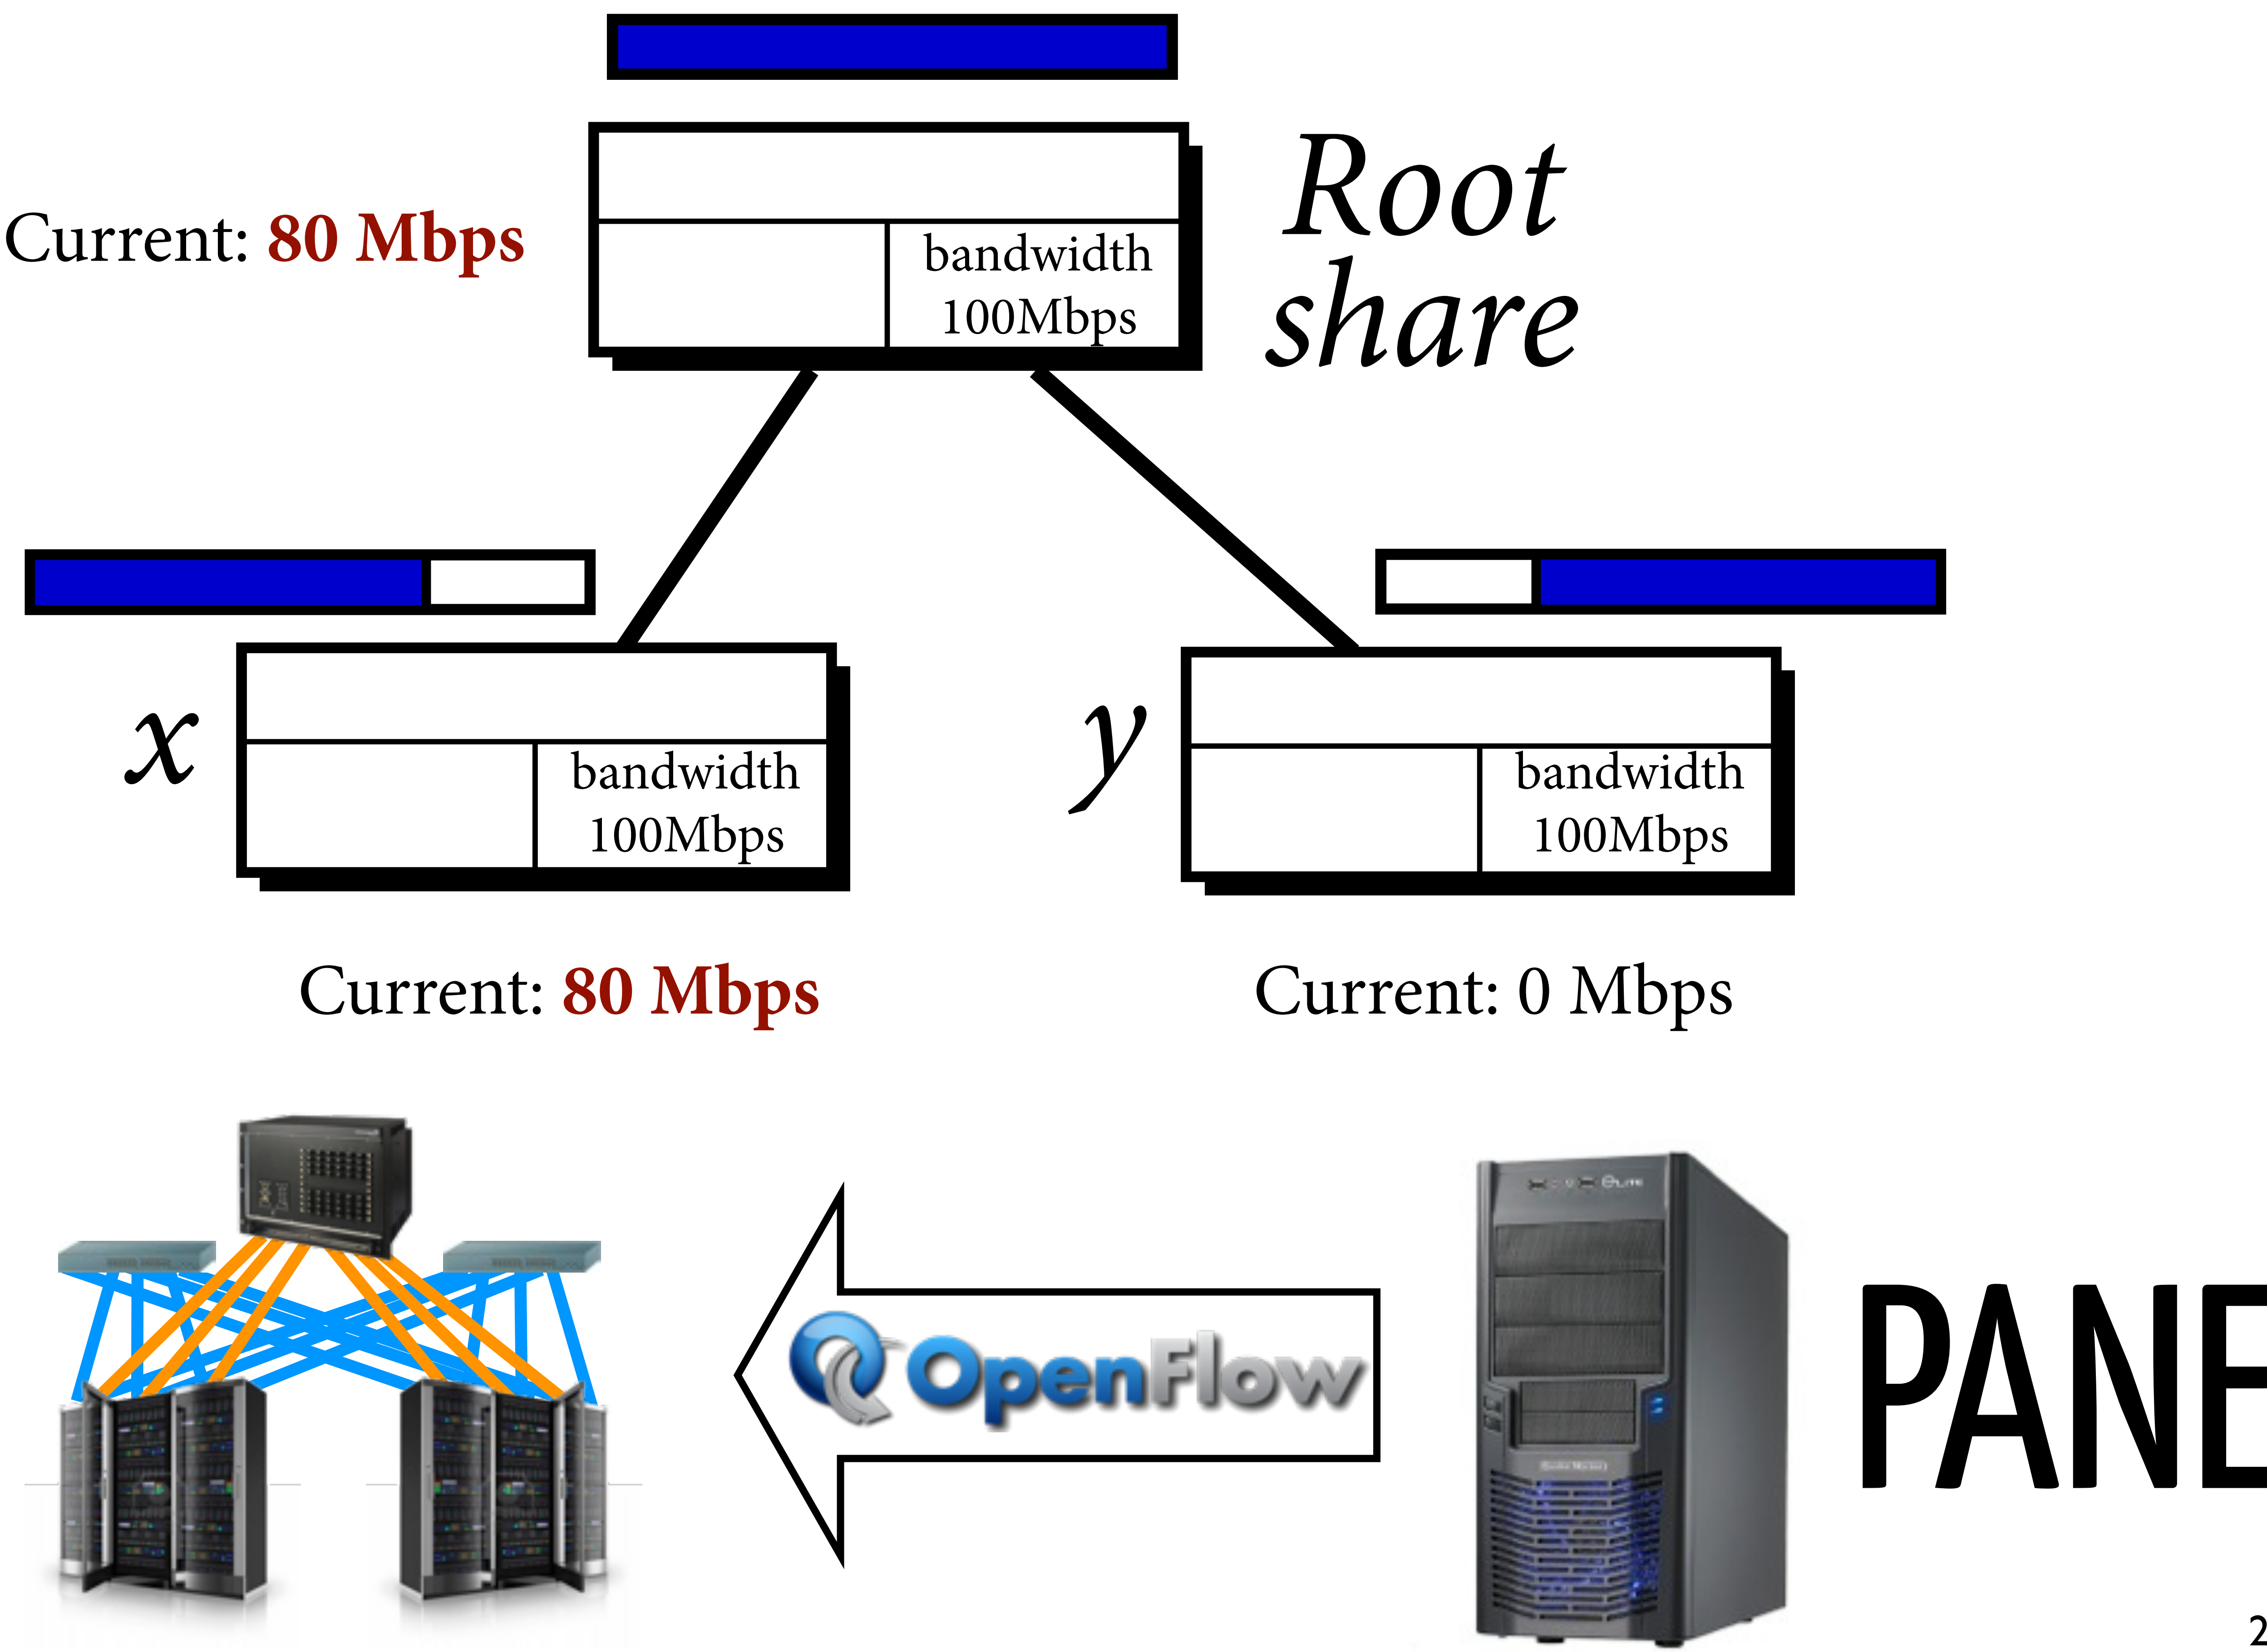
Root share
Current: 80 Mbps
bandwidth 100Mbps
x
bandwidth 100Mbps
Current: 80 Mbps
y
bandwidth 100Mbps
Current: 0 Mbps
OpenFlow
PANE

Protocol Sketch

PANE

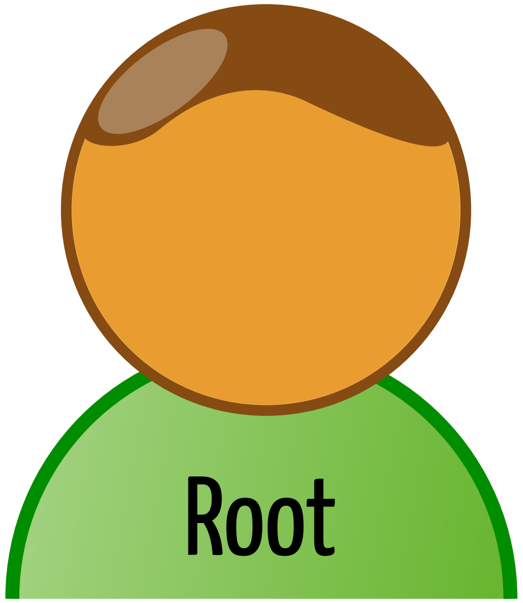

PANE

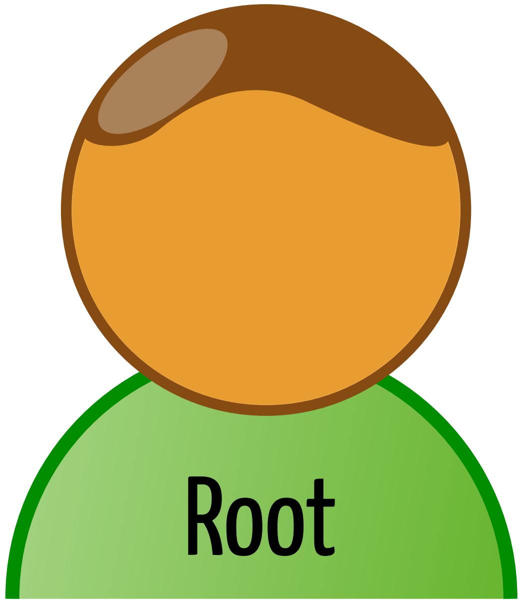

PANE

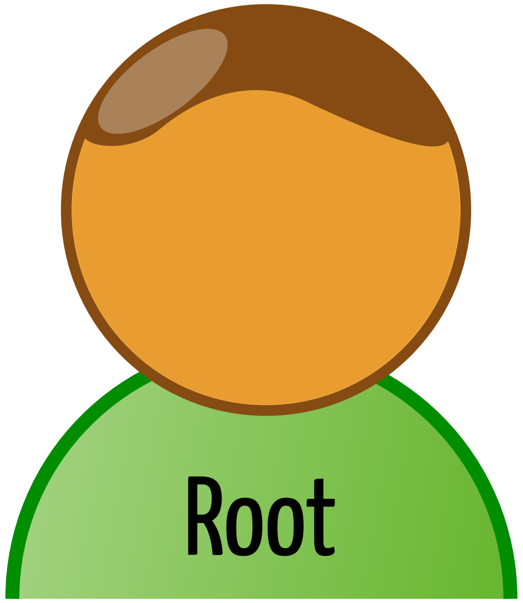


```
NewShare A for
(user=Alice) [reserve <= 10Mb]
on rootShare.
```

PANE

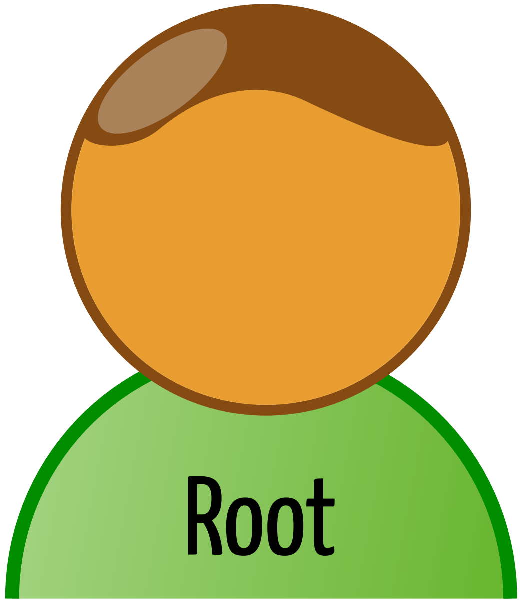

NewShare A for
(user=Alice) [reserve <= 10Mb]
on rootShare.

OK

PANE

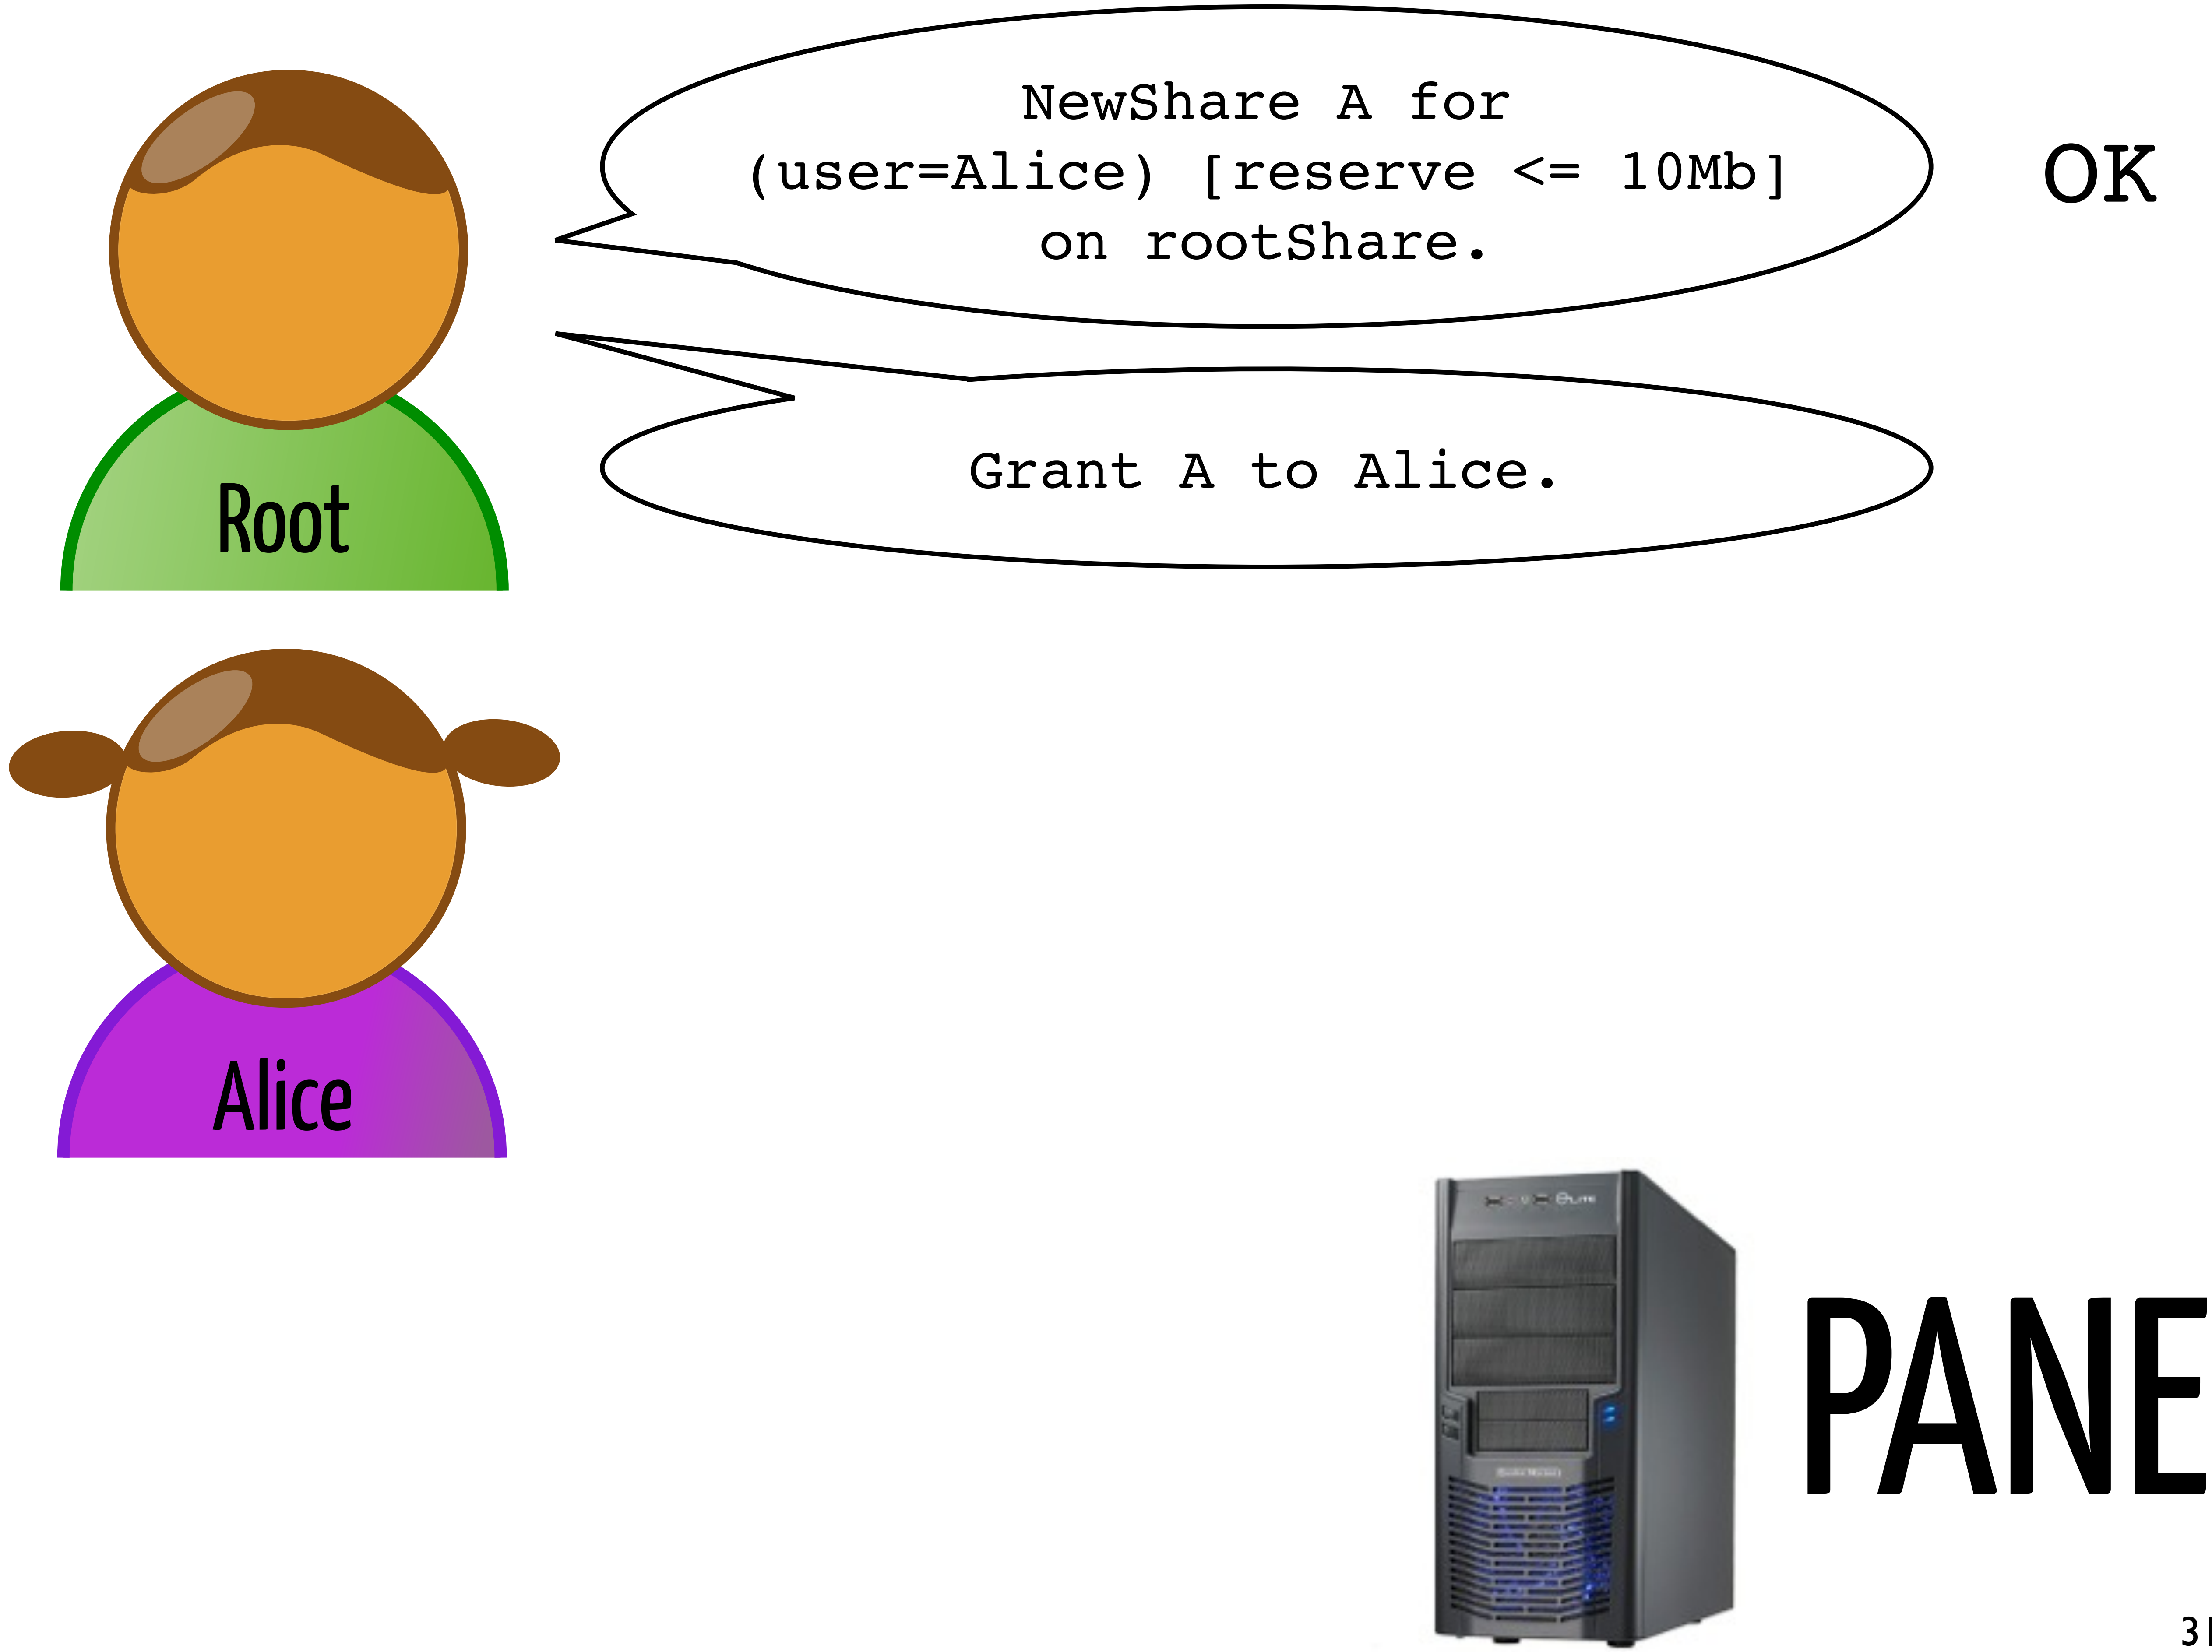
Root
NewShare A for
(user=Alice) [reserve <= 10Mb]
on rootShare.
OK
Grant A to Alice.
Alice
PANE

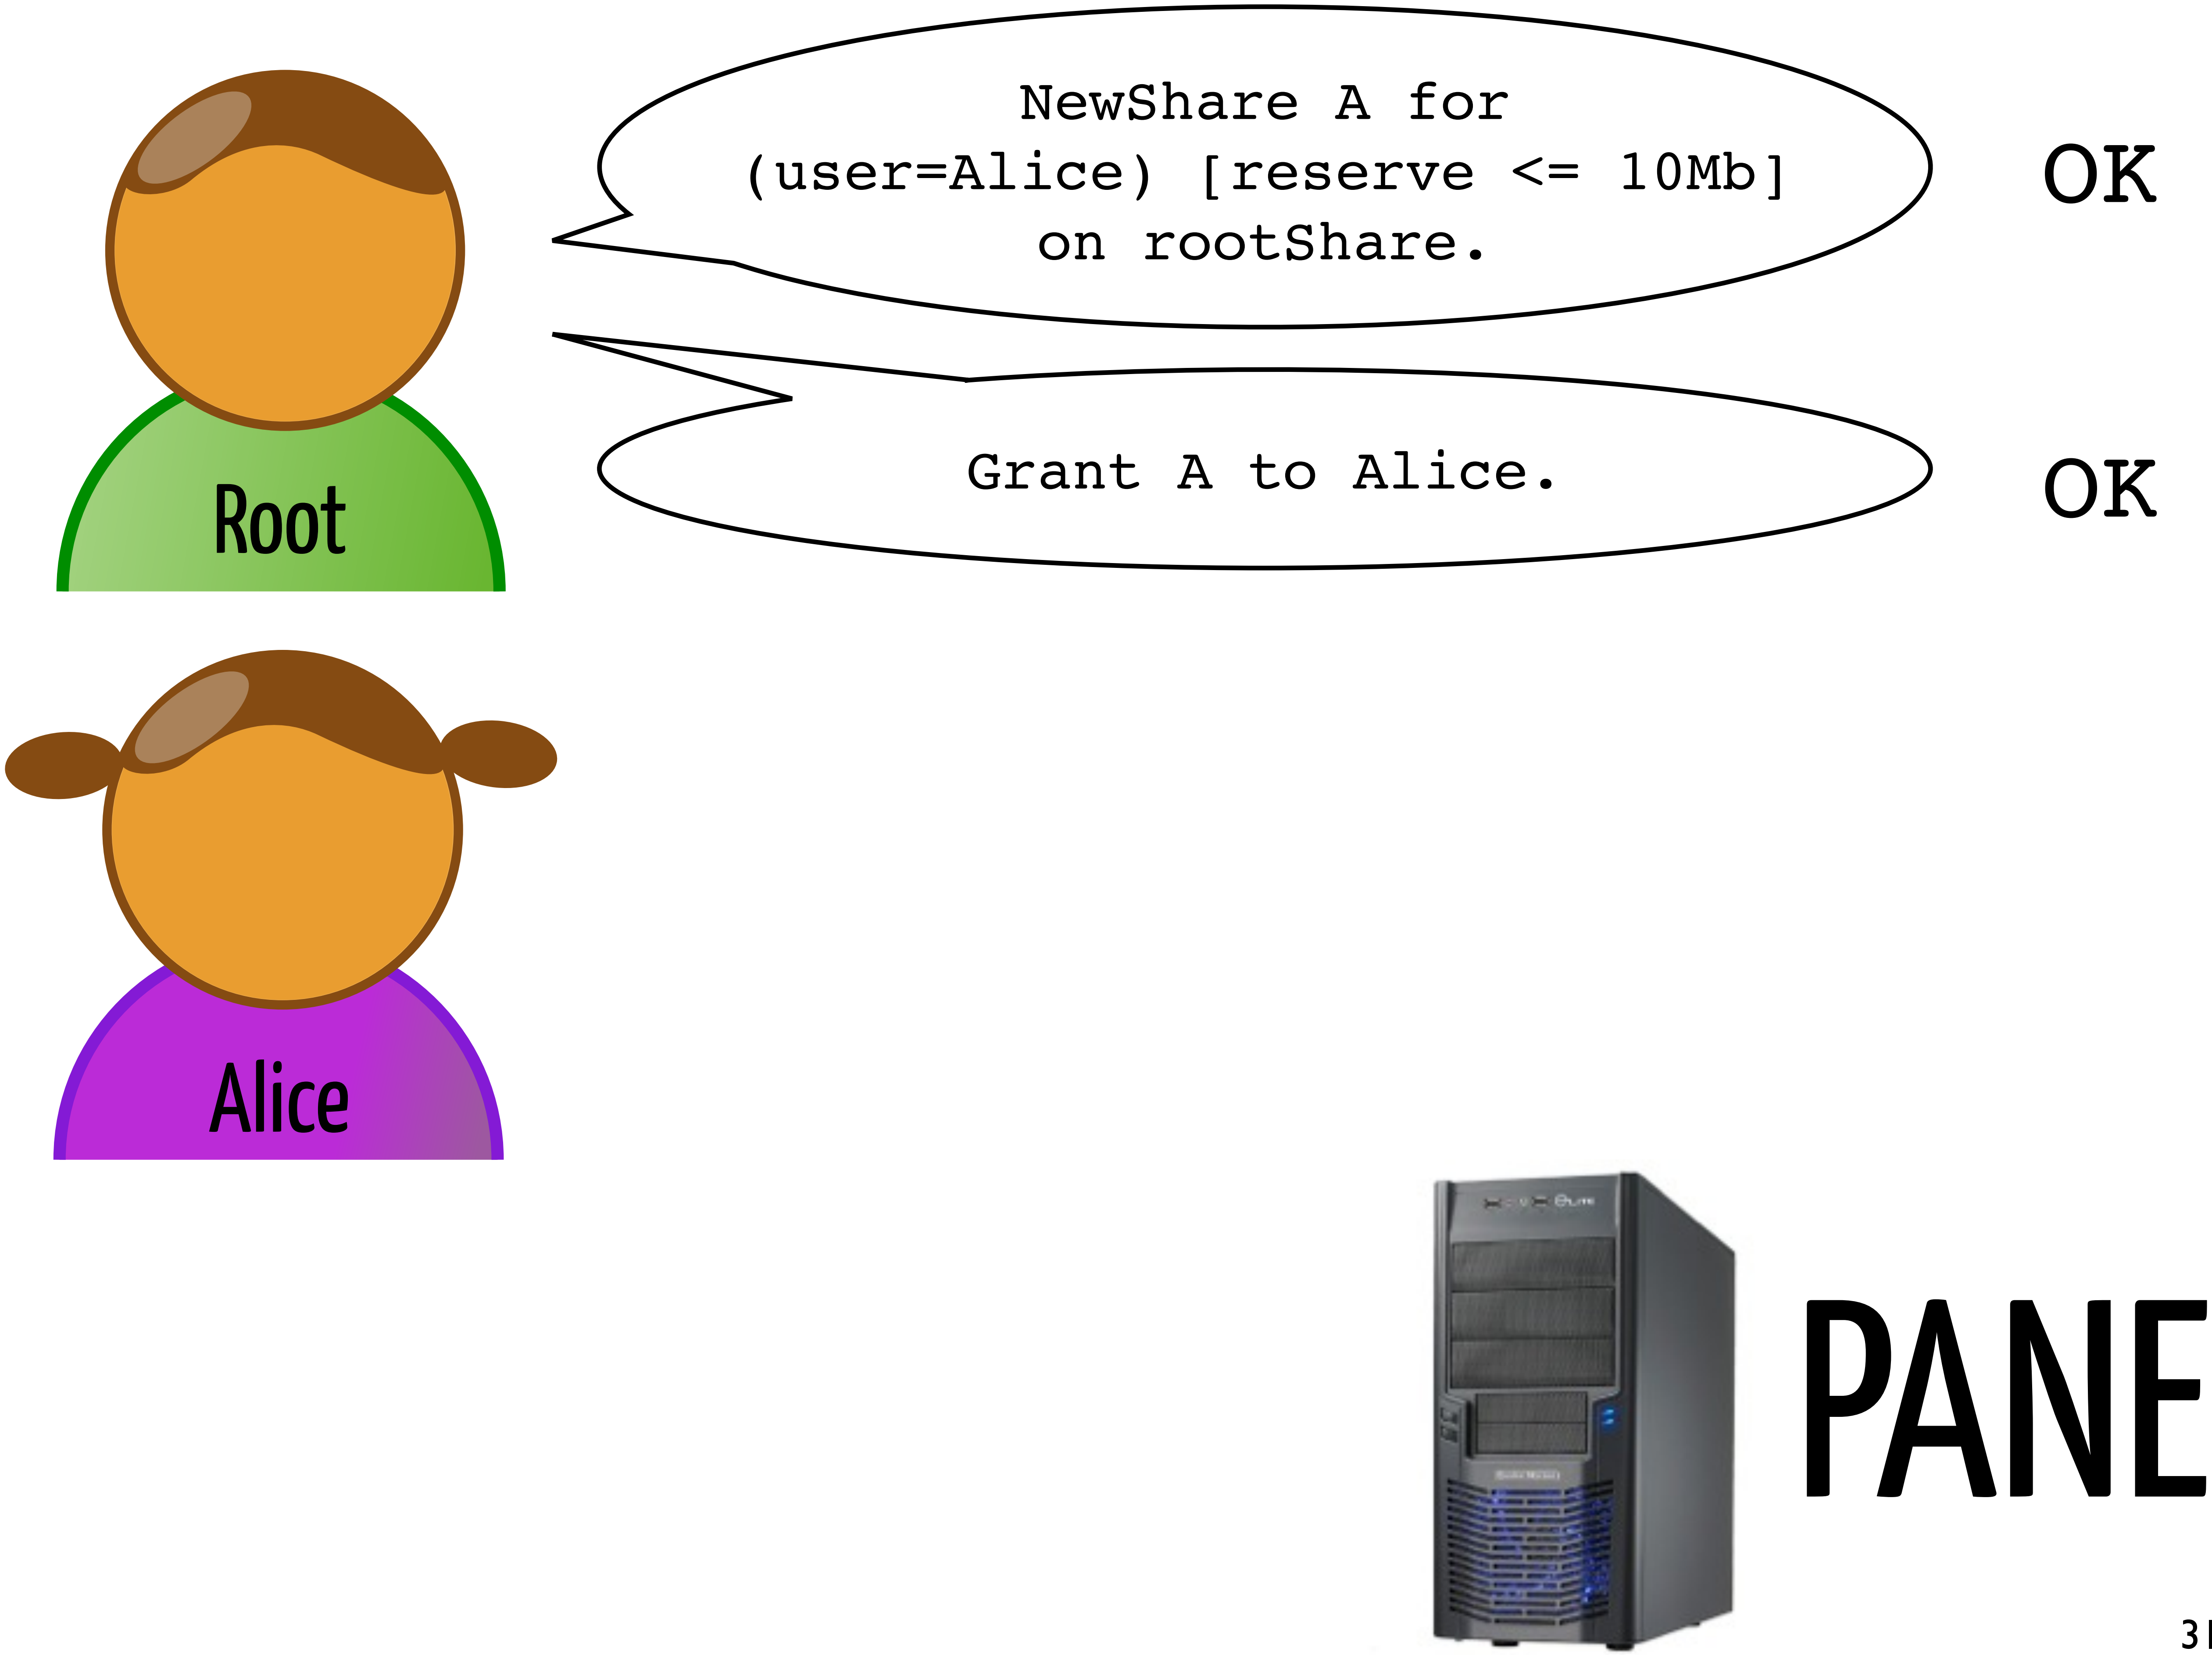
NewShare A for
(user=Alice) [reserve <= 10Mb]
on rootShare.
OK
Grant A to Alice.
OK
Root
Alice
PANE

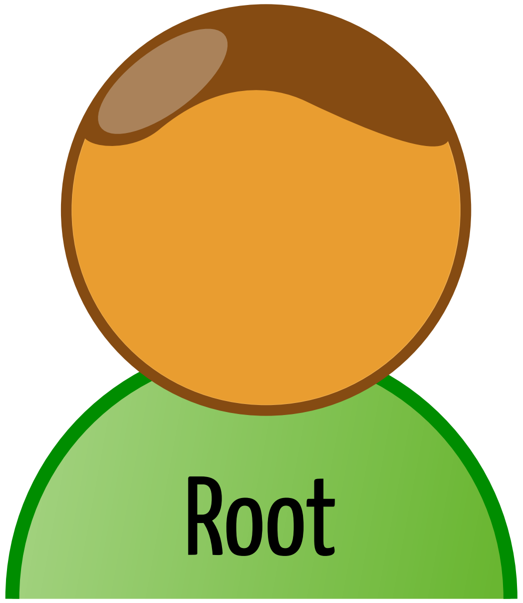

NewShare A for
(user=Alice) [reserve <= 10Mb]
on rootShare.

OK

Grant A to Alice.

OK

reserve(user=Alice,
dstPort=80) = 5Mb on A
from now to +10min.

PANE

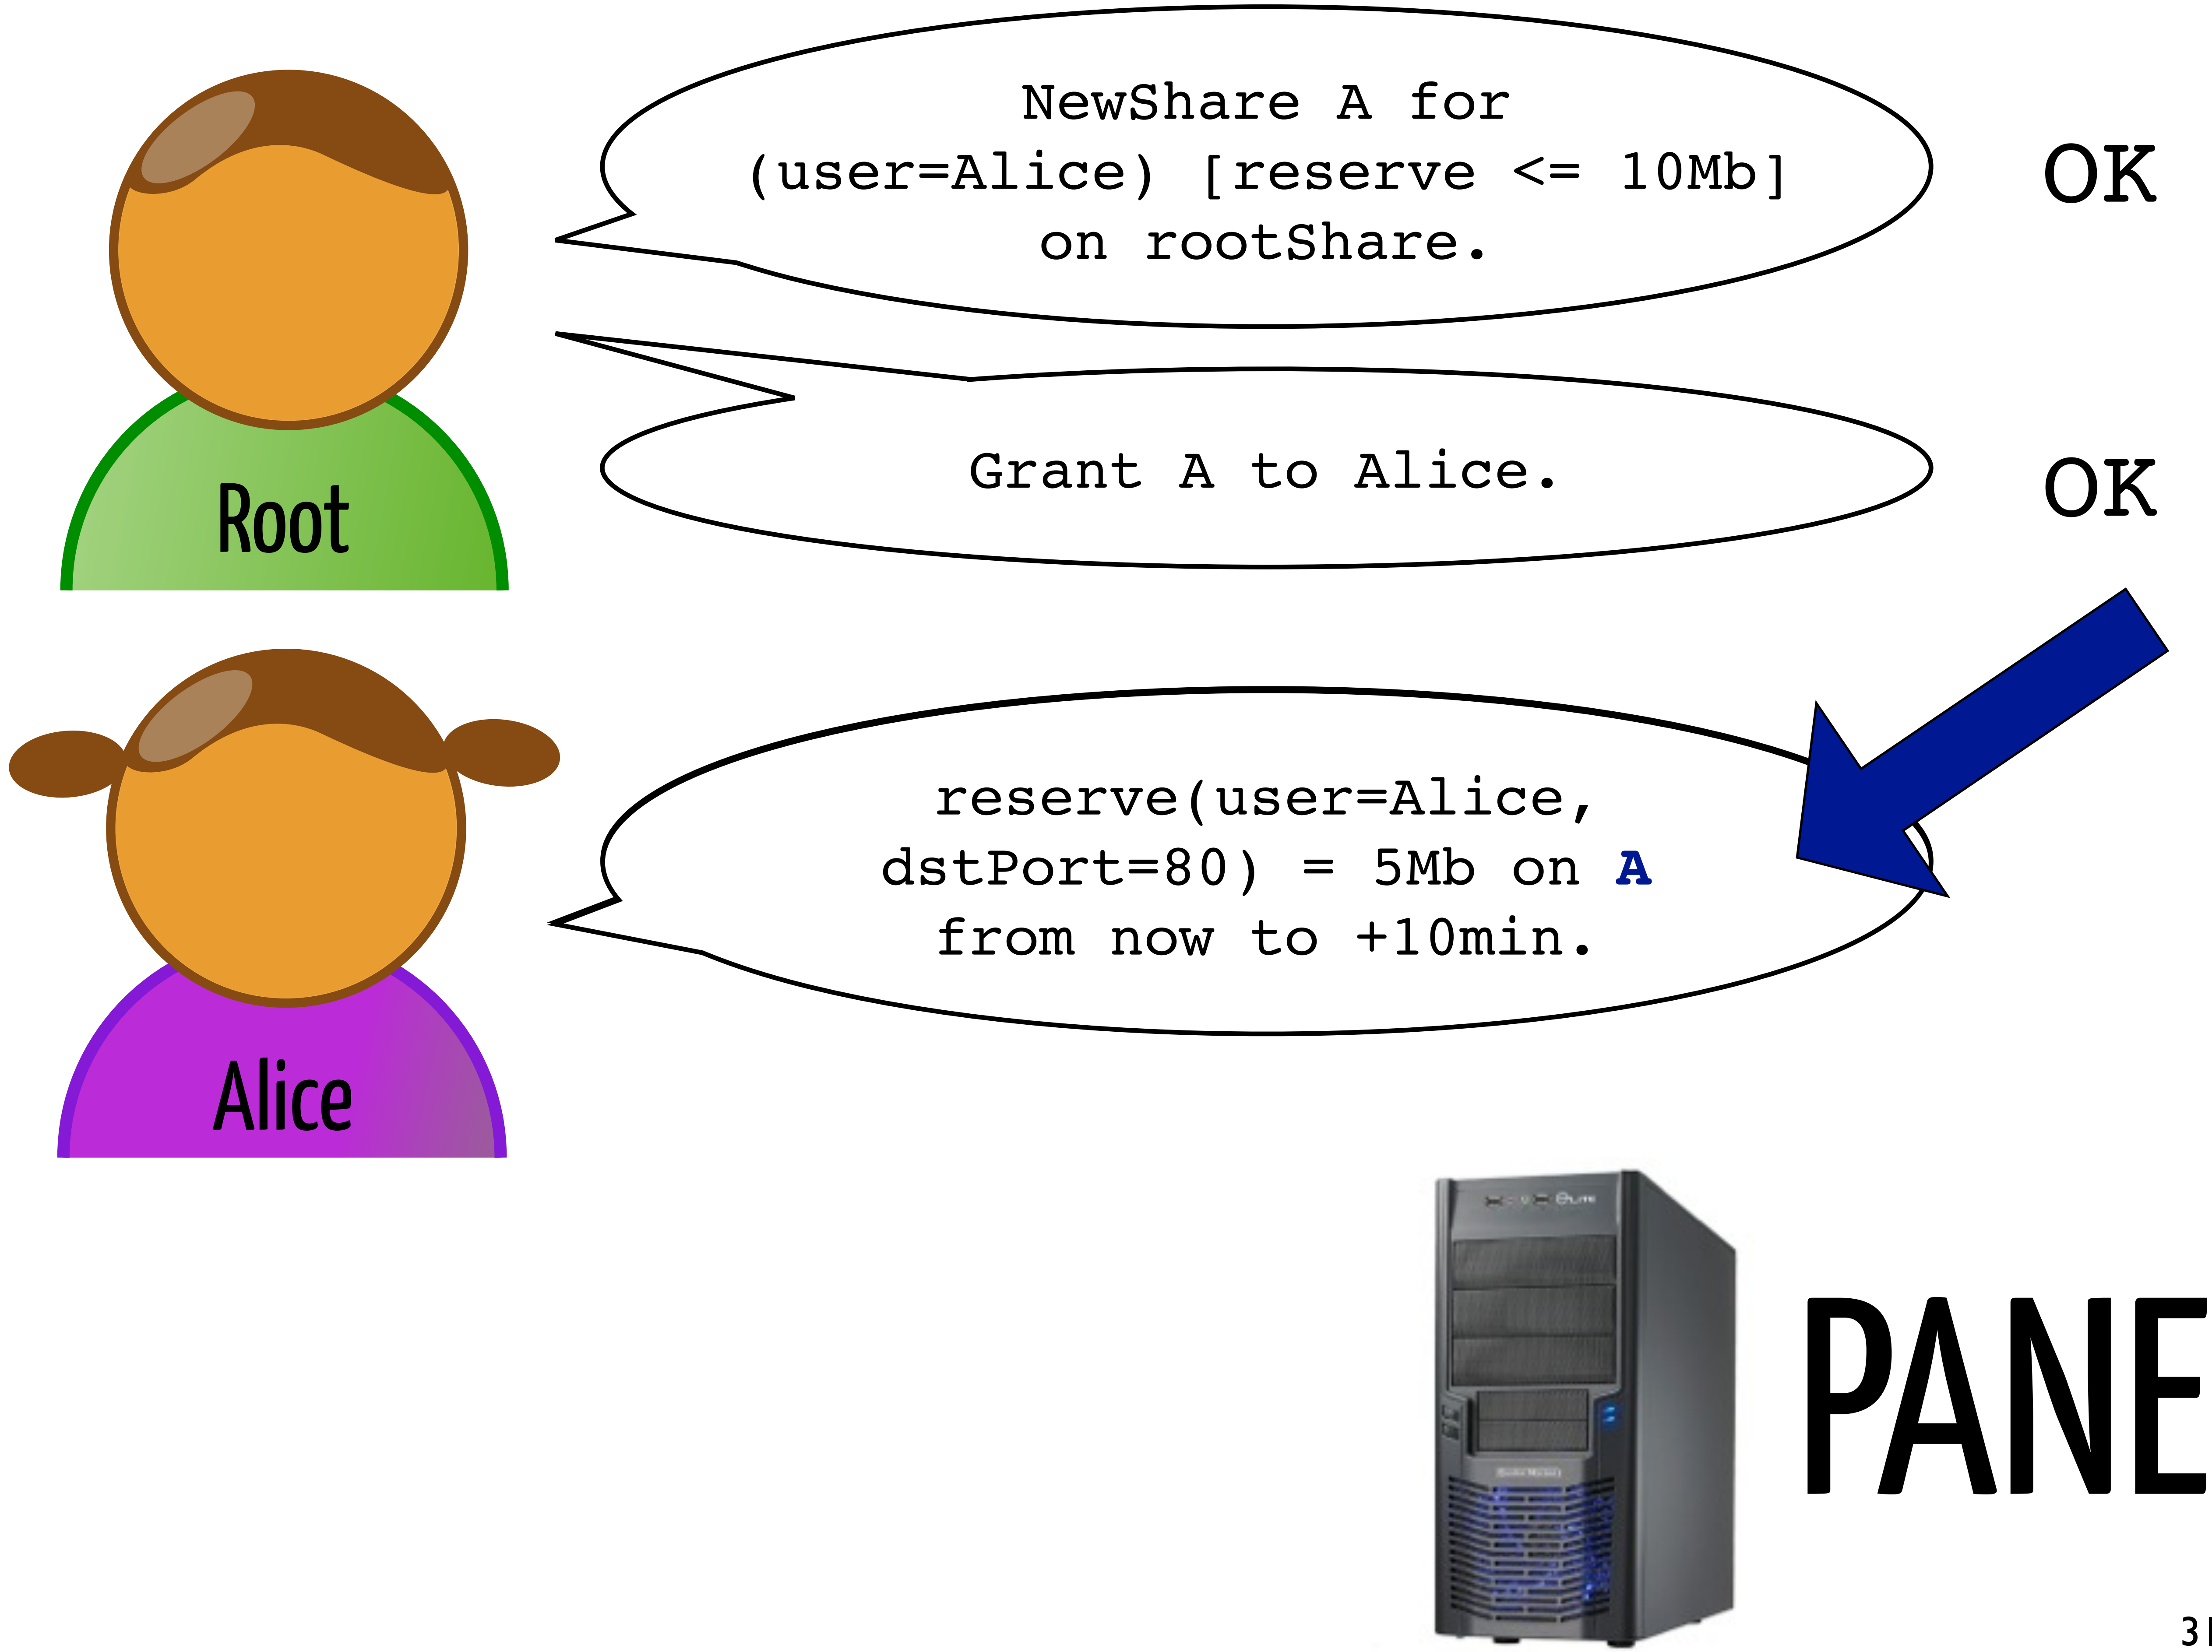
NewShare A for
(user=Alice) [reserve <= 10Mb]
on rootShare.
OK
Grant A to Alice.
OK
Root
reserve(user=Alice,
dstPort=80) = 5Mb on A
from now to +10min.
Alice
PANE

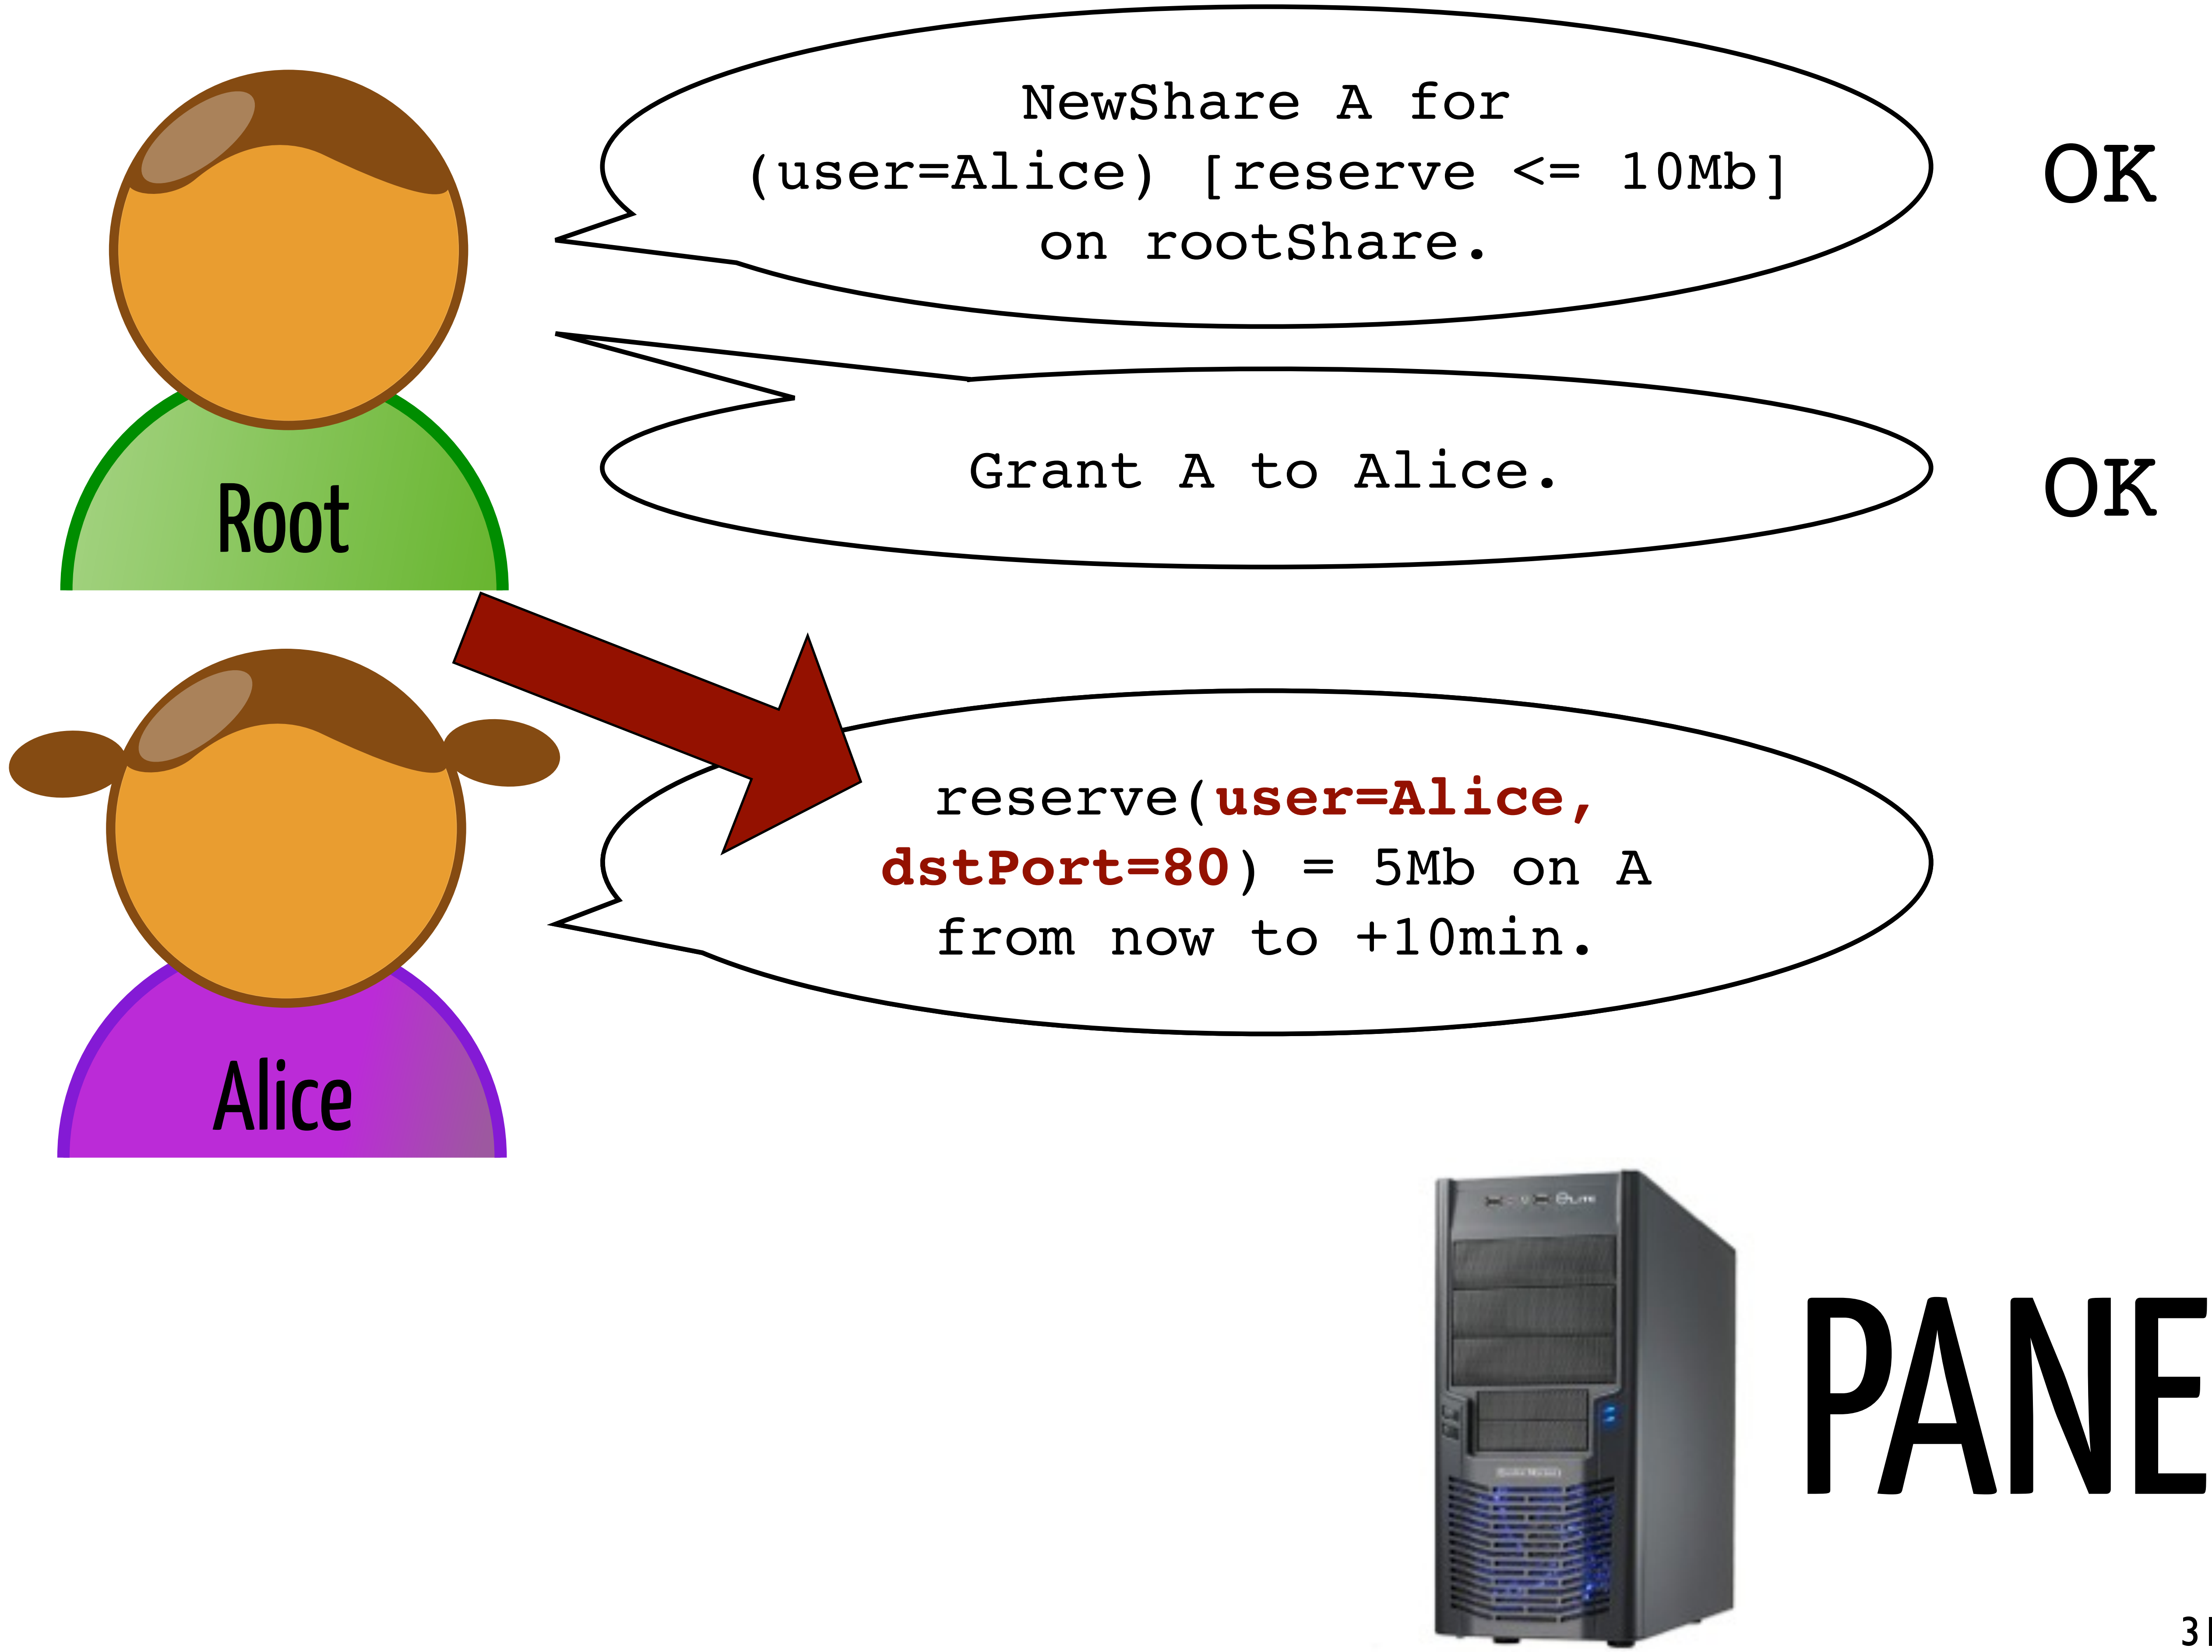
NewShare A for (user=Alice) [reserve <= 10Mb] on rootShare.
OK
Grant A to Alice.
OK
Root
reserve(user=Alice, dstPort=80) = 5Mb on A from now to +10min.
Alice
PANE

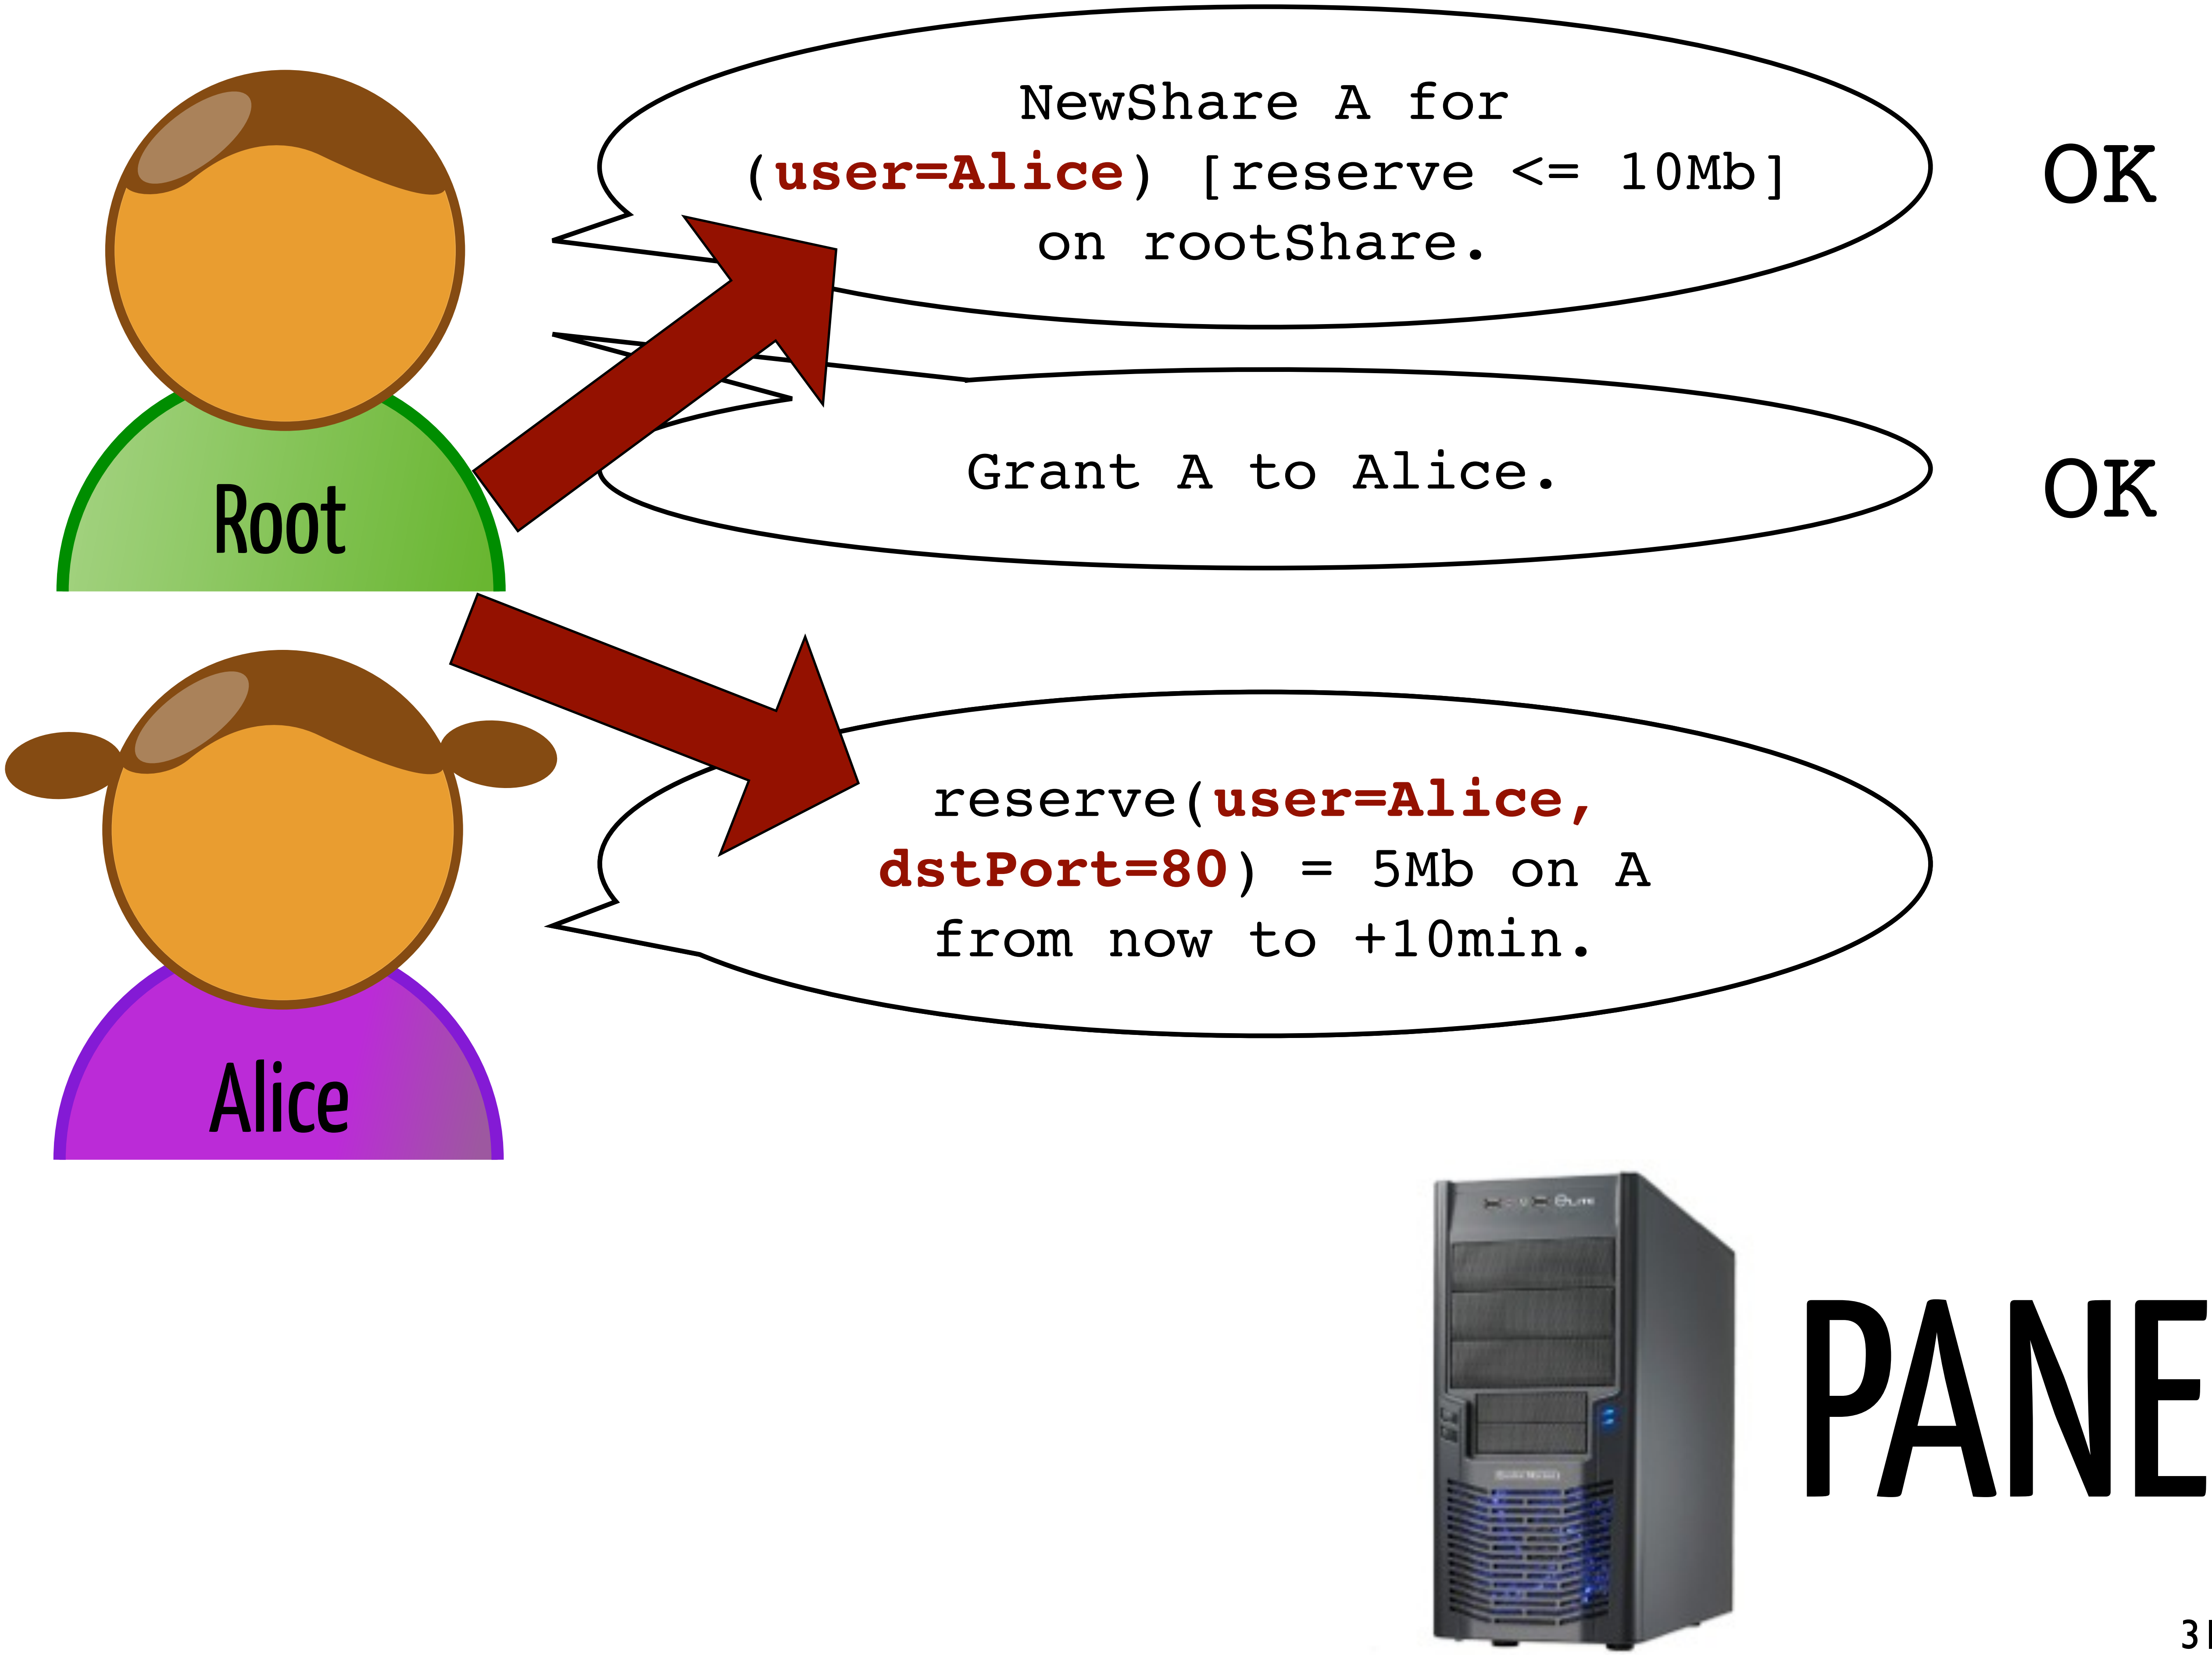

Root
NewShare A for (user=Alice) [reserve <= 10Mb] on rootShare.
OK
Grant A to Alice.
OK
Alice
reserve(user=Alice, dstPort=80) = 5Mb on A from now to +10min.
PANE

reserve(user=Alice,
dstPort=80) = 5Mb on A
from now to +10min.

PANE

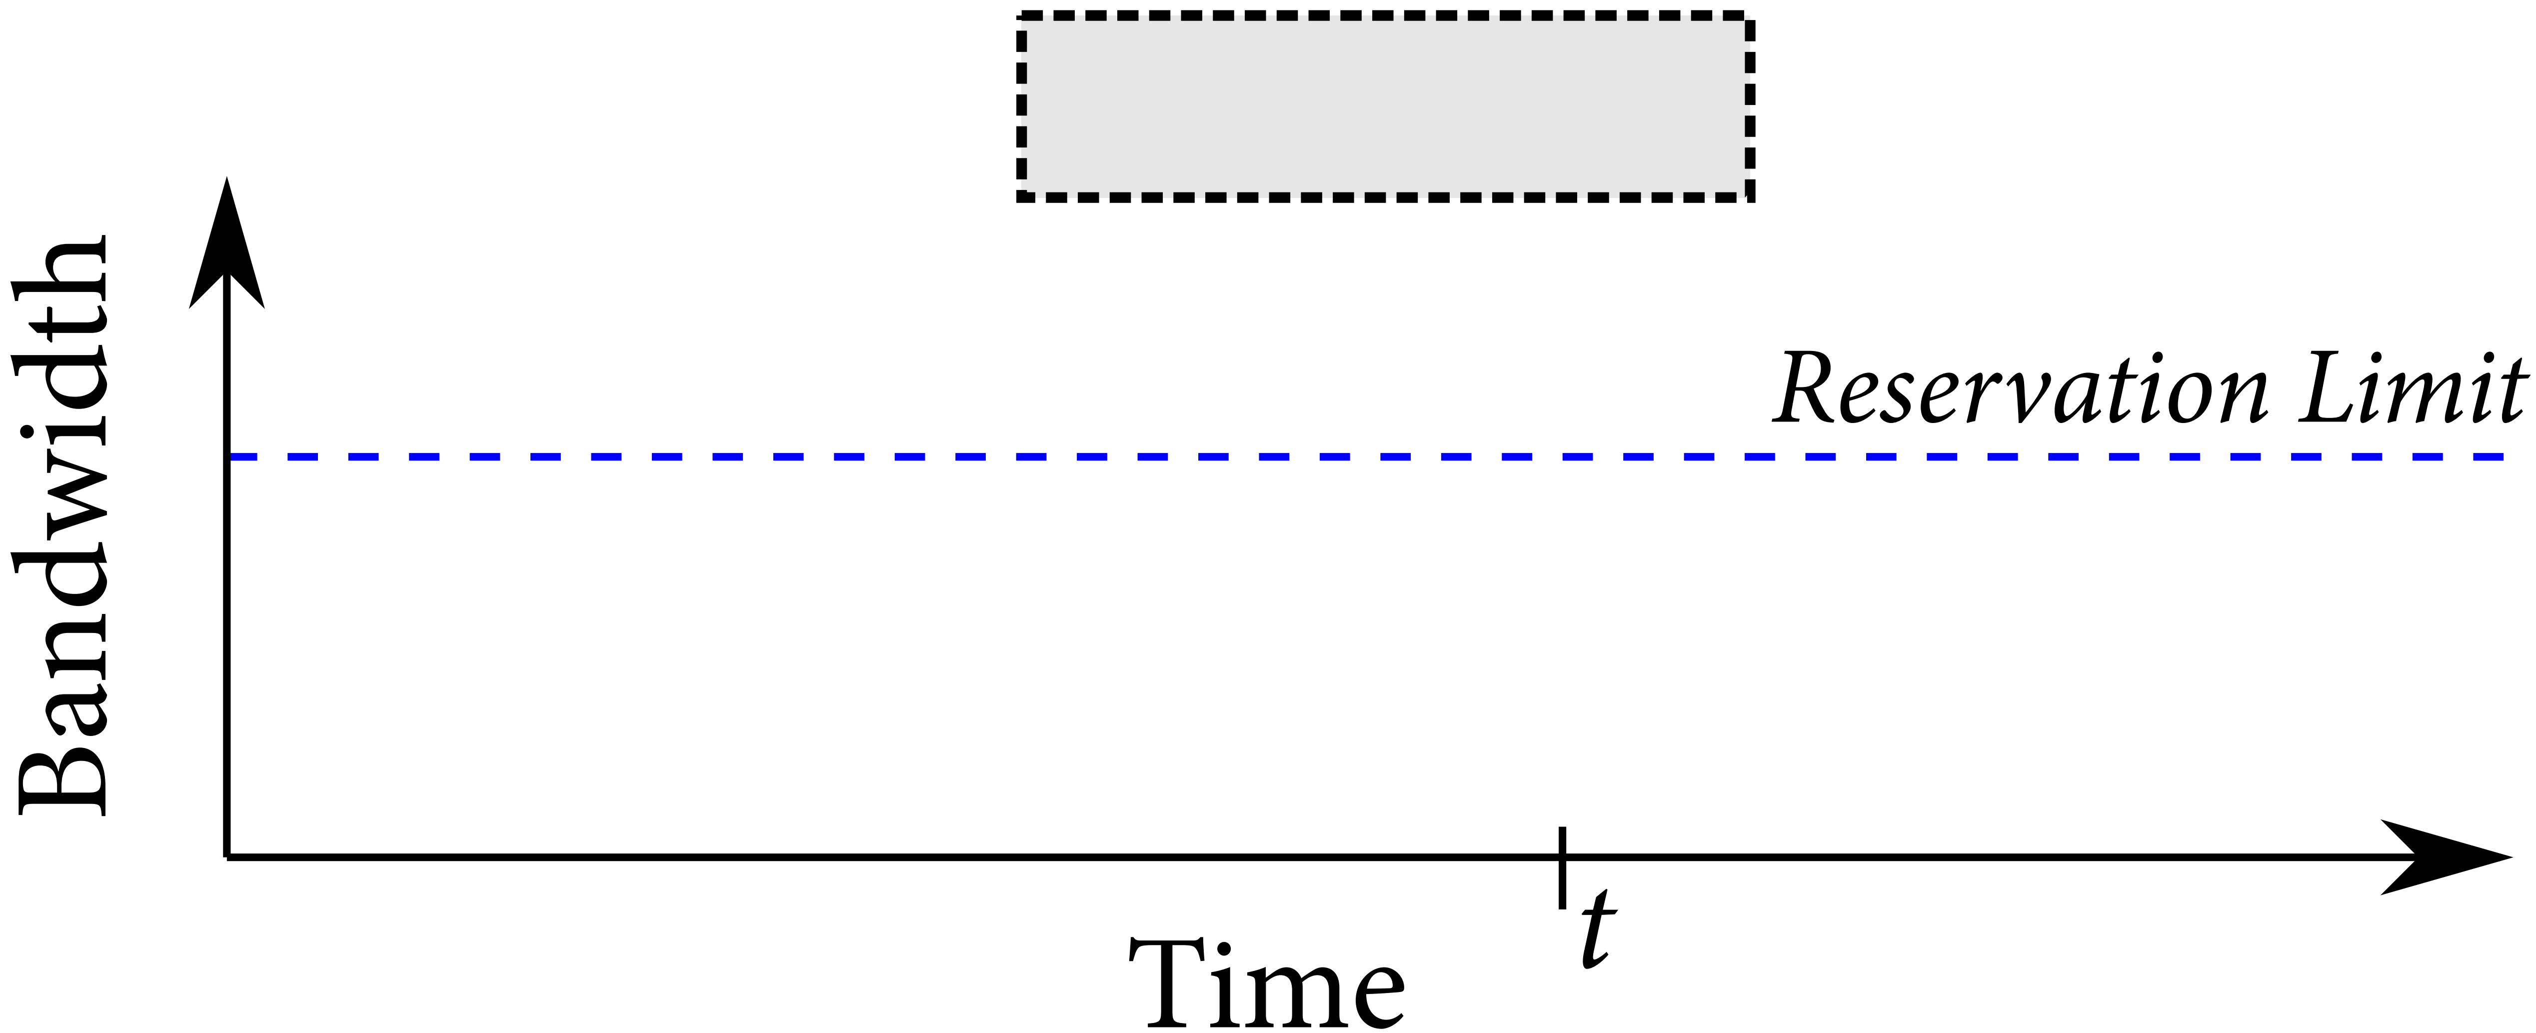

reserve(user=Alice,
dstPort=80) = 5Mb on A
from now to +10min.

PANE

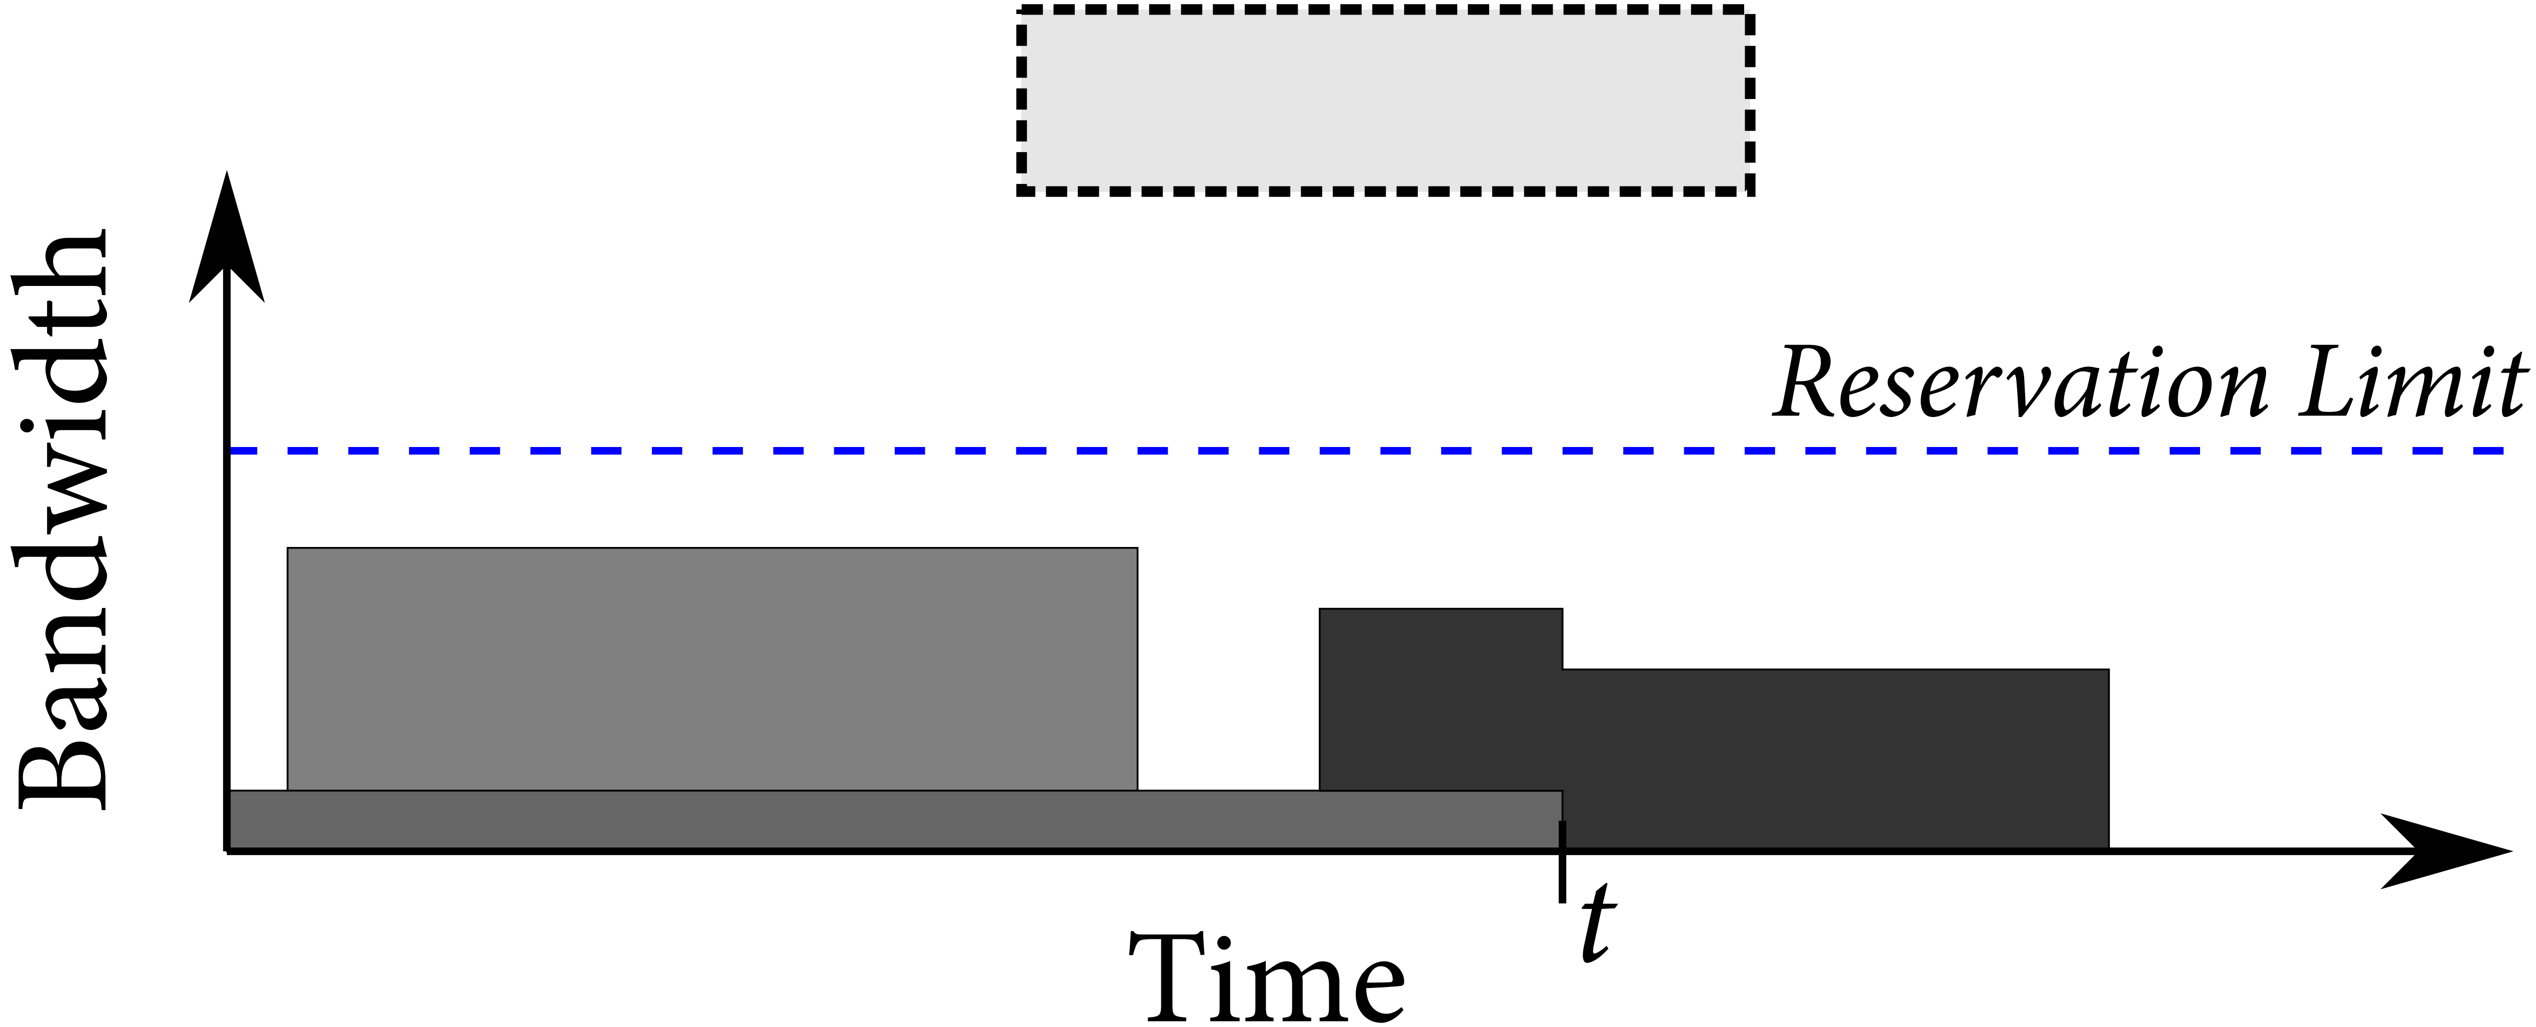

reserve(user=Alice, dstPort=80) = 5Mb on A from now to +10min.

PANE

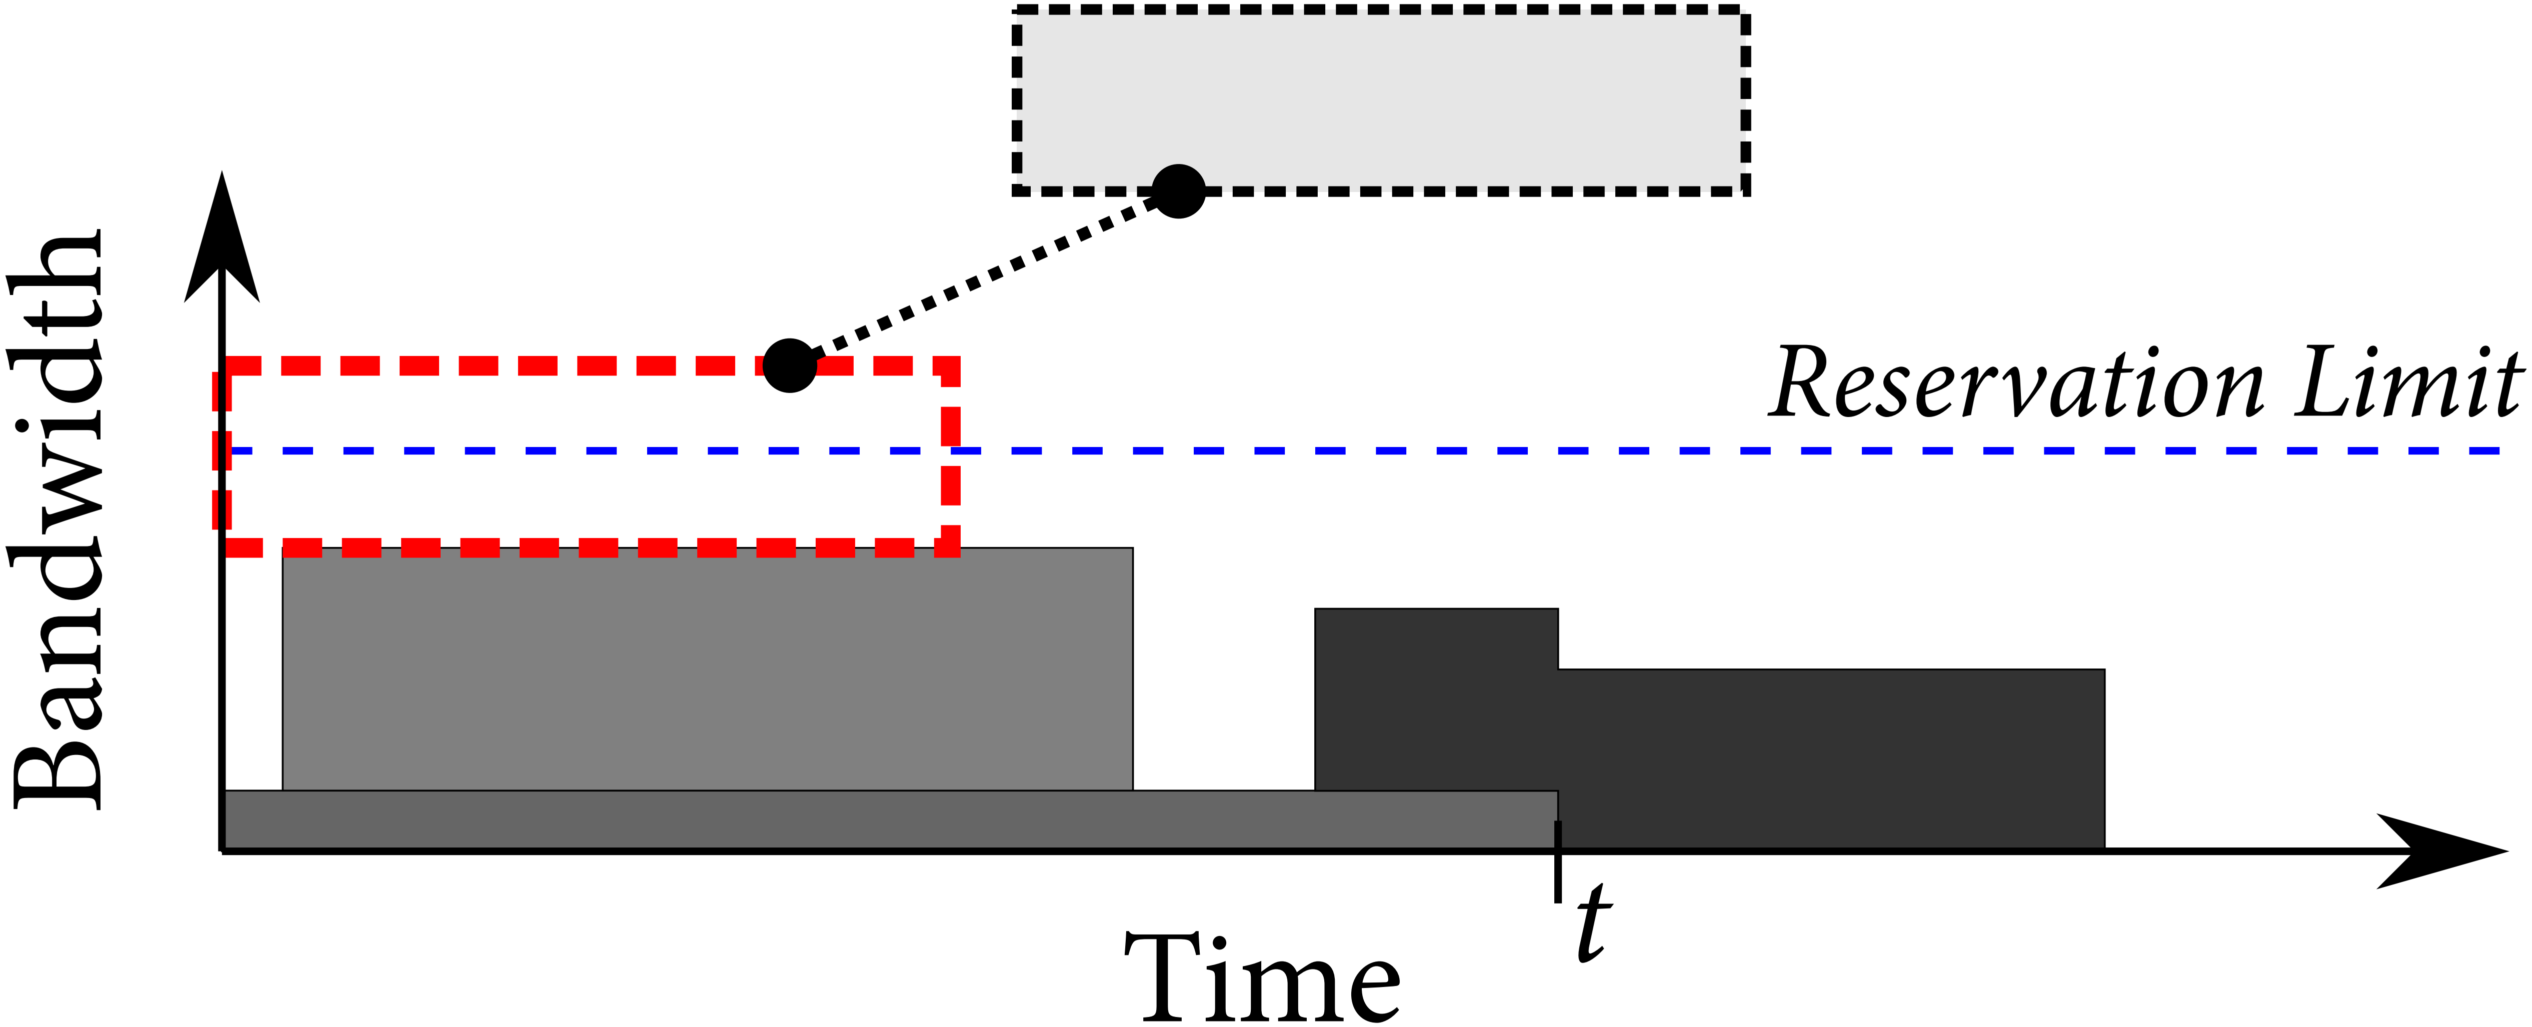

reserve(user=Alice,
dstPort=80) = 5Mb on A
from now to +10min.

PANE

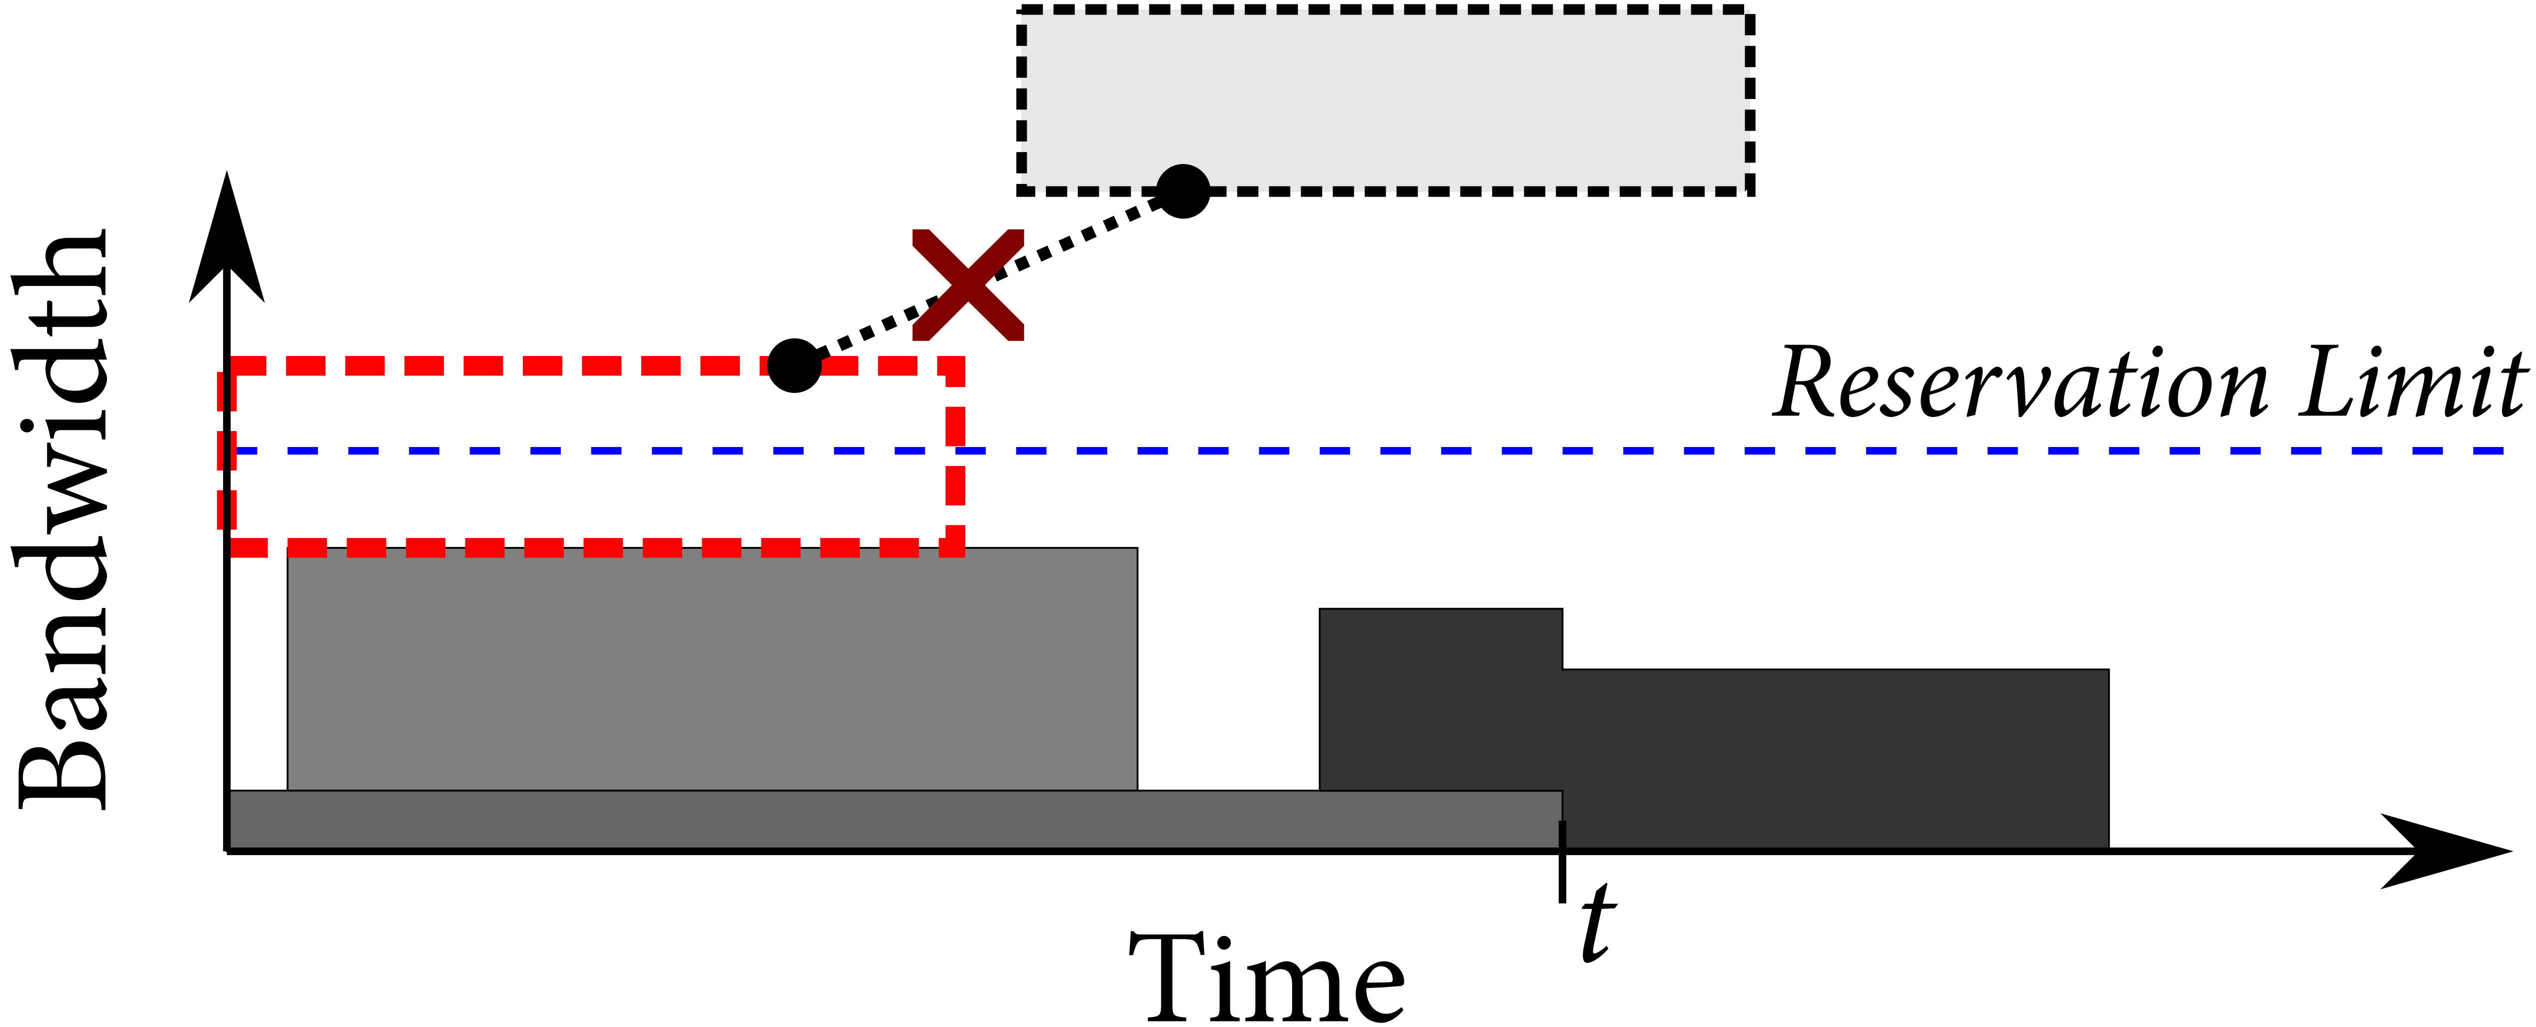
Bandwidth
Reservation Limit
t
Time

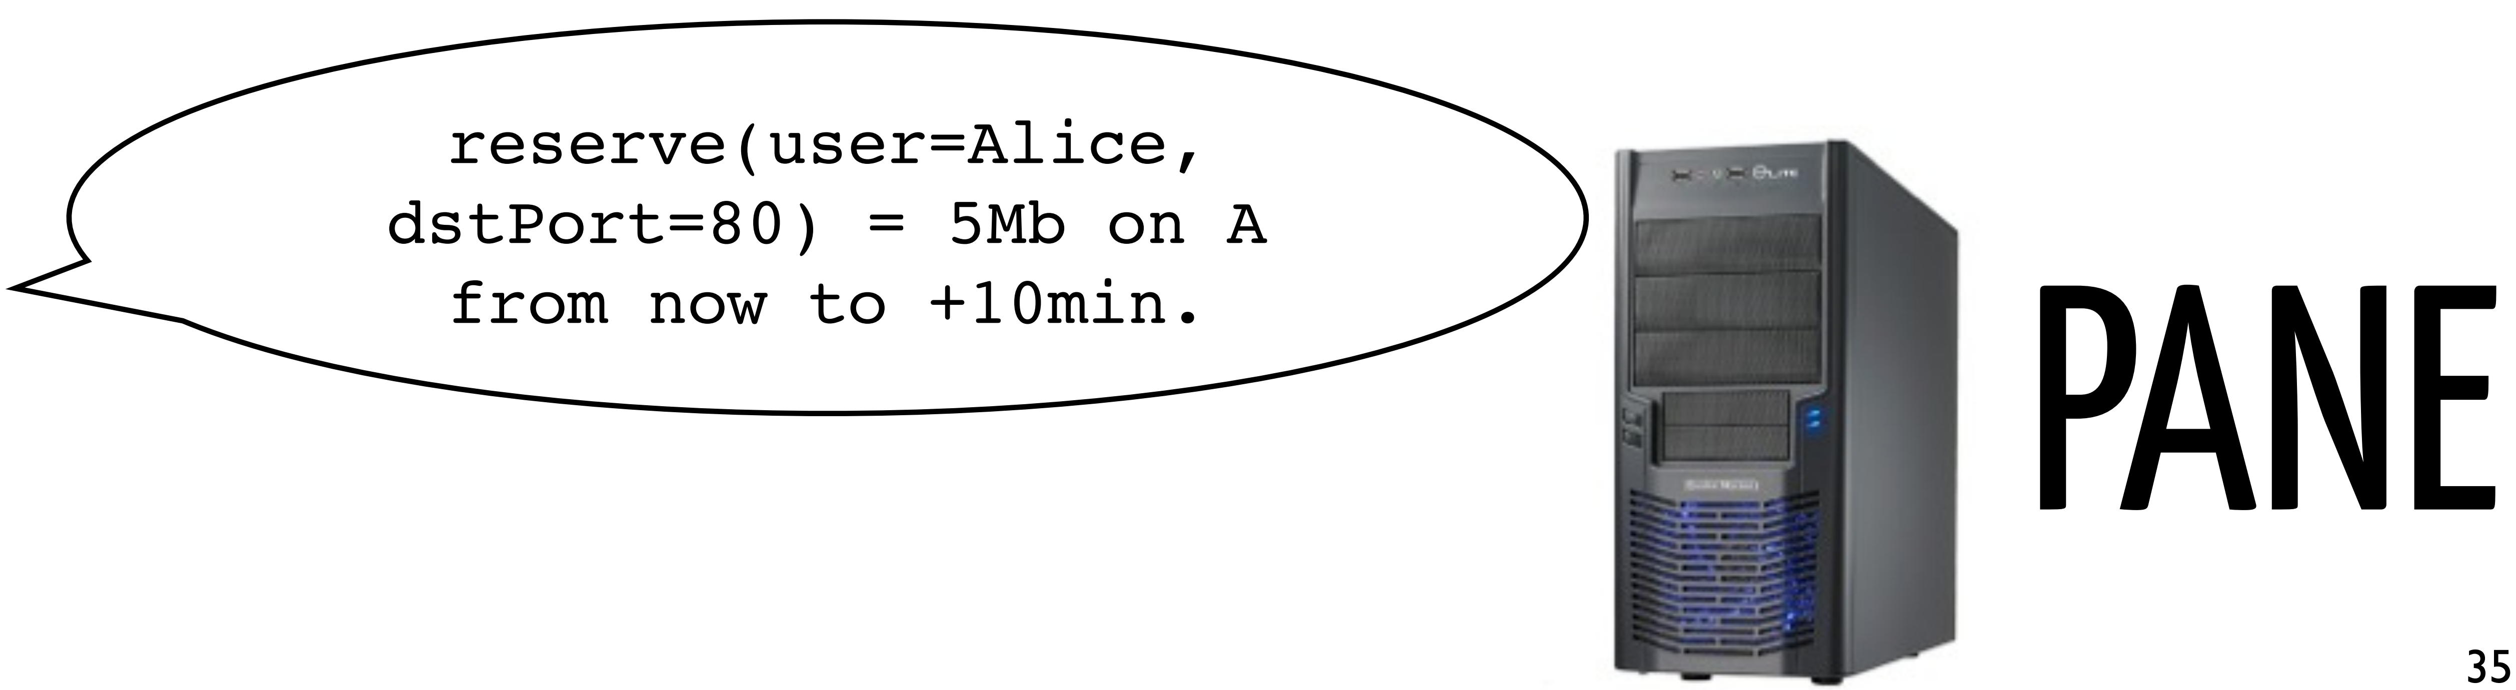
reserve(user=Alice,
dstPort=80) = 5Mb on A
from now to +10min.
PANE

reserve(user=Alice,
dstPort=80) = 5Mb on A
from now to +10min.

NO

PANE


reserve(user=Alice,
dstPort=80) = 5Mb on A
from now to +10min.
Alice

NO

PANE

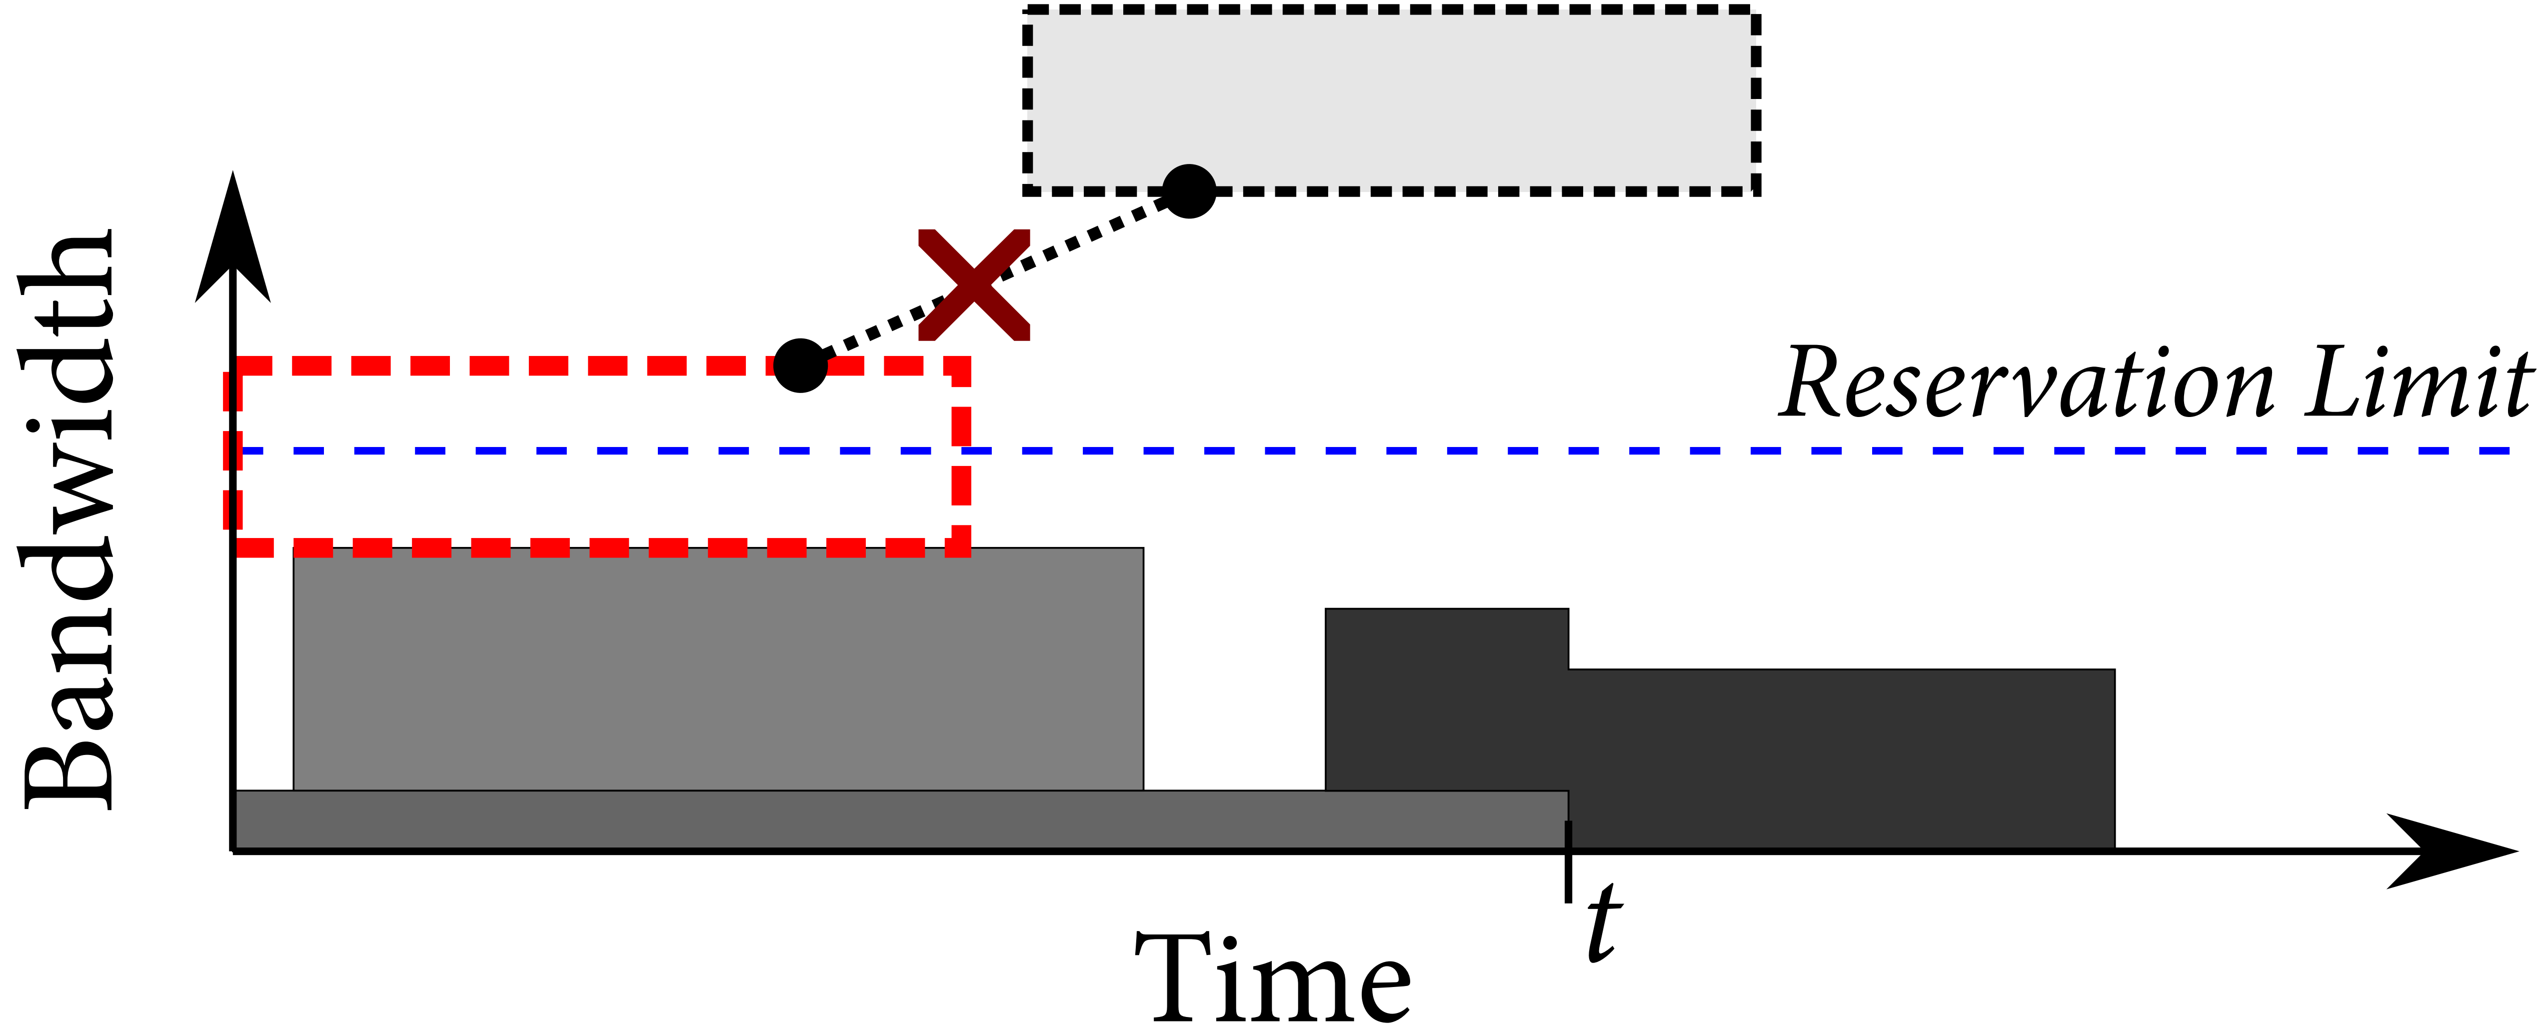
Bandwidth
Reservation Limit
Time
t

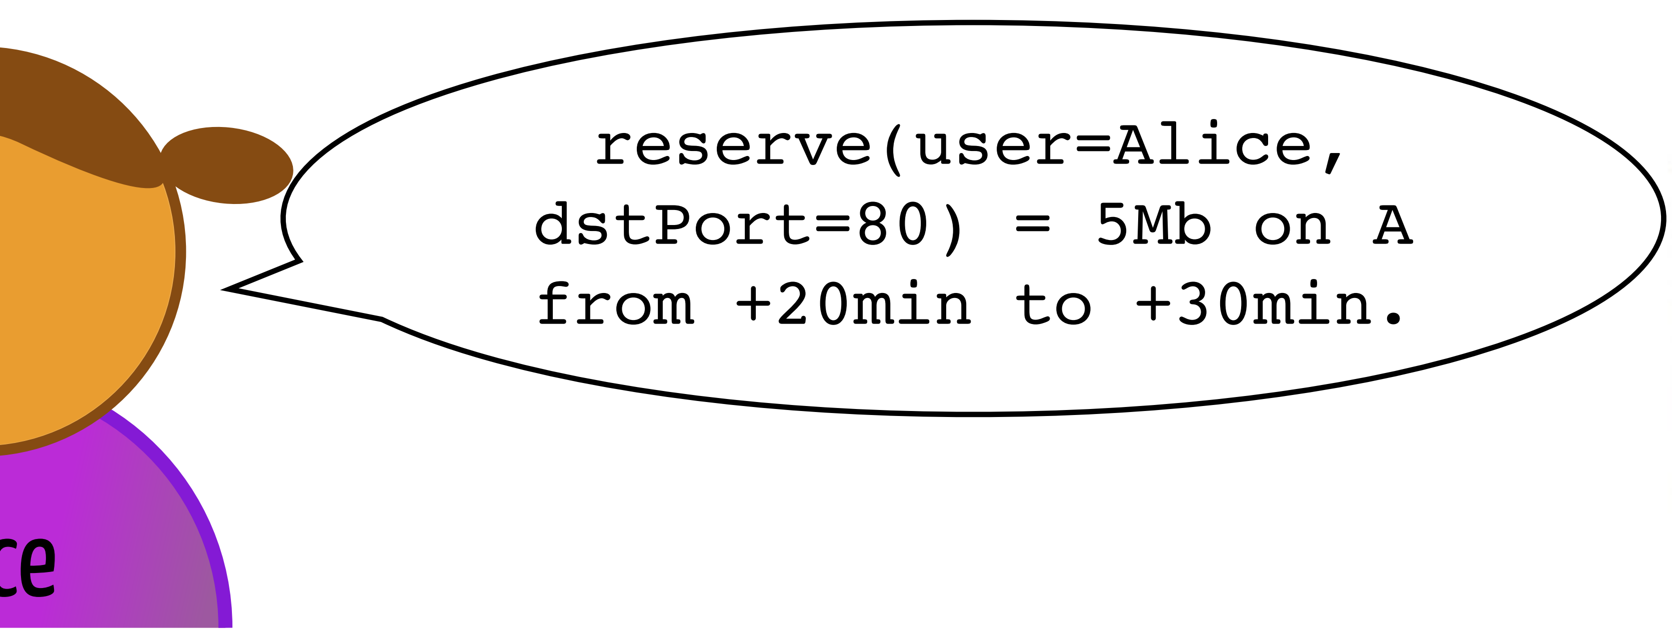
reserve(user=Alice,
dstPort=80) = 5Mb on A
from +20min to +30min.

PANE

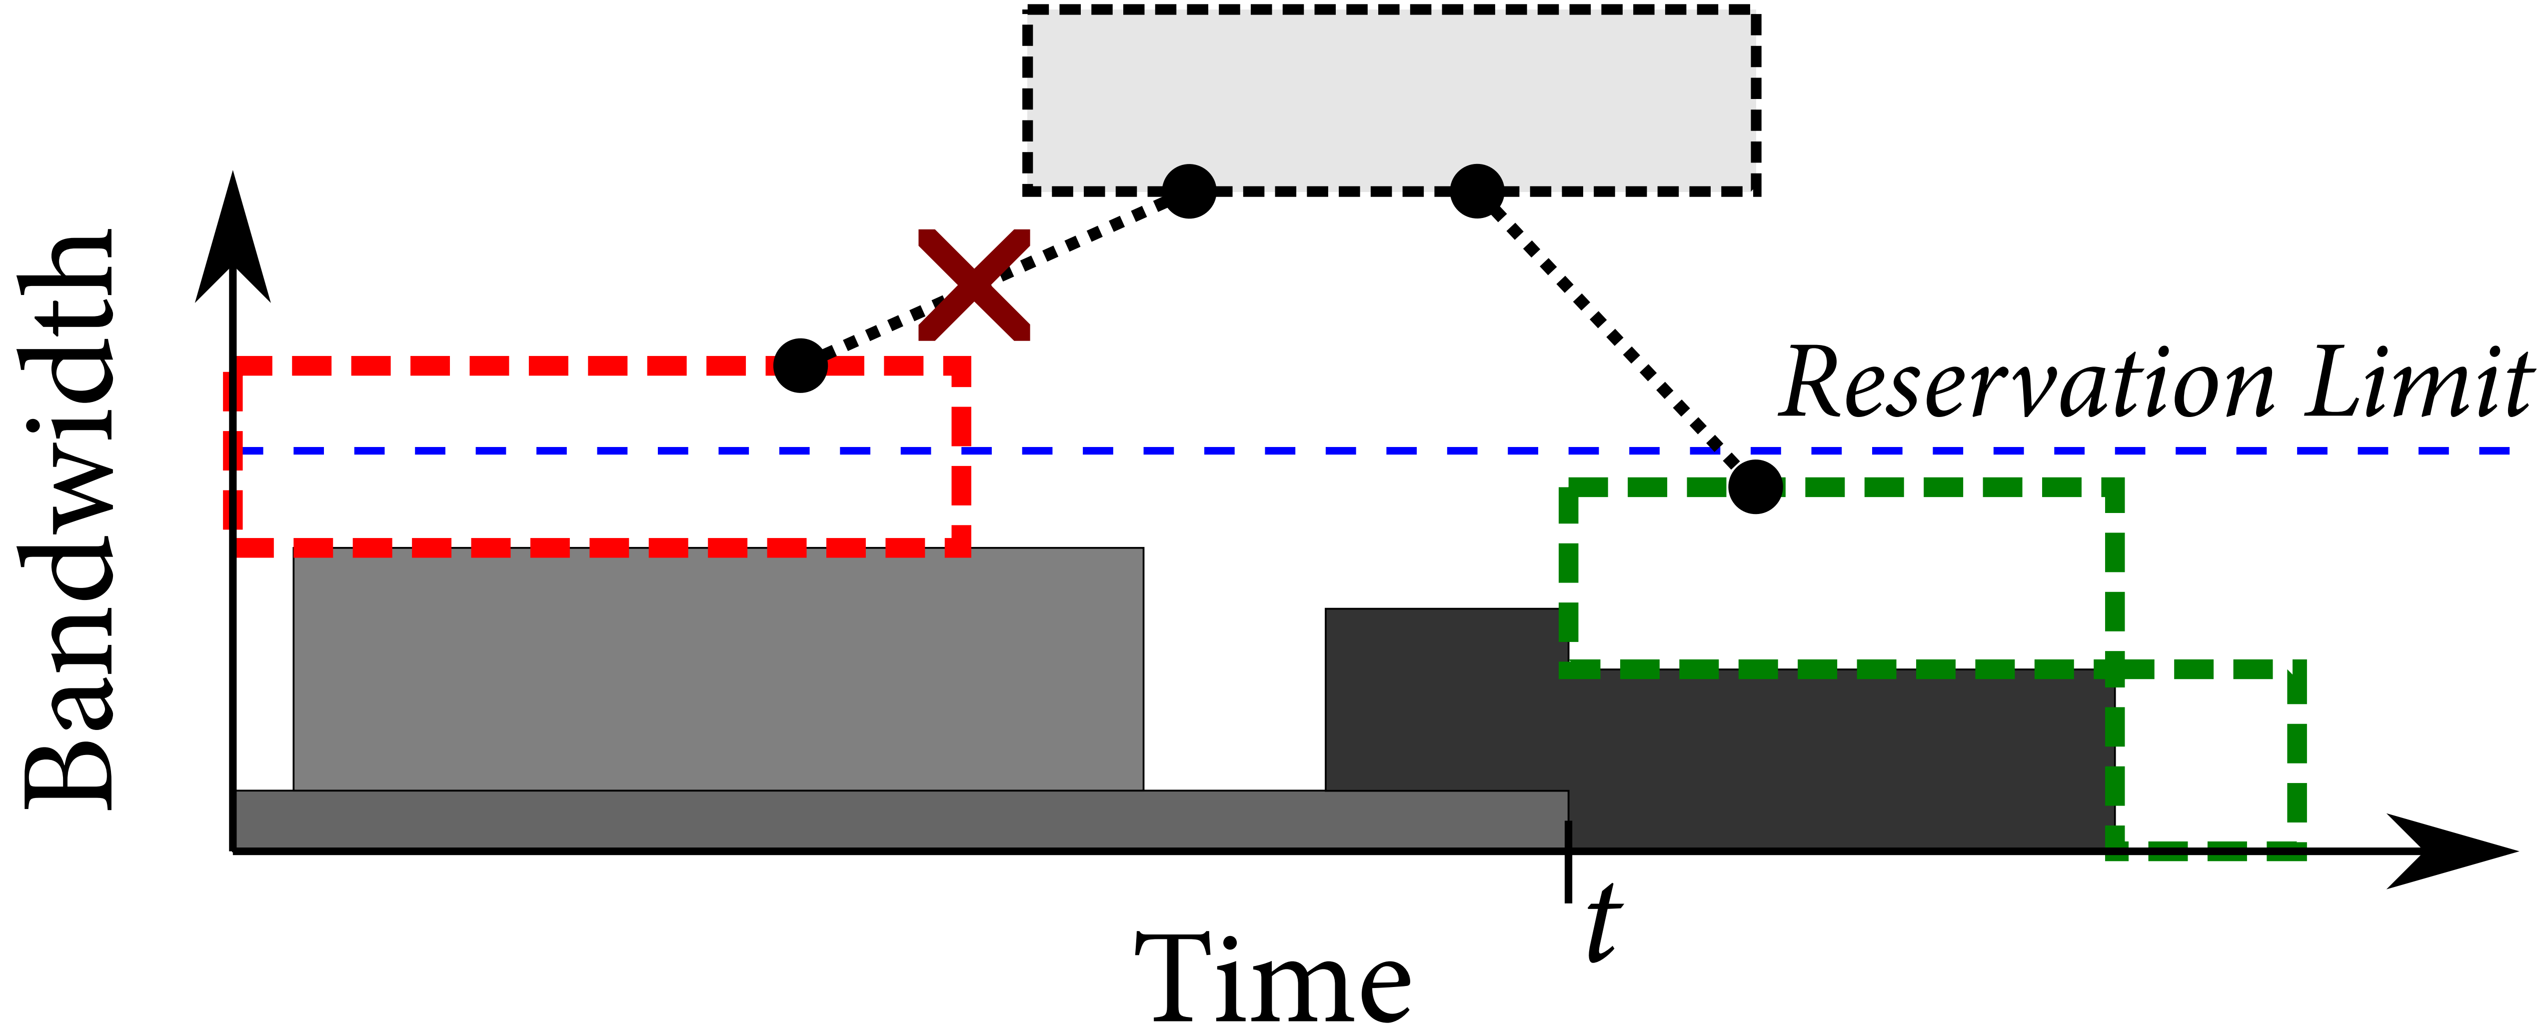
Bandwidth
Reservation Limit
t
Time

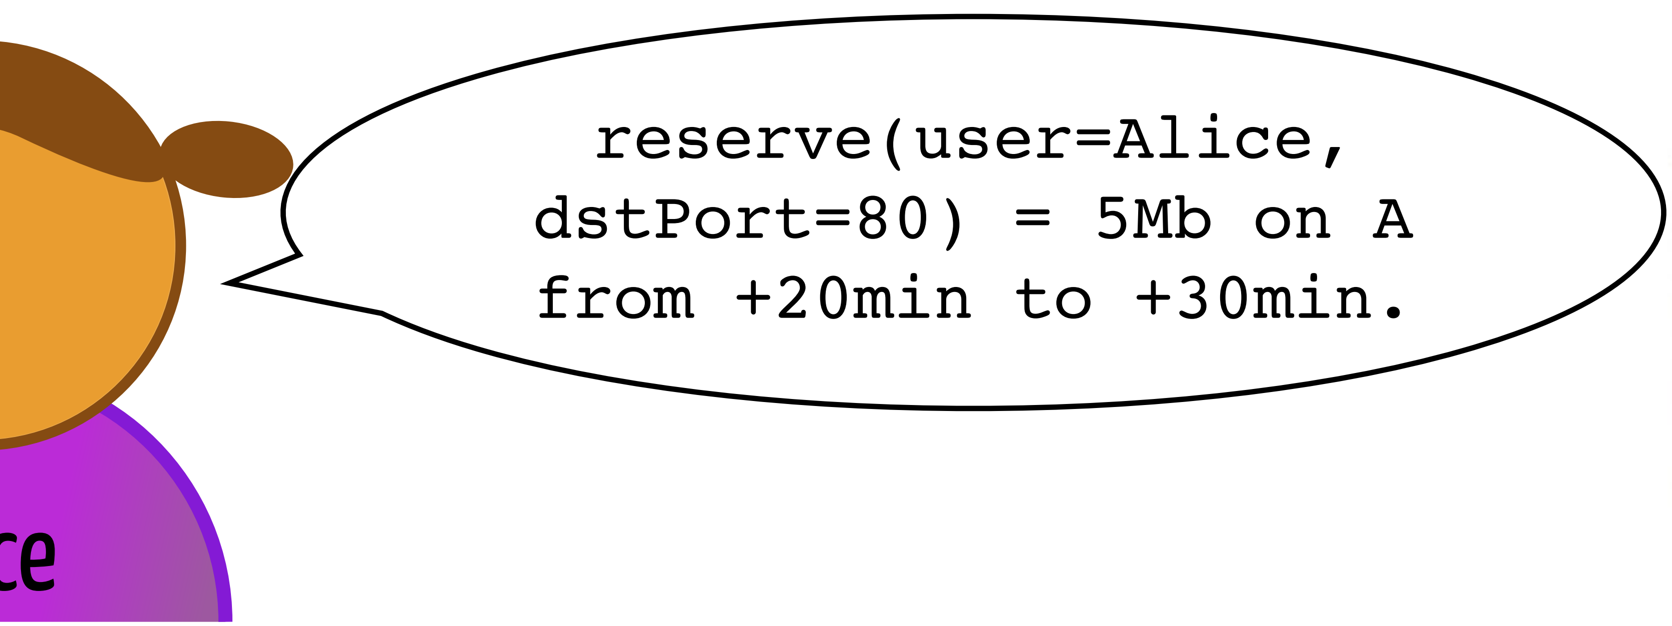
reserve(user=Alice,
dstPort=80) = 5Mb on A
from +20min to +30min.

PANE

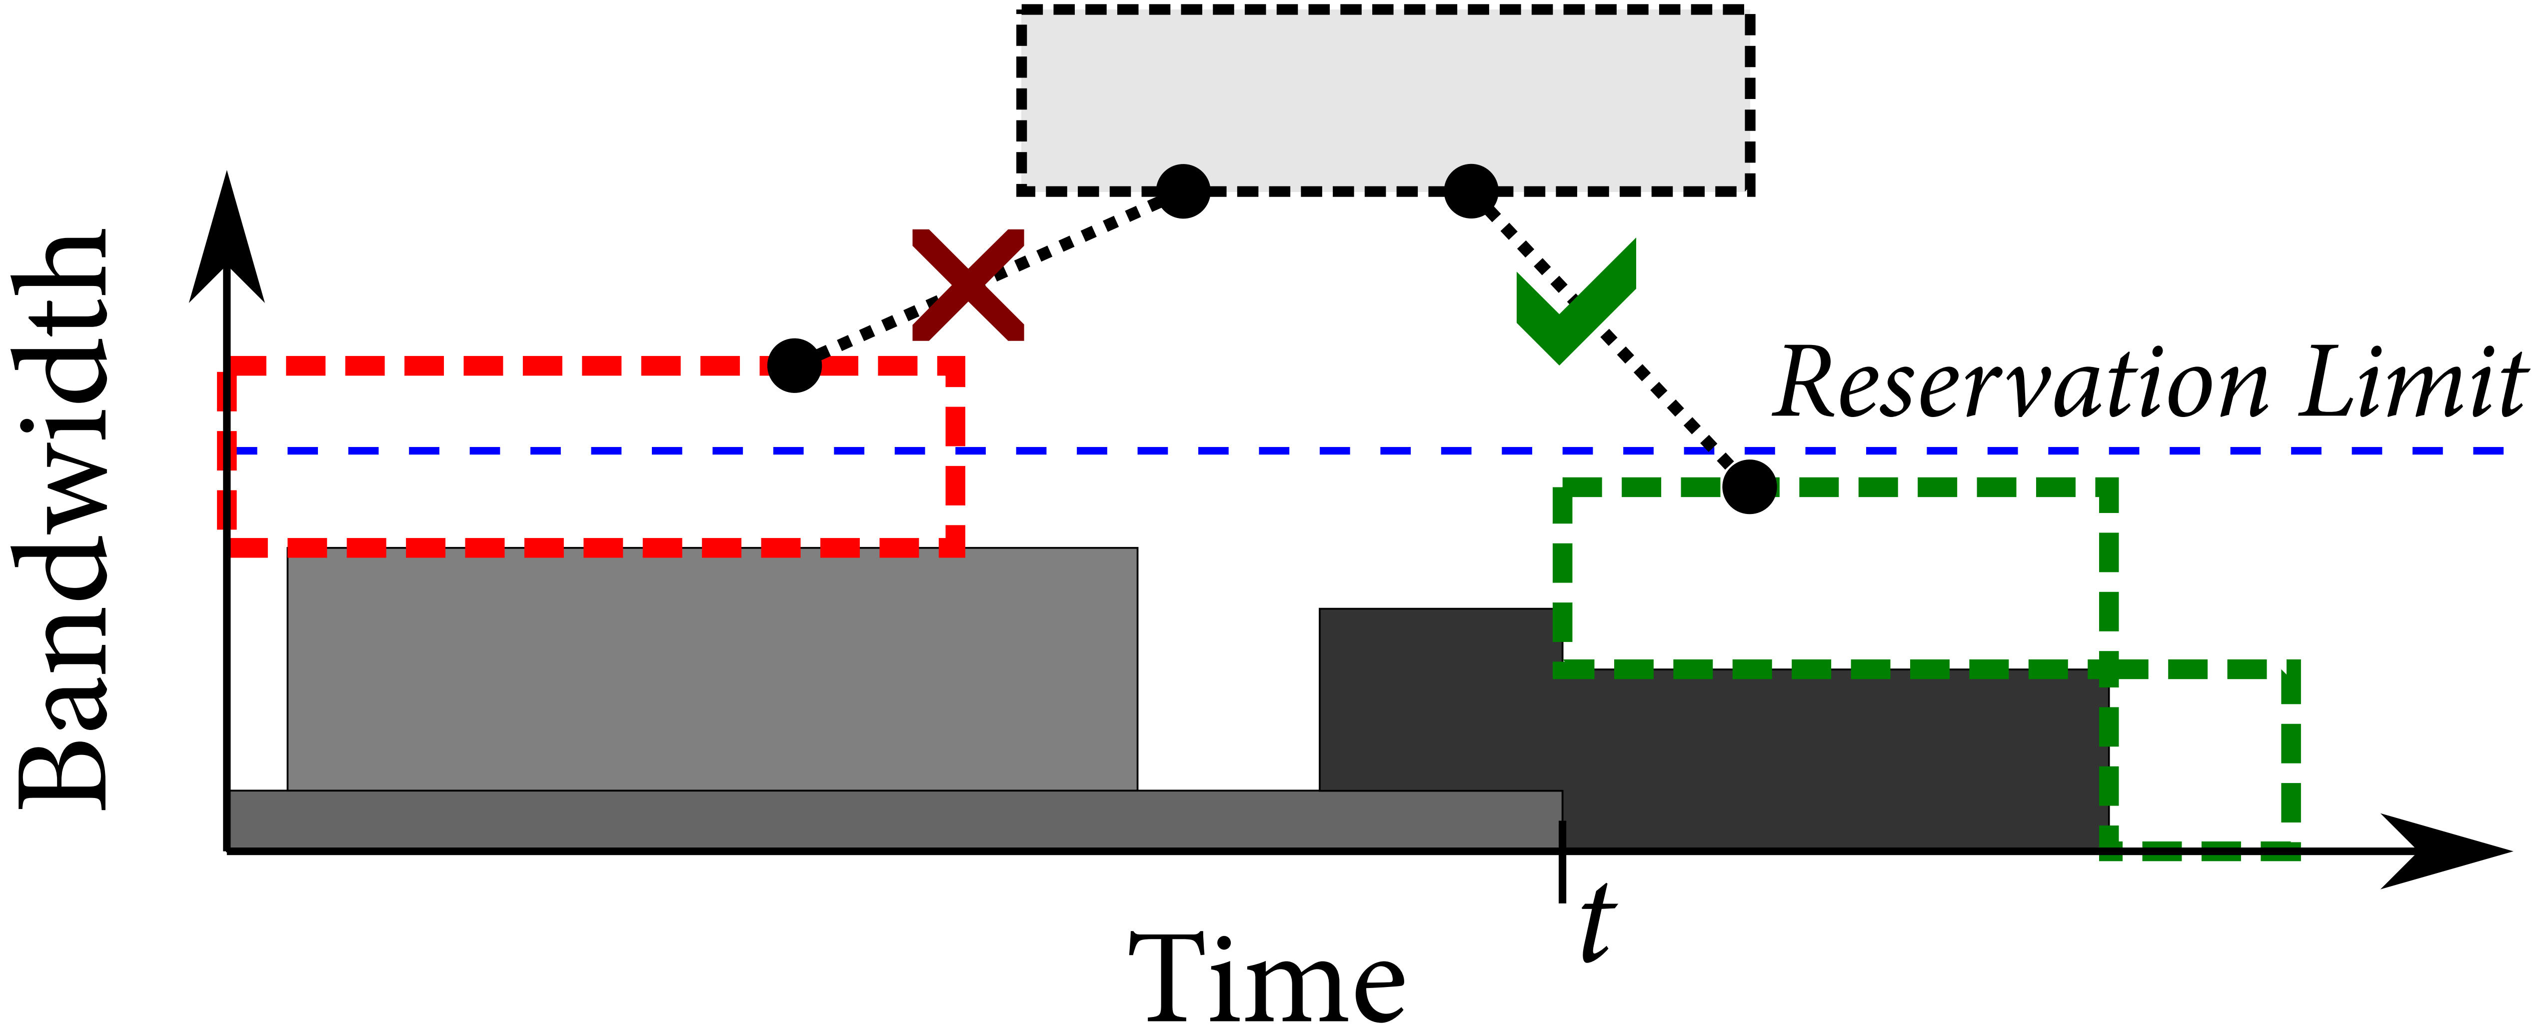
Bandwidth
Reservation Limit
t
Time

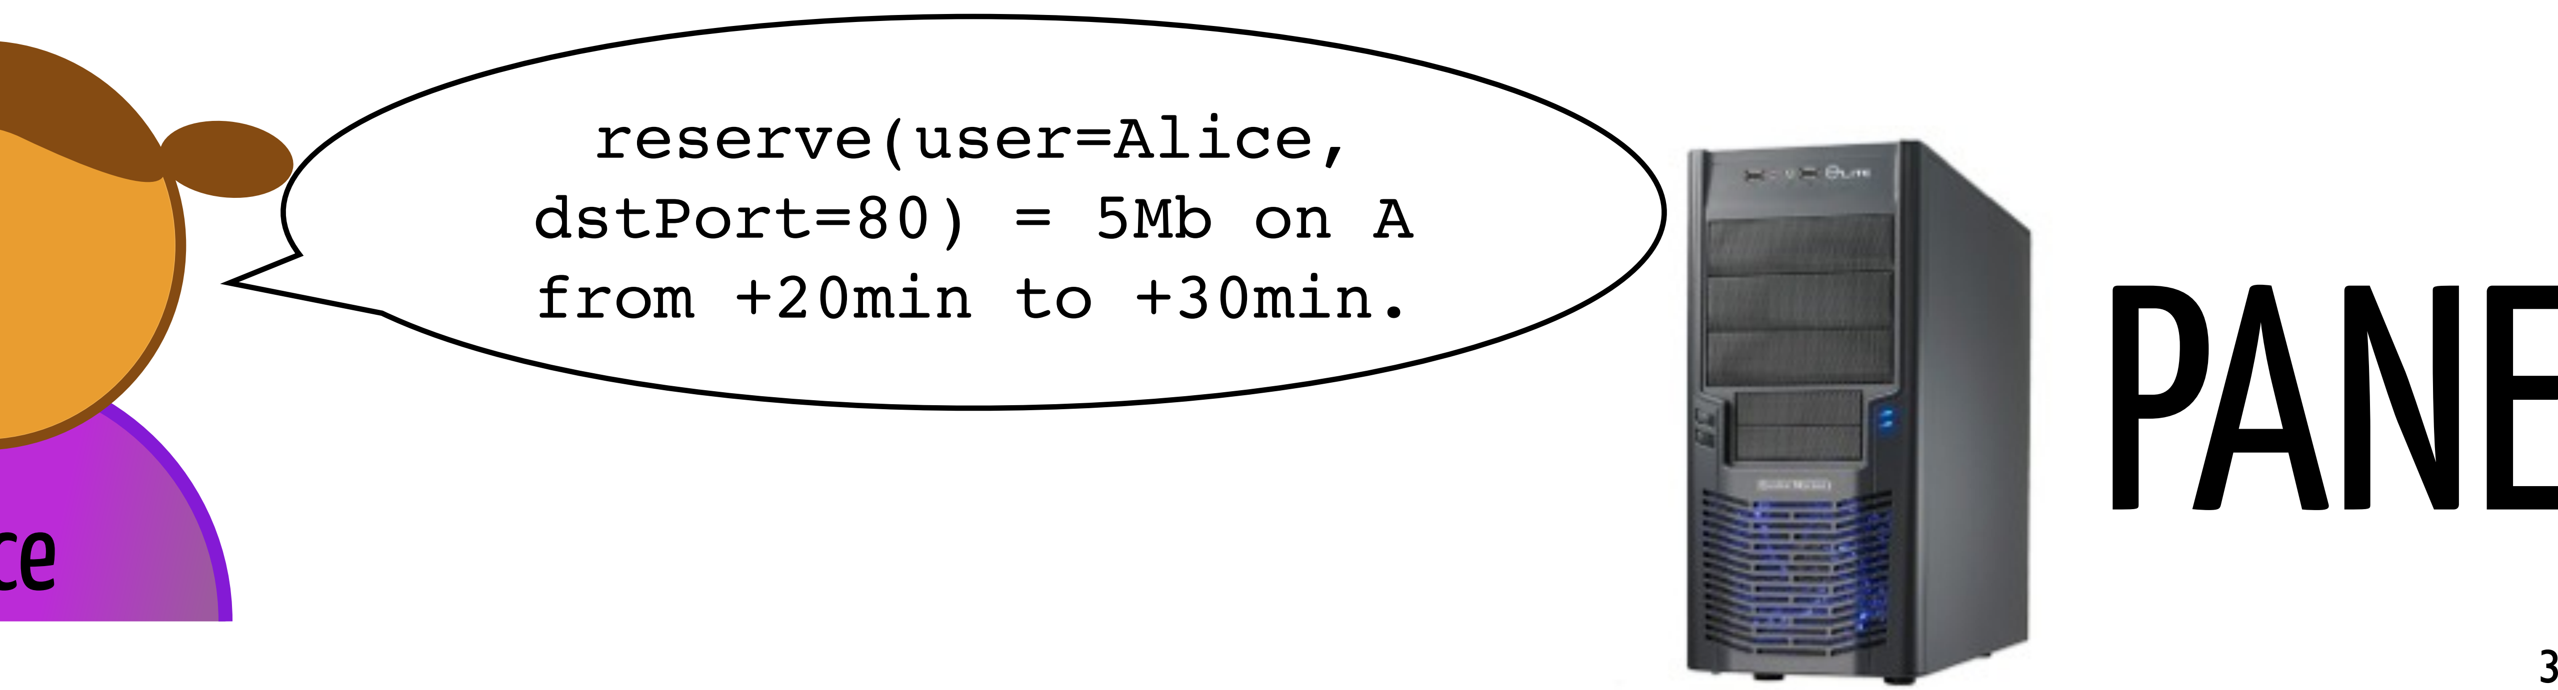
reserve(user=Alice,
dstPort=80) = 5Mb on A
from +20min to +30min.
PANE


```
reserve(user=Alice,
dstPort=80) = 5Mb on A
from now to +10min.
```

NO

```
reserve(user=Alice,
dstPort=80) = 5Mb on A
from +20min to +30min.
```

OK

PANE


```
reserve(user=Alice,
dstPort=80) = 5Mb on A
from now to +10min.
```

NO

```
reserve(user=Alice,
dstPort=80) = 5Mb on A
from +20min to +30min.
```

OK

PANE

PANE

PANE

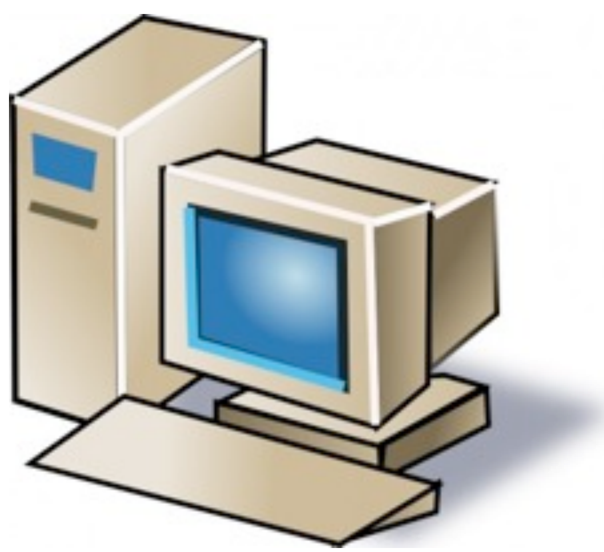

10.0.0.2

PANE

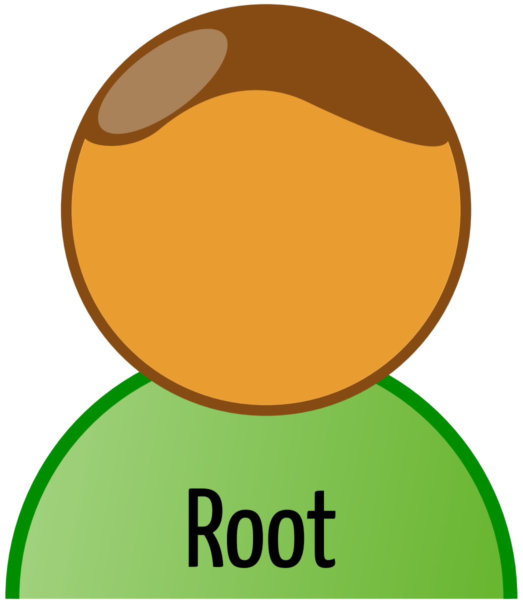

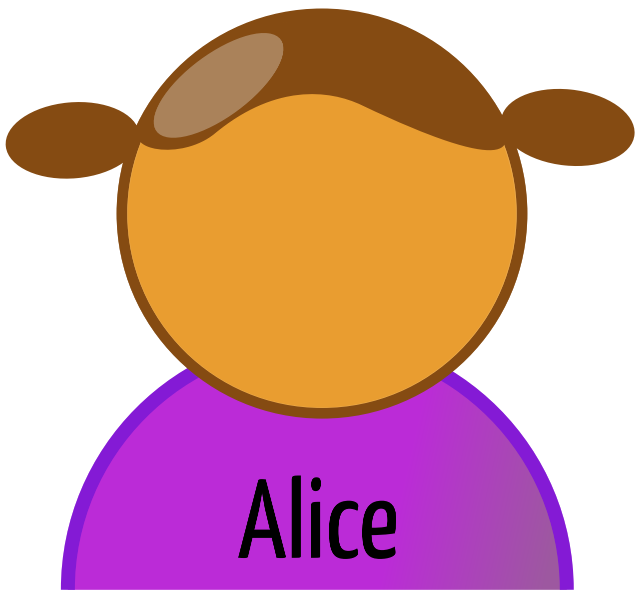

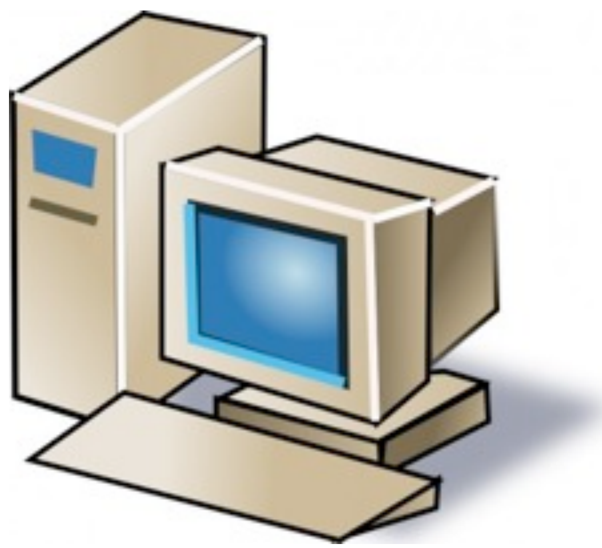

10.0.0.2

PANE

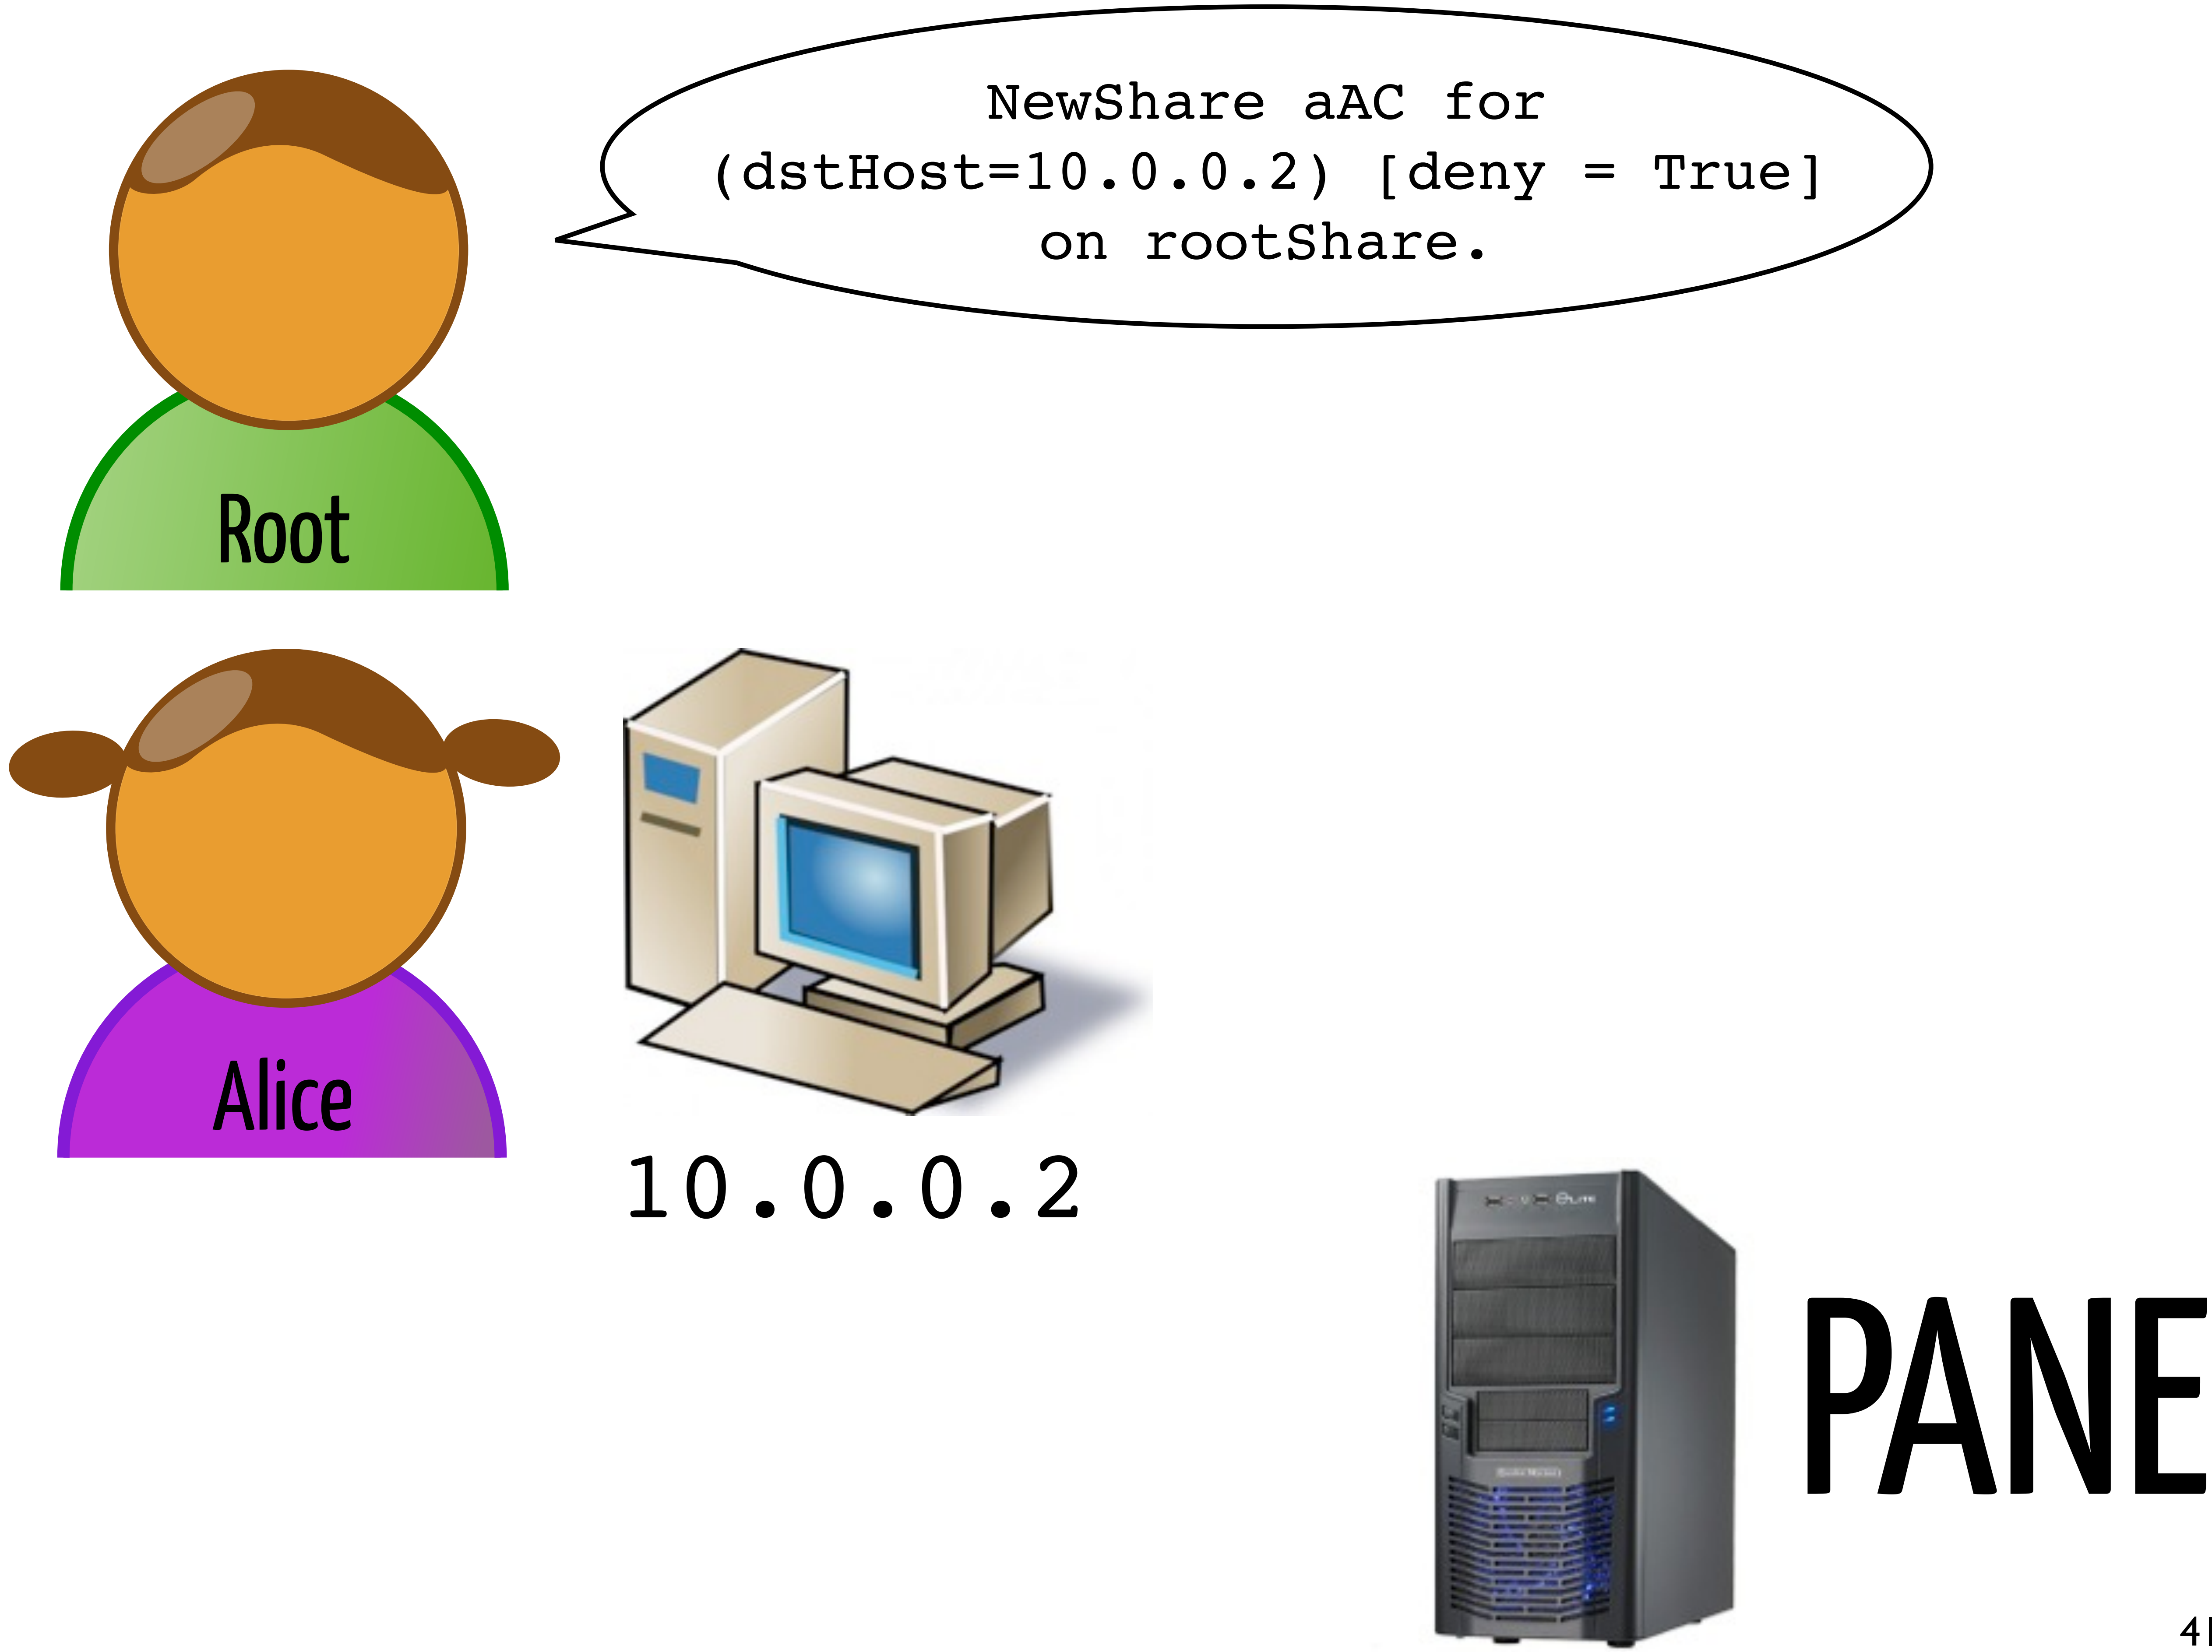
NewShare aAC for
(dstHost=10.0.0.2) [deny = True]
on rootShare.
Root
Alice
10.0.0.2
PANE

NewShare aAC for
(dstHost=10.0.0.2) [deny = True]
on rootShare.

OK

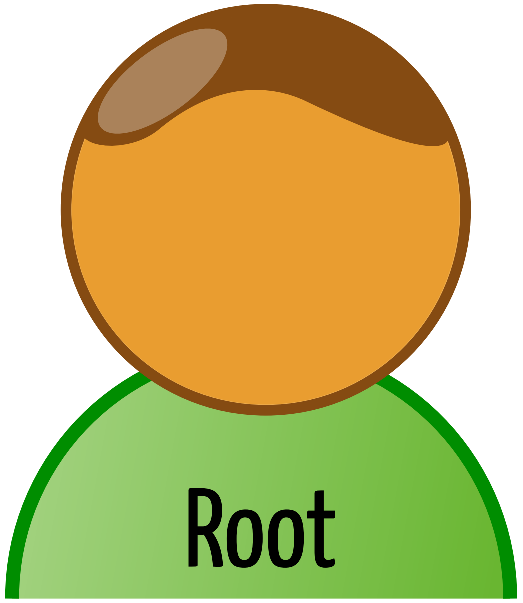

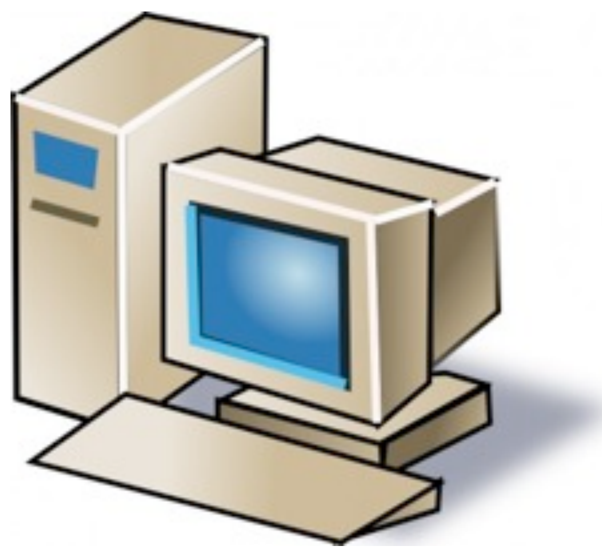

10.0.0.2

PANE

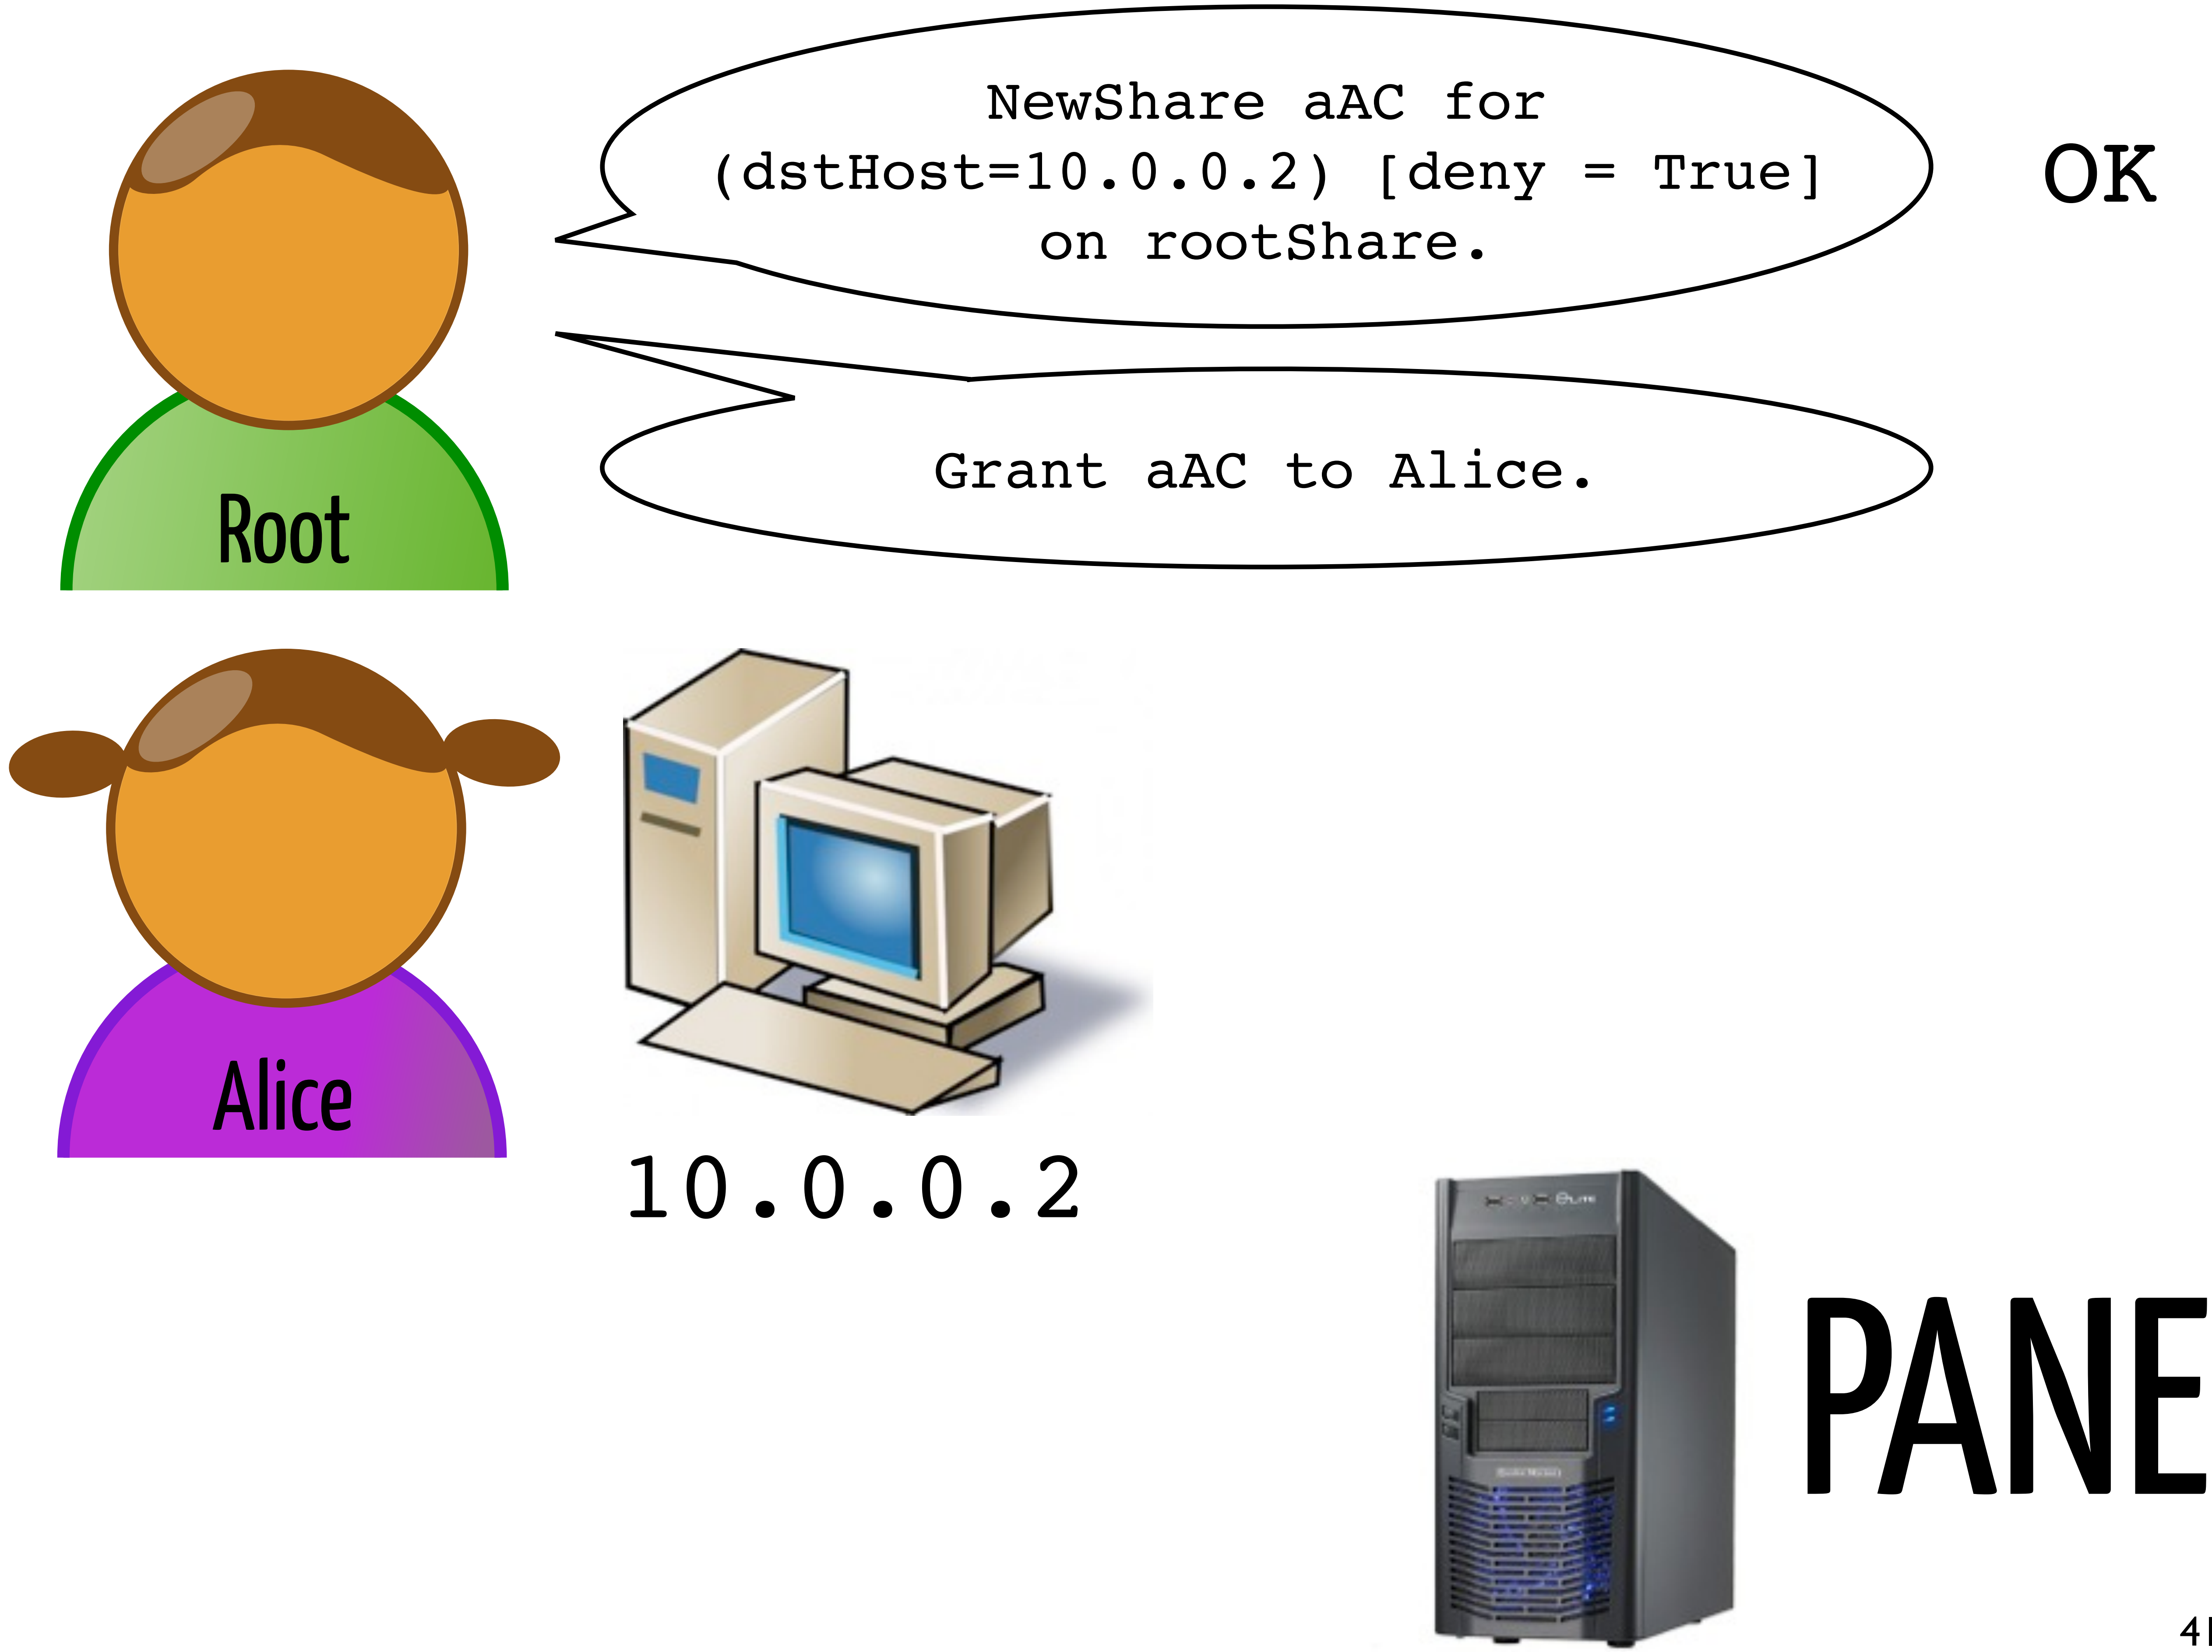

NewShare aAC for (dstHost=10.0.0.2) [deny = True] on rootShare.
OK
Grant aAC to Alice.
Root
Alice
10.0.0.2
PANE

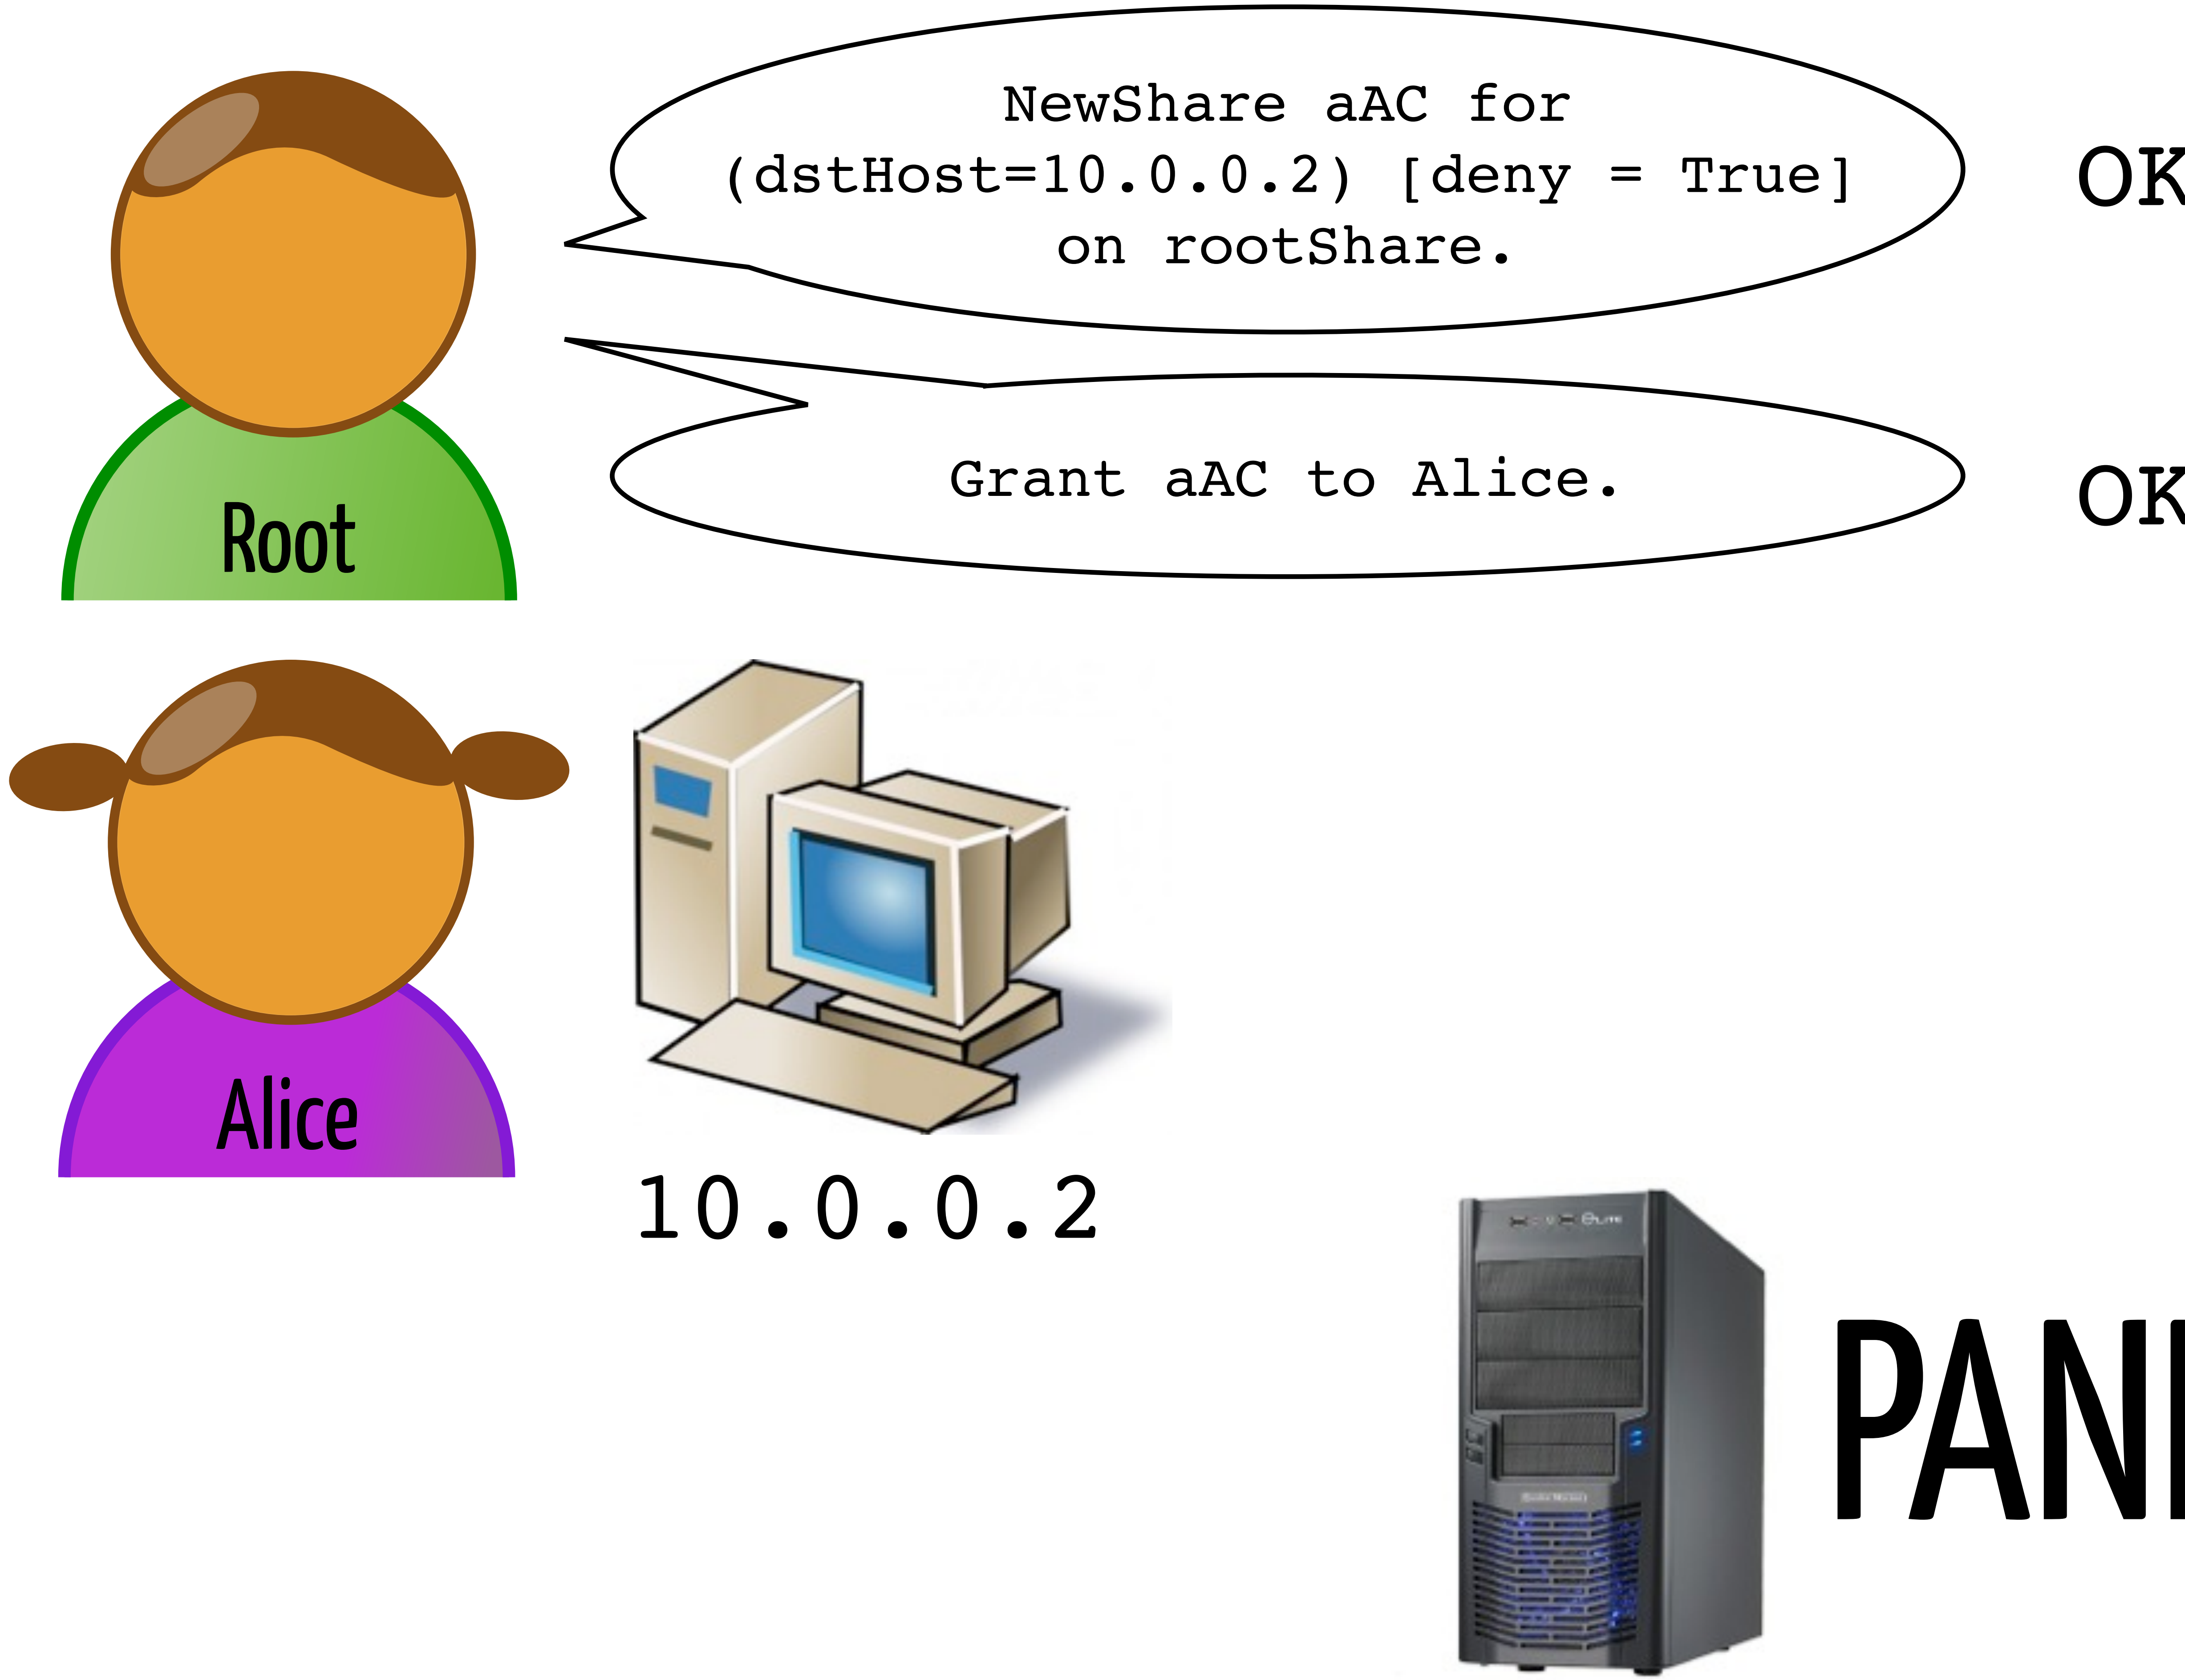
NewShare aAC for
(dstHost=10.0.0.2) [deny = True]
on rootShare.
OK
Grant aAC to Alice.
OK
Root
Alice
10.0.0.2
PANE

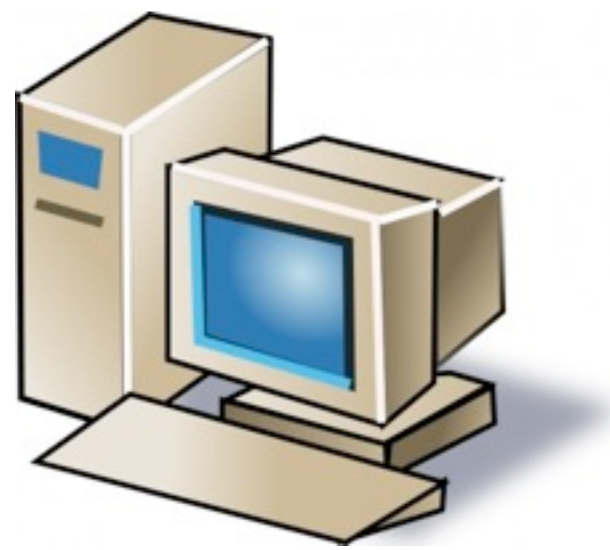

10.0.0.2

PANE

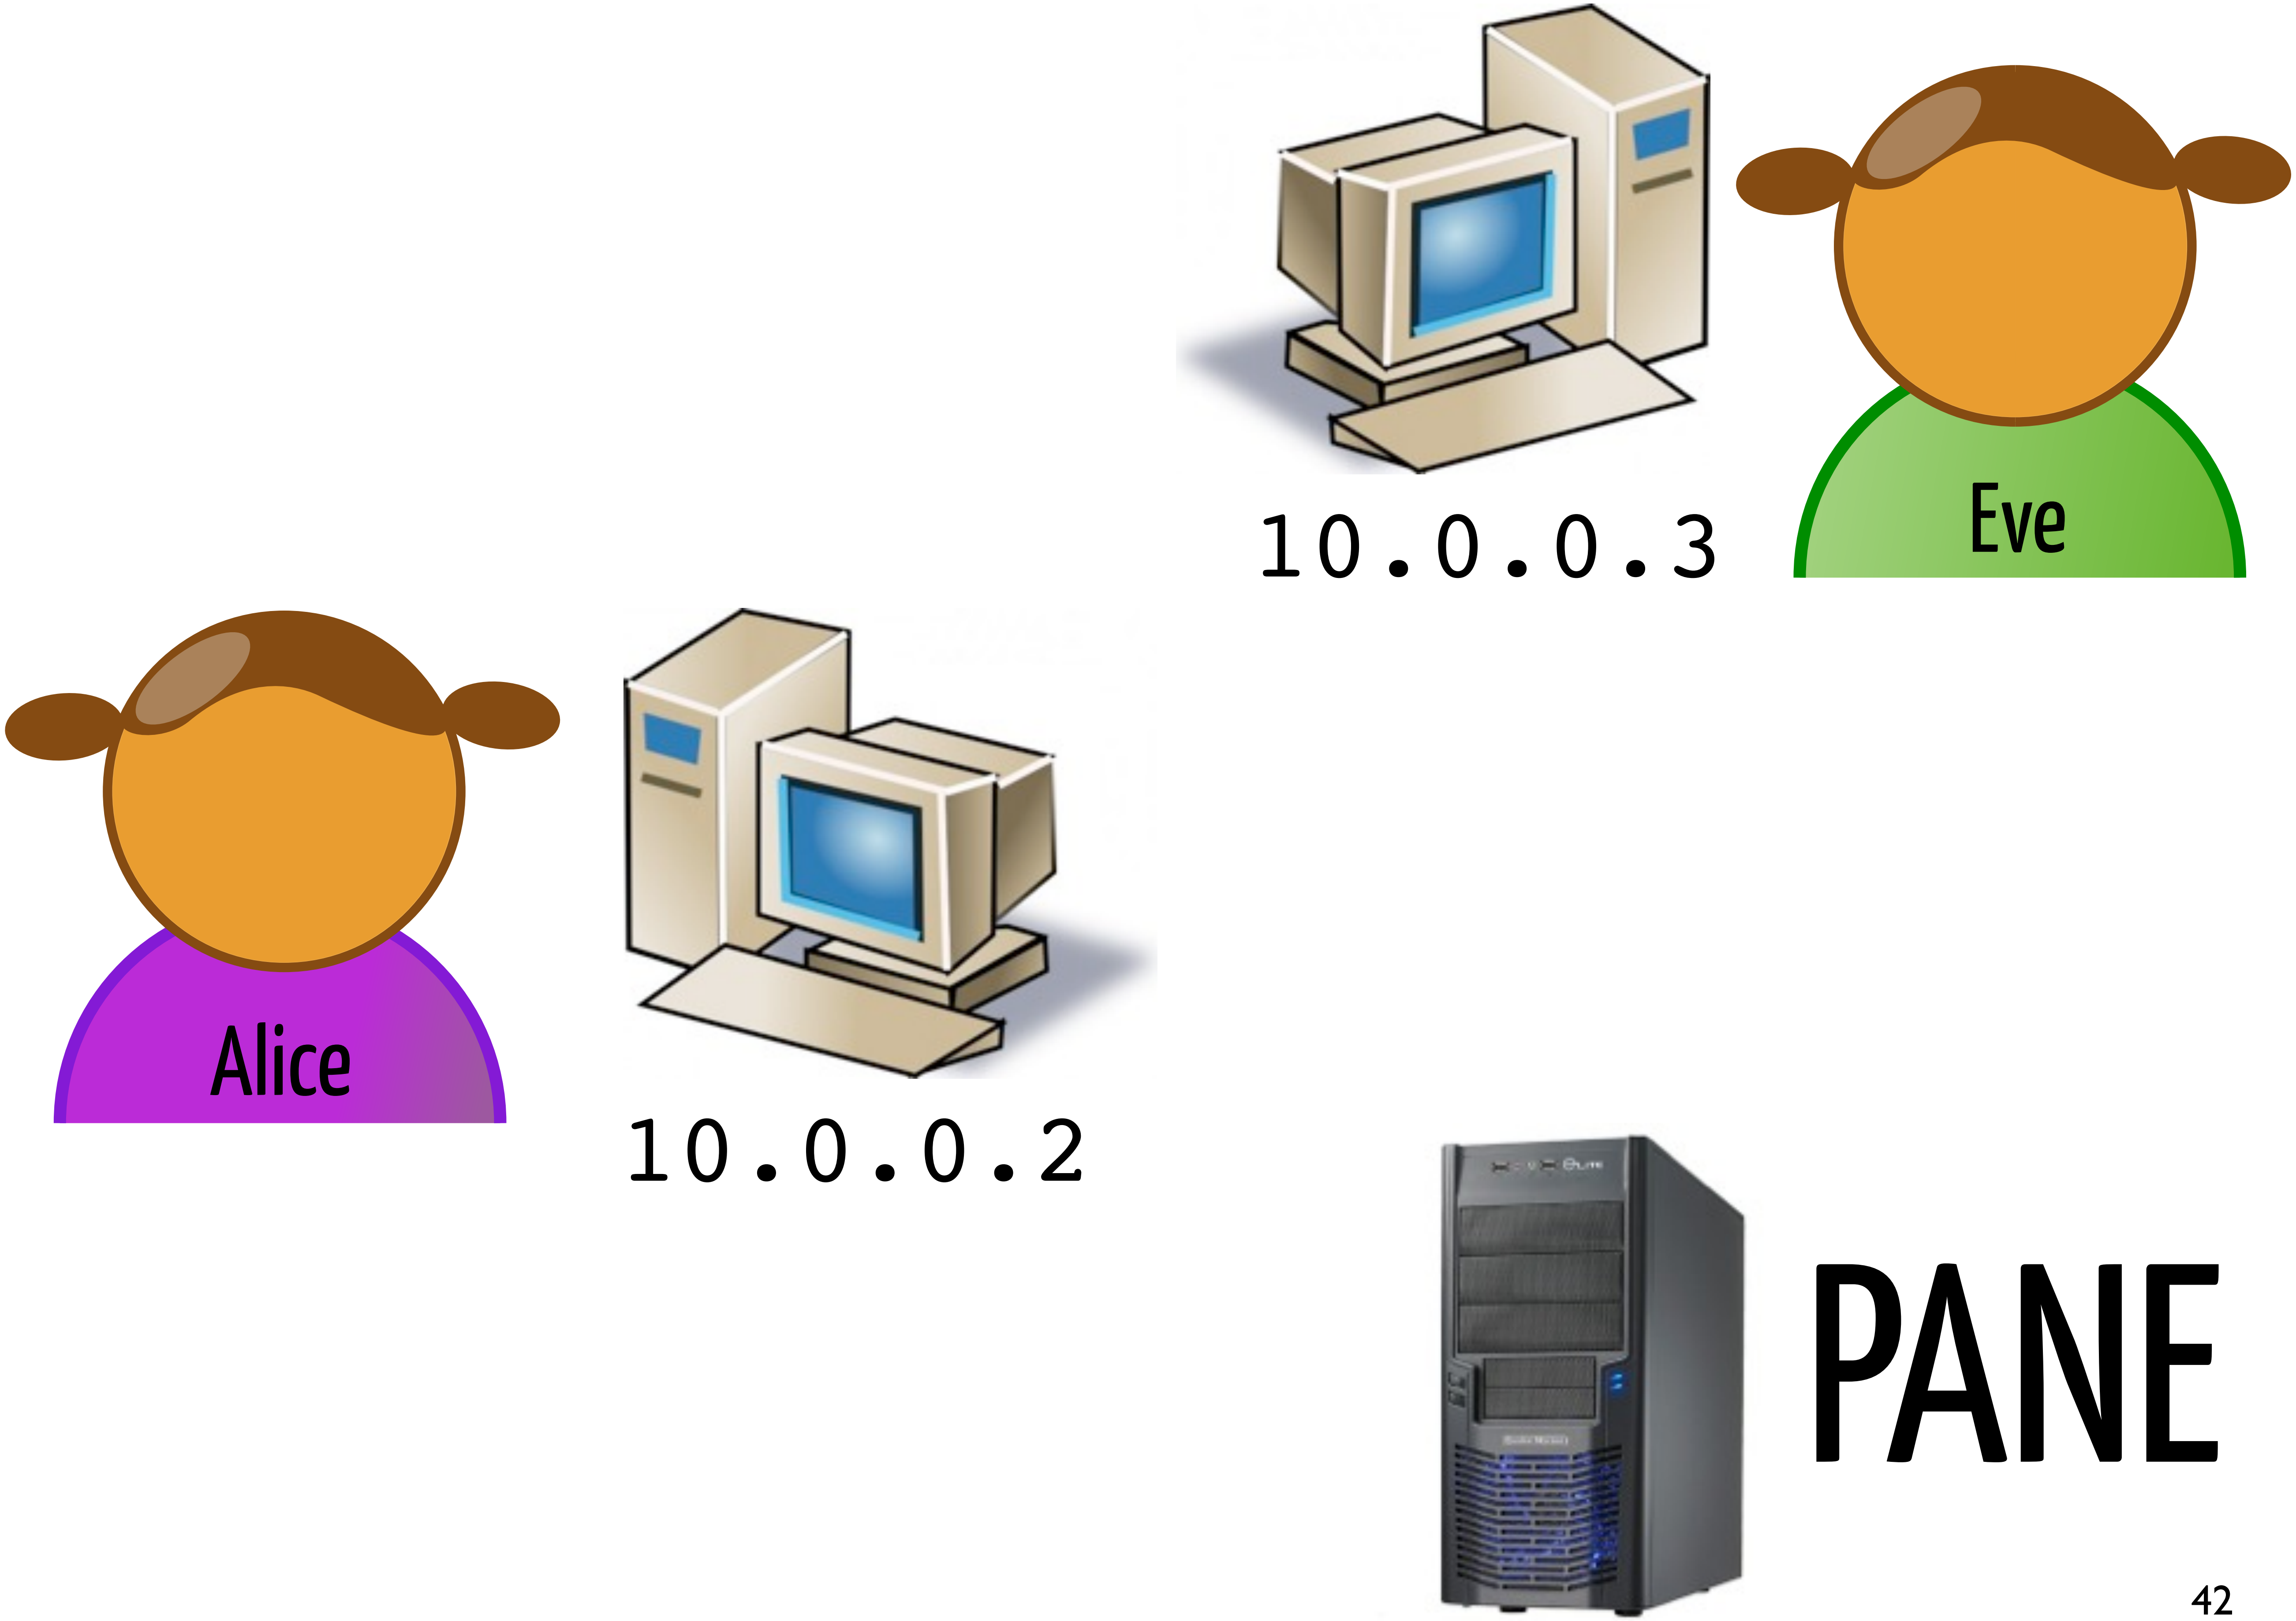
10.0.0.3
Eve
Alice
10.0.0.2
PANE

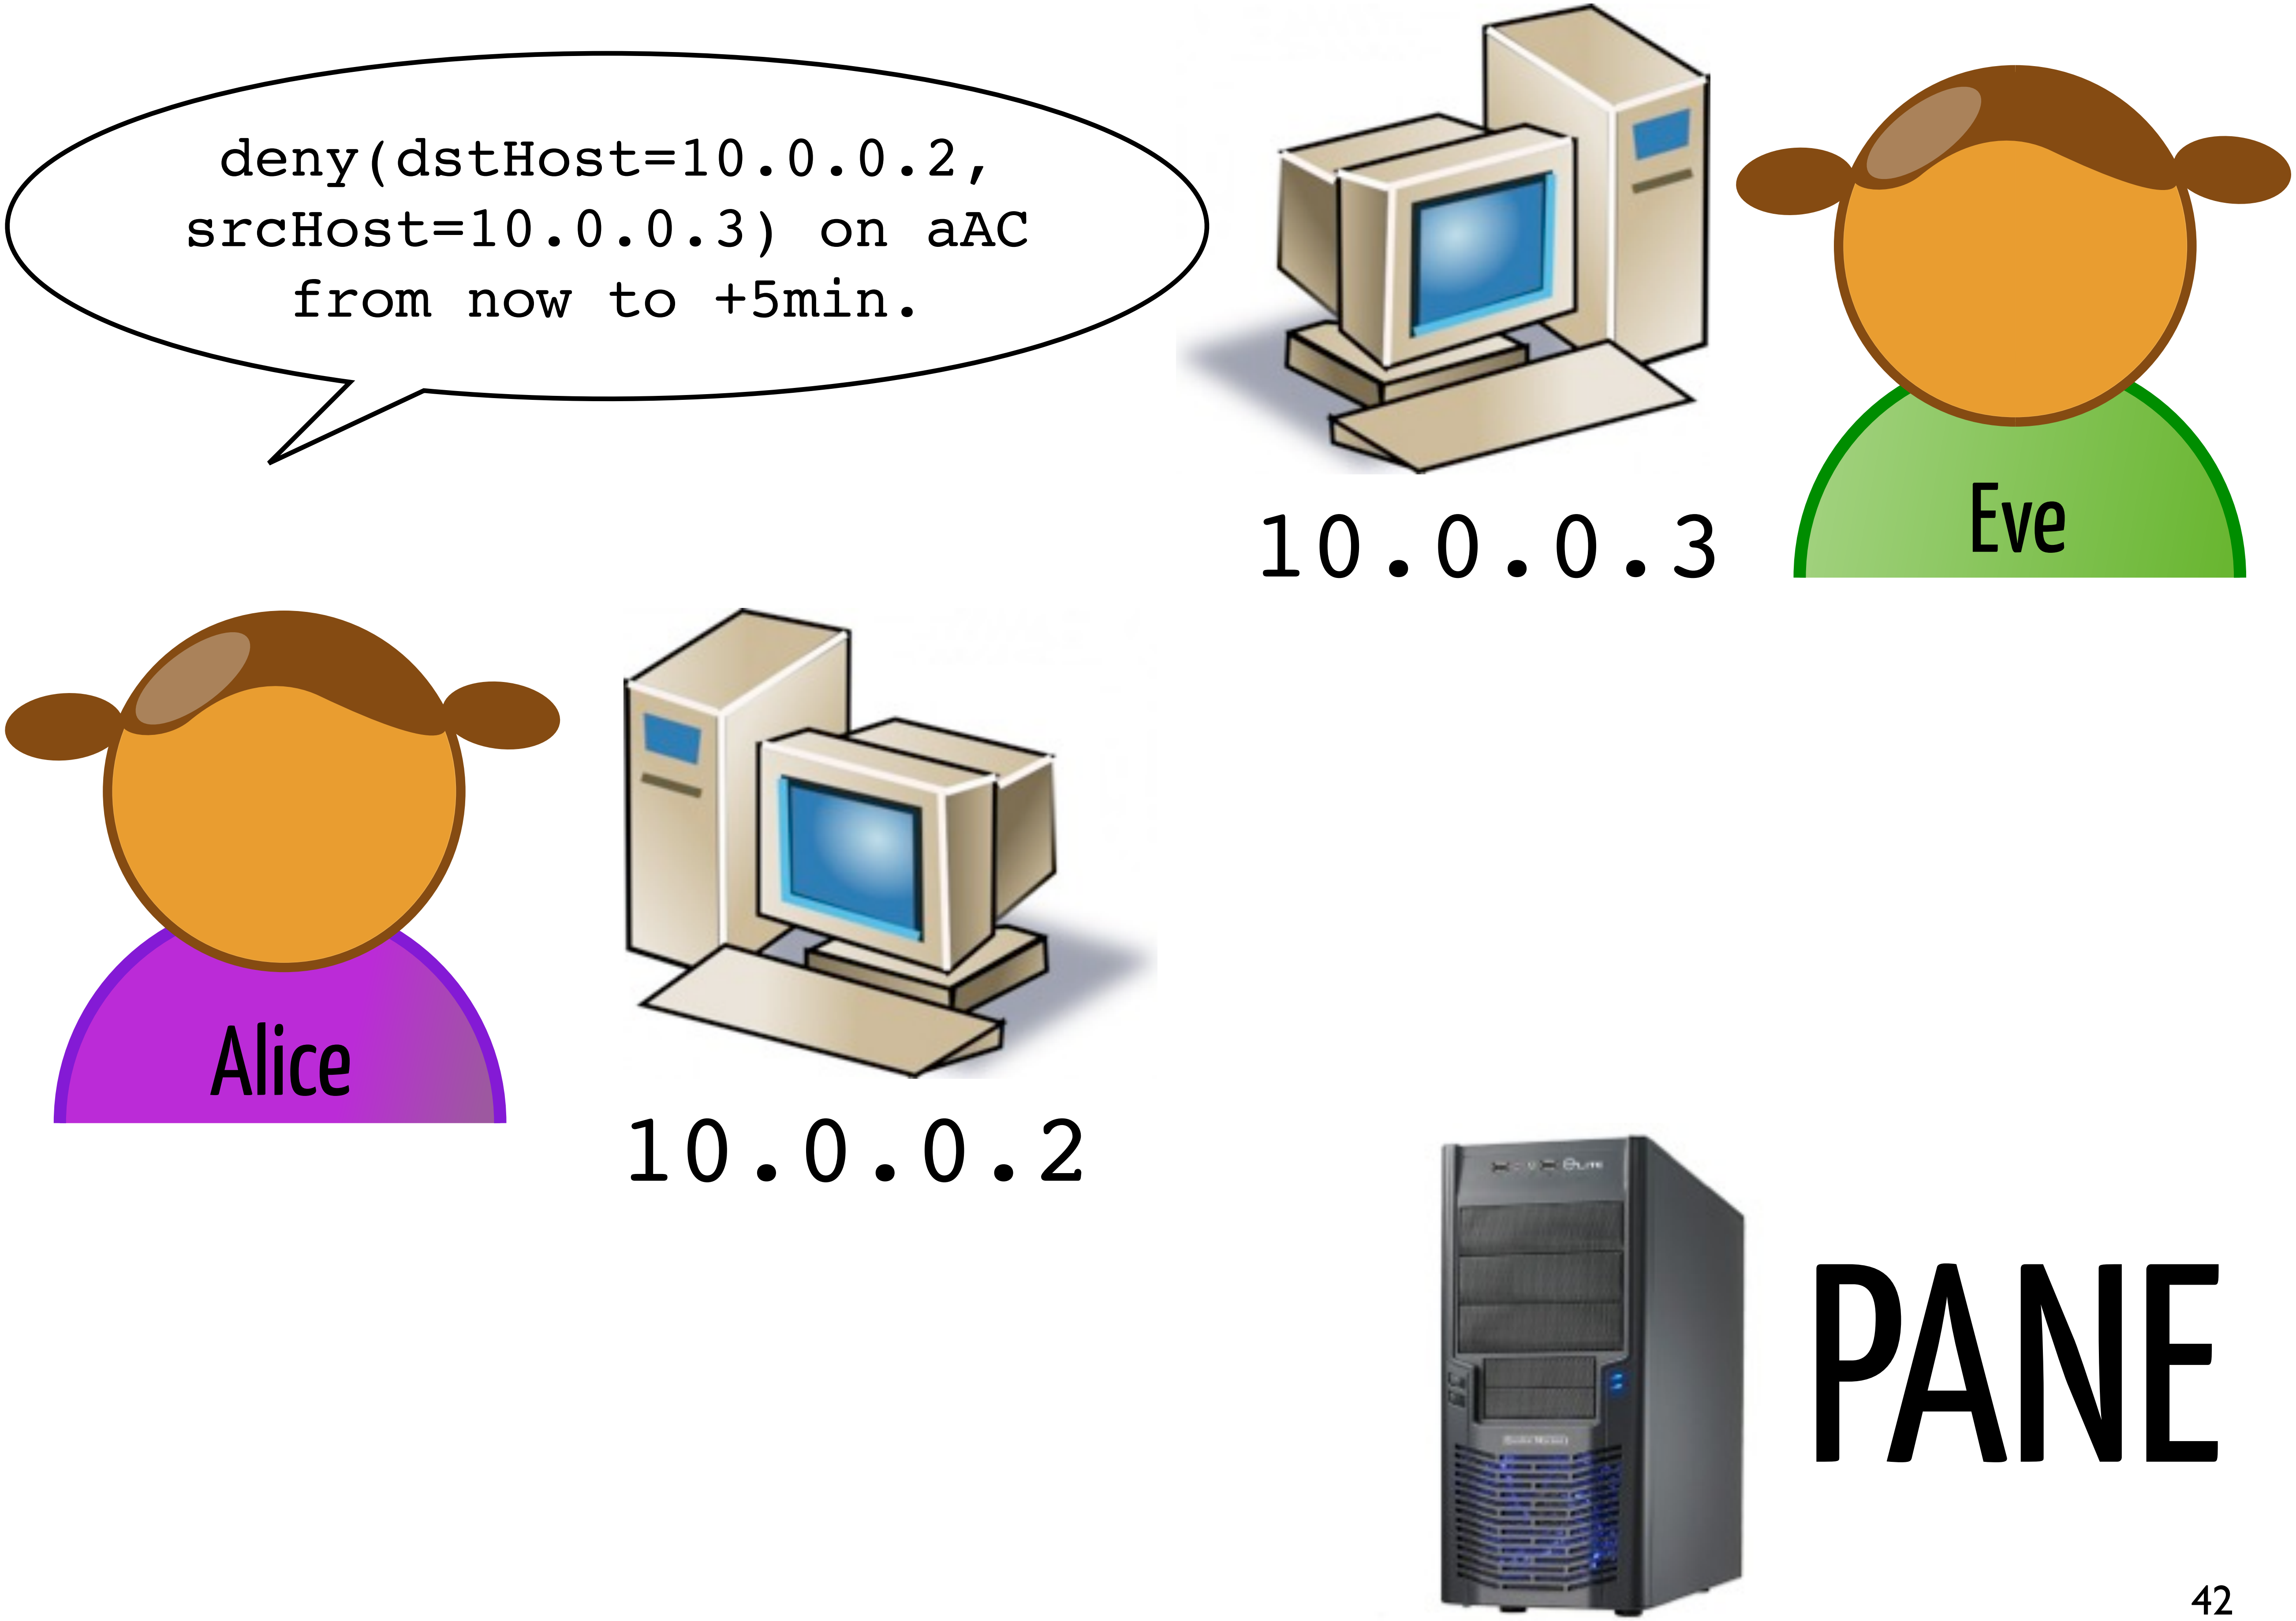
deny(dstHost=10.0.0.2,
srcHost=10.0.0.3) on aAC
from now to +5min.
10.0.0.3
Eve
Alice
10.0.0.2
PANE

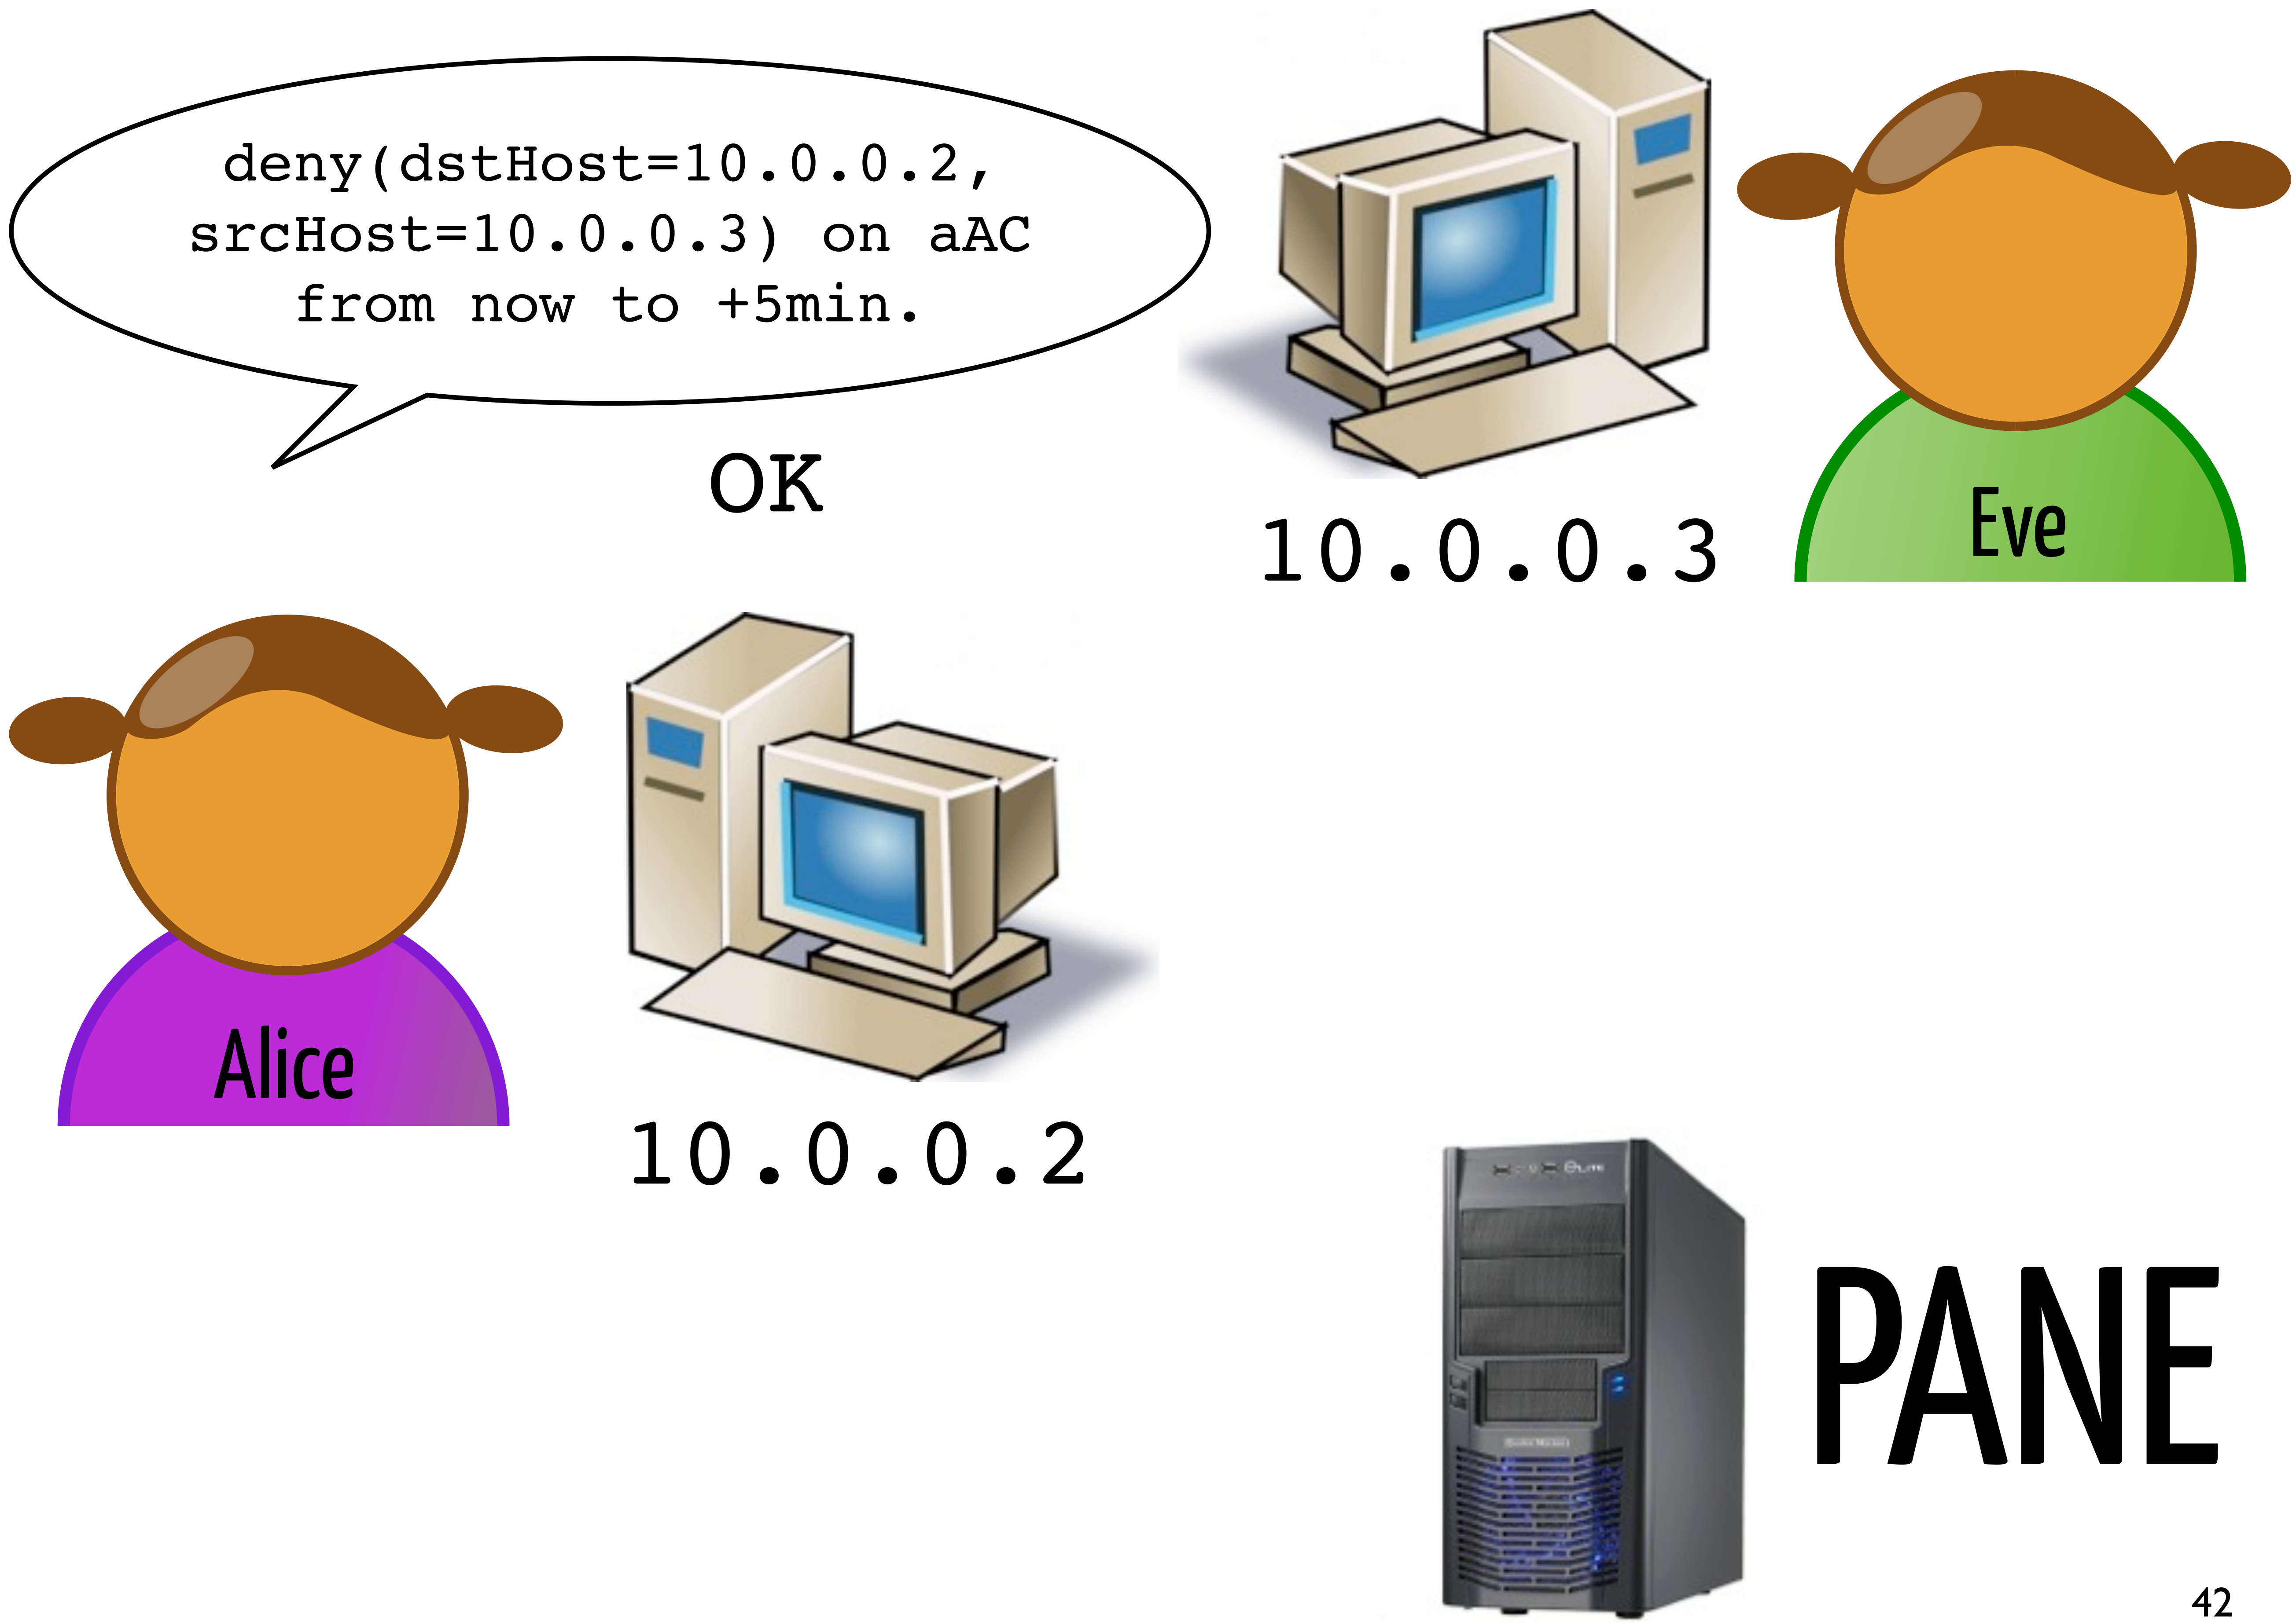
deny(dstHost=10.0.0.2,
srcHost=10.0.0.3) on aAC
from now to +5min.
OK
10.0.0.3
Eve
Alice
10.0.0.2
PANE

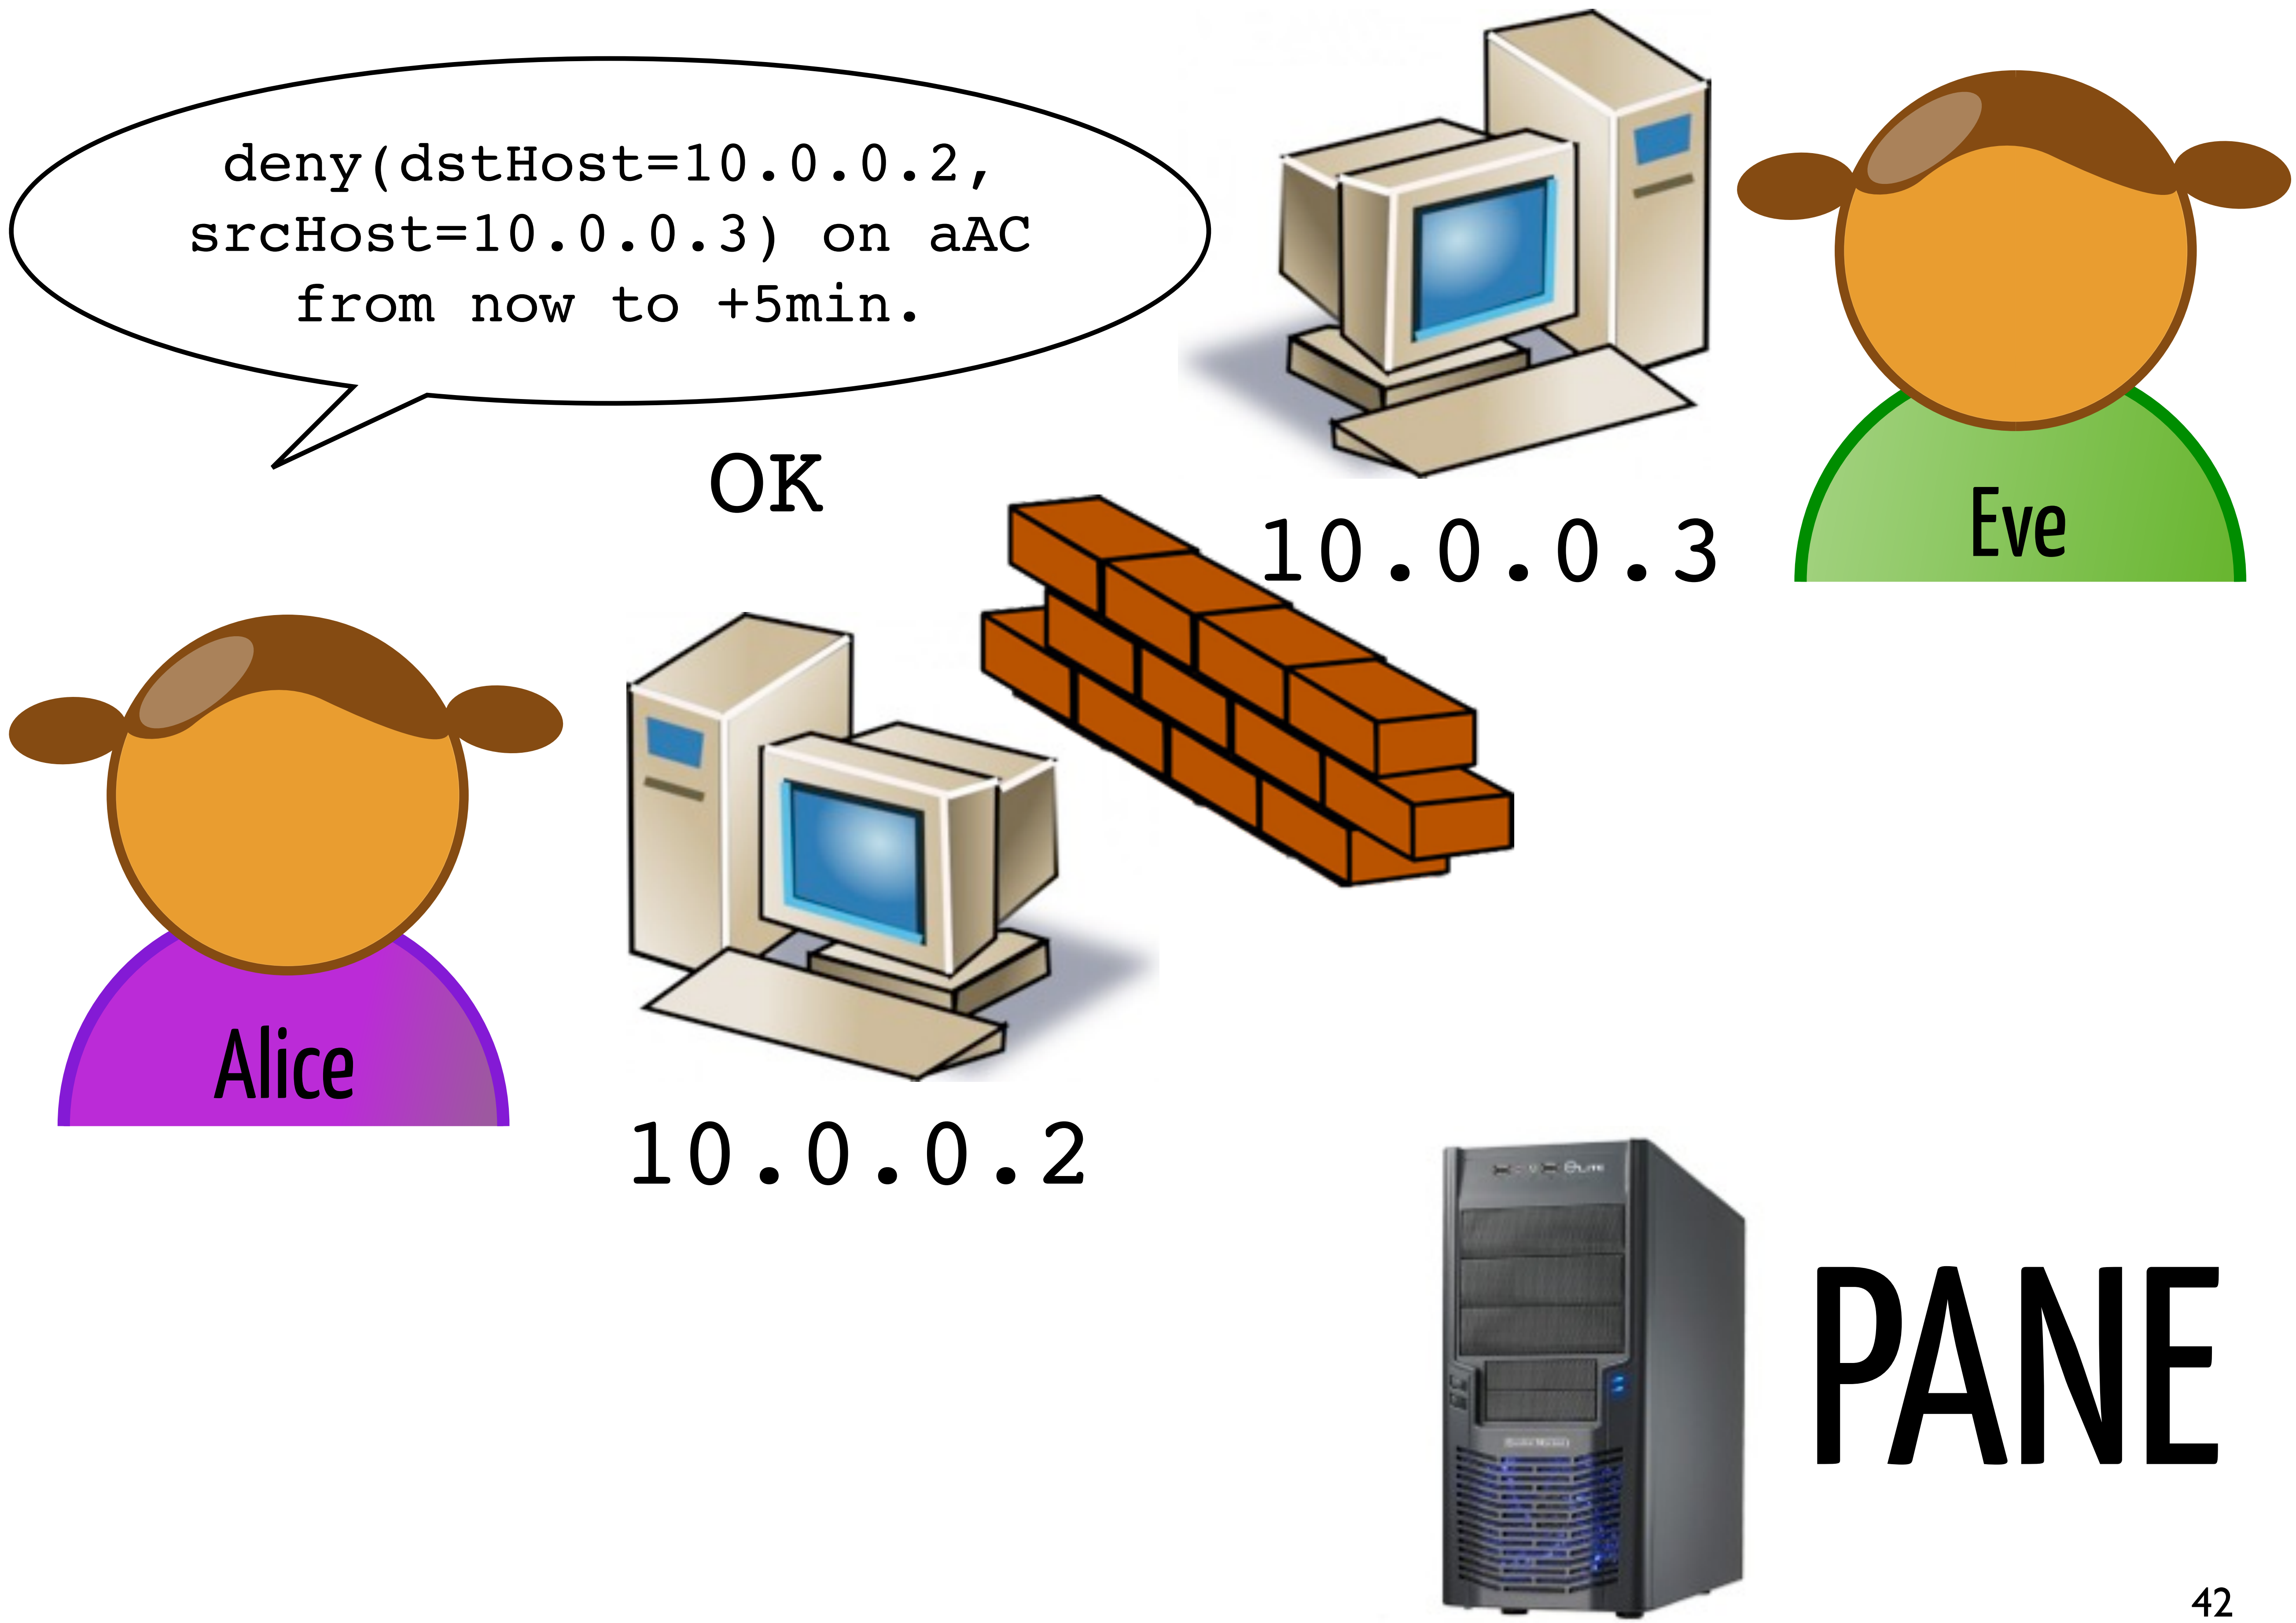

deny(dstHost=10.0.0.2,
srcHost=10.0.0.3) on aAC
from now to +5min.
OK
10.0.0.3
Eve
Alice
10.0.0.2
PANE

Dynamic Flow Processing

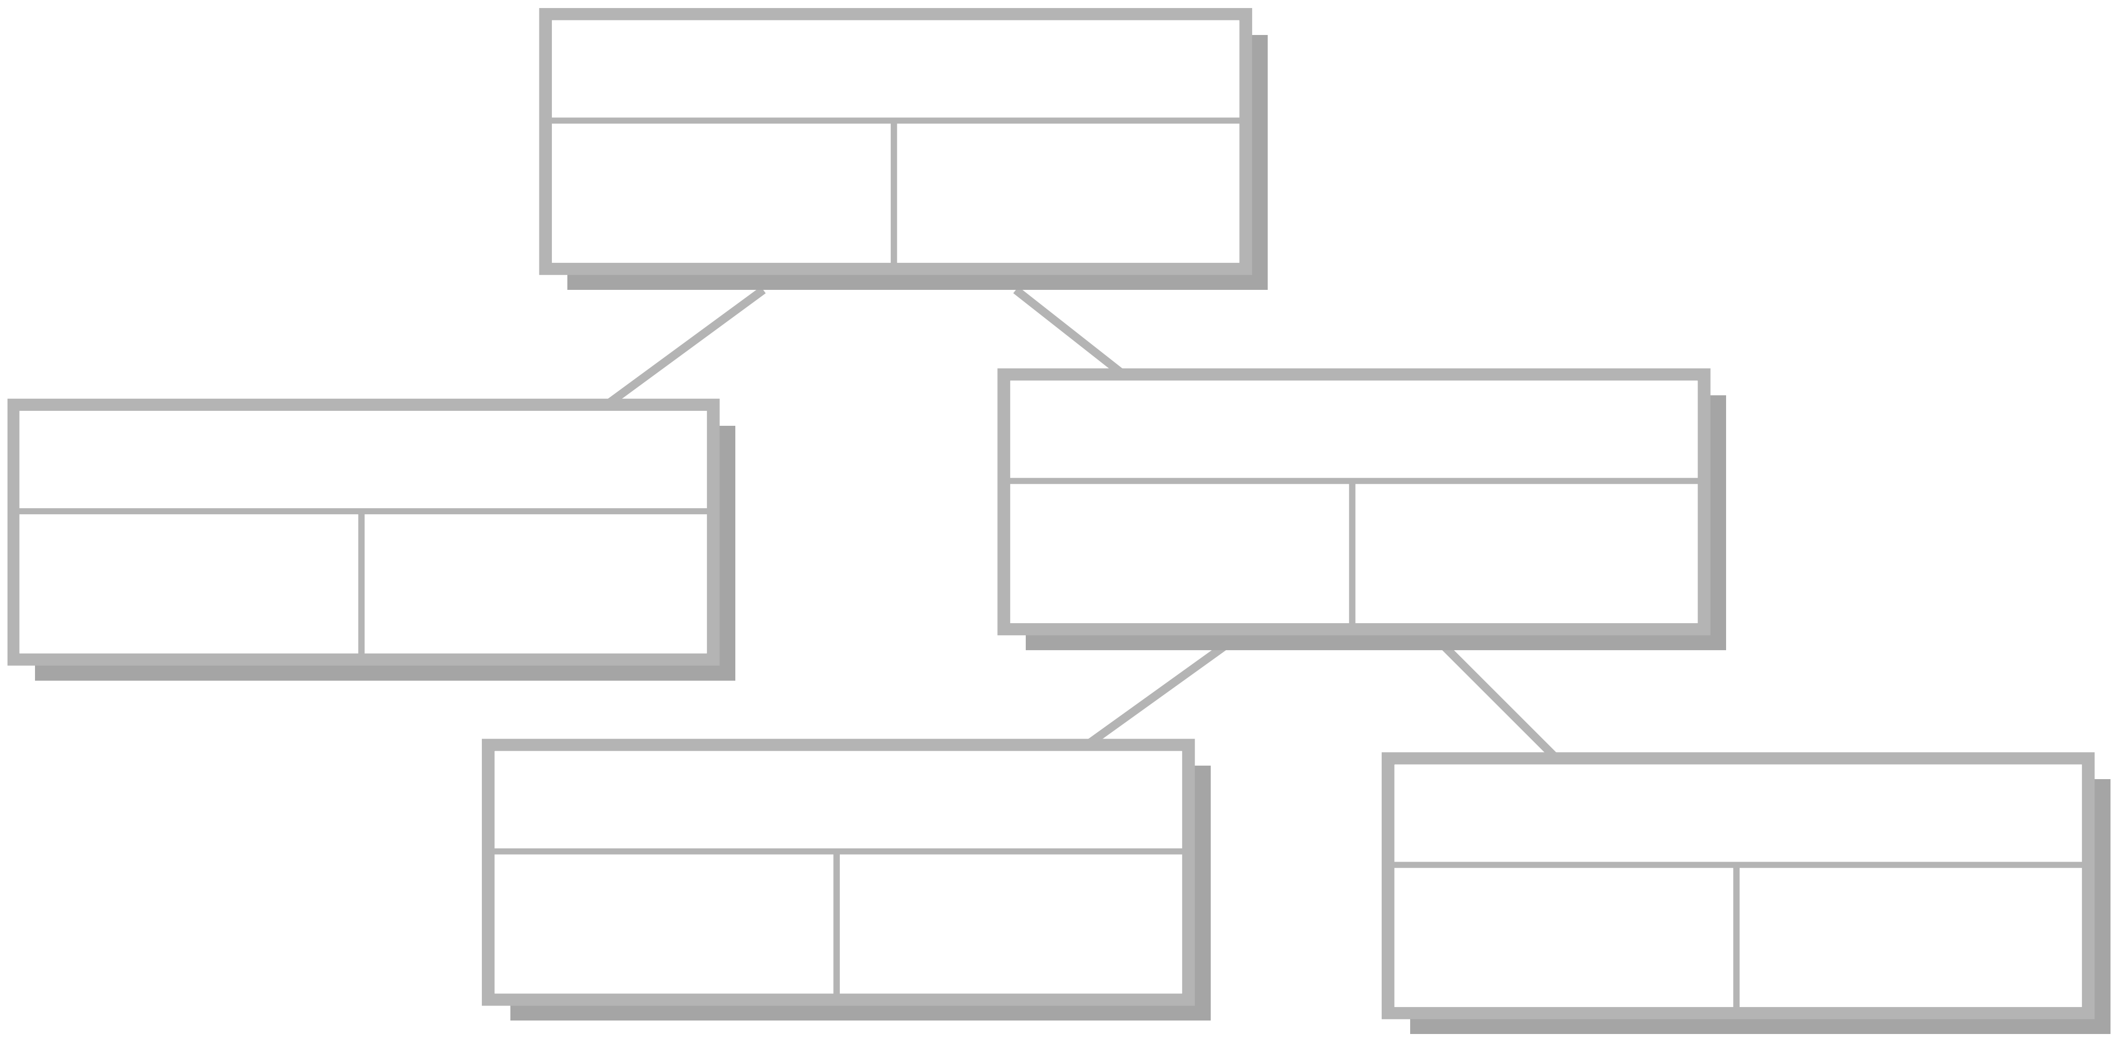

Hierarchy of Policies

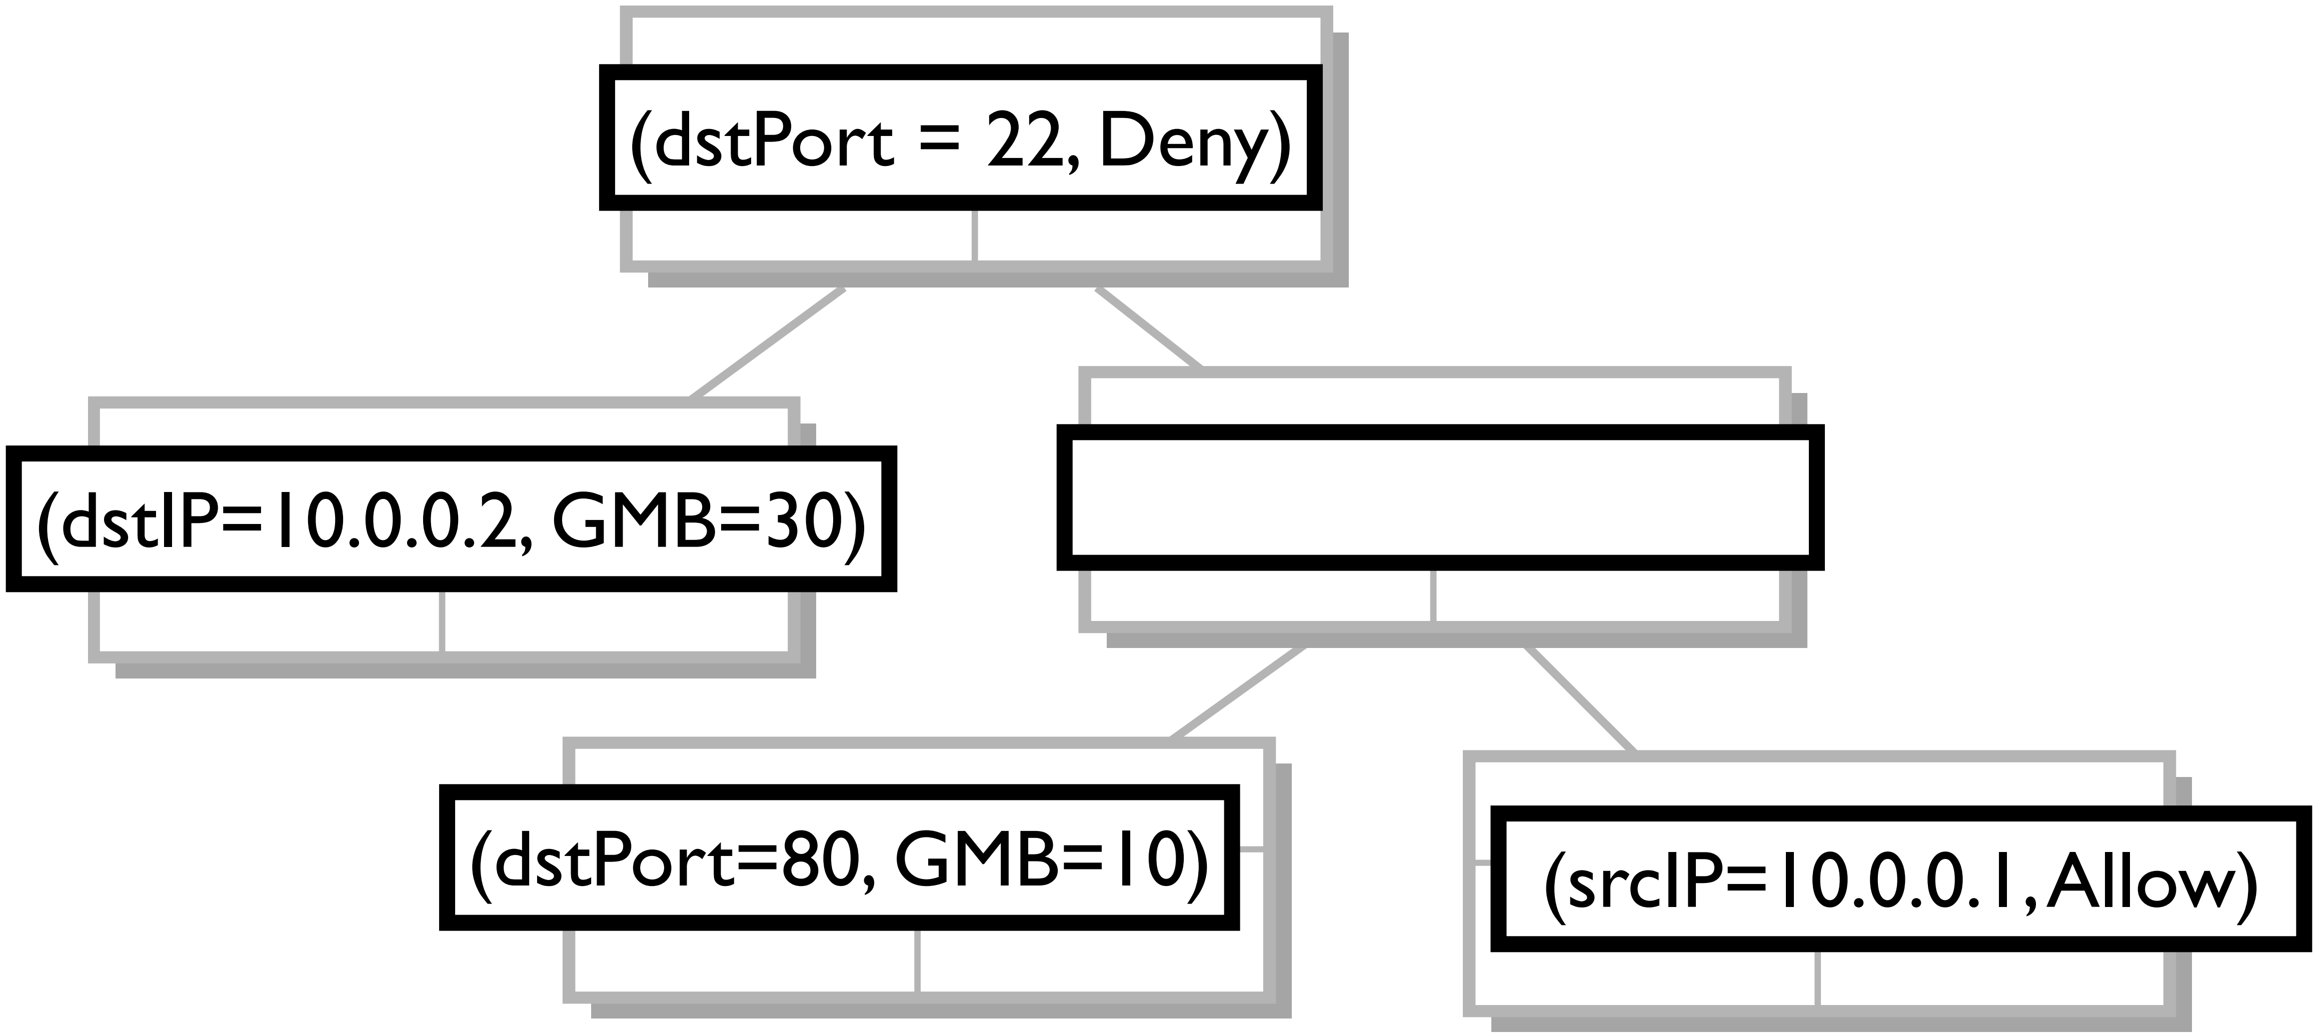

Hierarchy of Policies

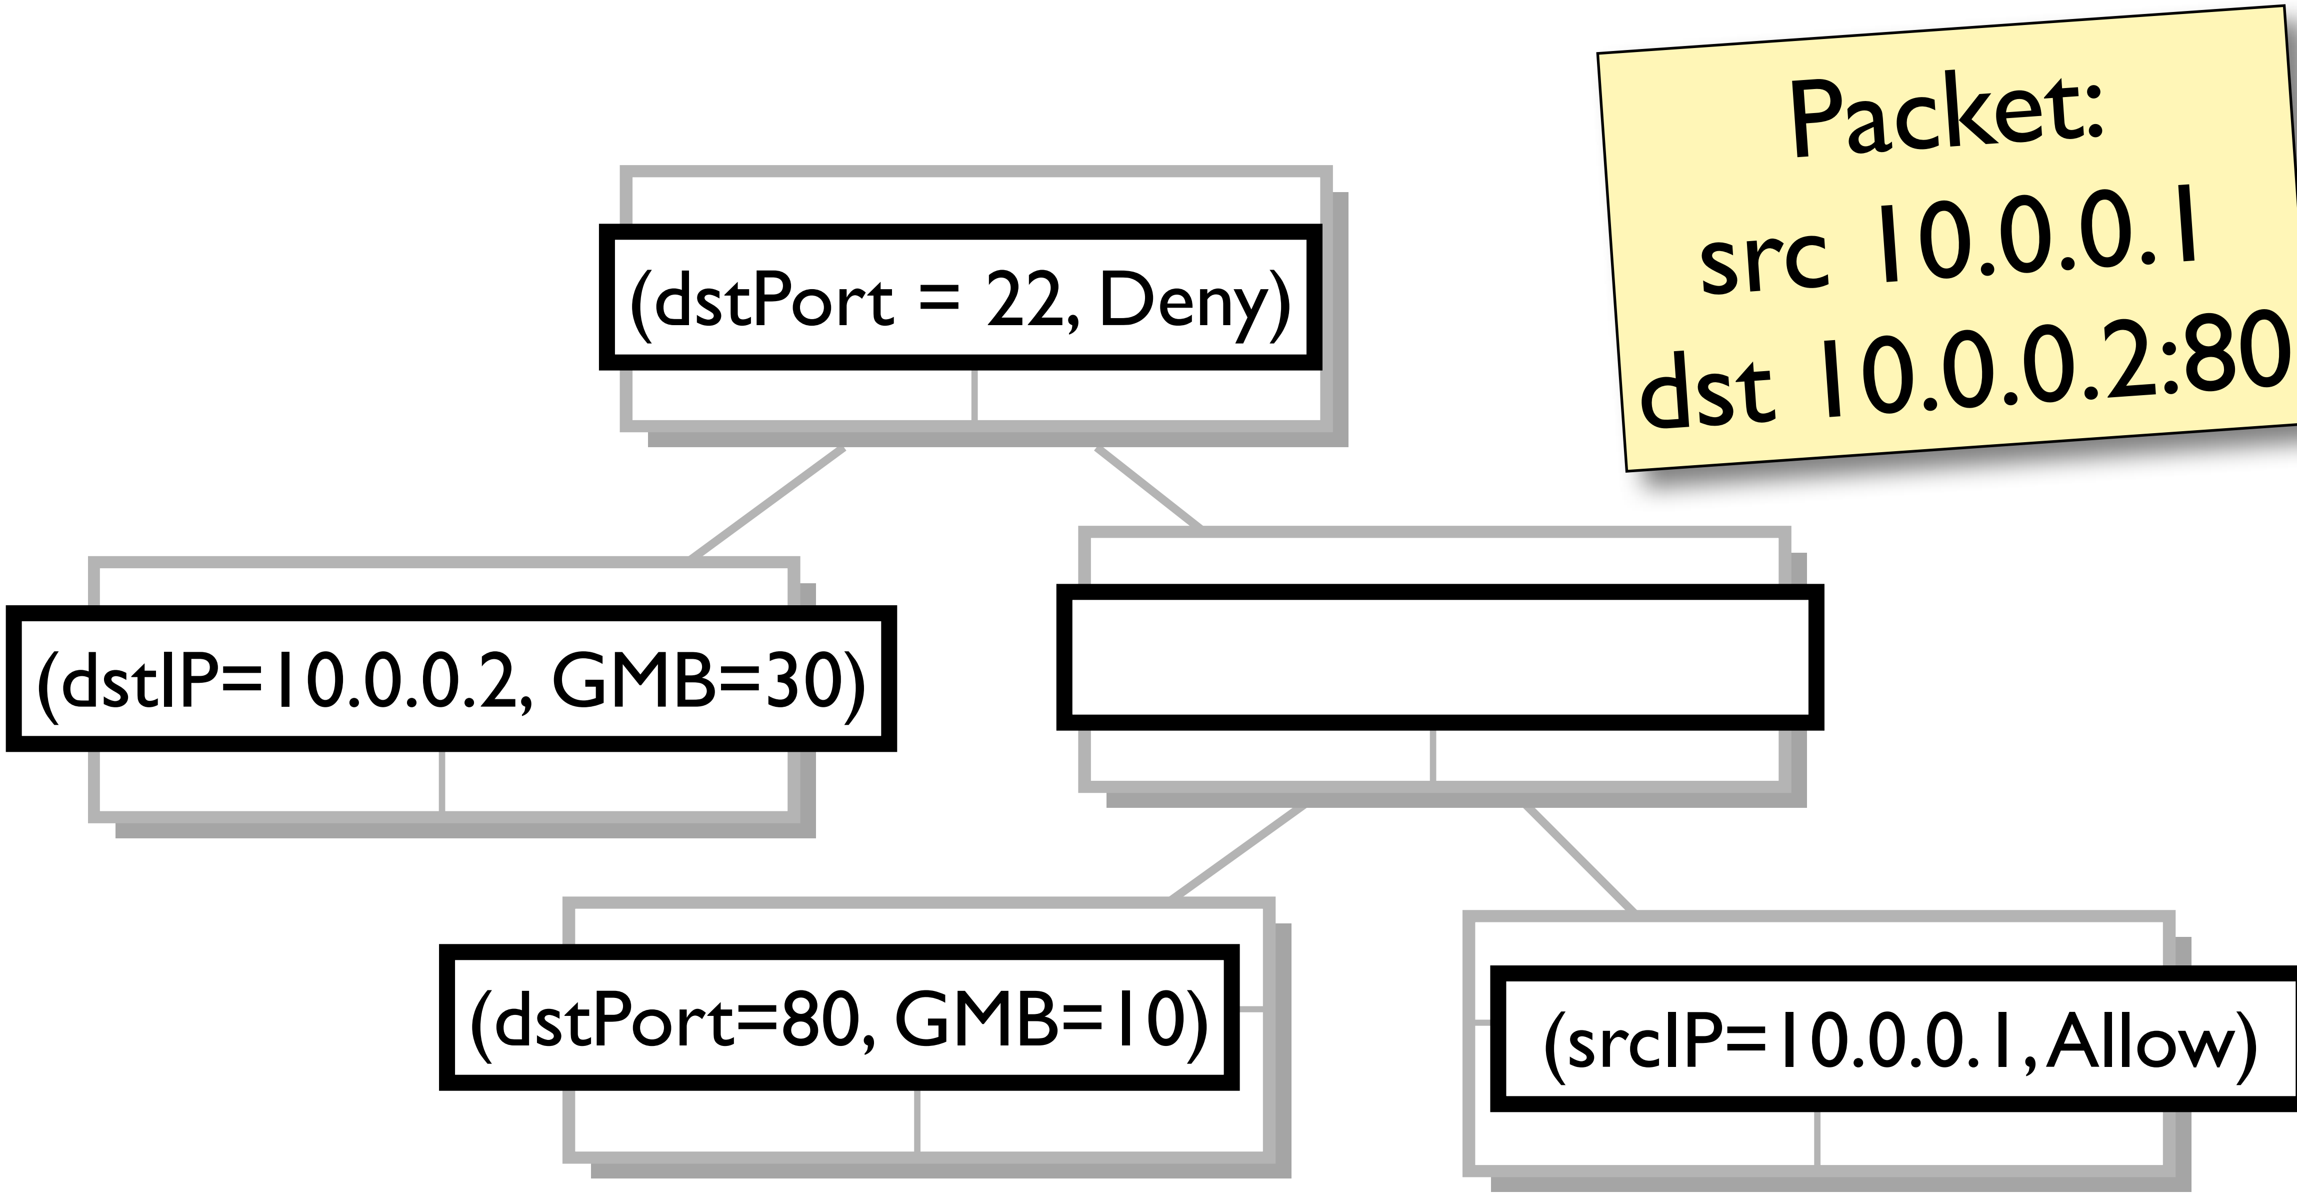

Hierarchy of Policies

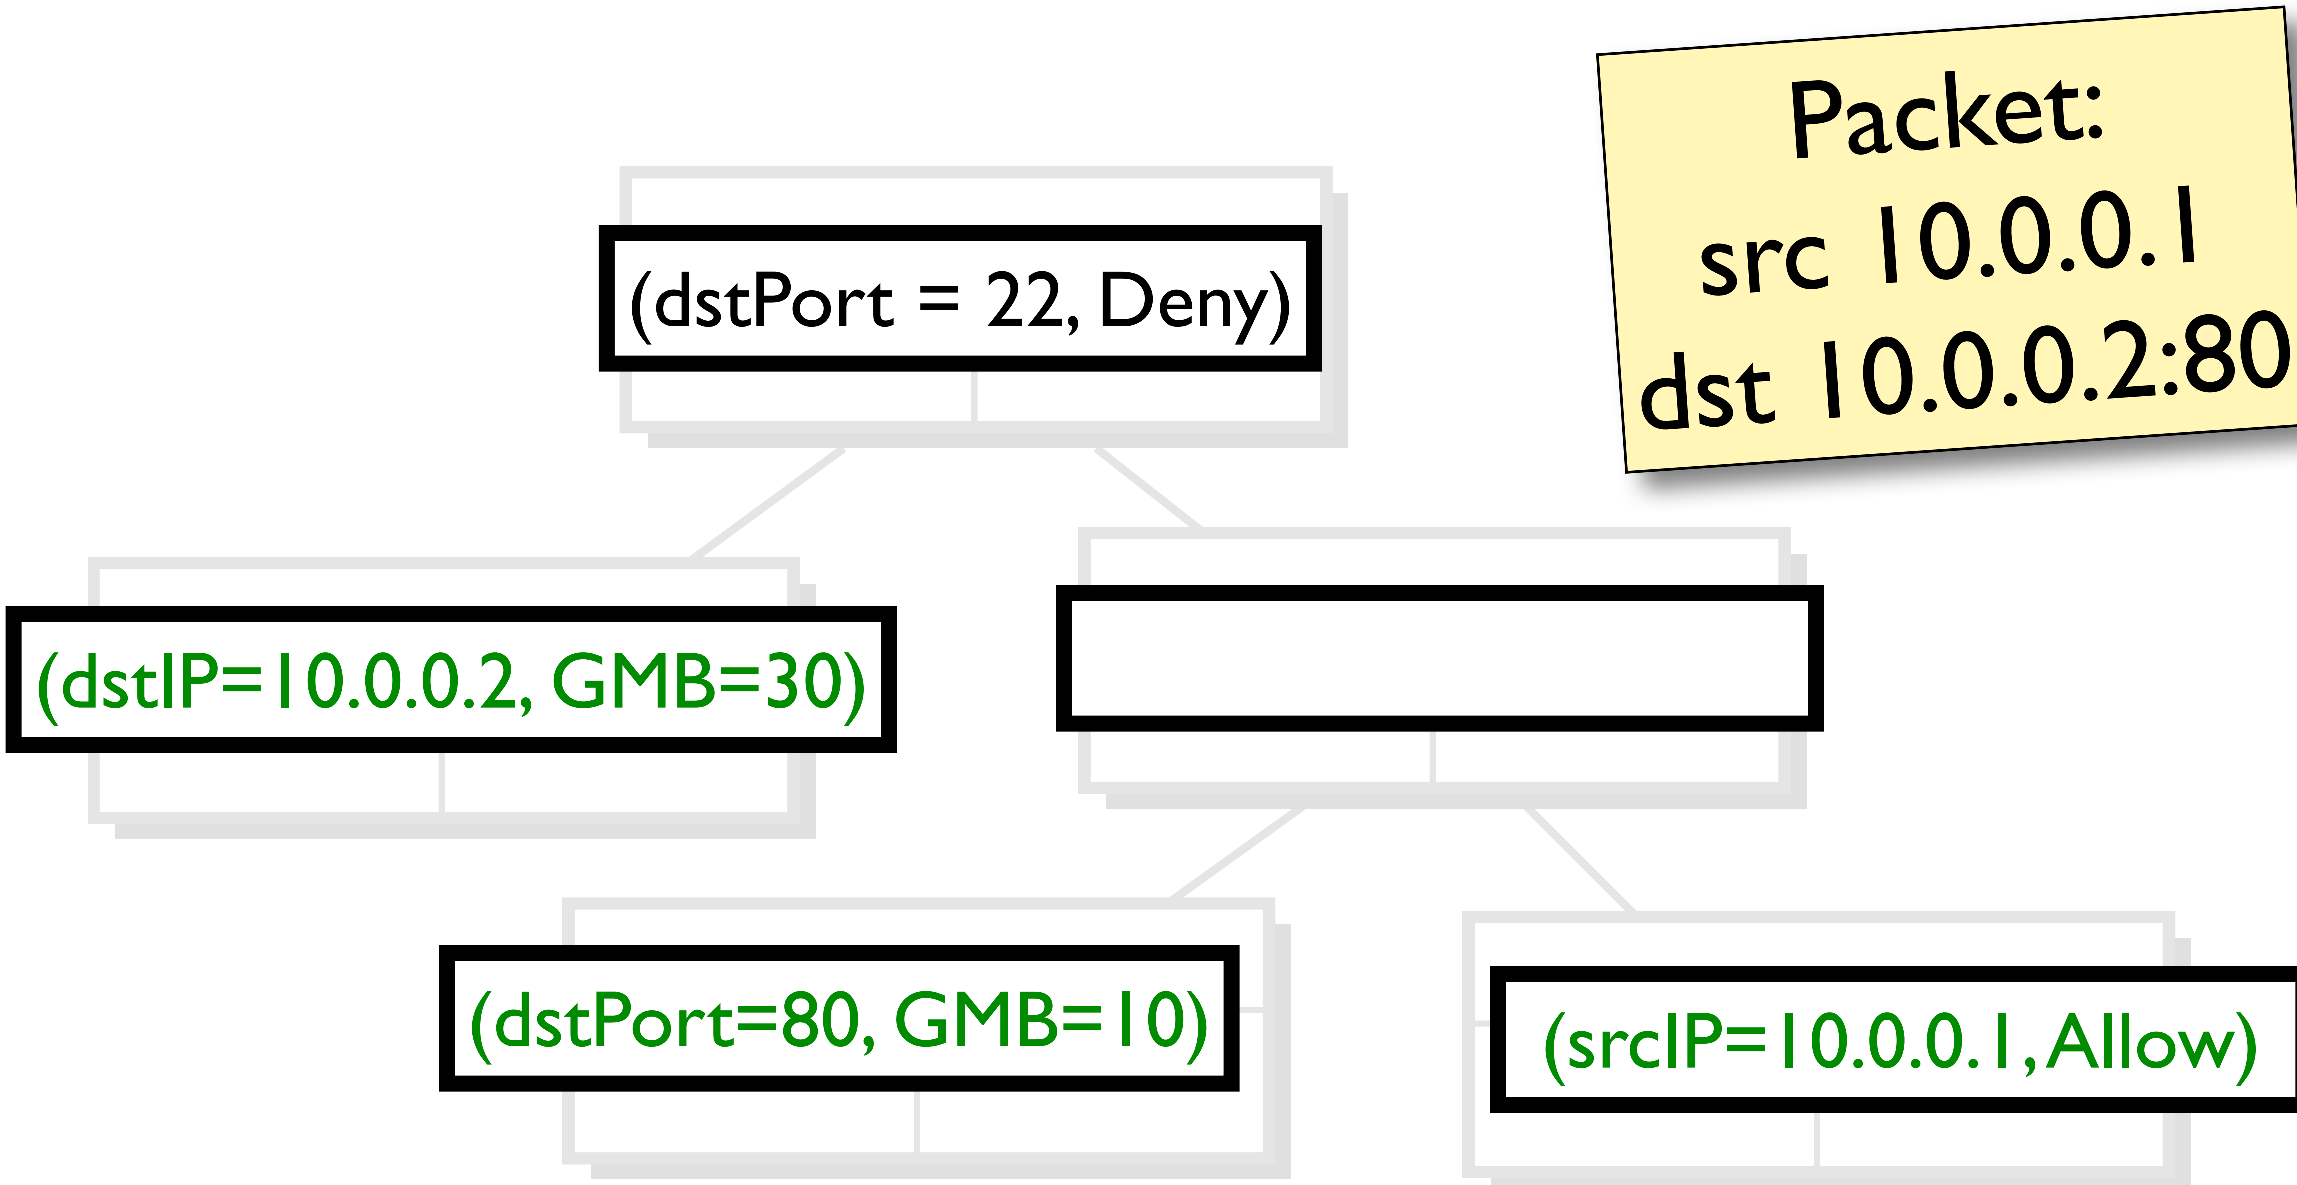

Hierarchical Flow Table

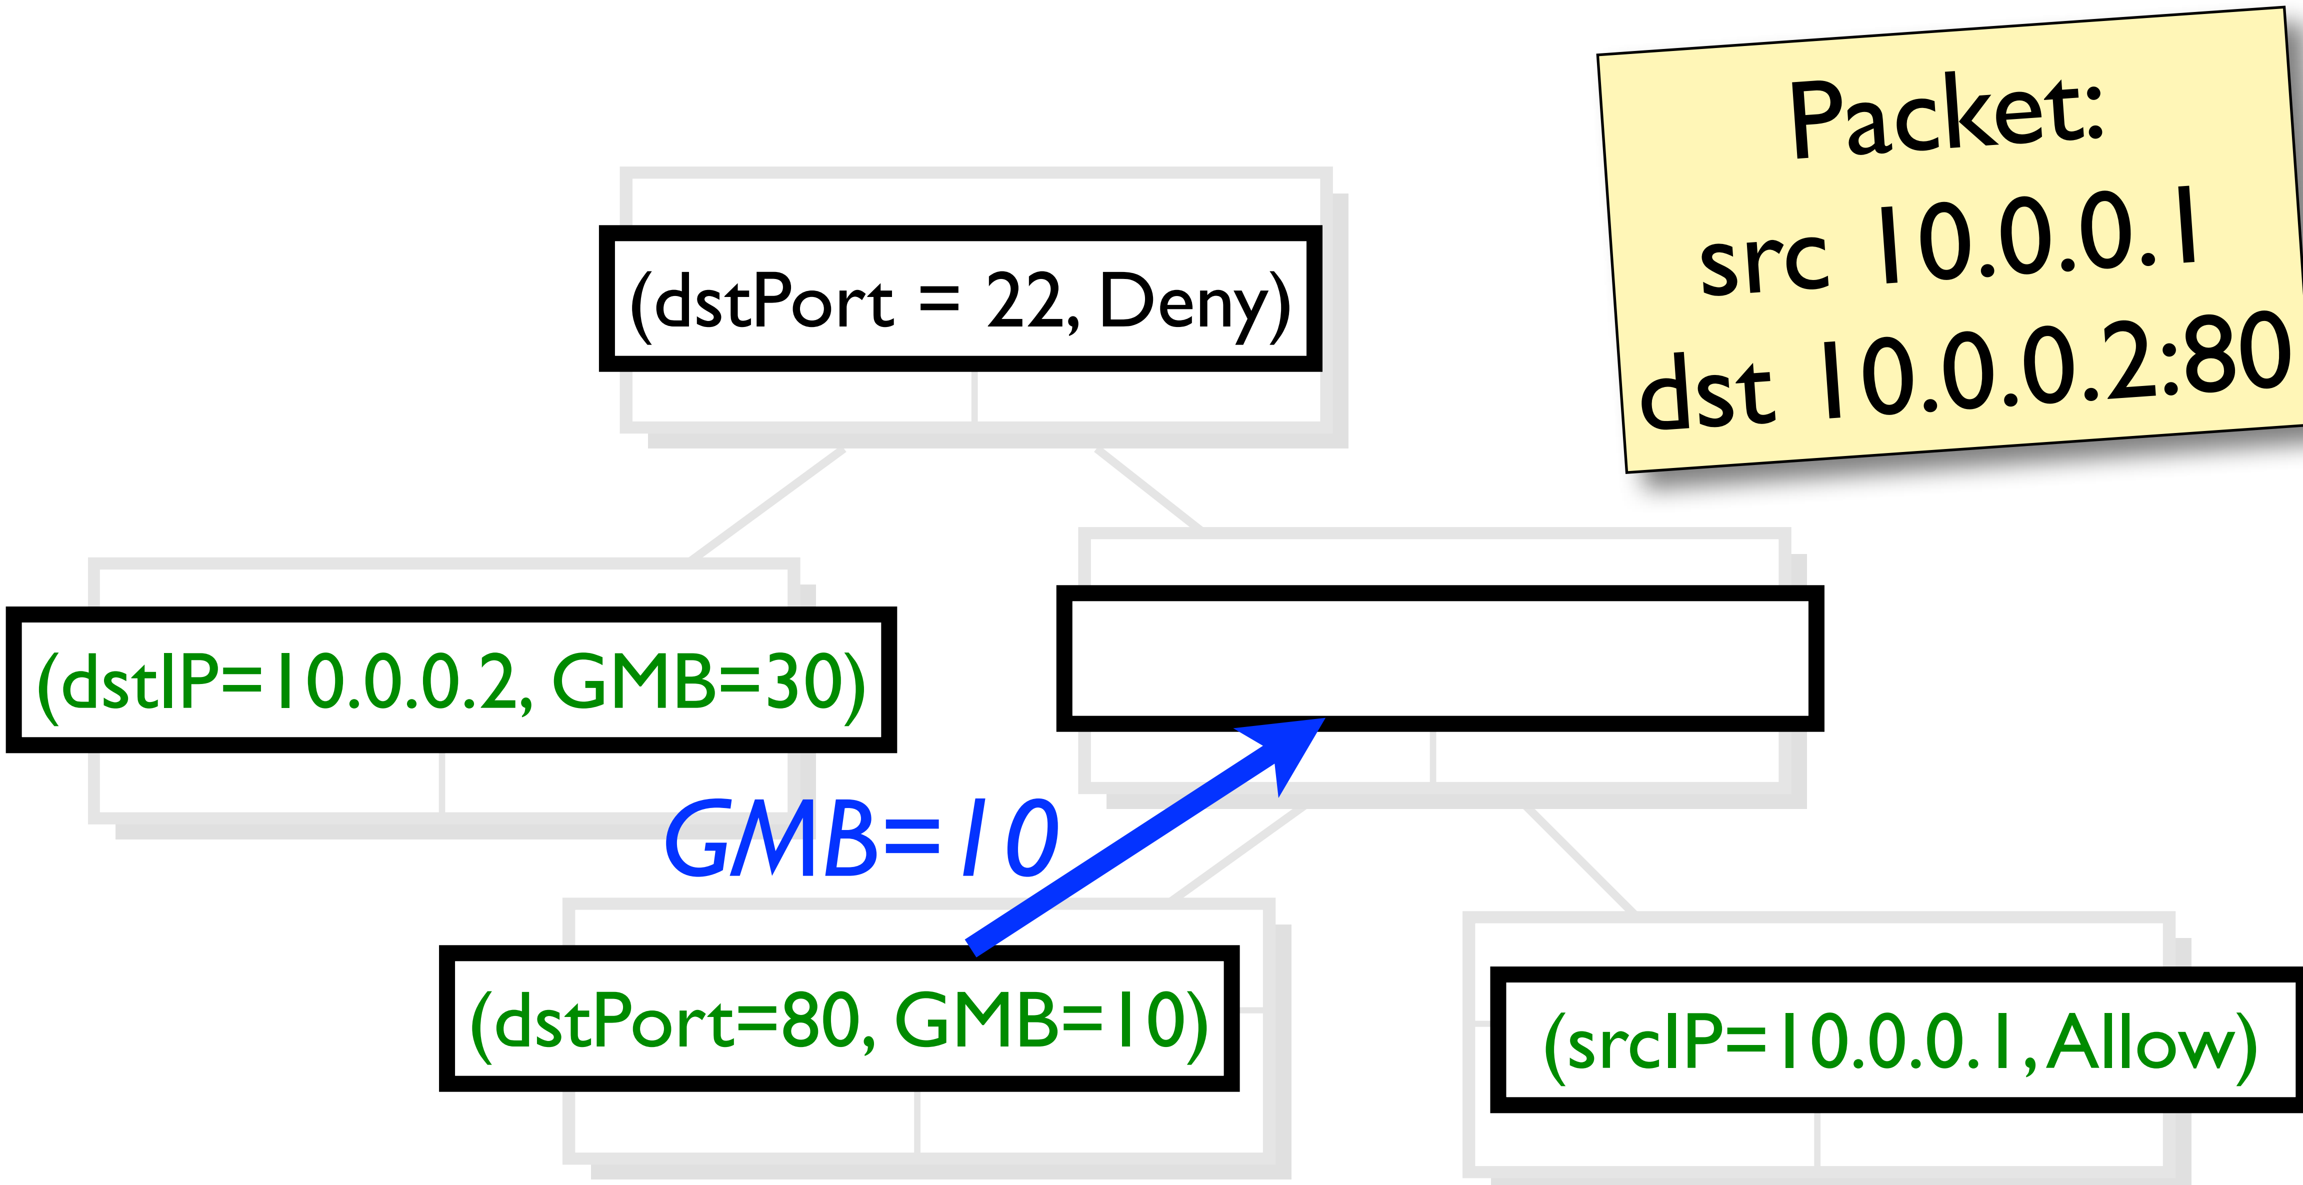

Hierarchical Flow Table

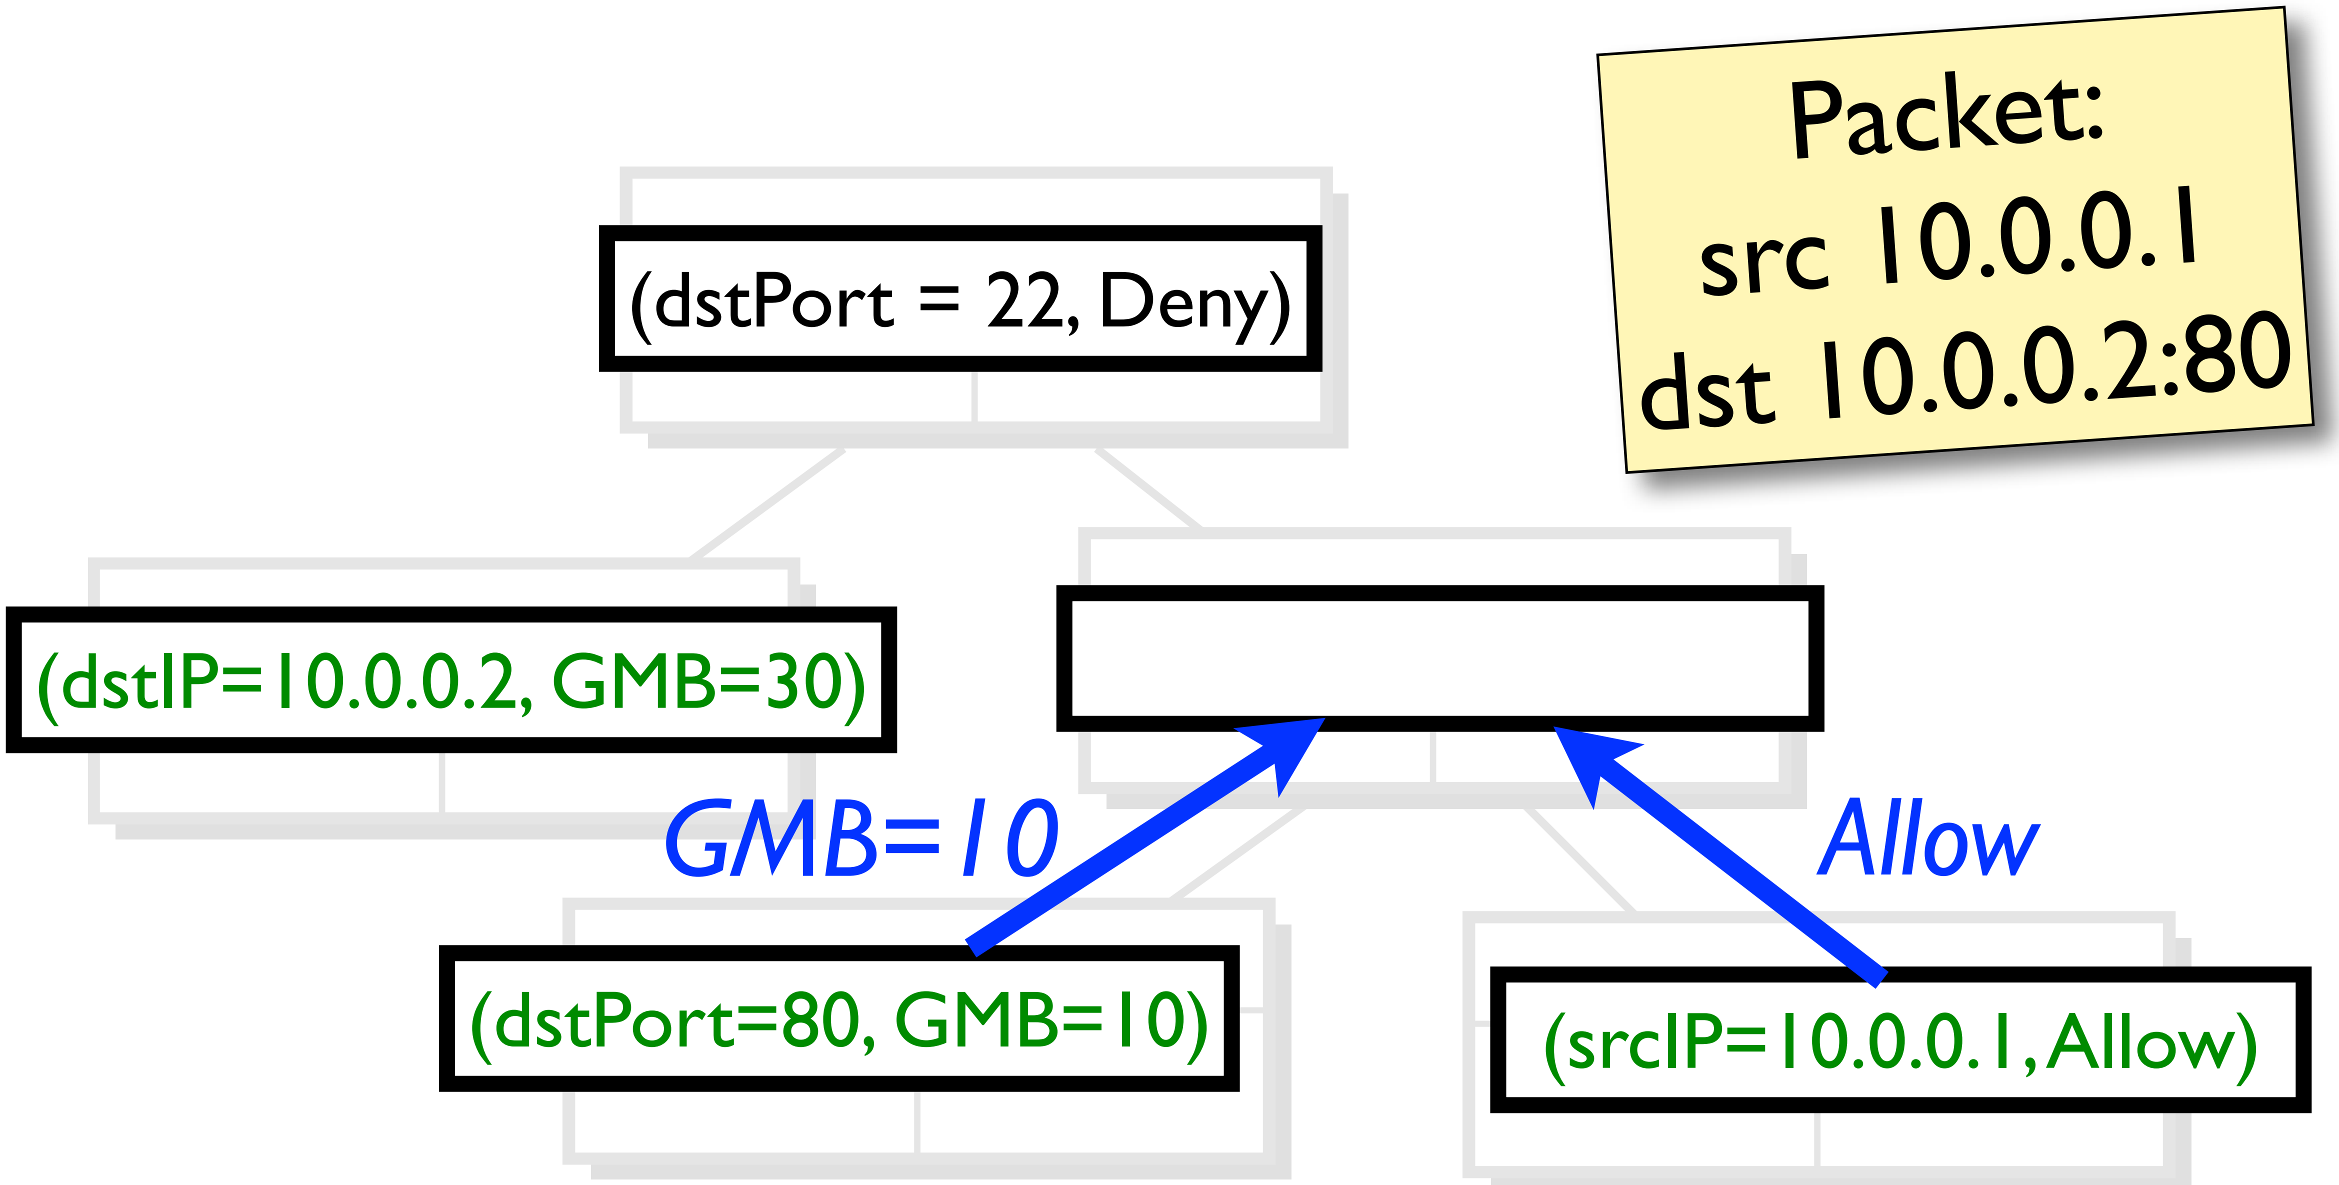

Hierarchical Flow Table

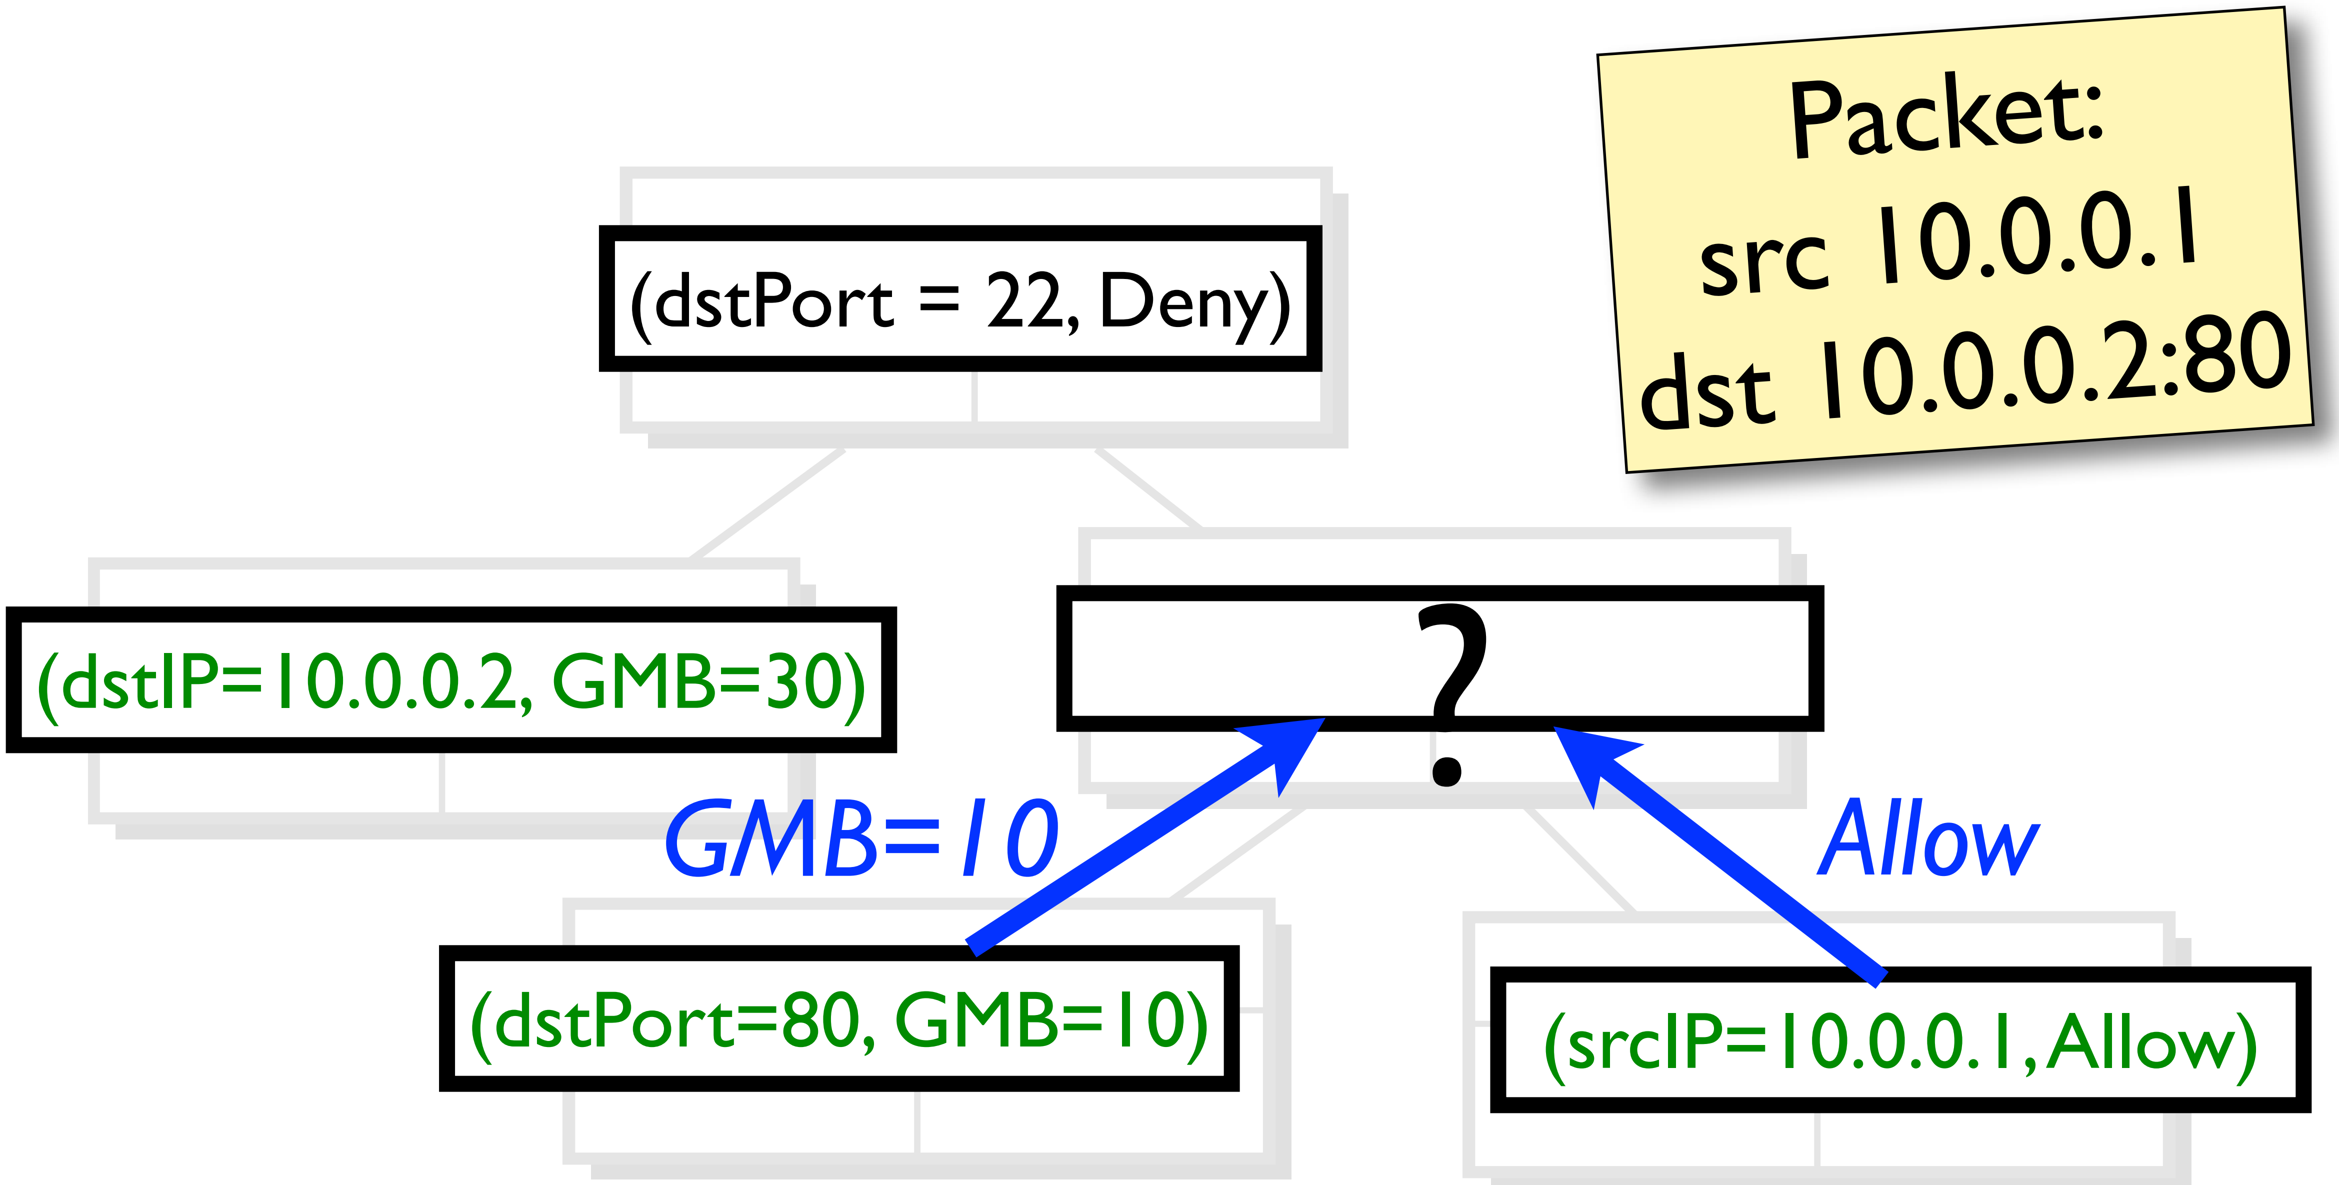

Hierarchical Flow Table

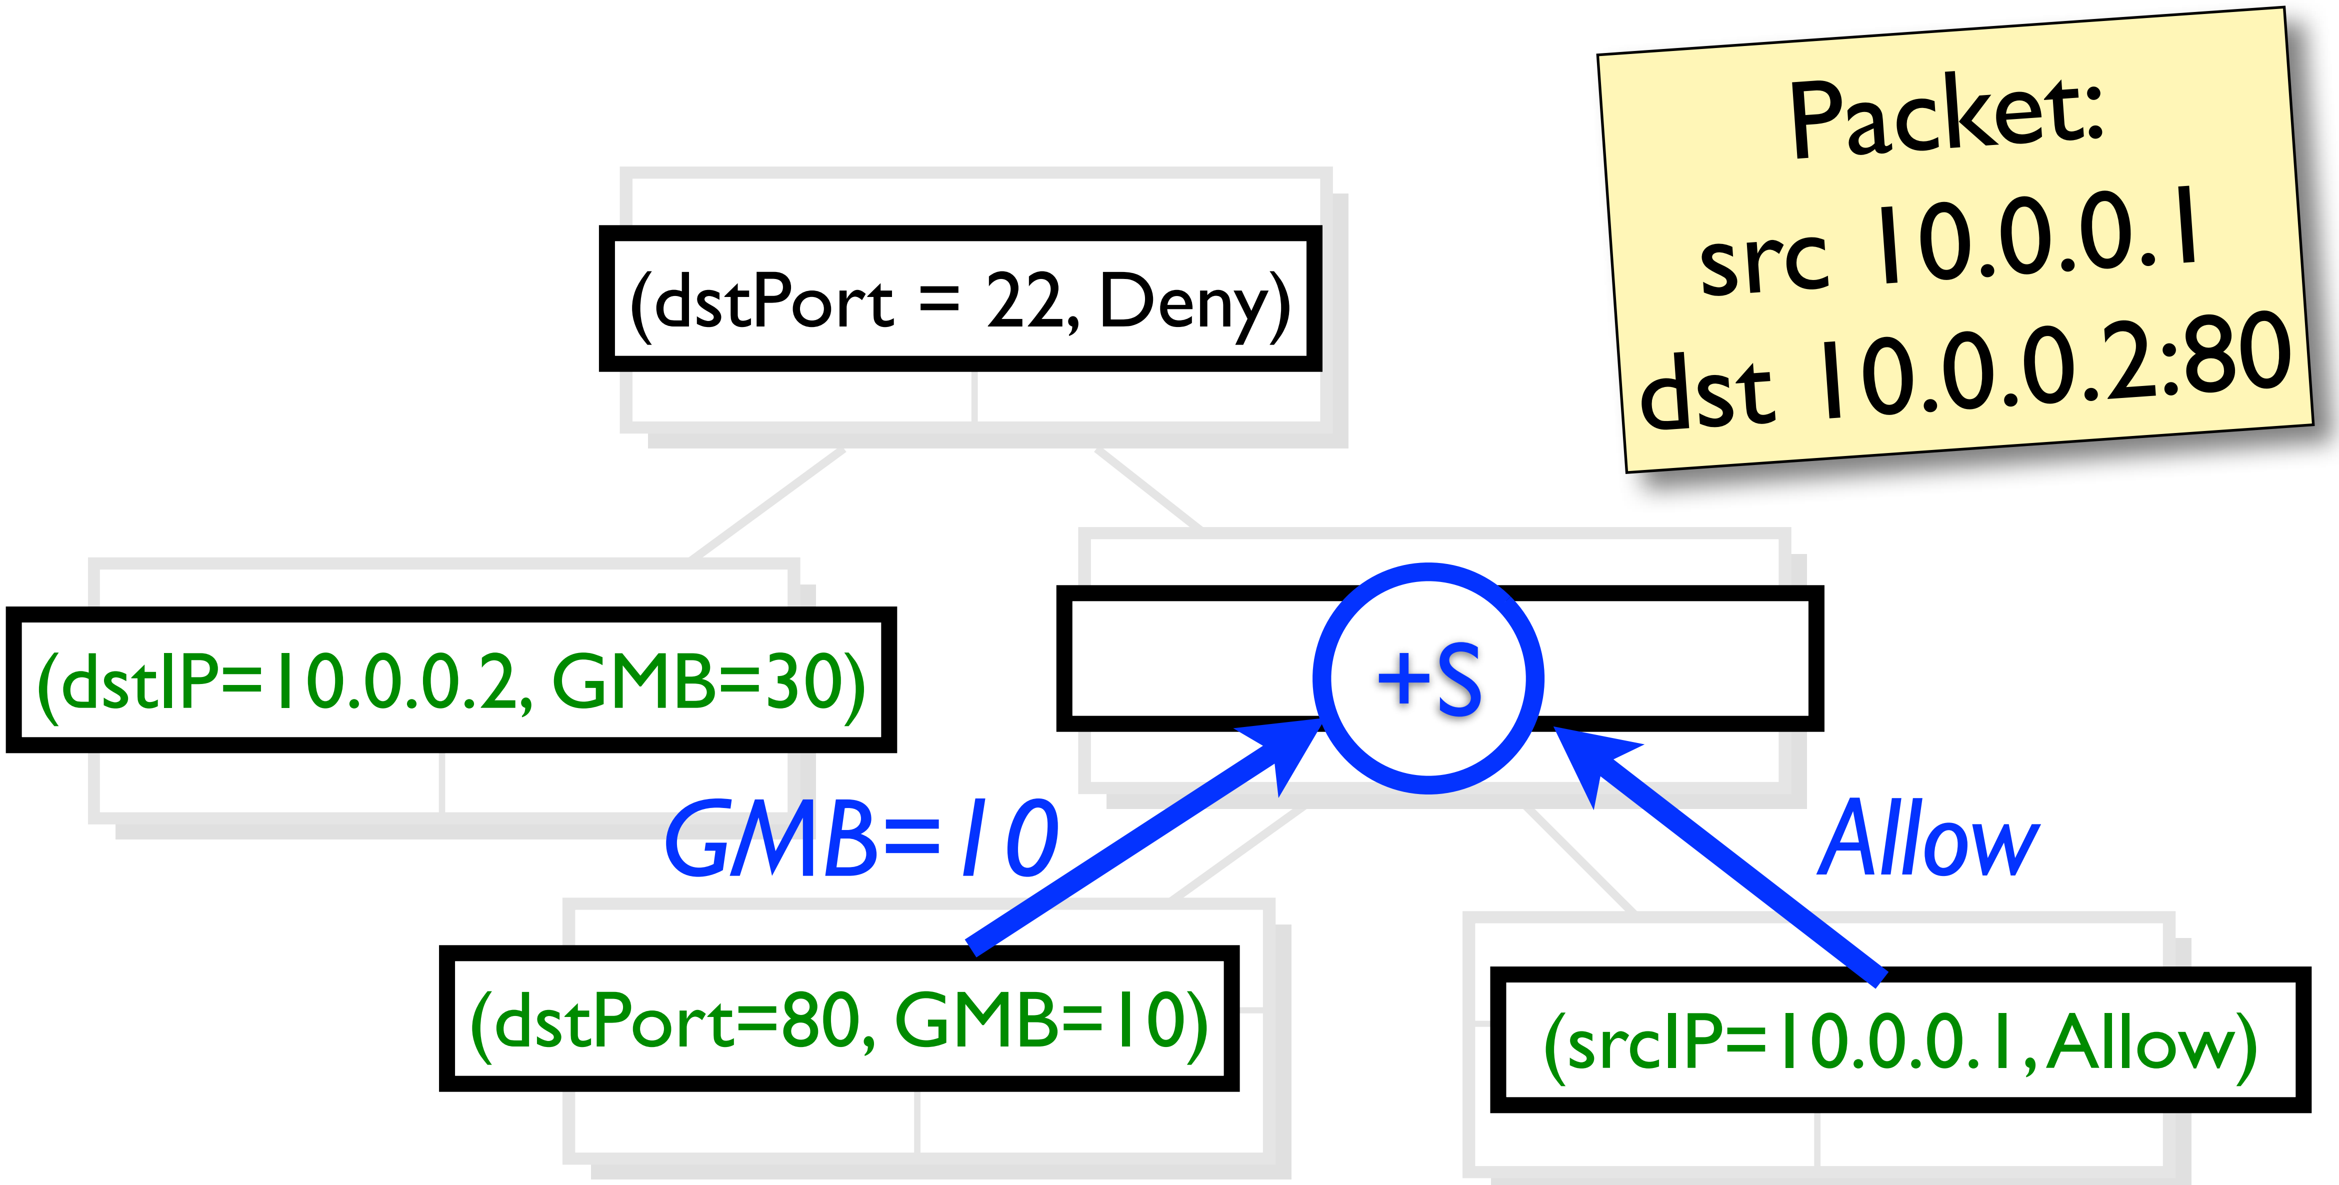

Hierarchical Flow Table

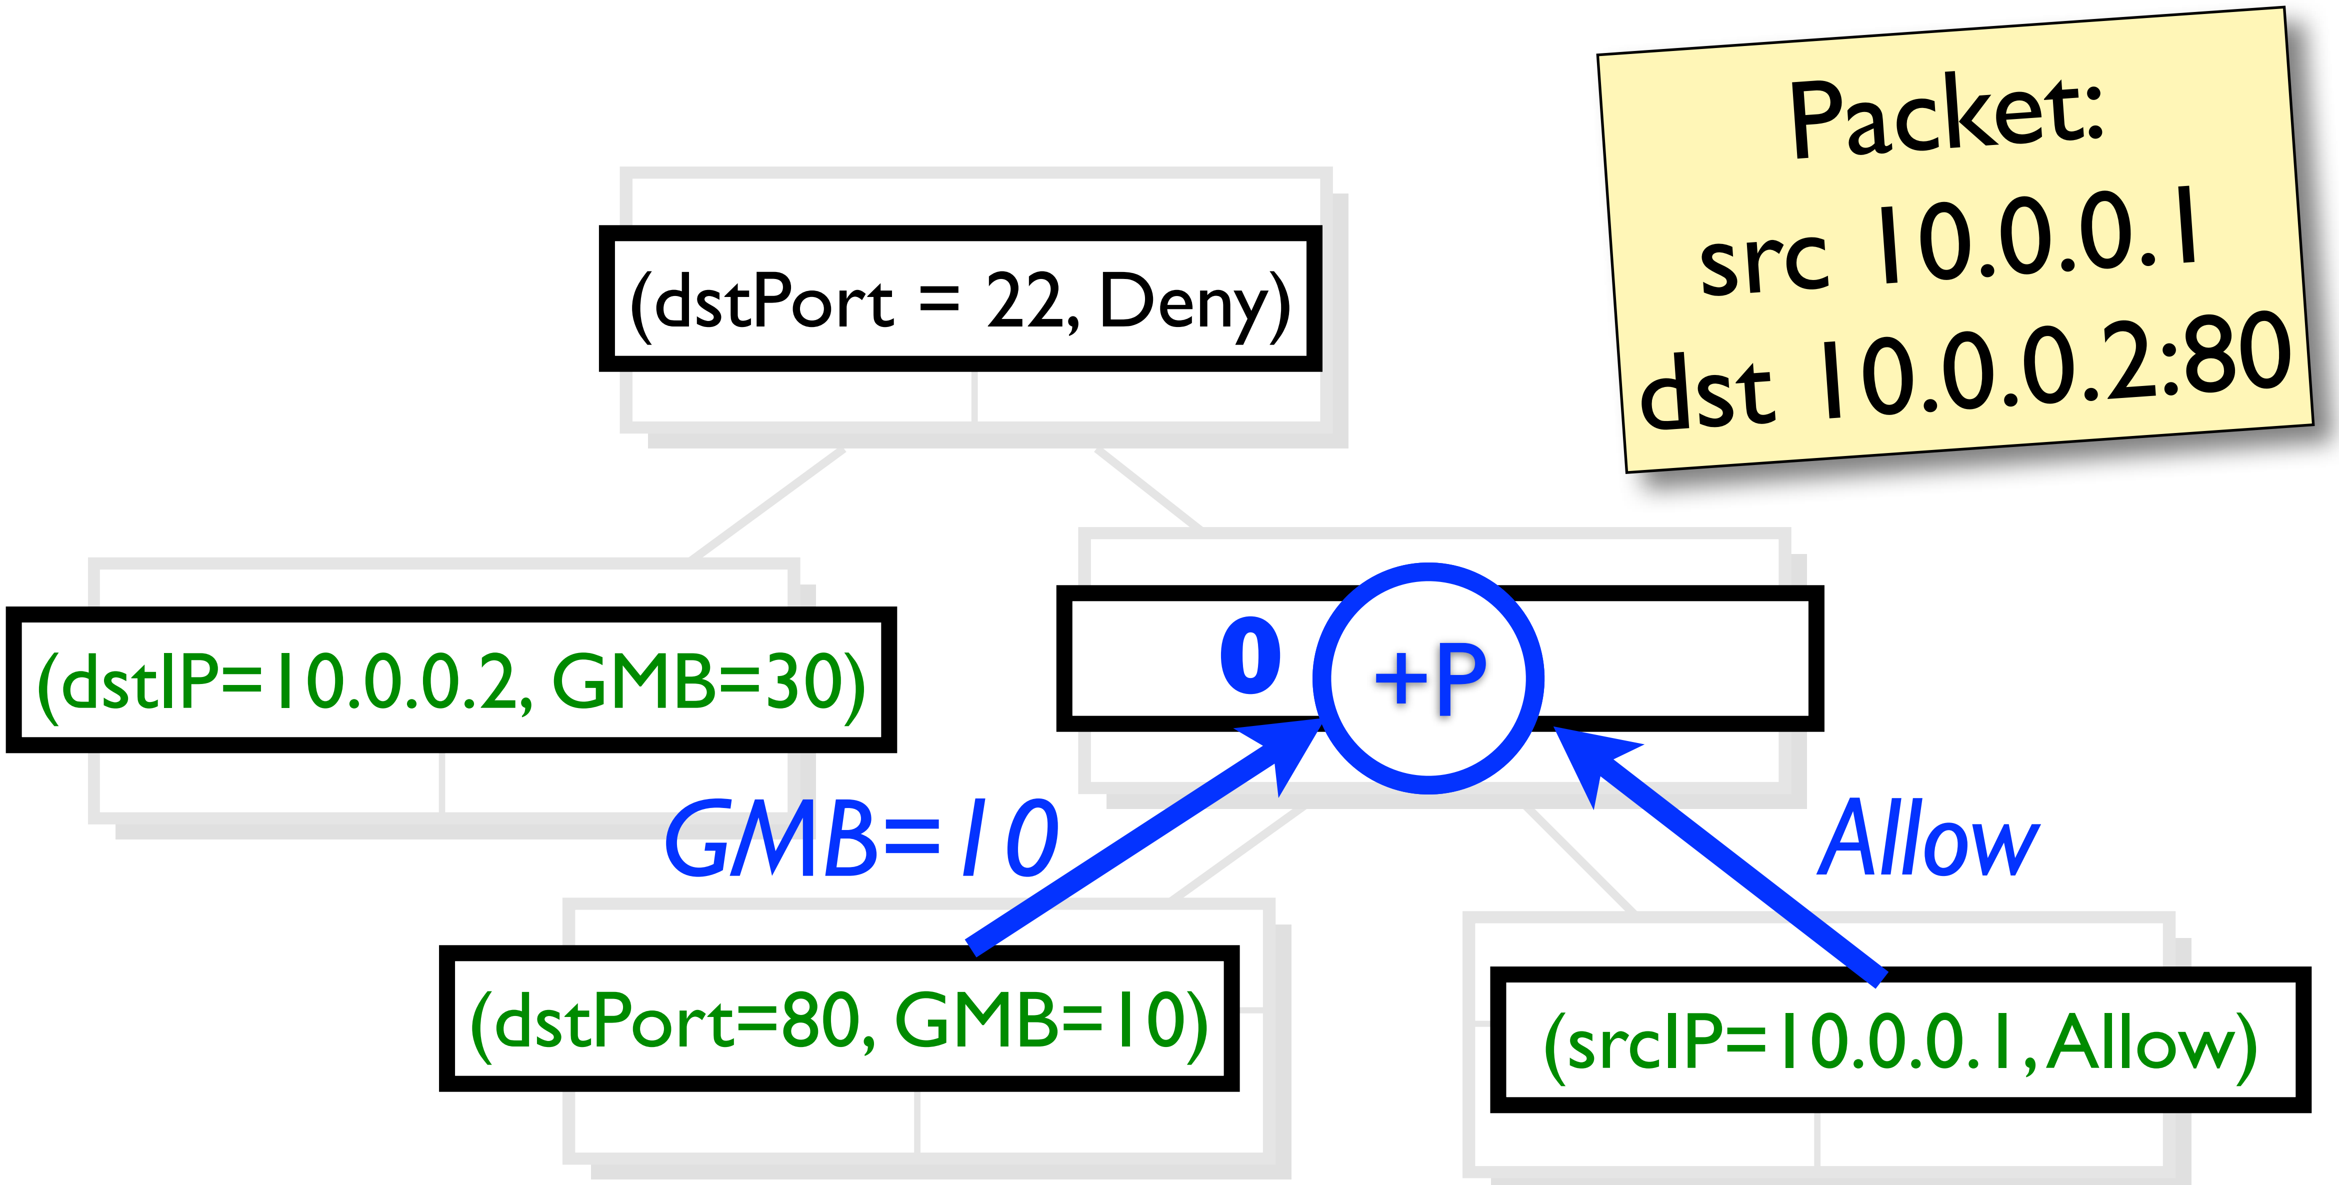

Hierarchical Flow Table

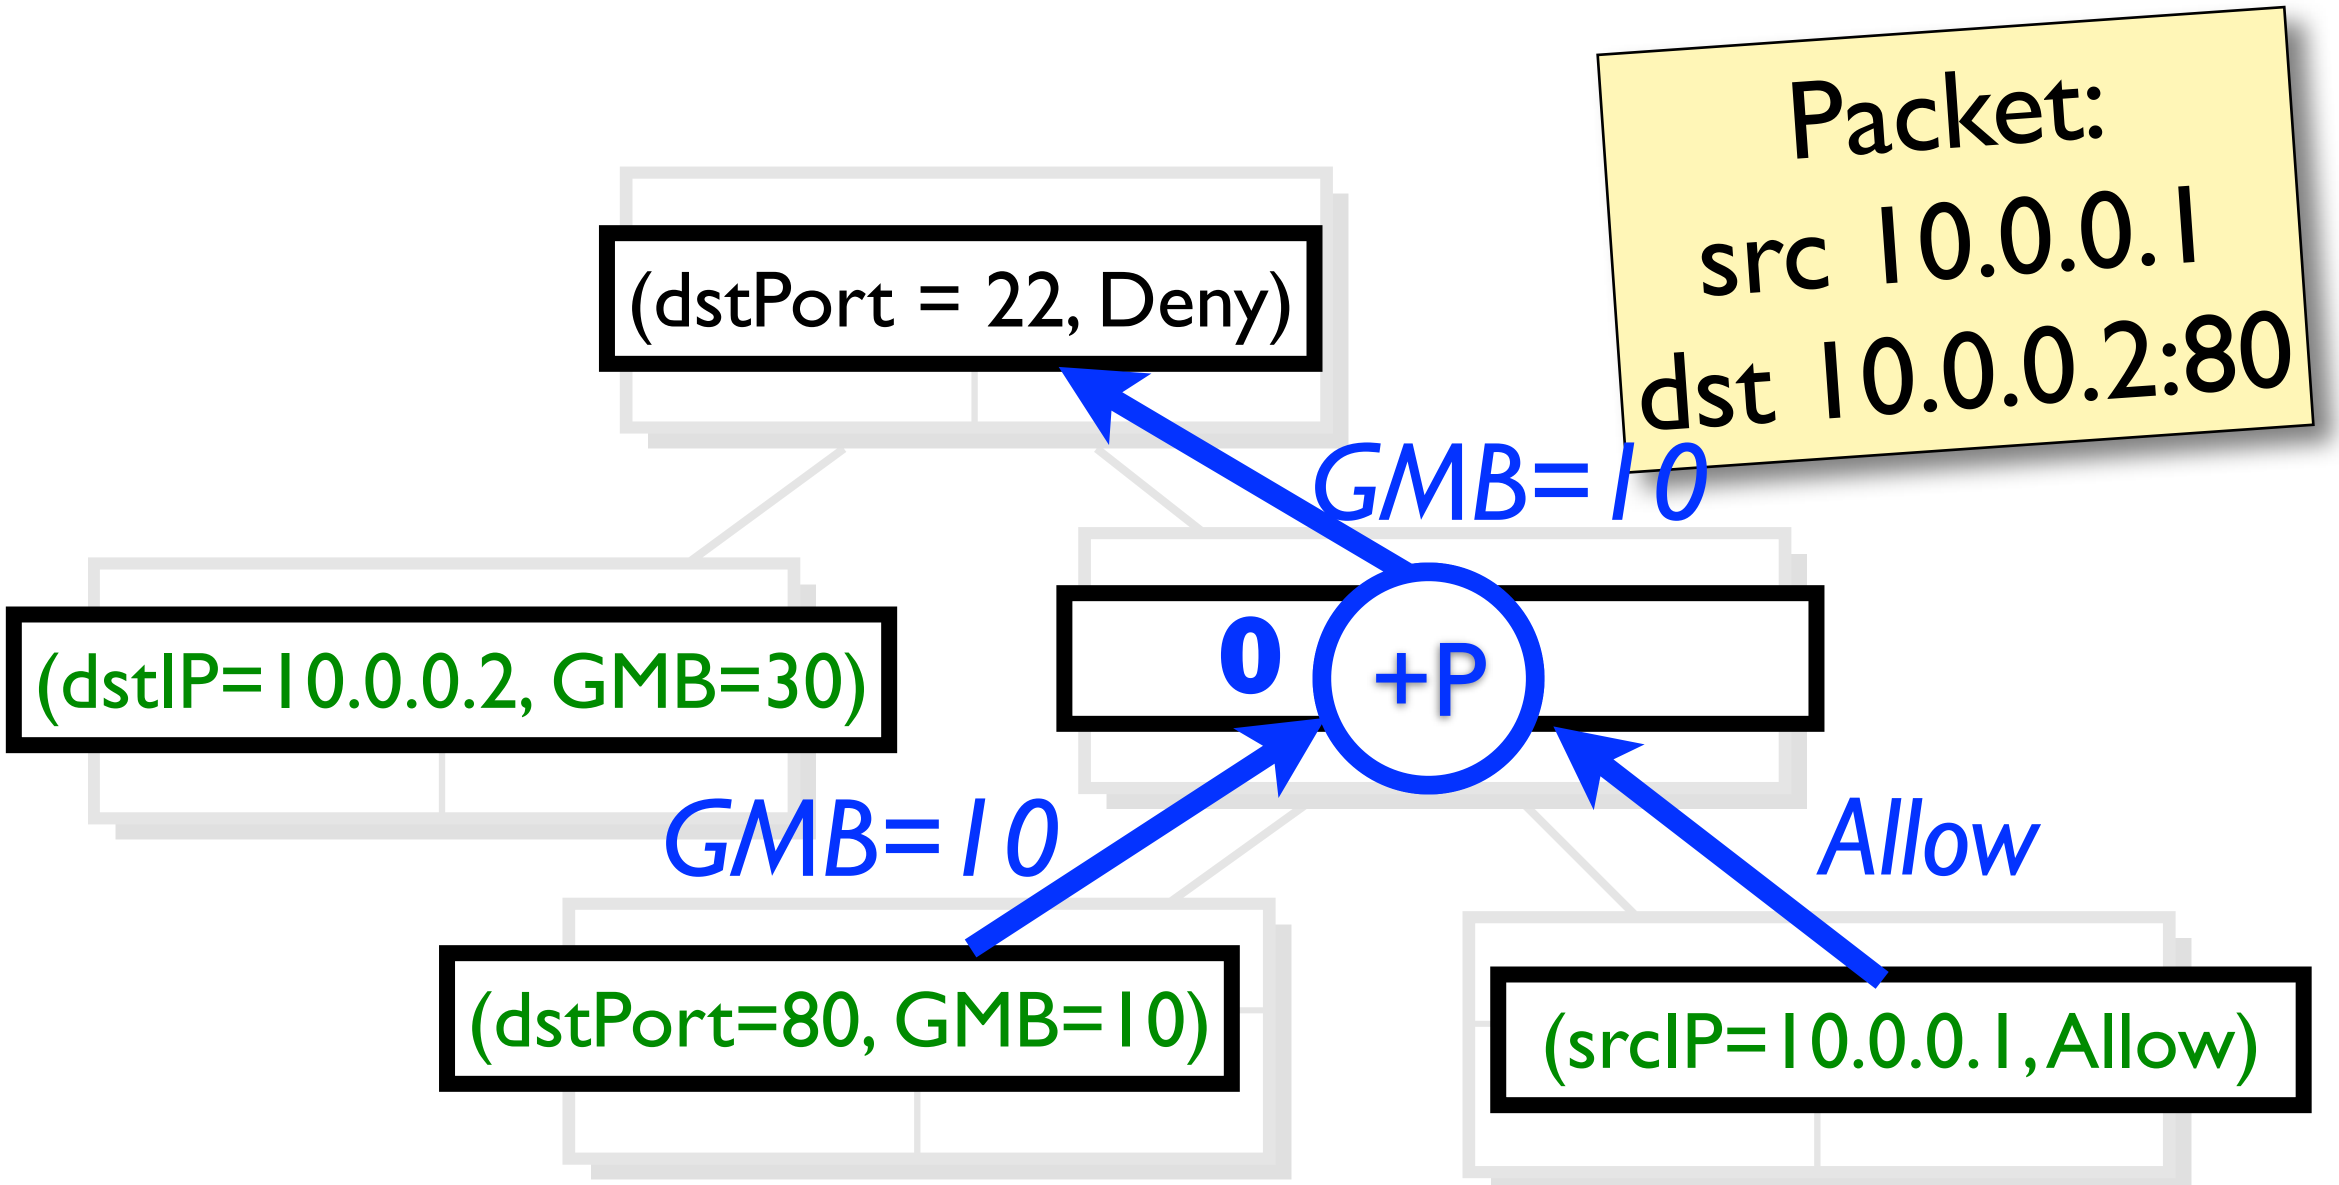

Hierarchical Flow Table

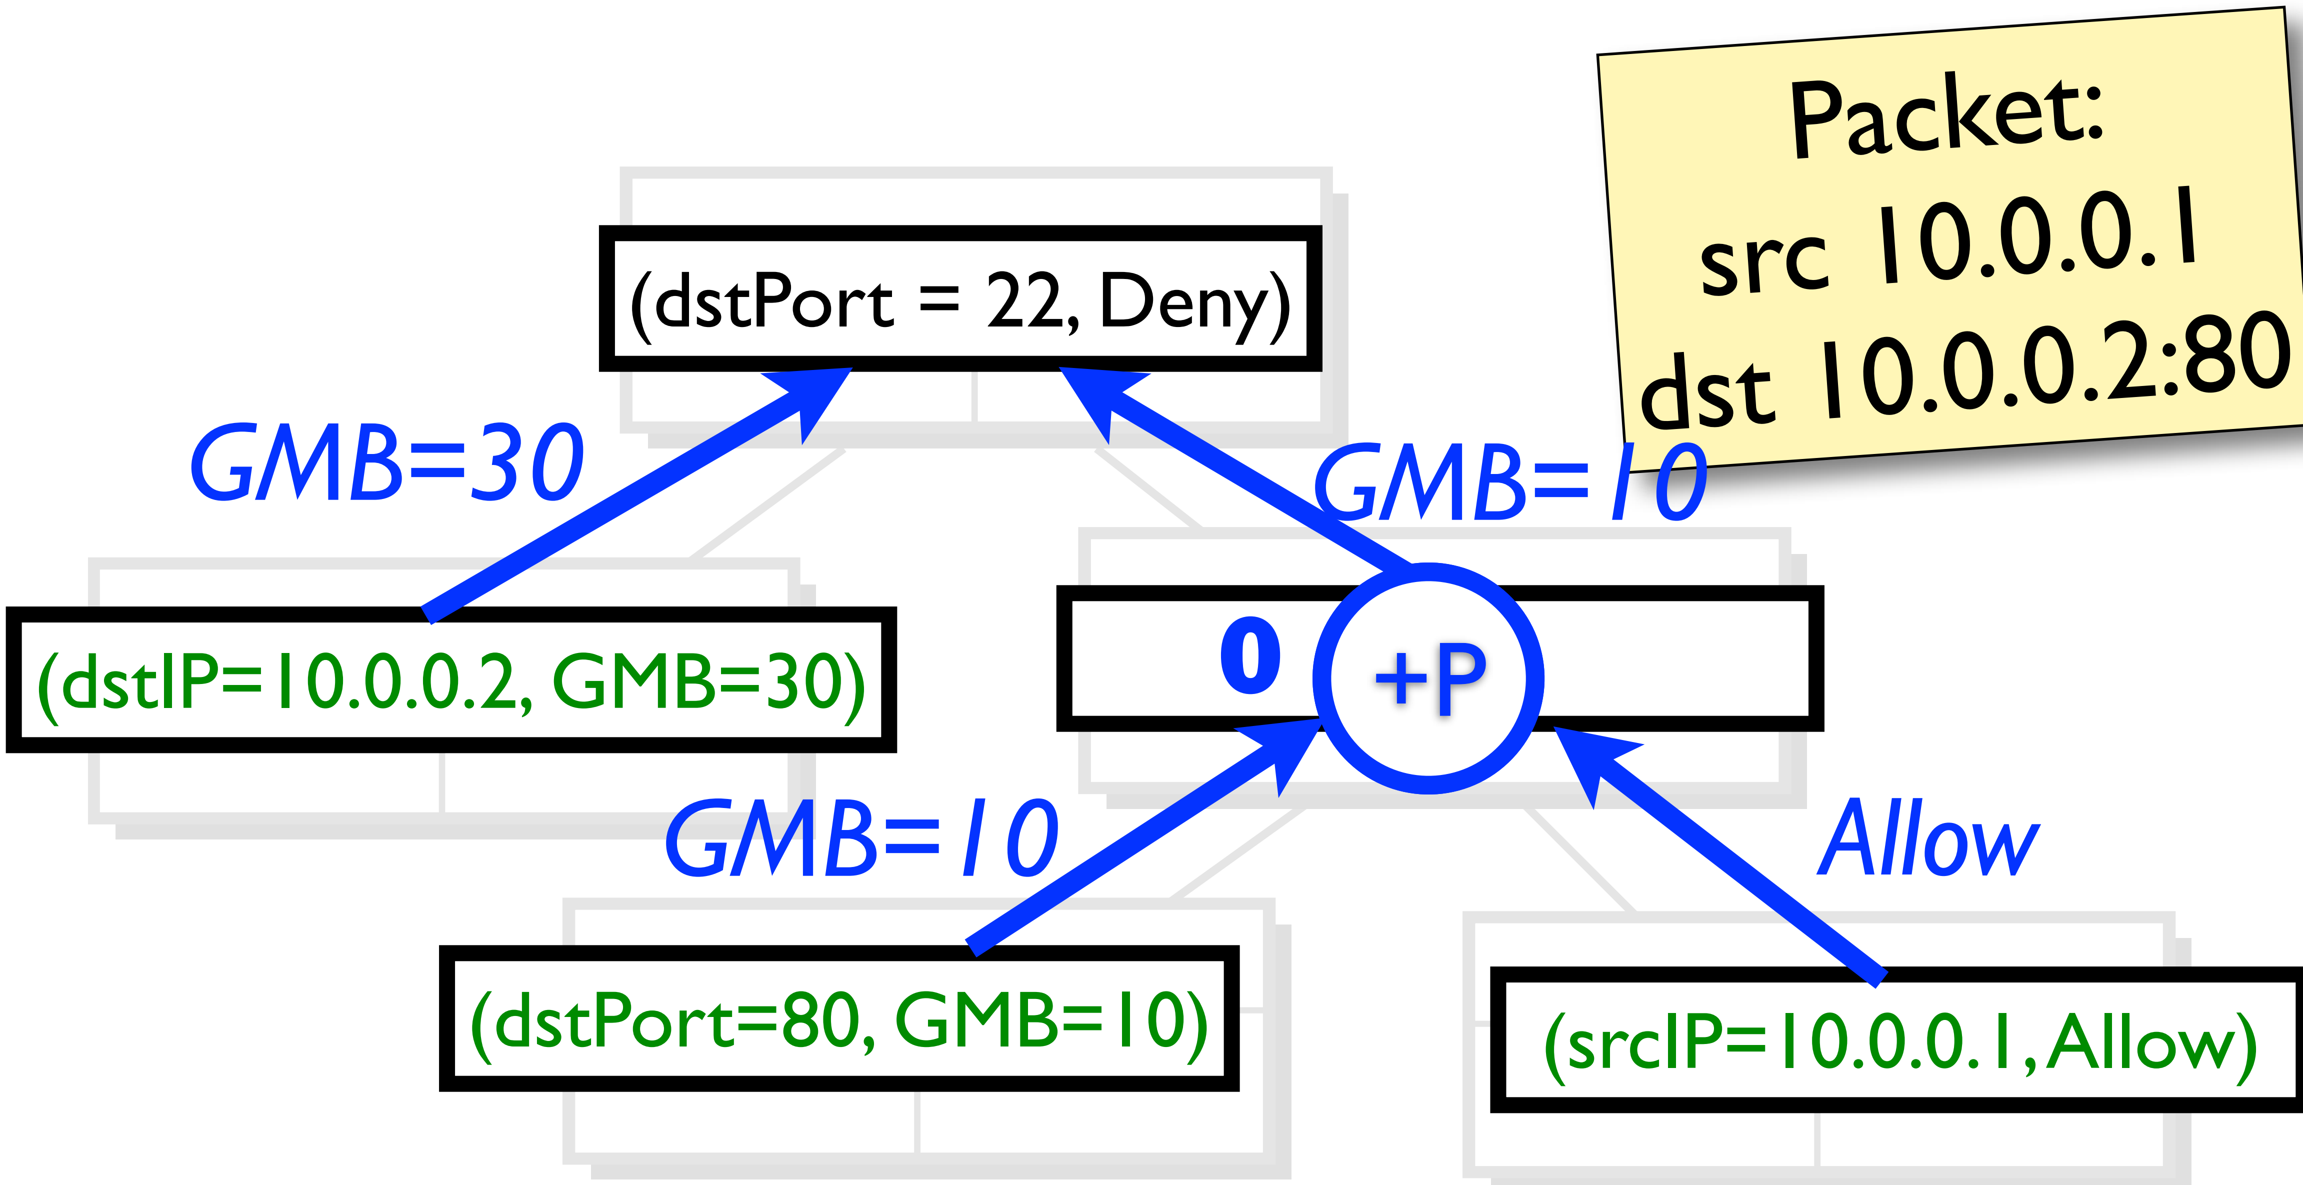

Hierarchical Flow Table

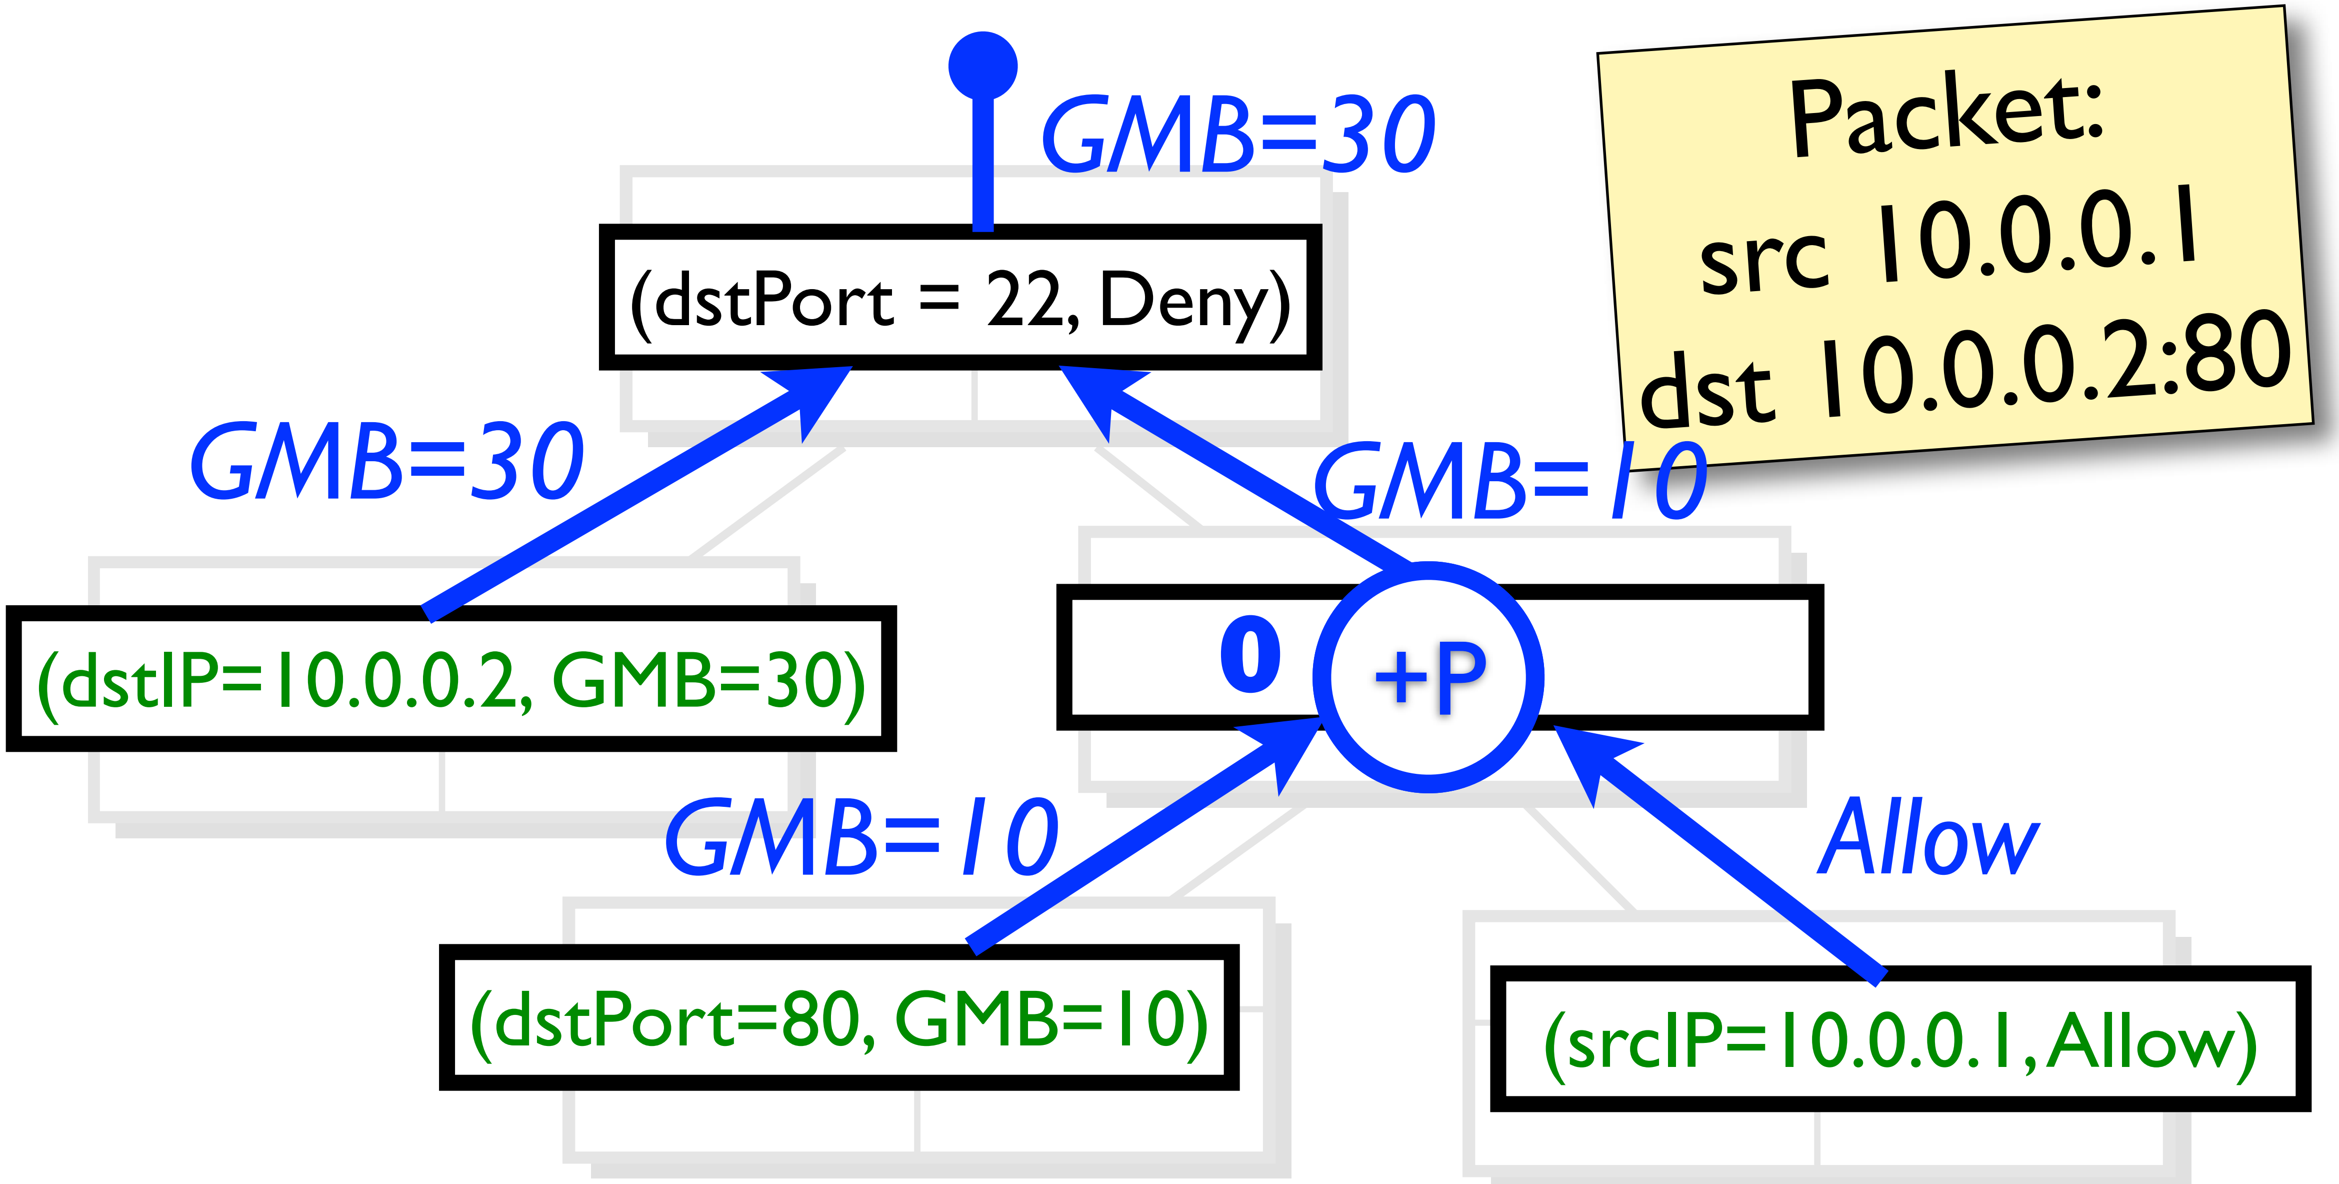

Hierarchical Flow Table

Requirements

*Associative, **0**-identity*

Commutative

+D *In node*

+S *Sibling*

+P *Parent-Sibling*

HFT operators

Requirements		In PANE
*Associative, **0**-identity* — *Commutative*	+D *In node* +S *Sibling*	D and S identical. Deny overrides Allow. GMB combines as **max**
	+P *Parent-Sibling*	Child overrides Parent for Access Control GMB combines as **max**

HFT operators

There’s a slight problem...

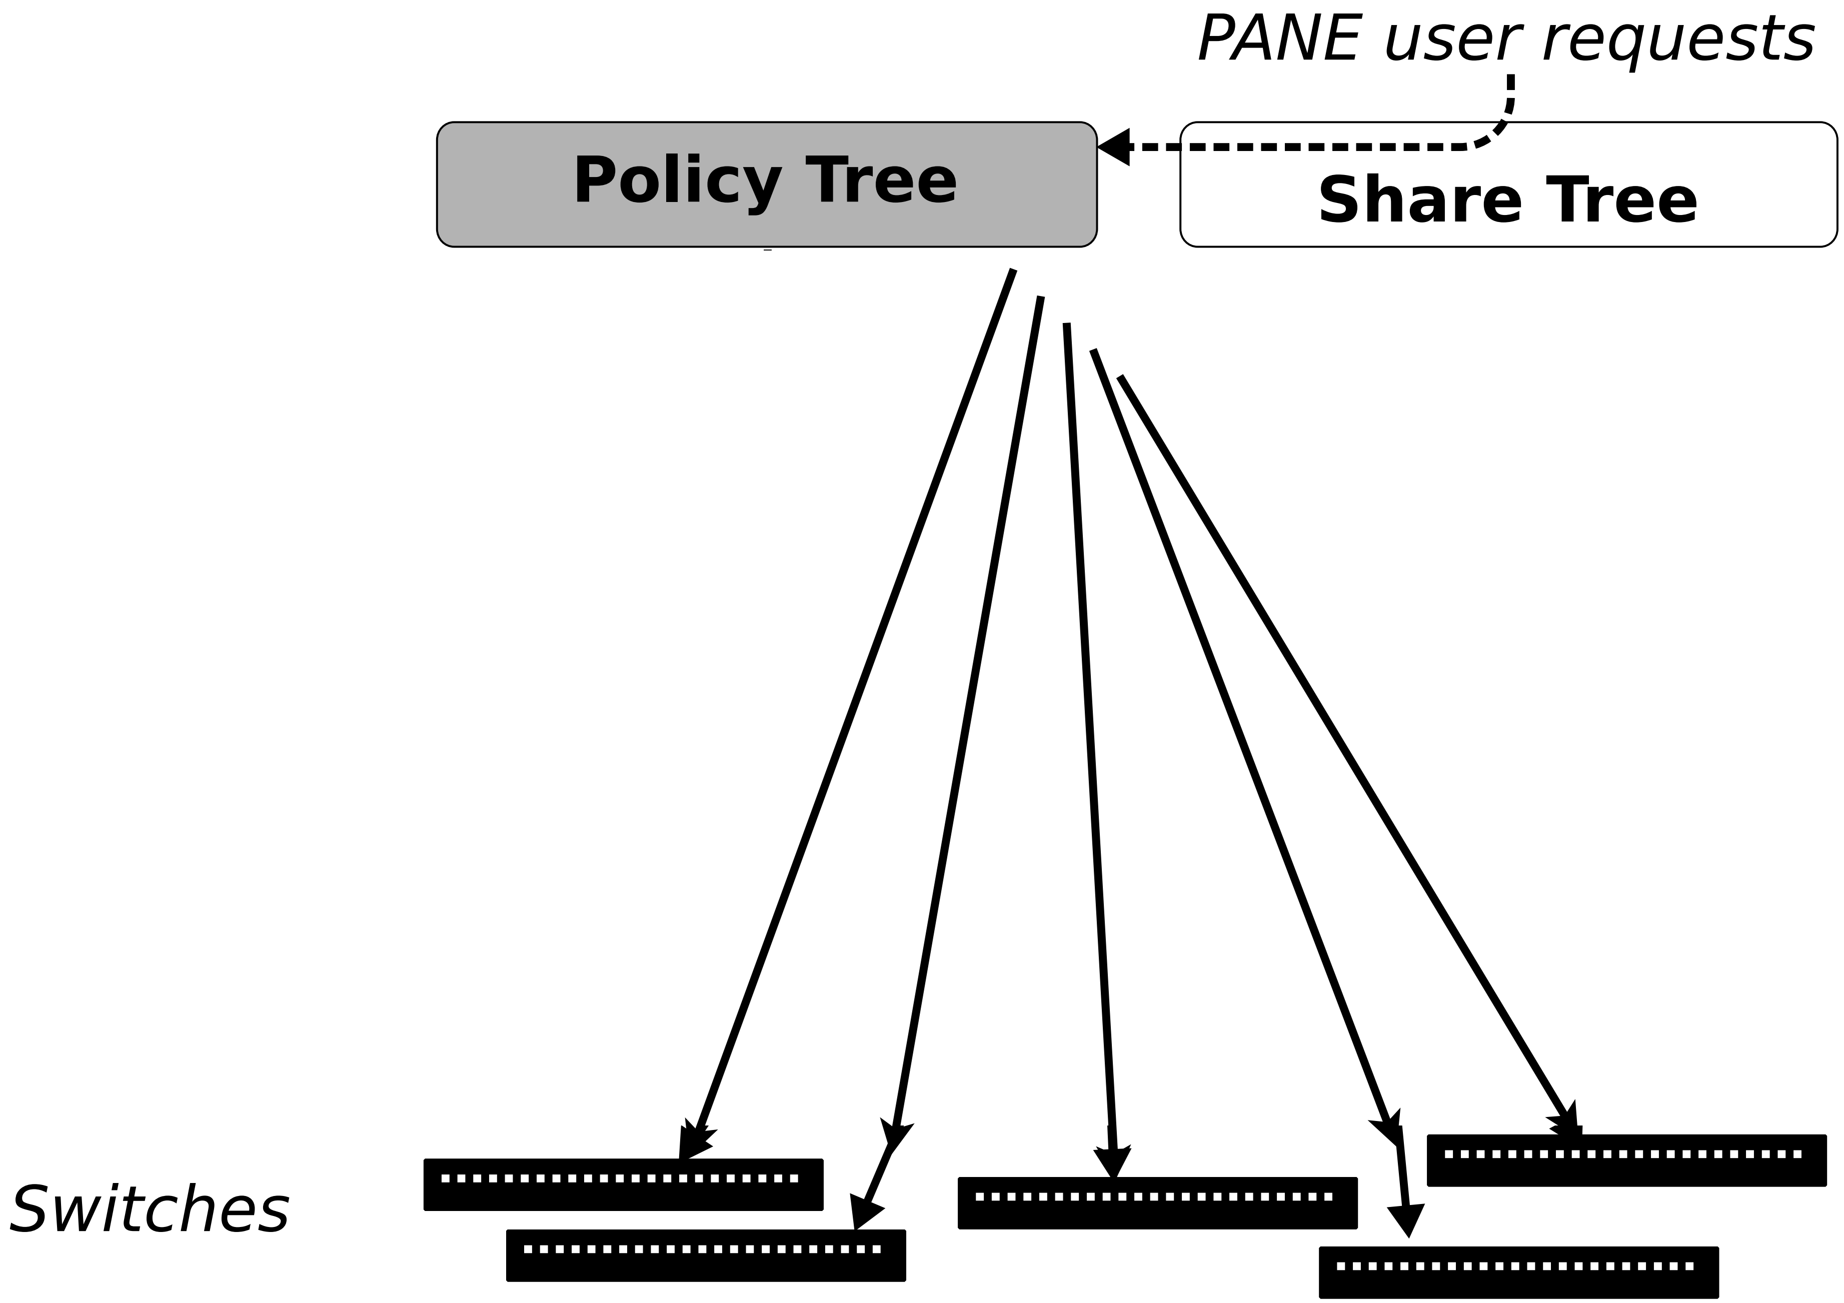
PANE user requests
Policy Tree
Share Tree
Switches

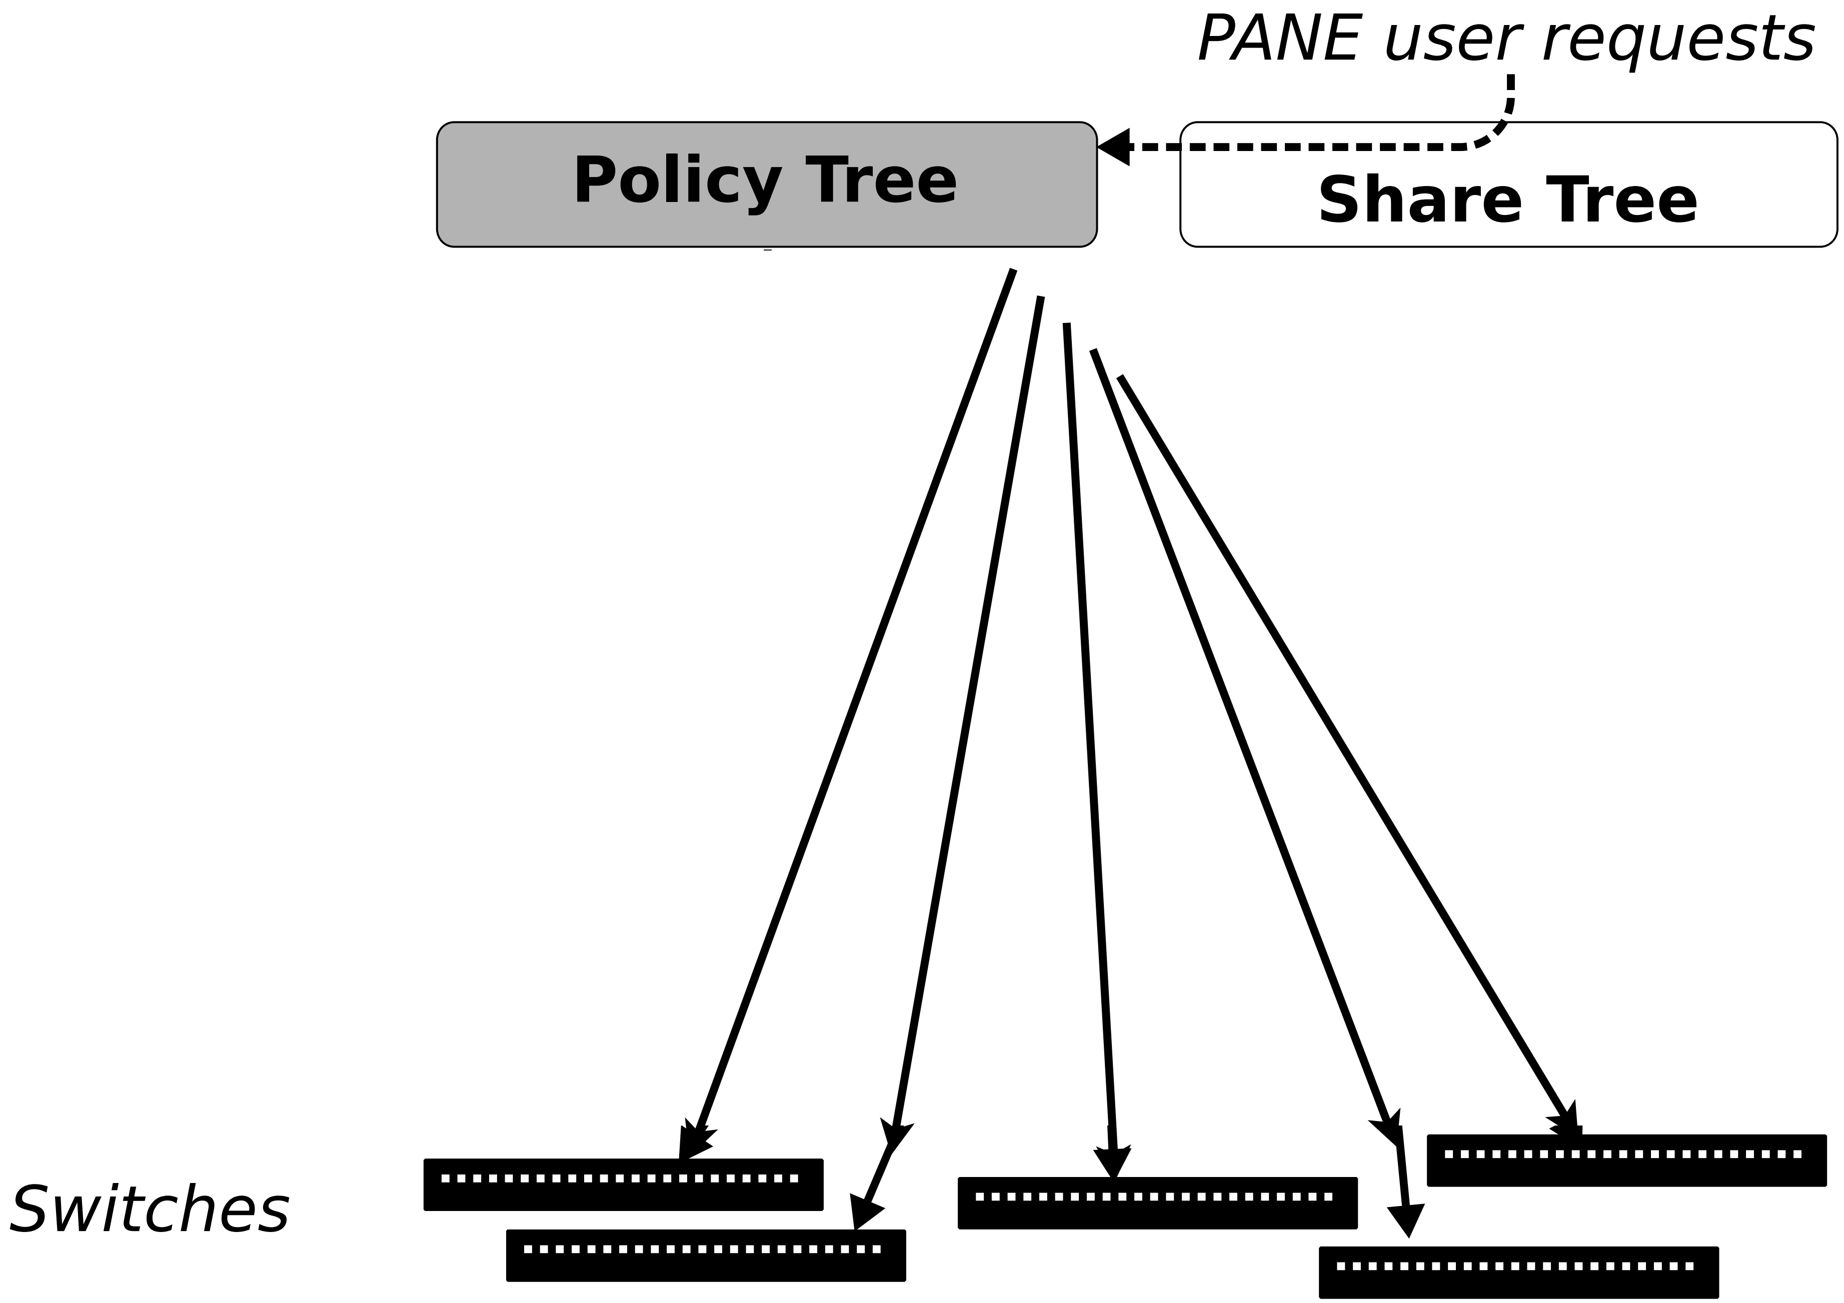

Switches don't grow trees

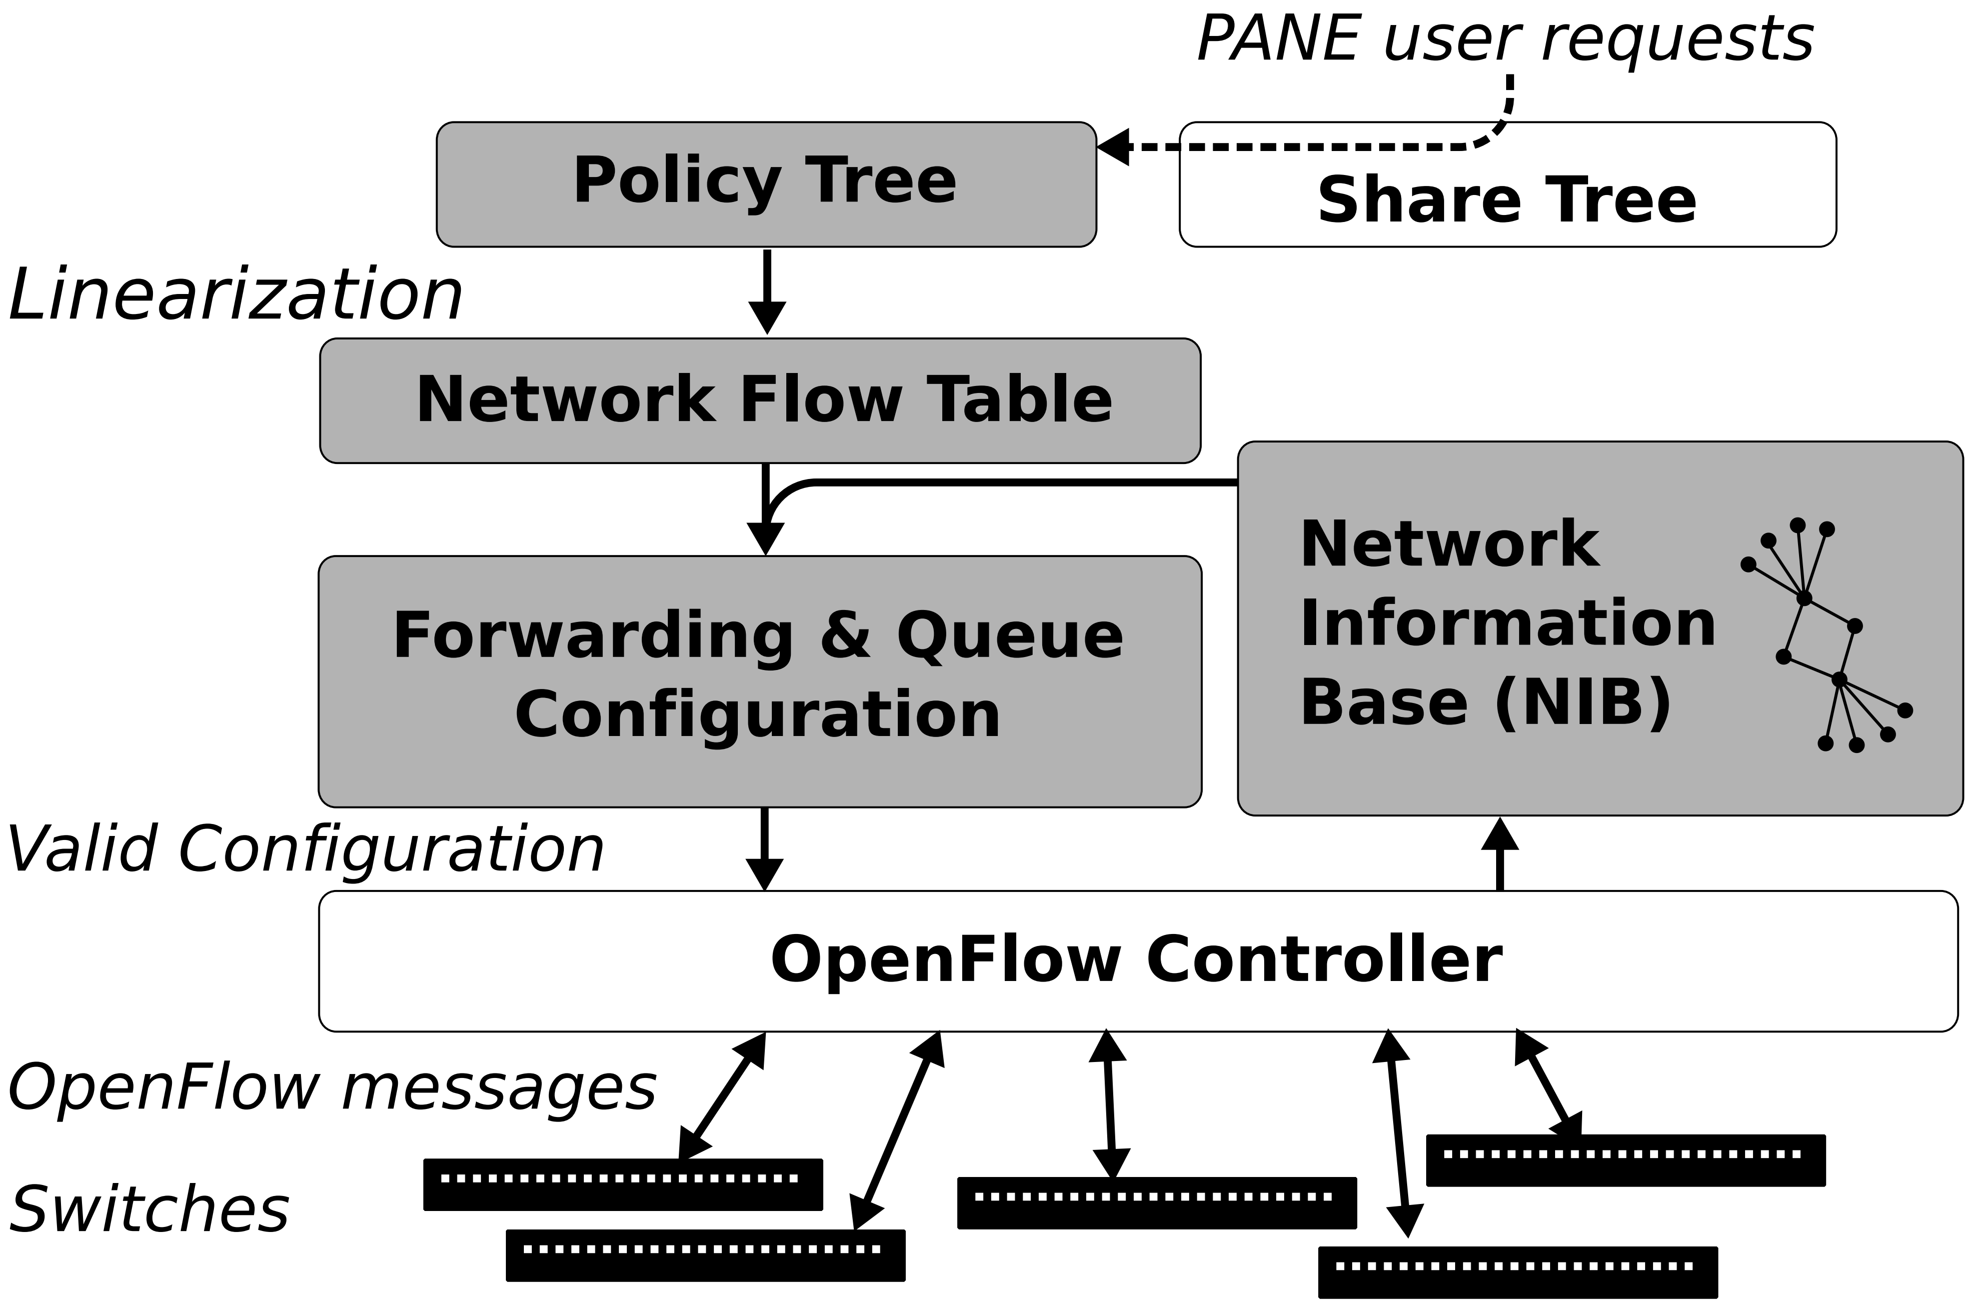

Switches don't grow trees

Current Status

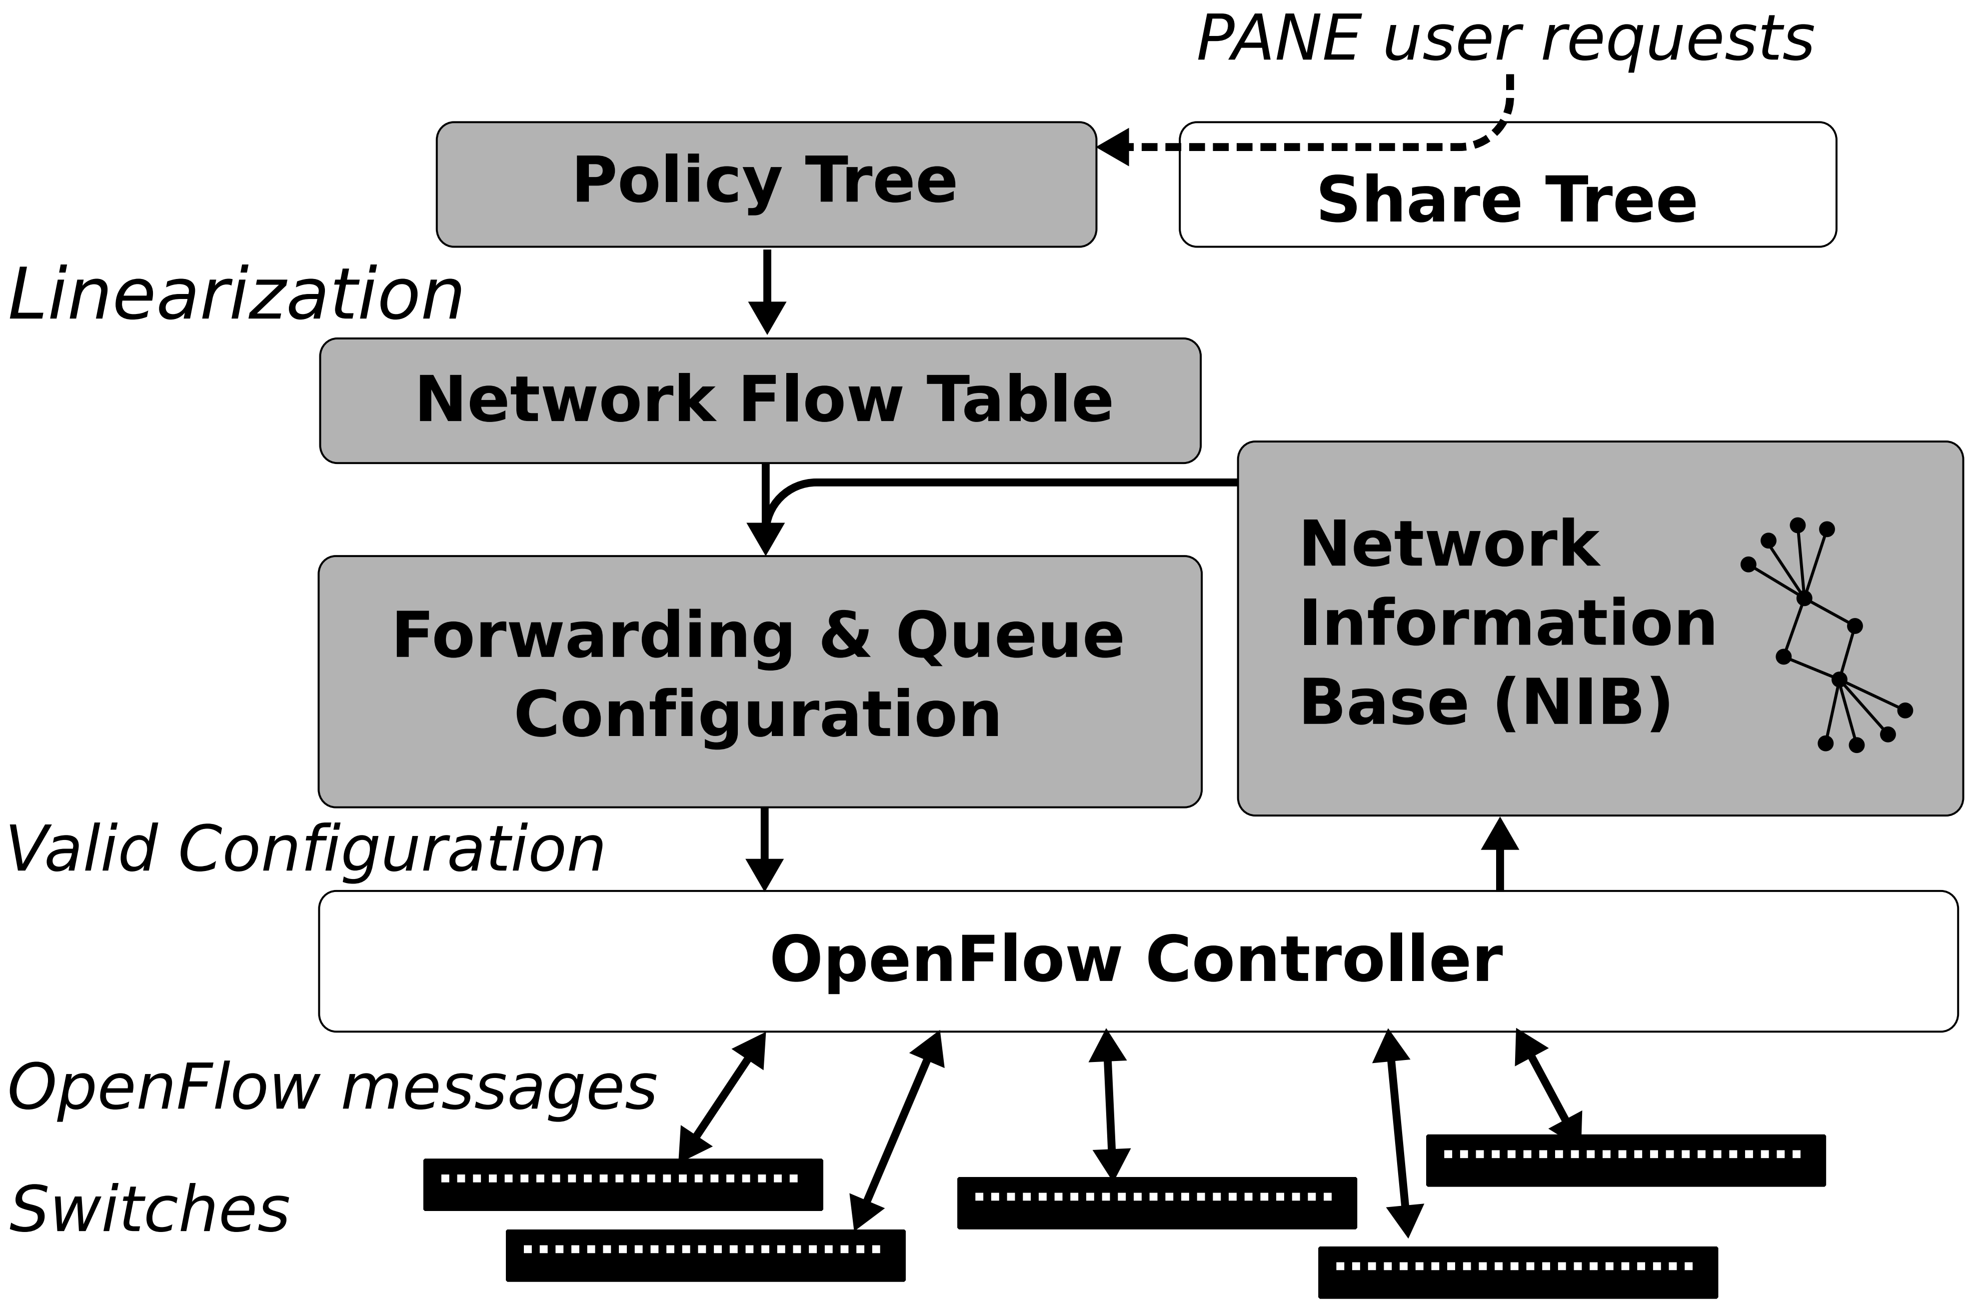
PANE user requests
Policy Tree
Share Tree
Linearization
Network Flow Table
Forwarding & Queue Configuration
Network Information Base (NIB)
Valid Configuration
OpenFlow Controller
OpenFlow messages
Switches

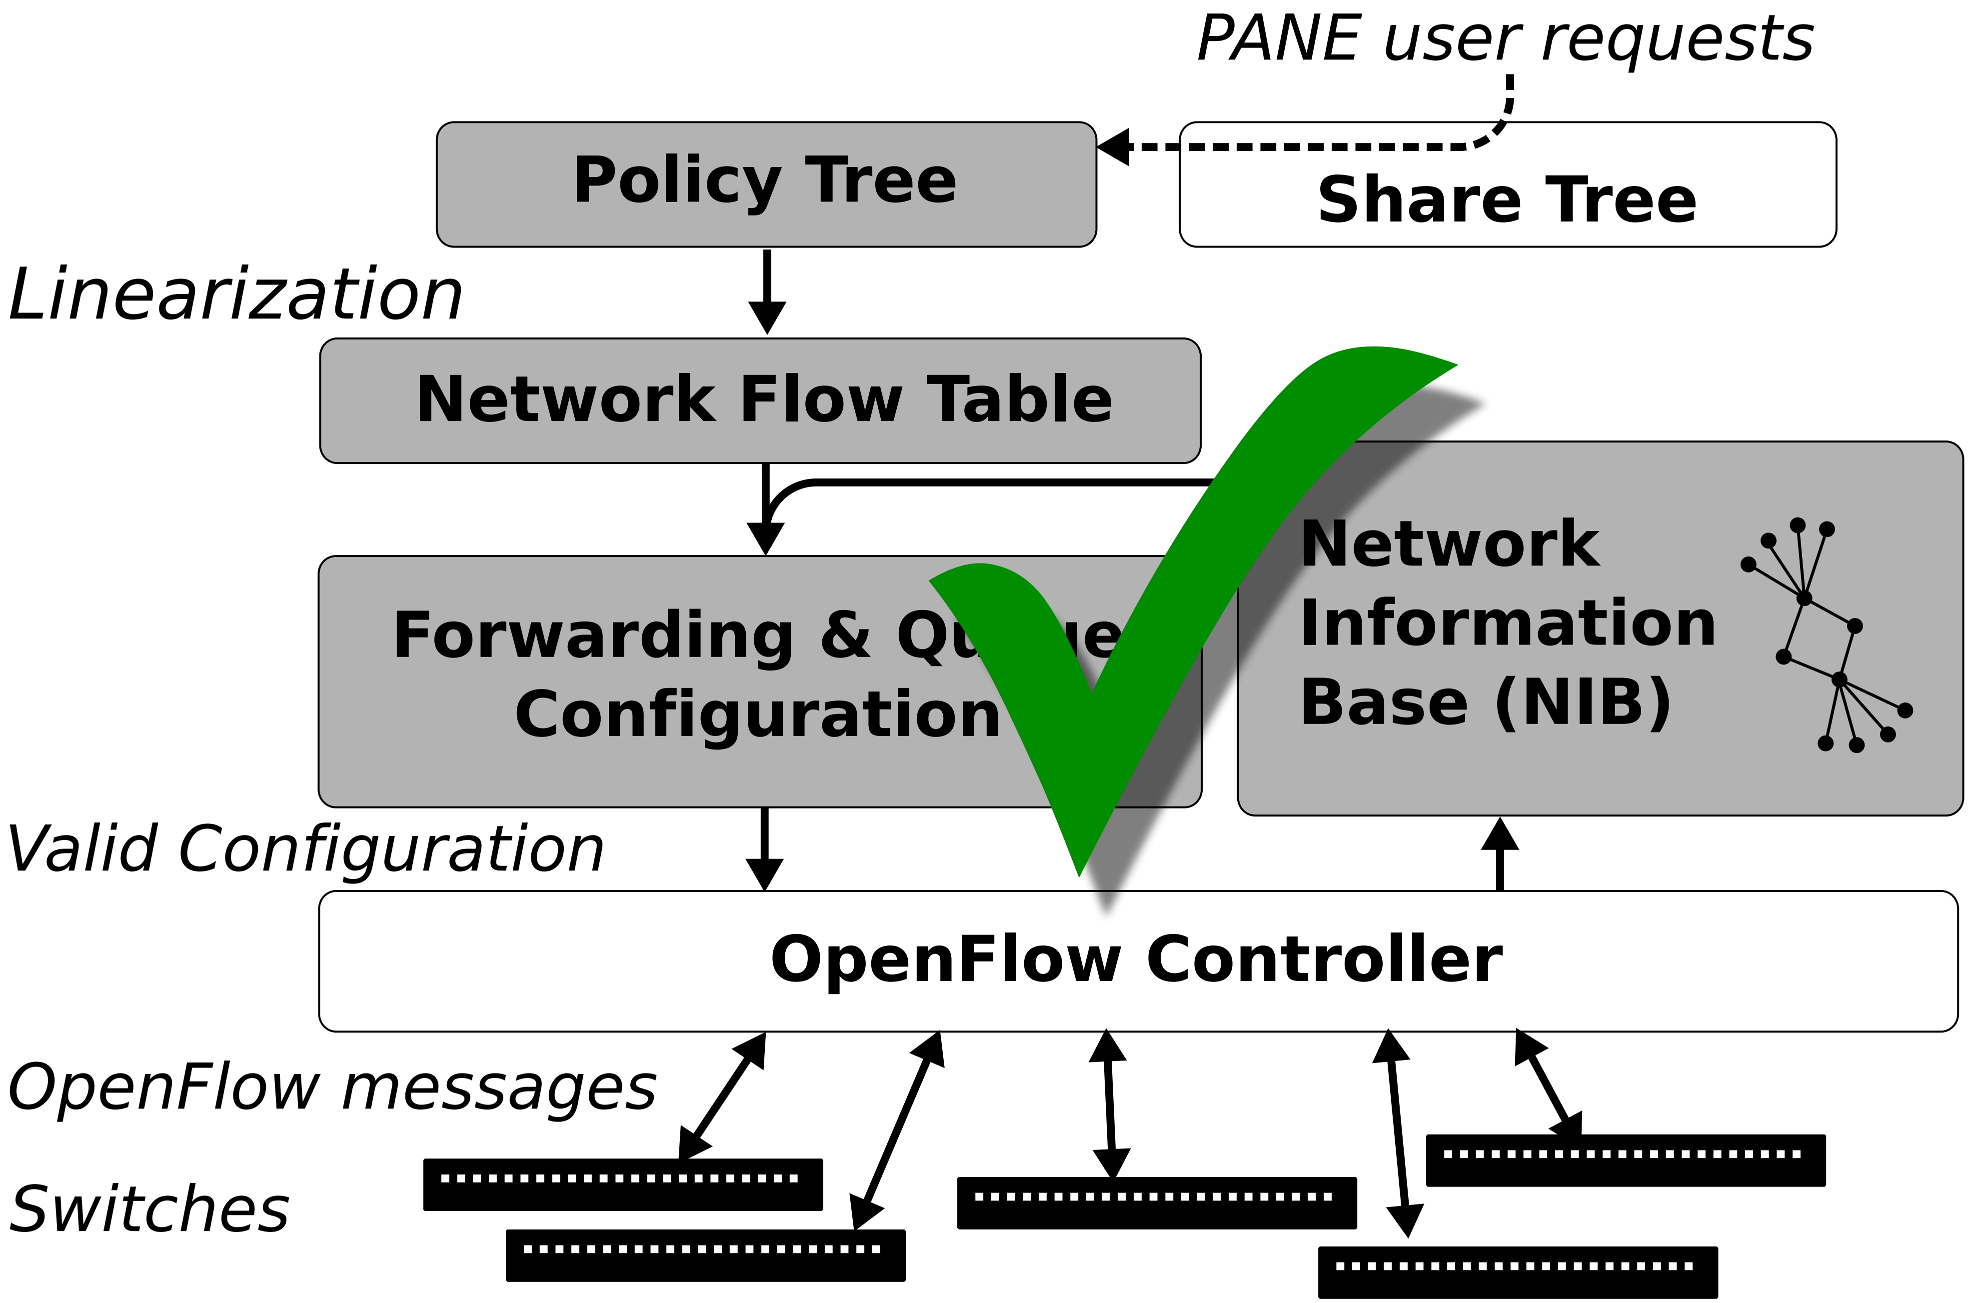
PANE user requests
Policy Tree
Share Tree
Linearization
Network Flow Table
Forwarding & Queue Configuration
Network Information Base (NIB)
Valid Configuration
OpenFlow Controller
OpenFlow messages
Switches

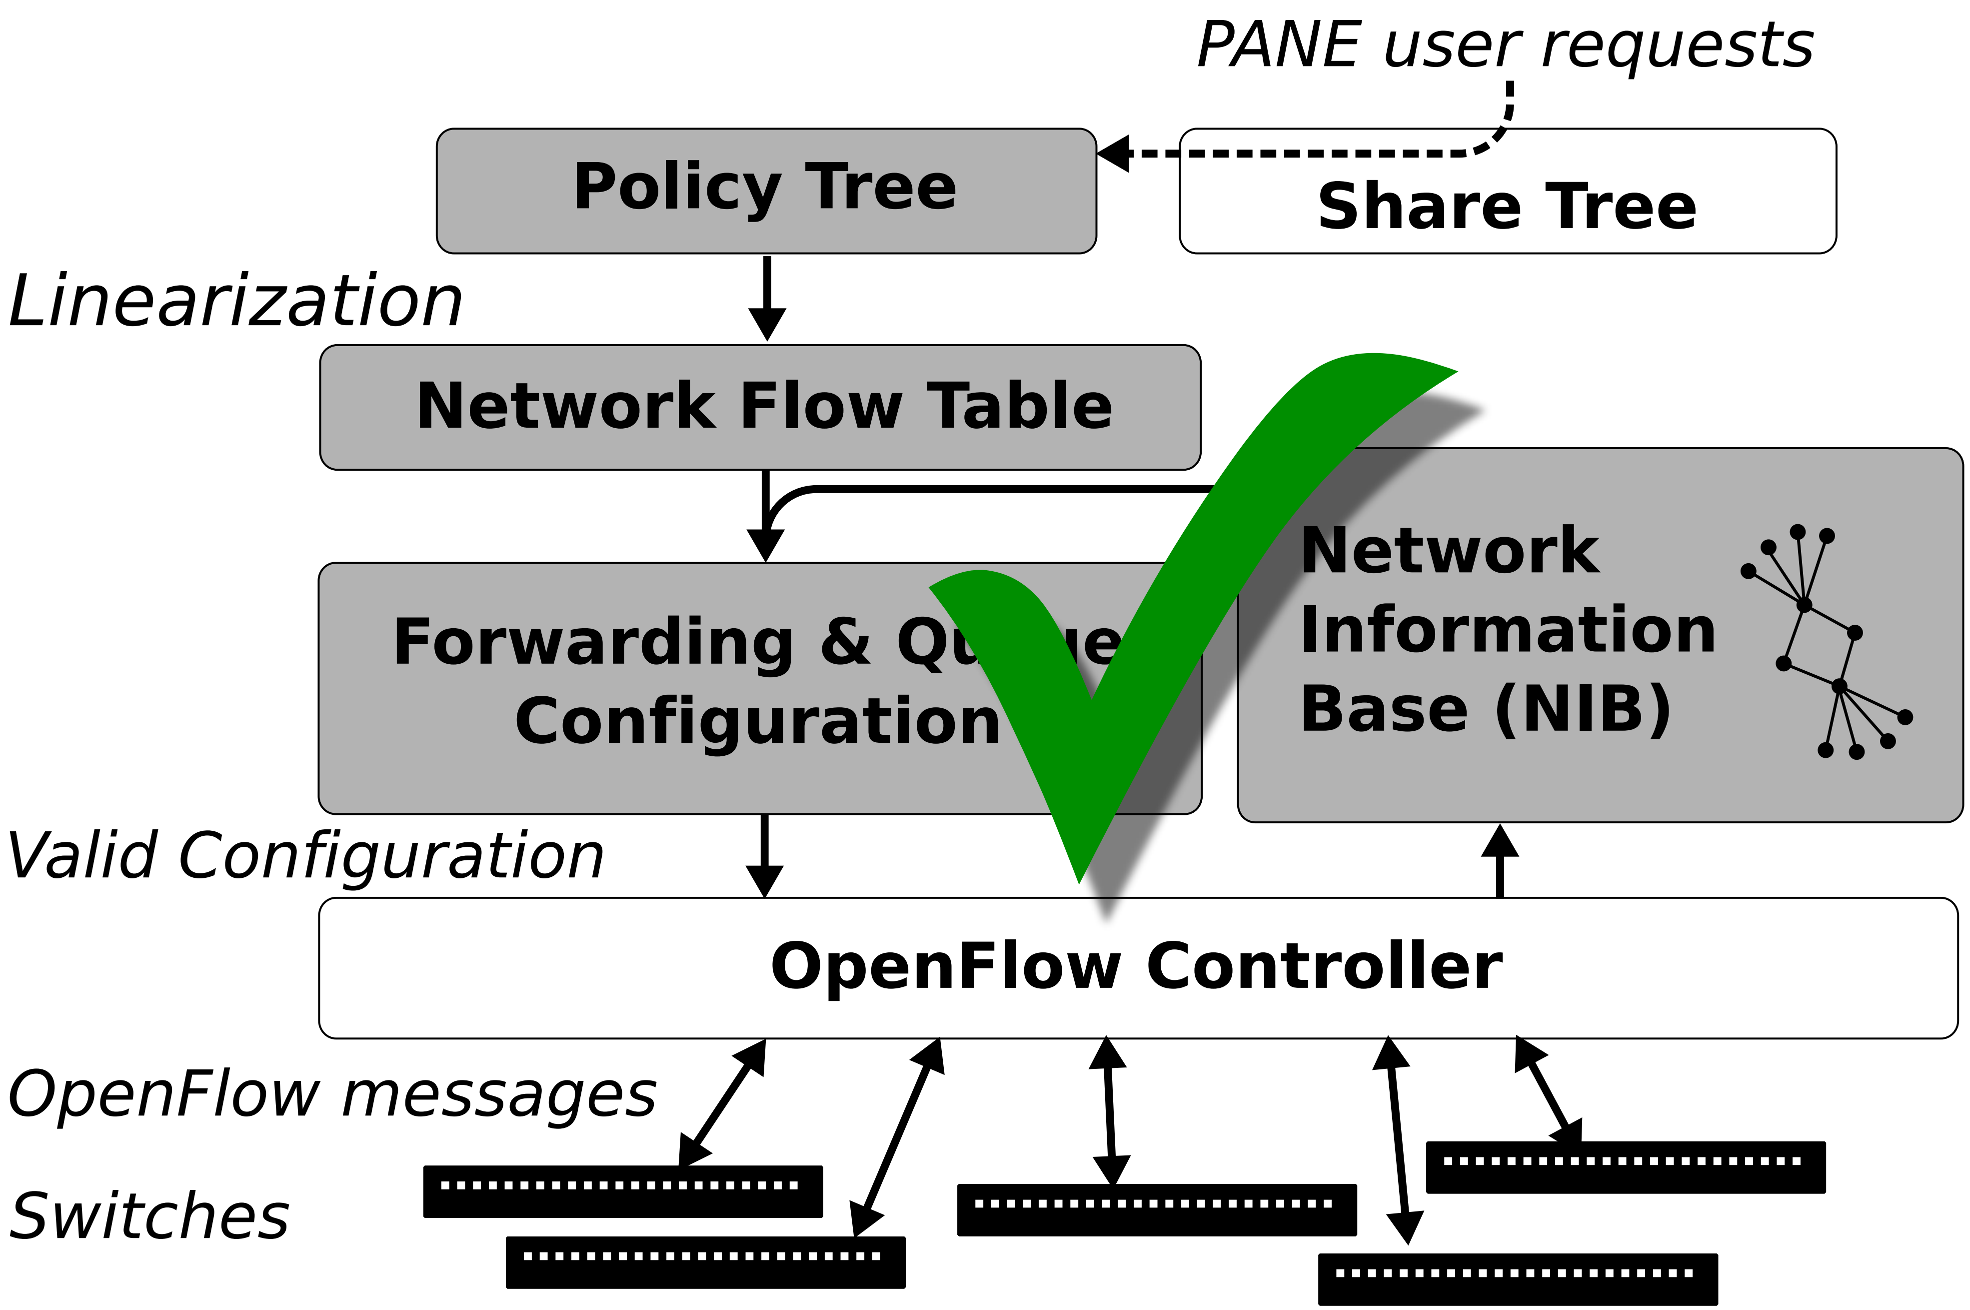

We've built all these components
Currently have Access Control and GMB

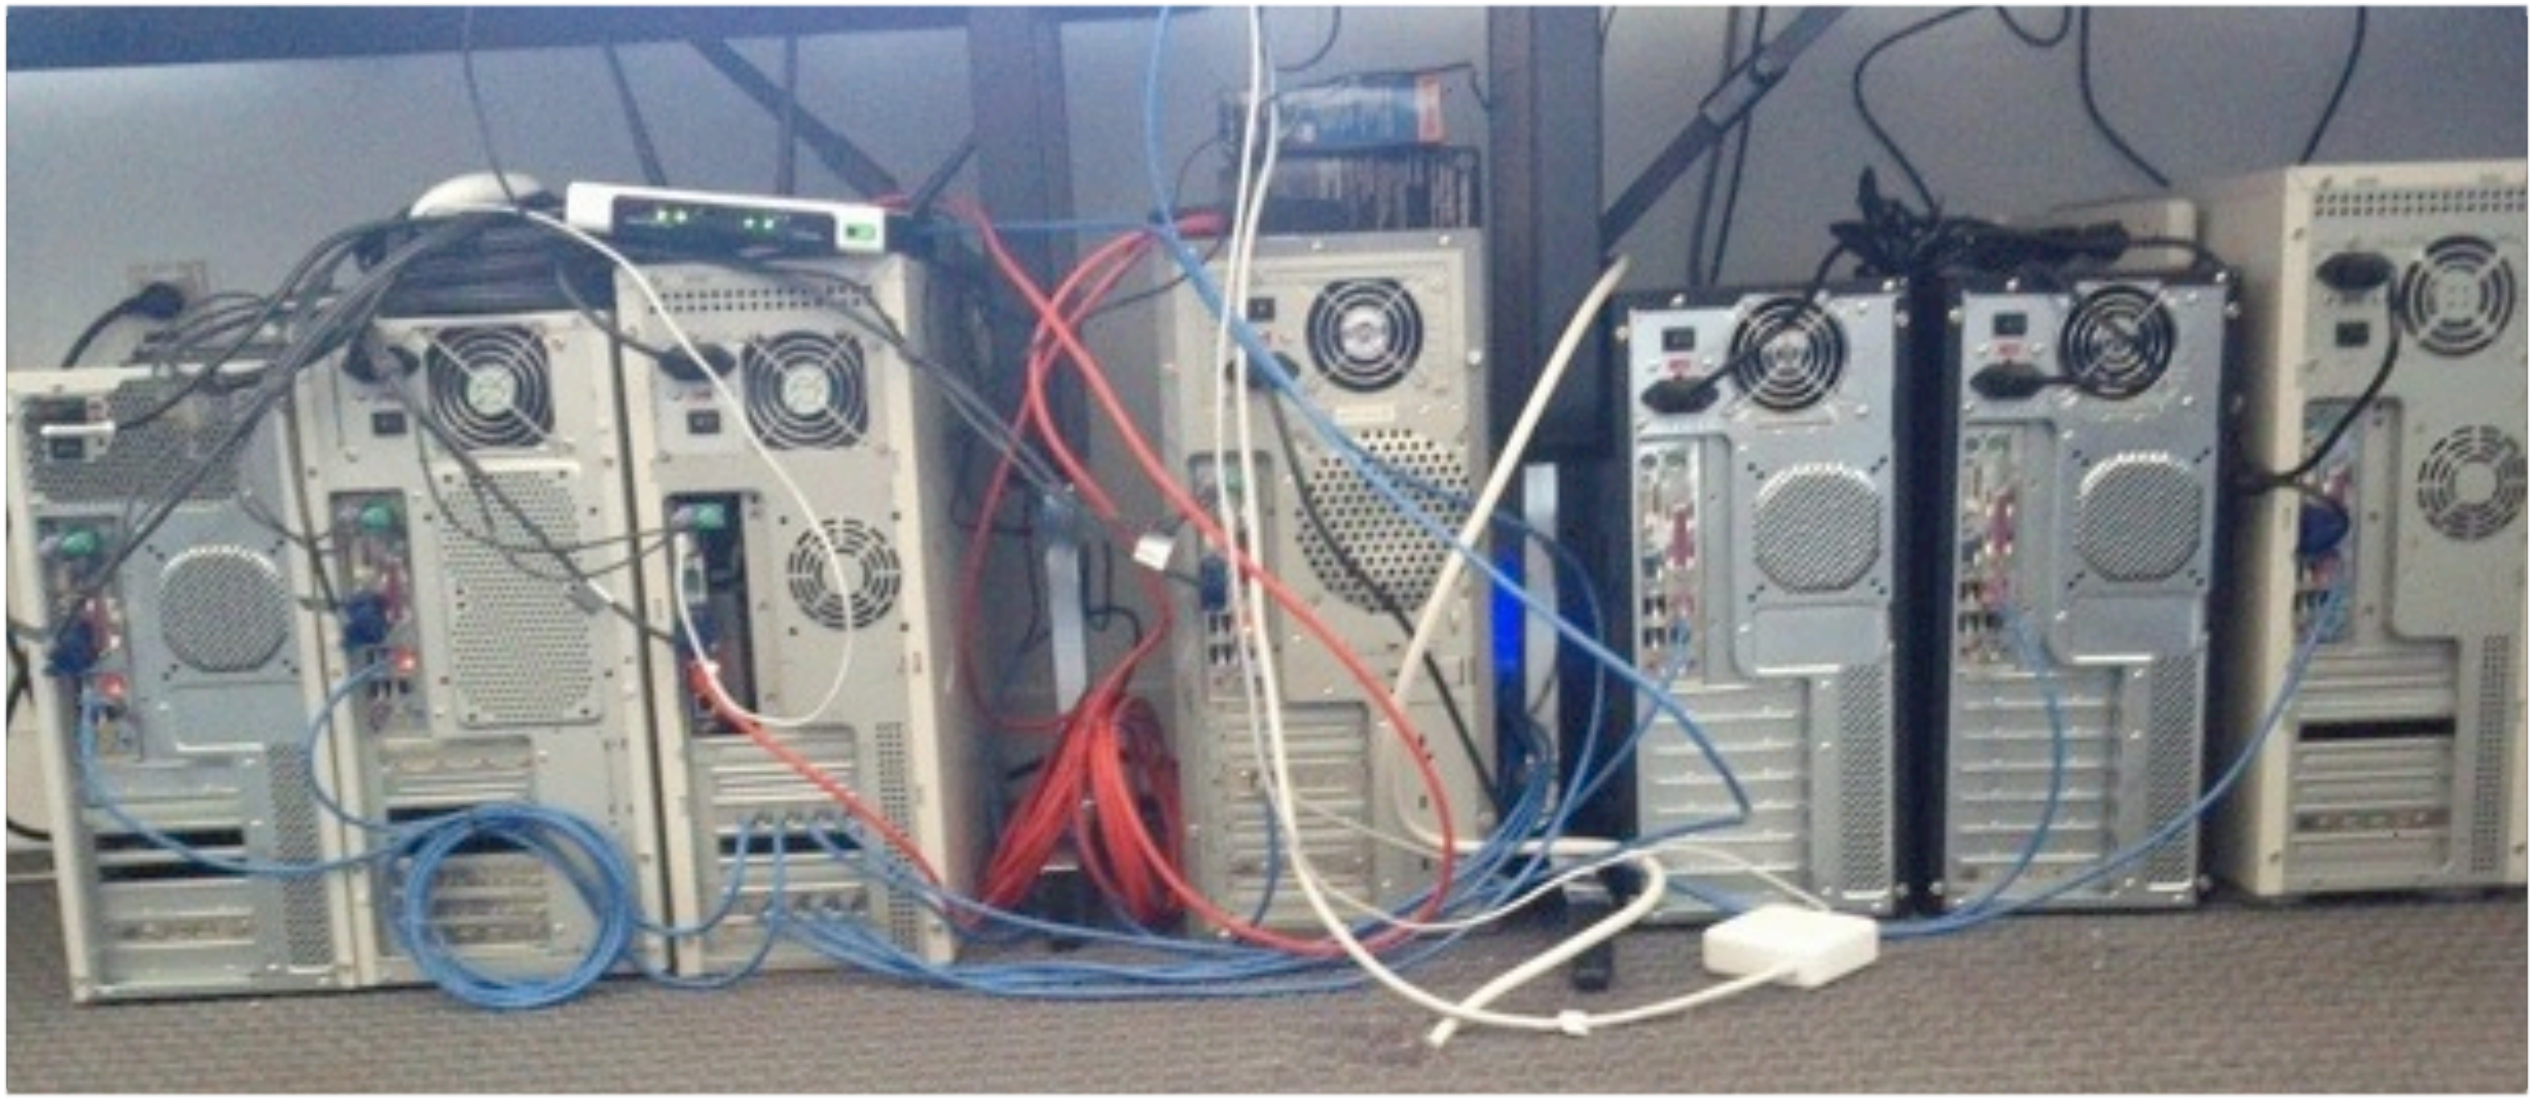

Toy Evaluation

Protecting Zookeeper

5 PANE-enabled Zookeper servers
1 client
Connected via single OpenVSwitch (3.3Gbps)
iPerf generating load on all links

Protecting Zookeeper

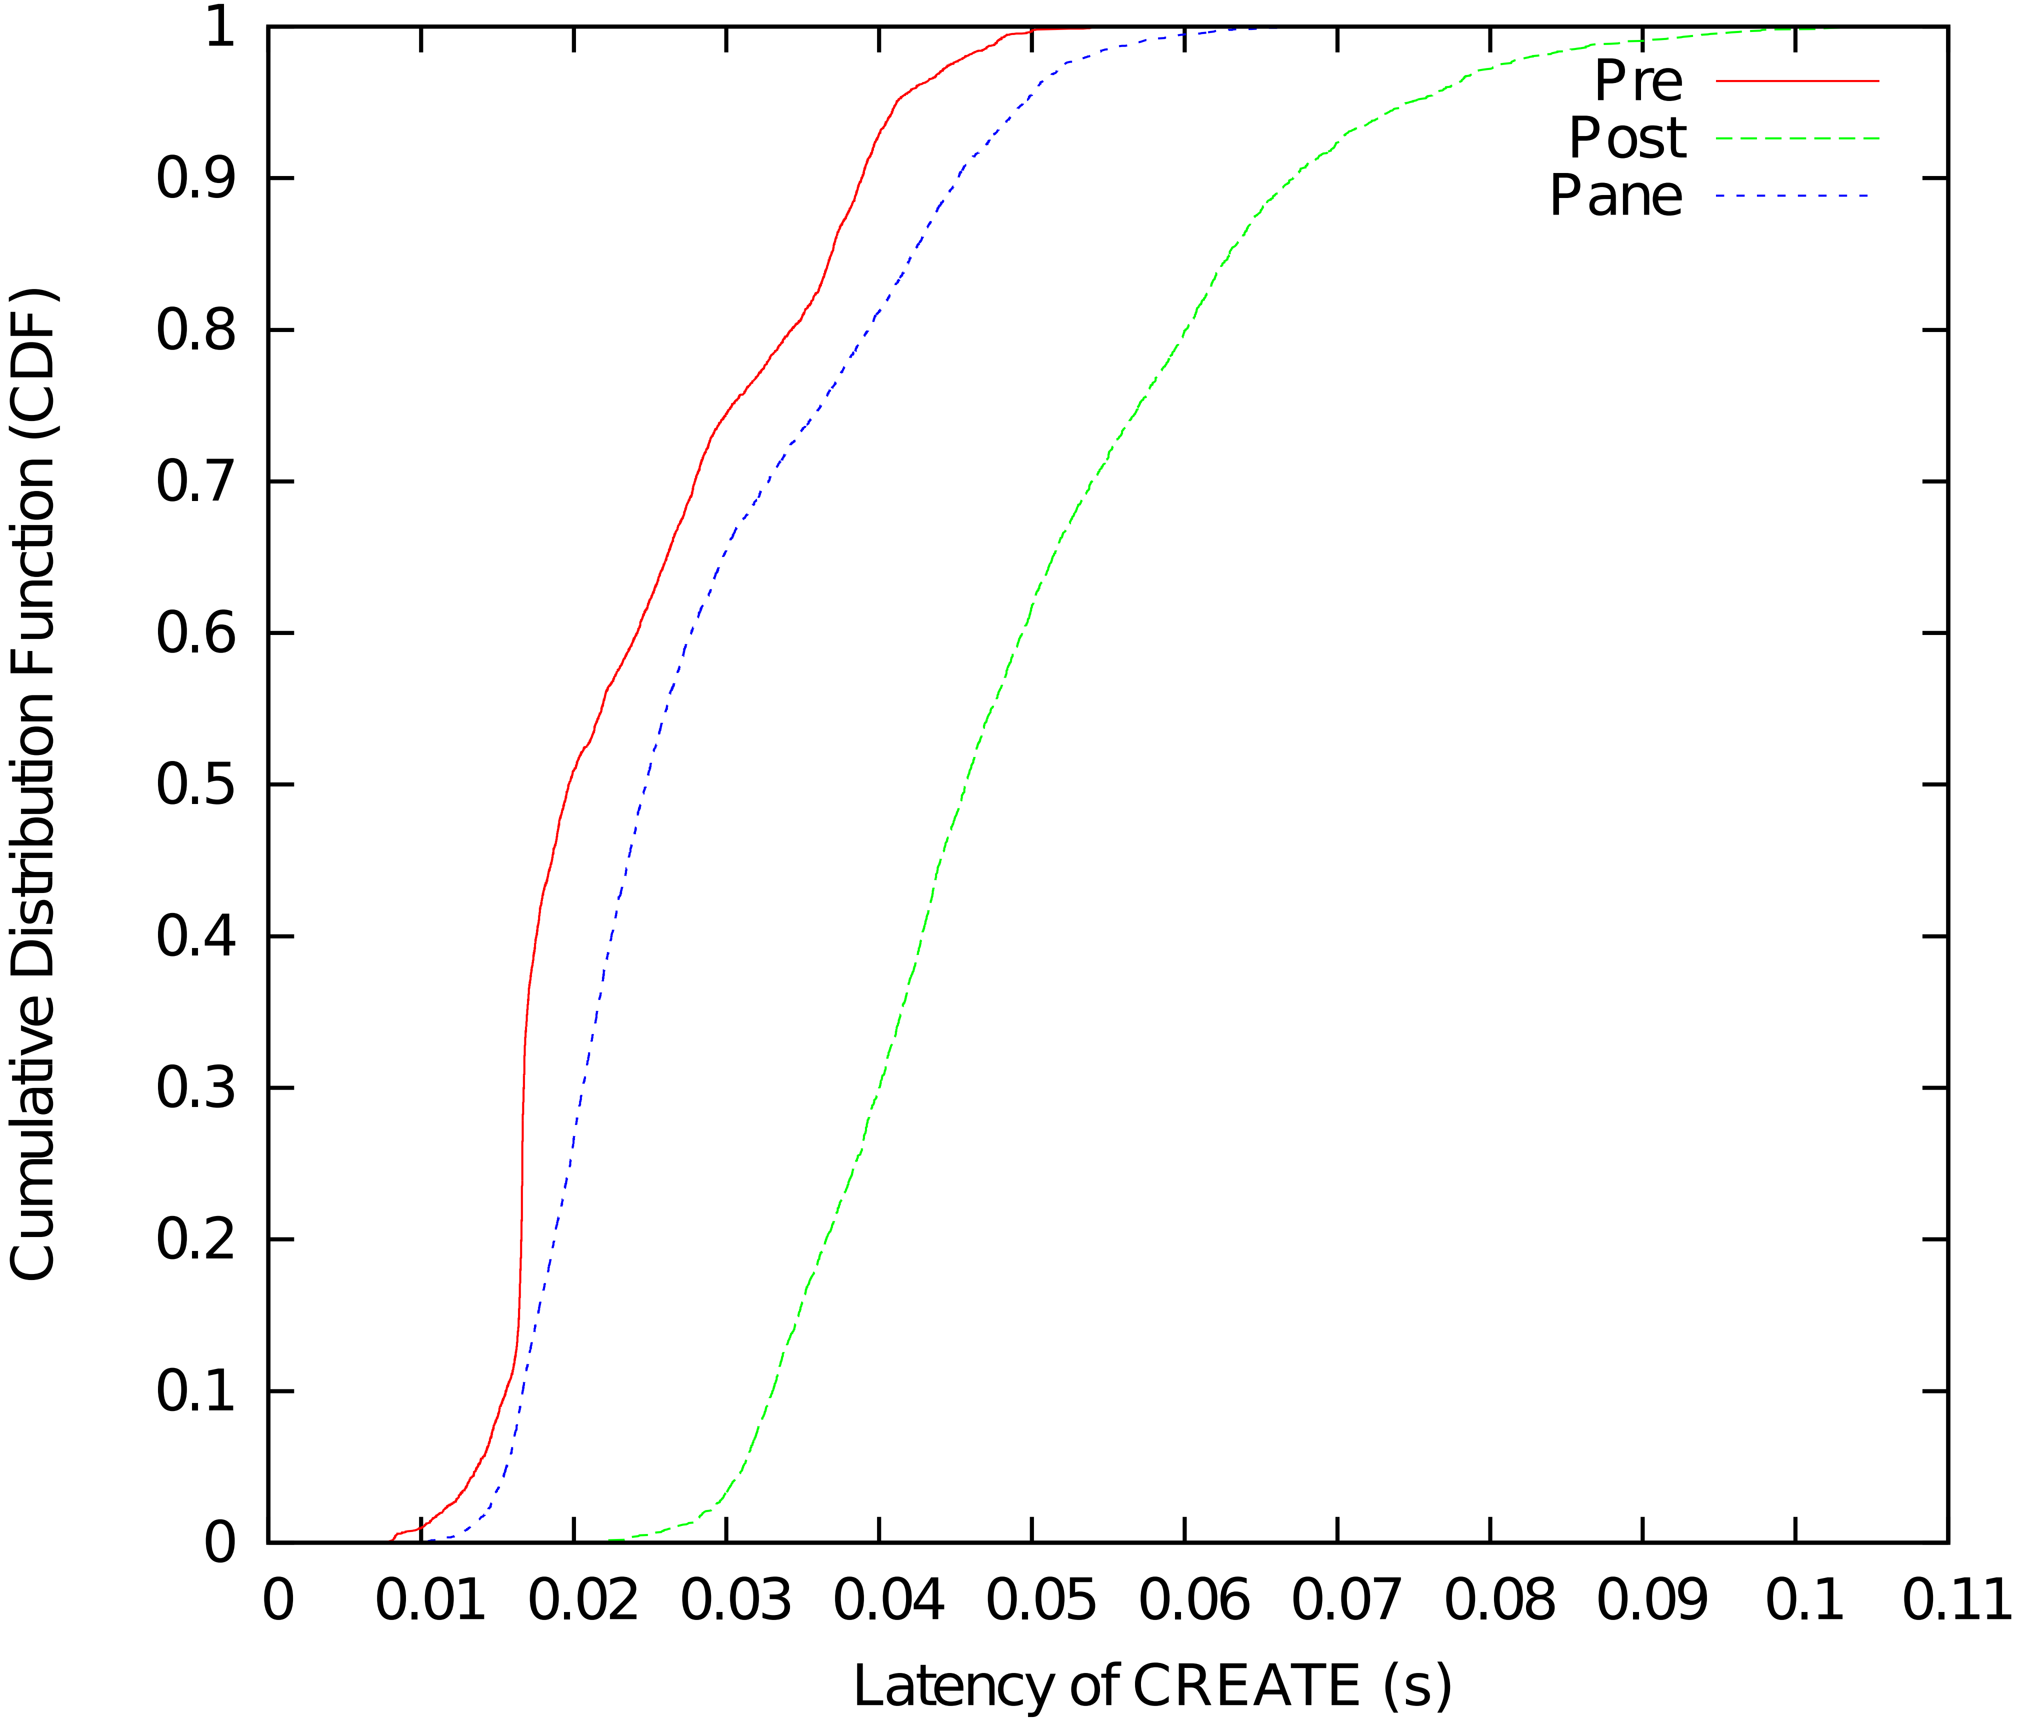

Cumulative Distribution Function (CDF)
Latency of CREATE (s)
Pre
Post
Pane
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0
0.01
0.02
0.03
0.04
0.05
0.06
0.07
0.08
0.09
0.1
0.11

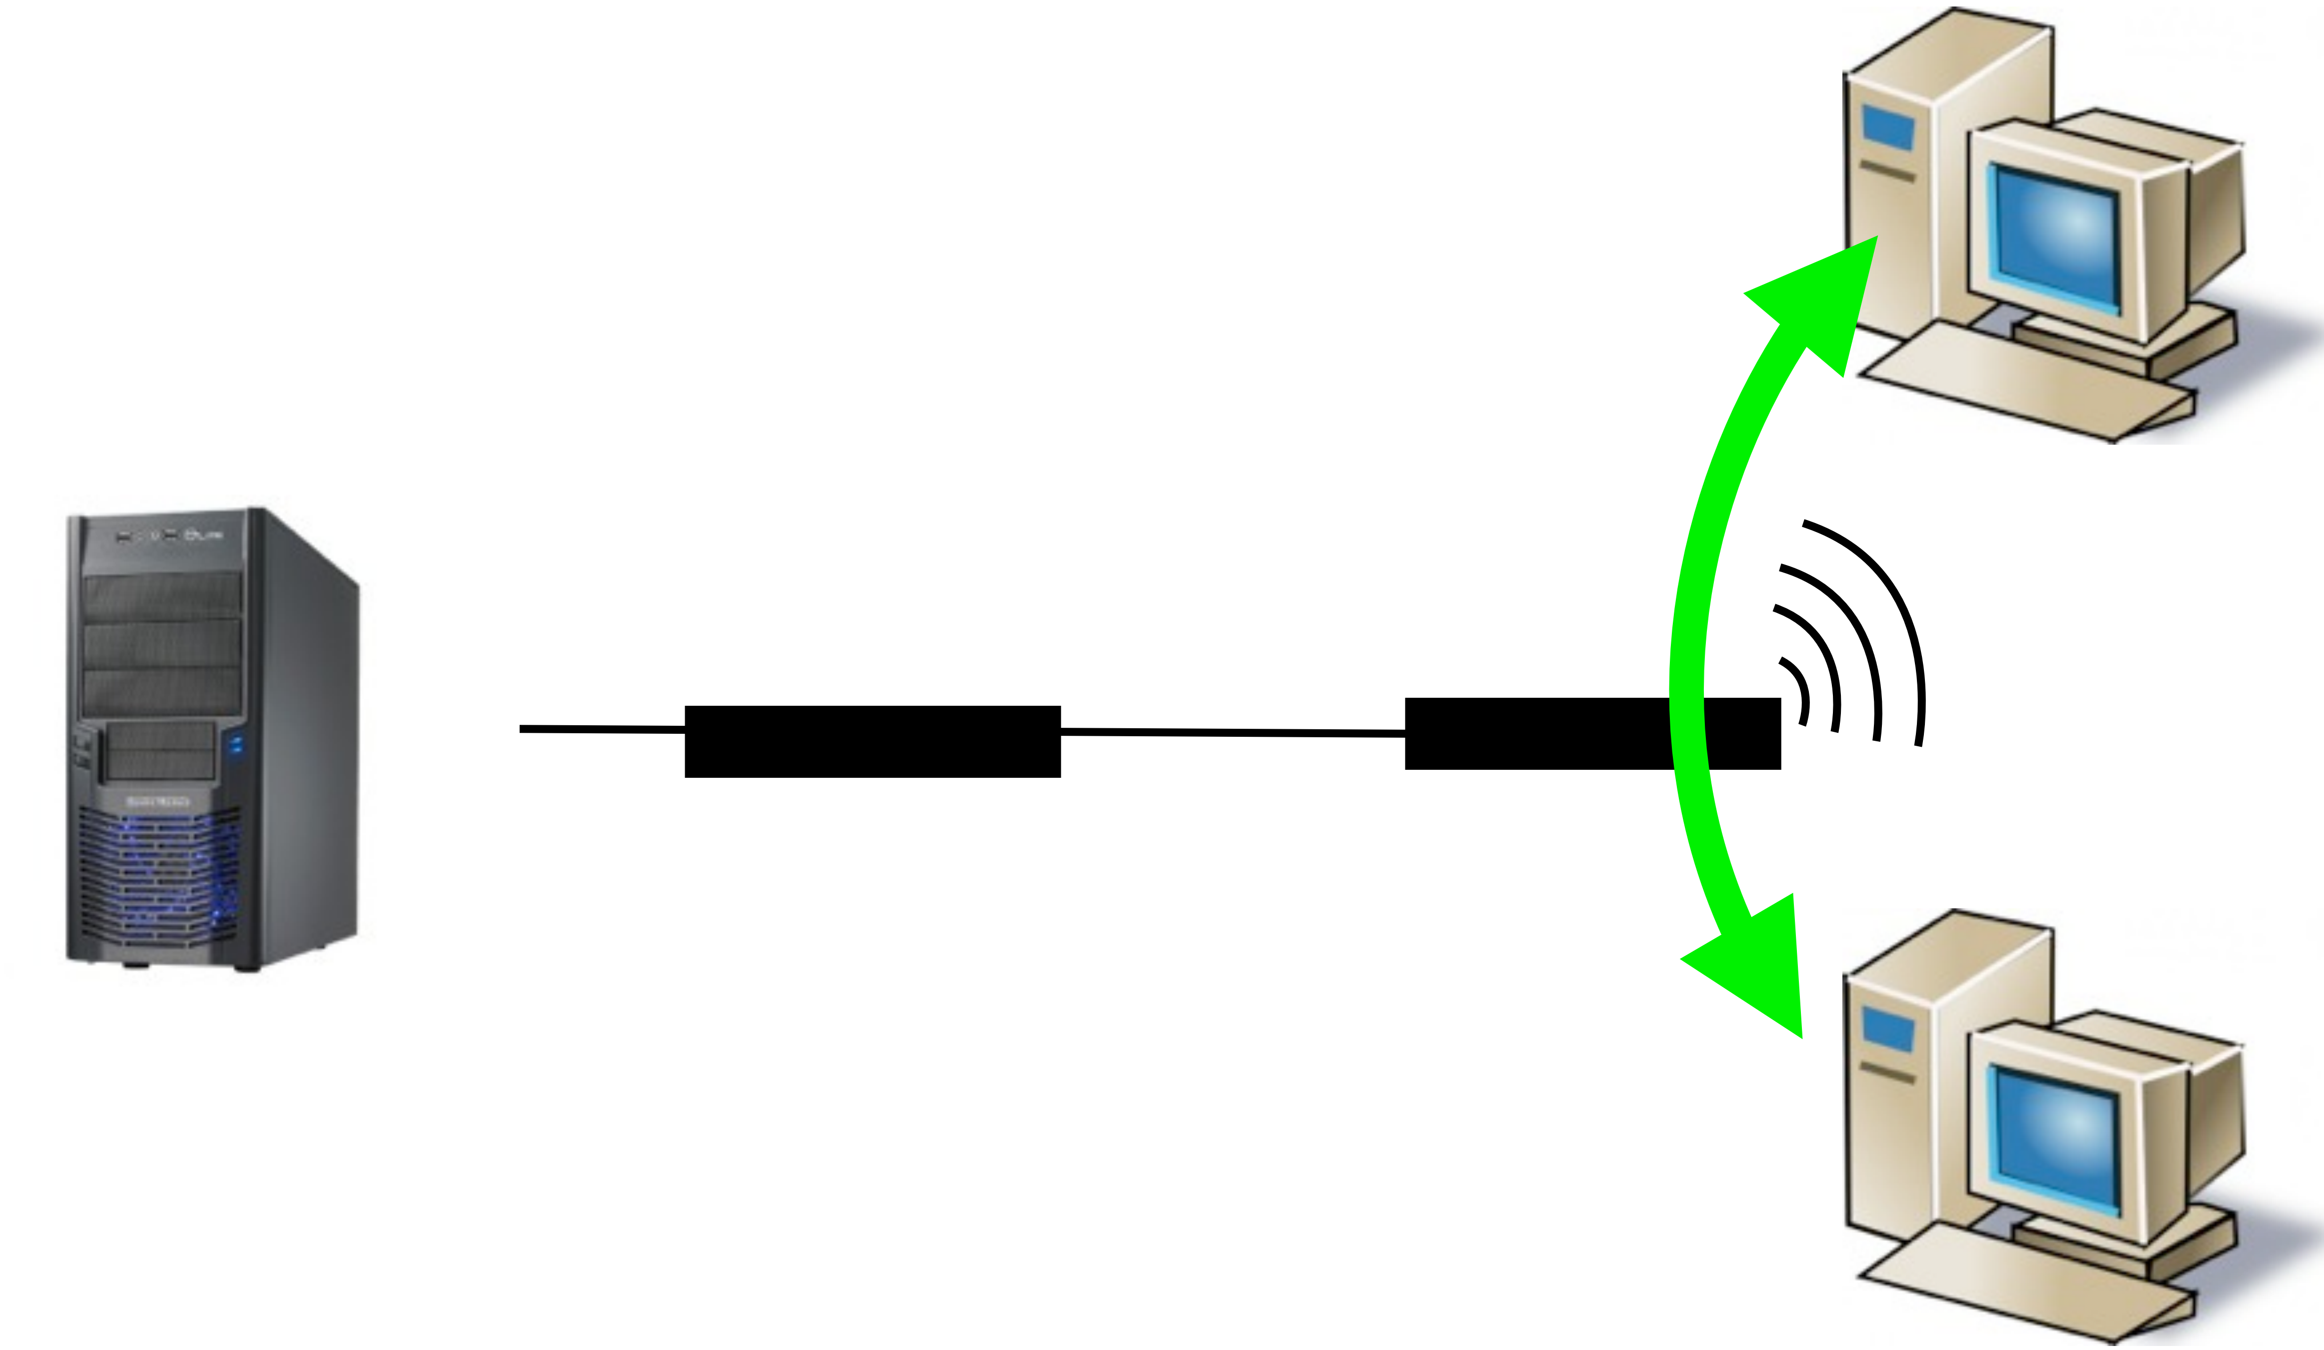

Denial-of-service

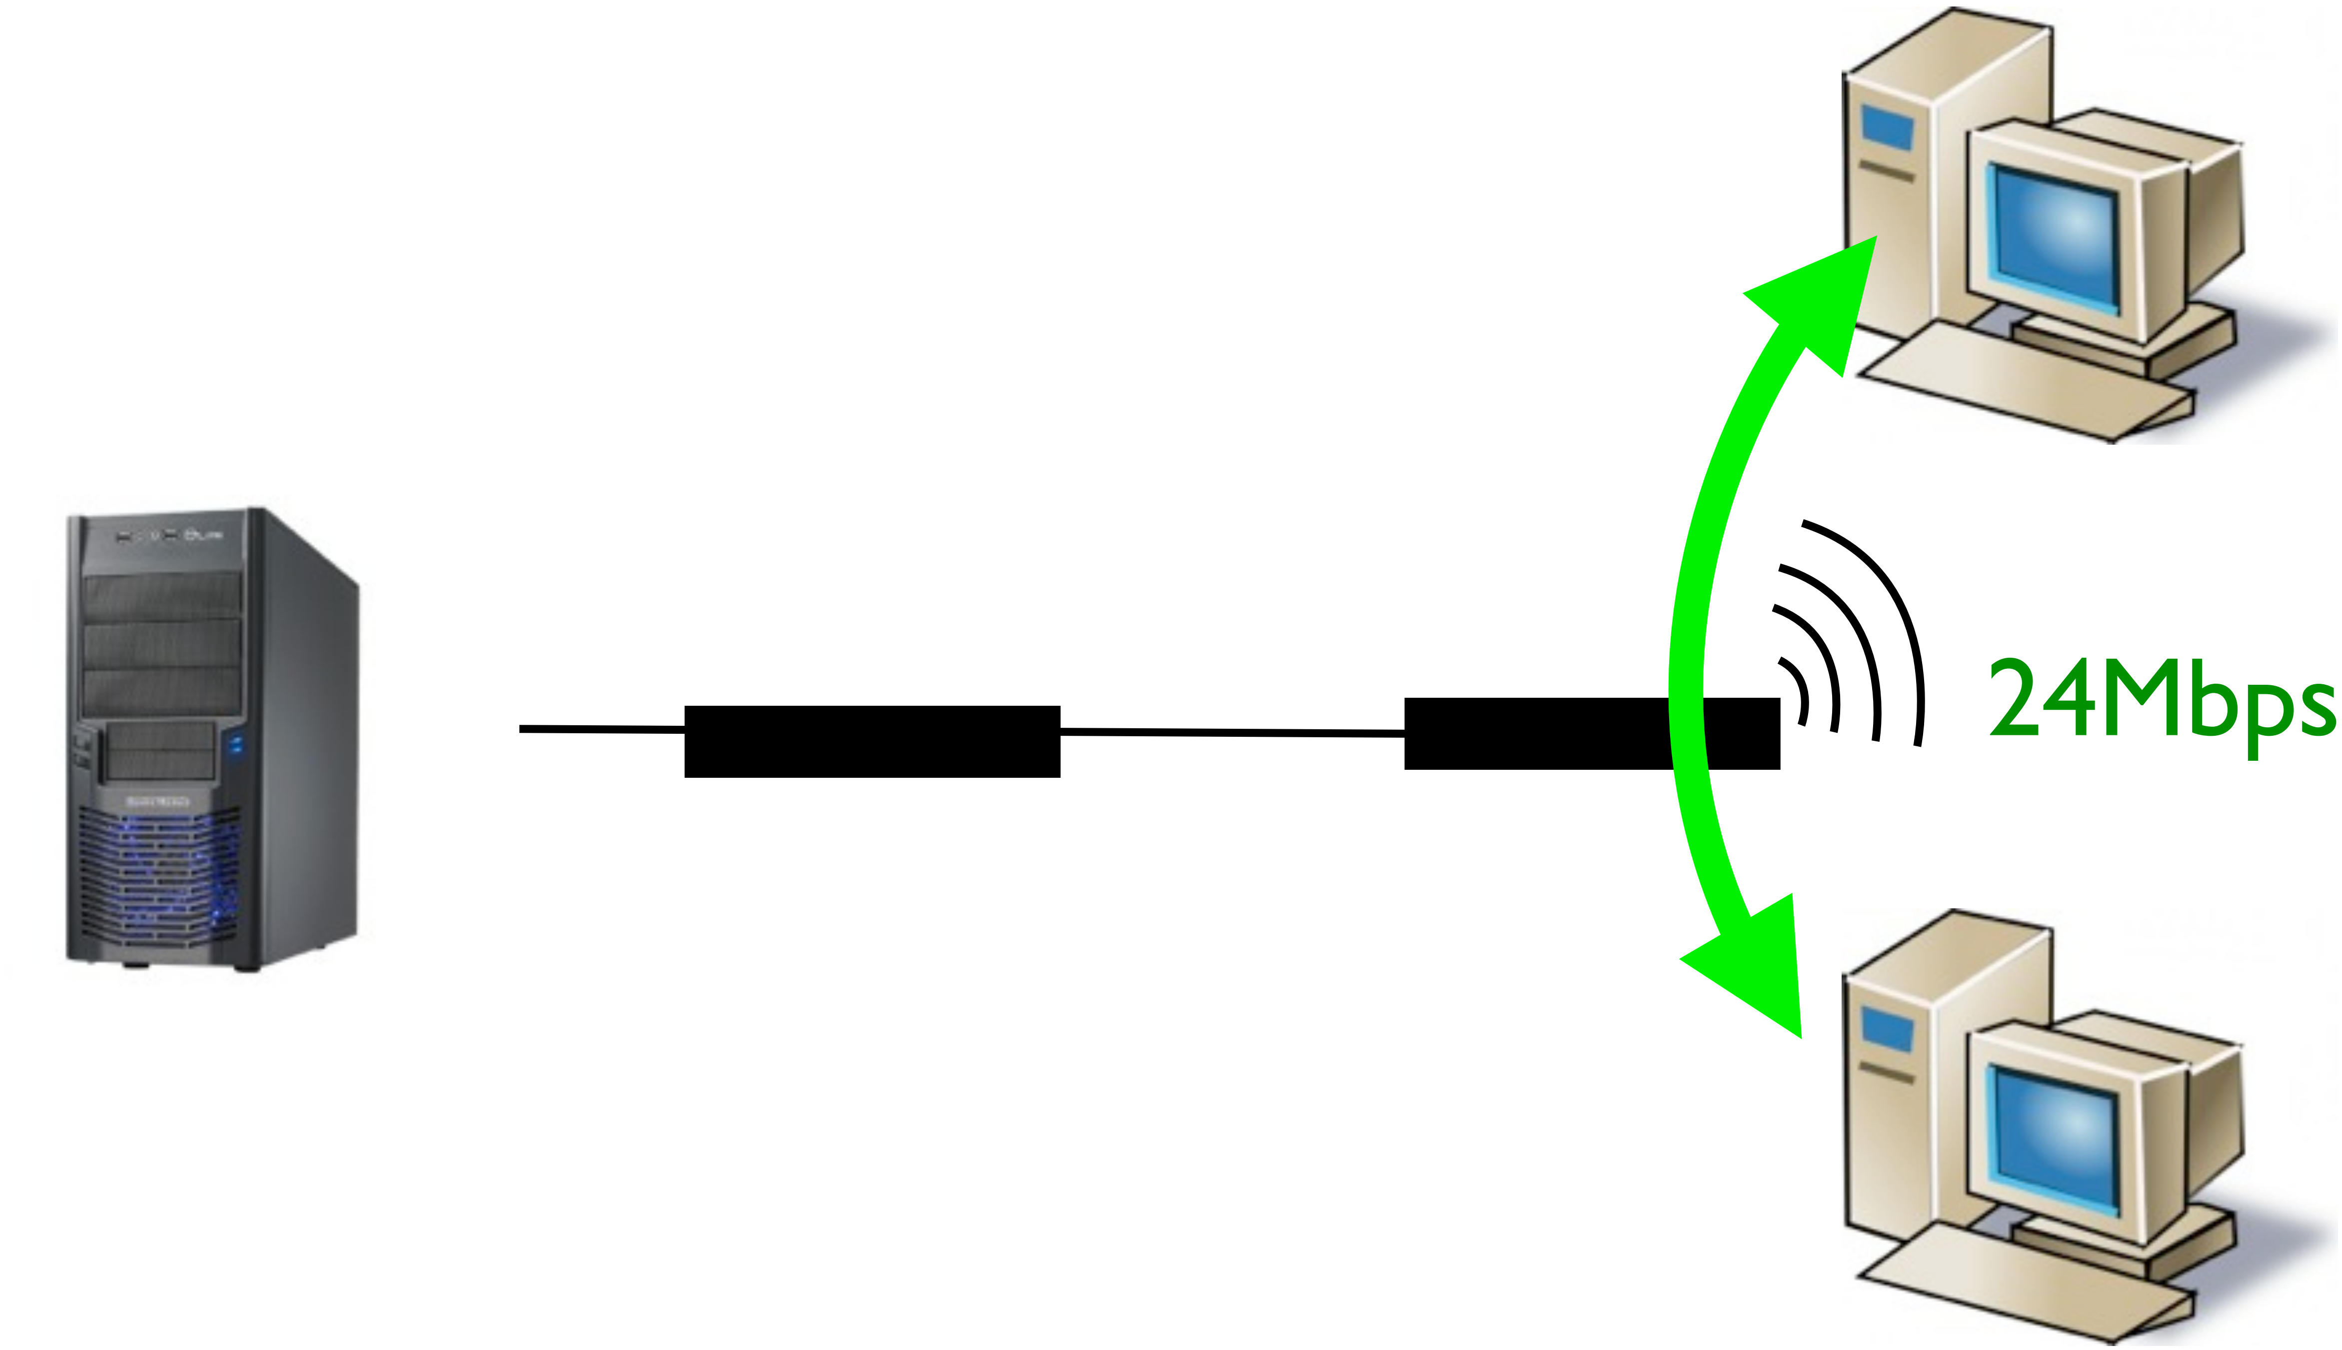

Denial-of-service

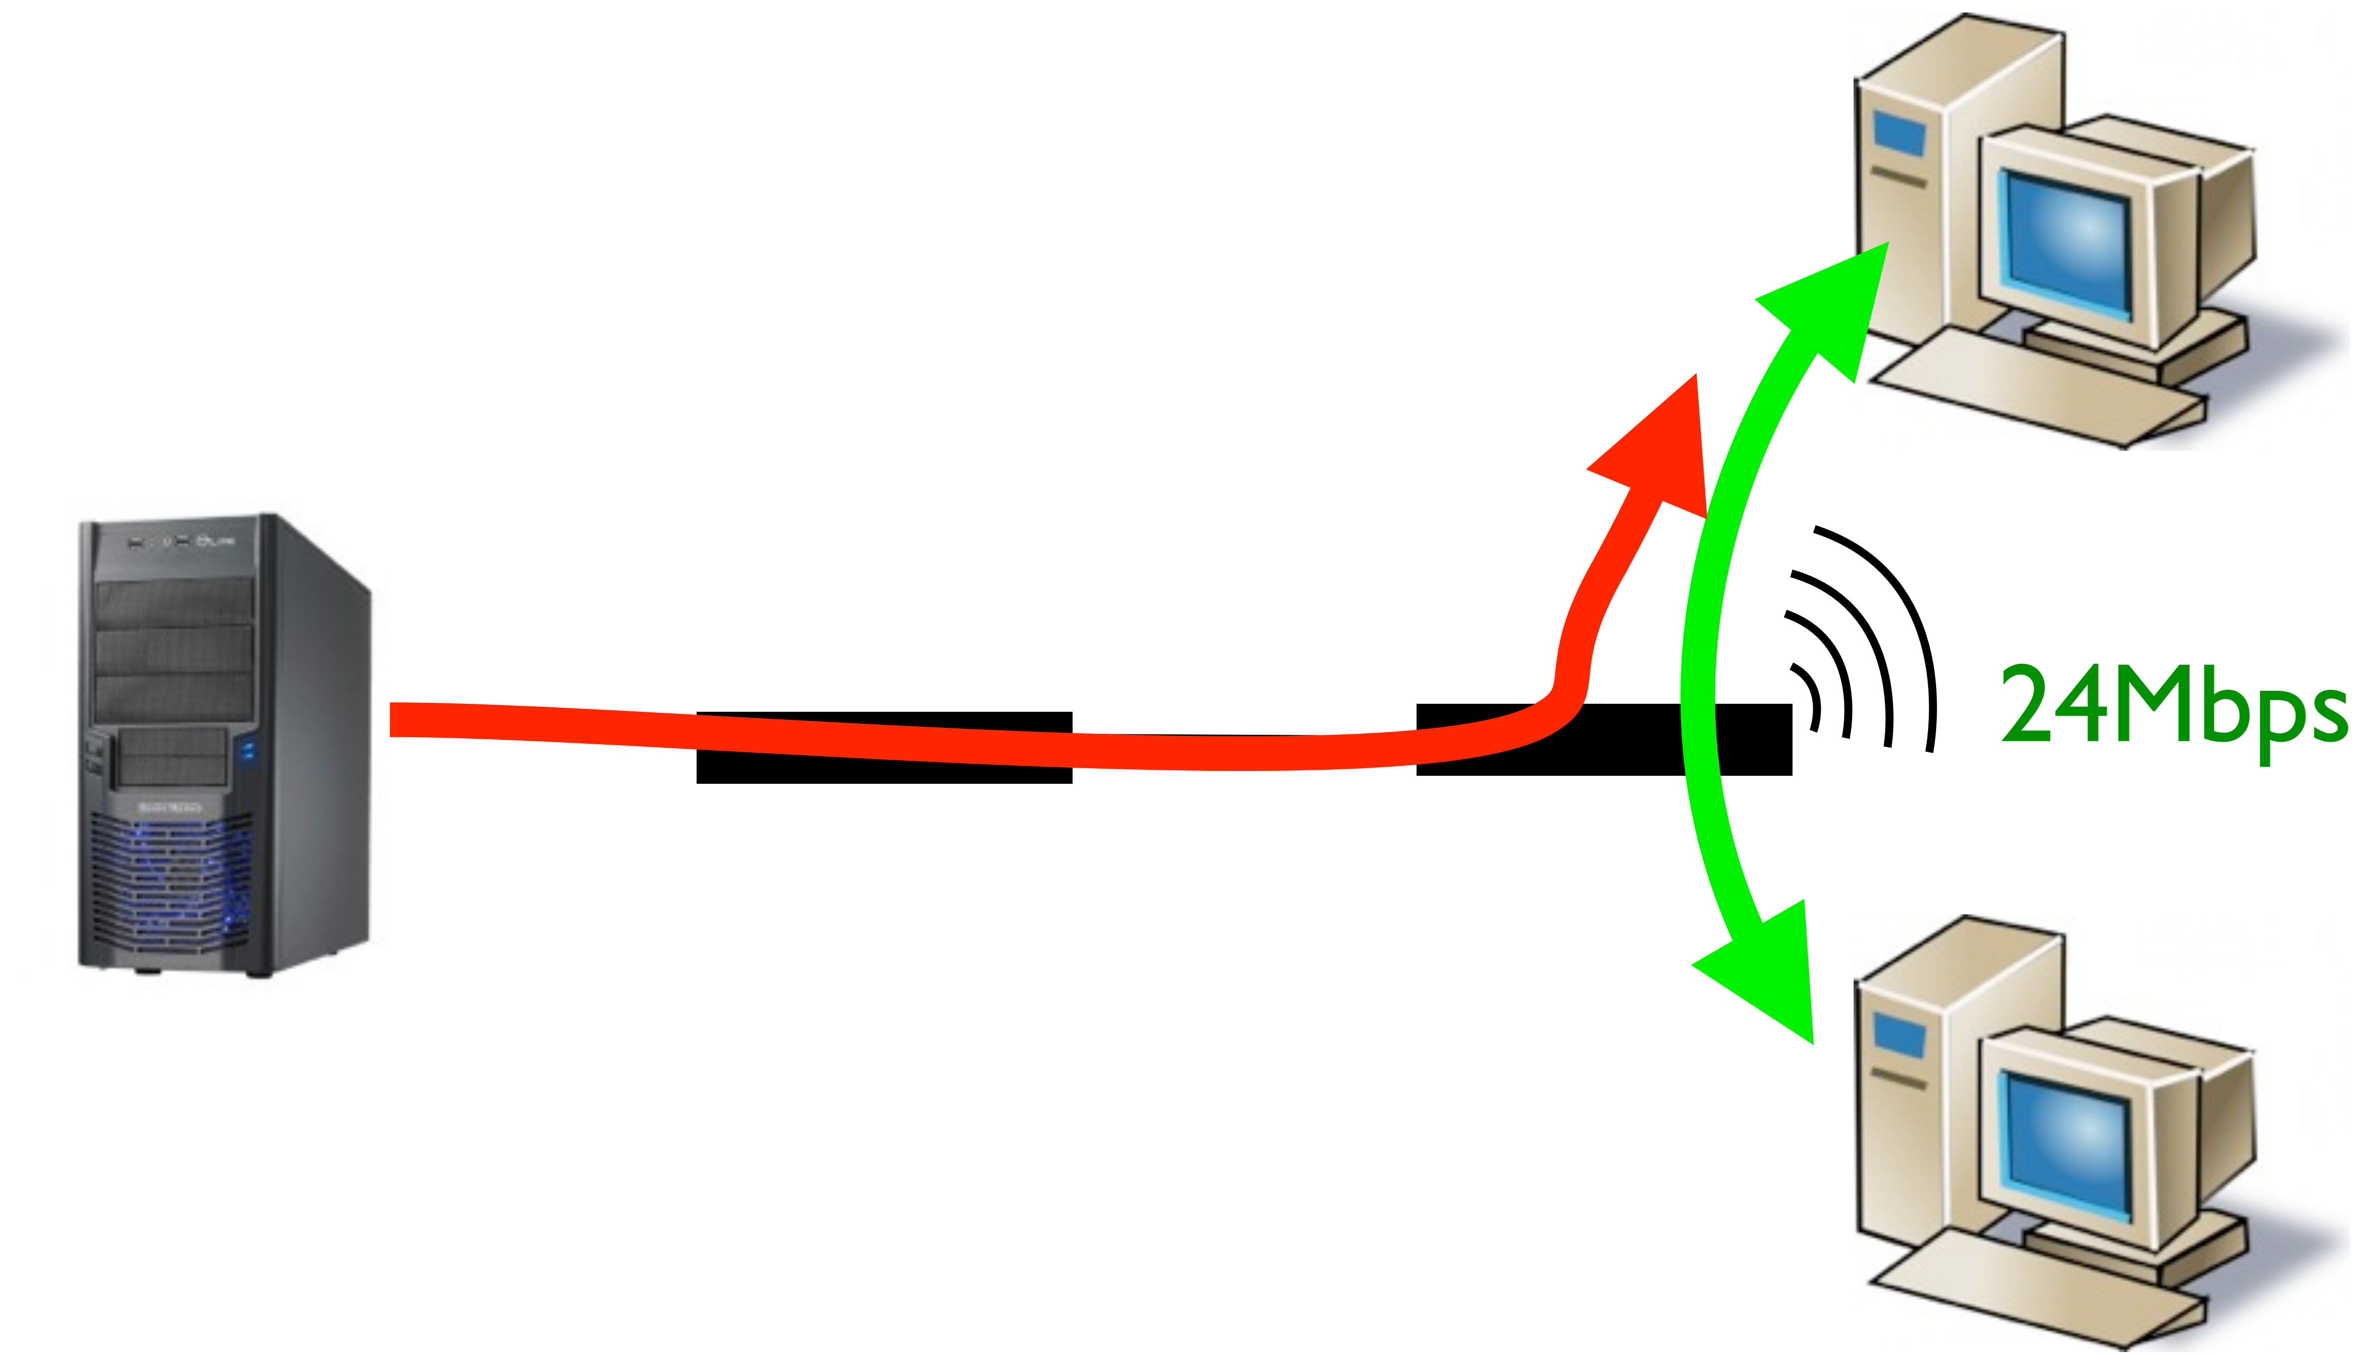

Denial-of-service

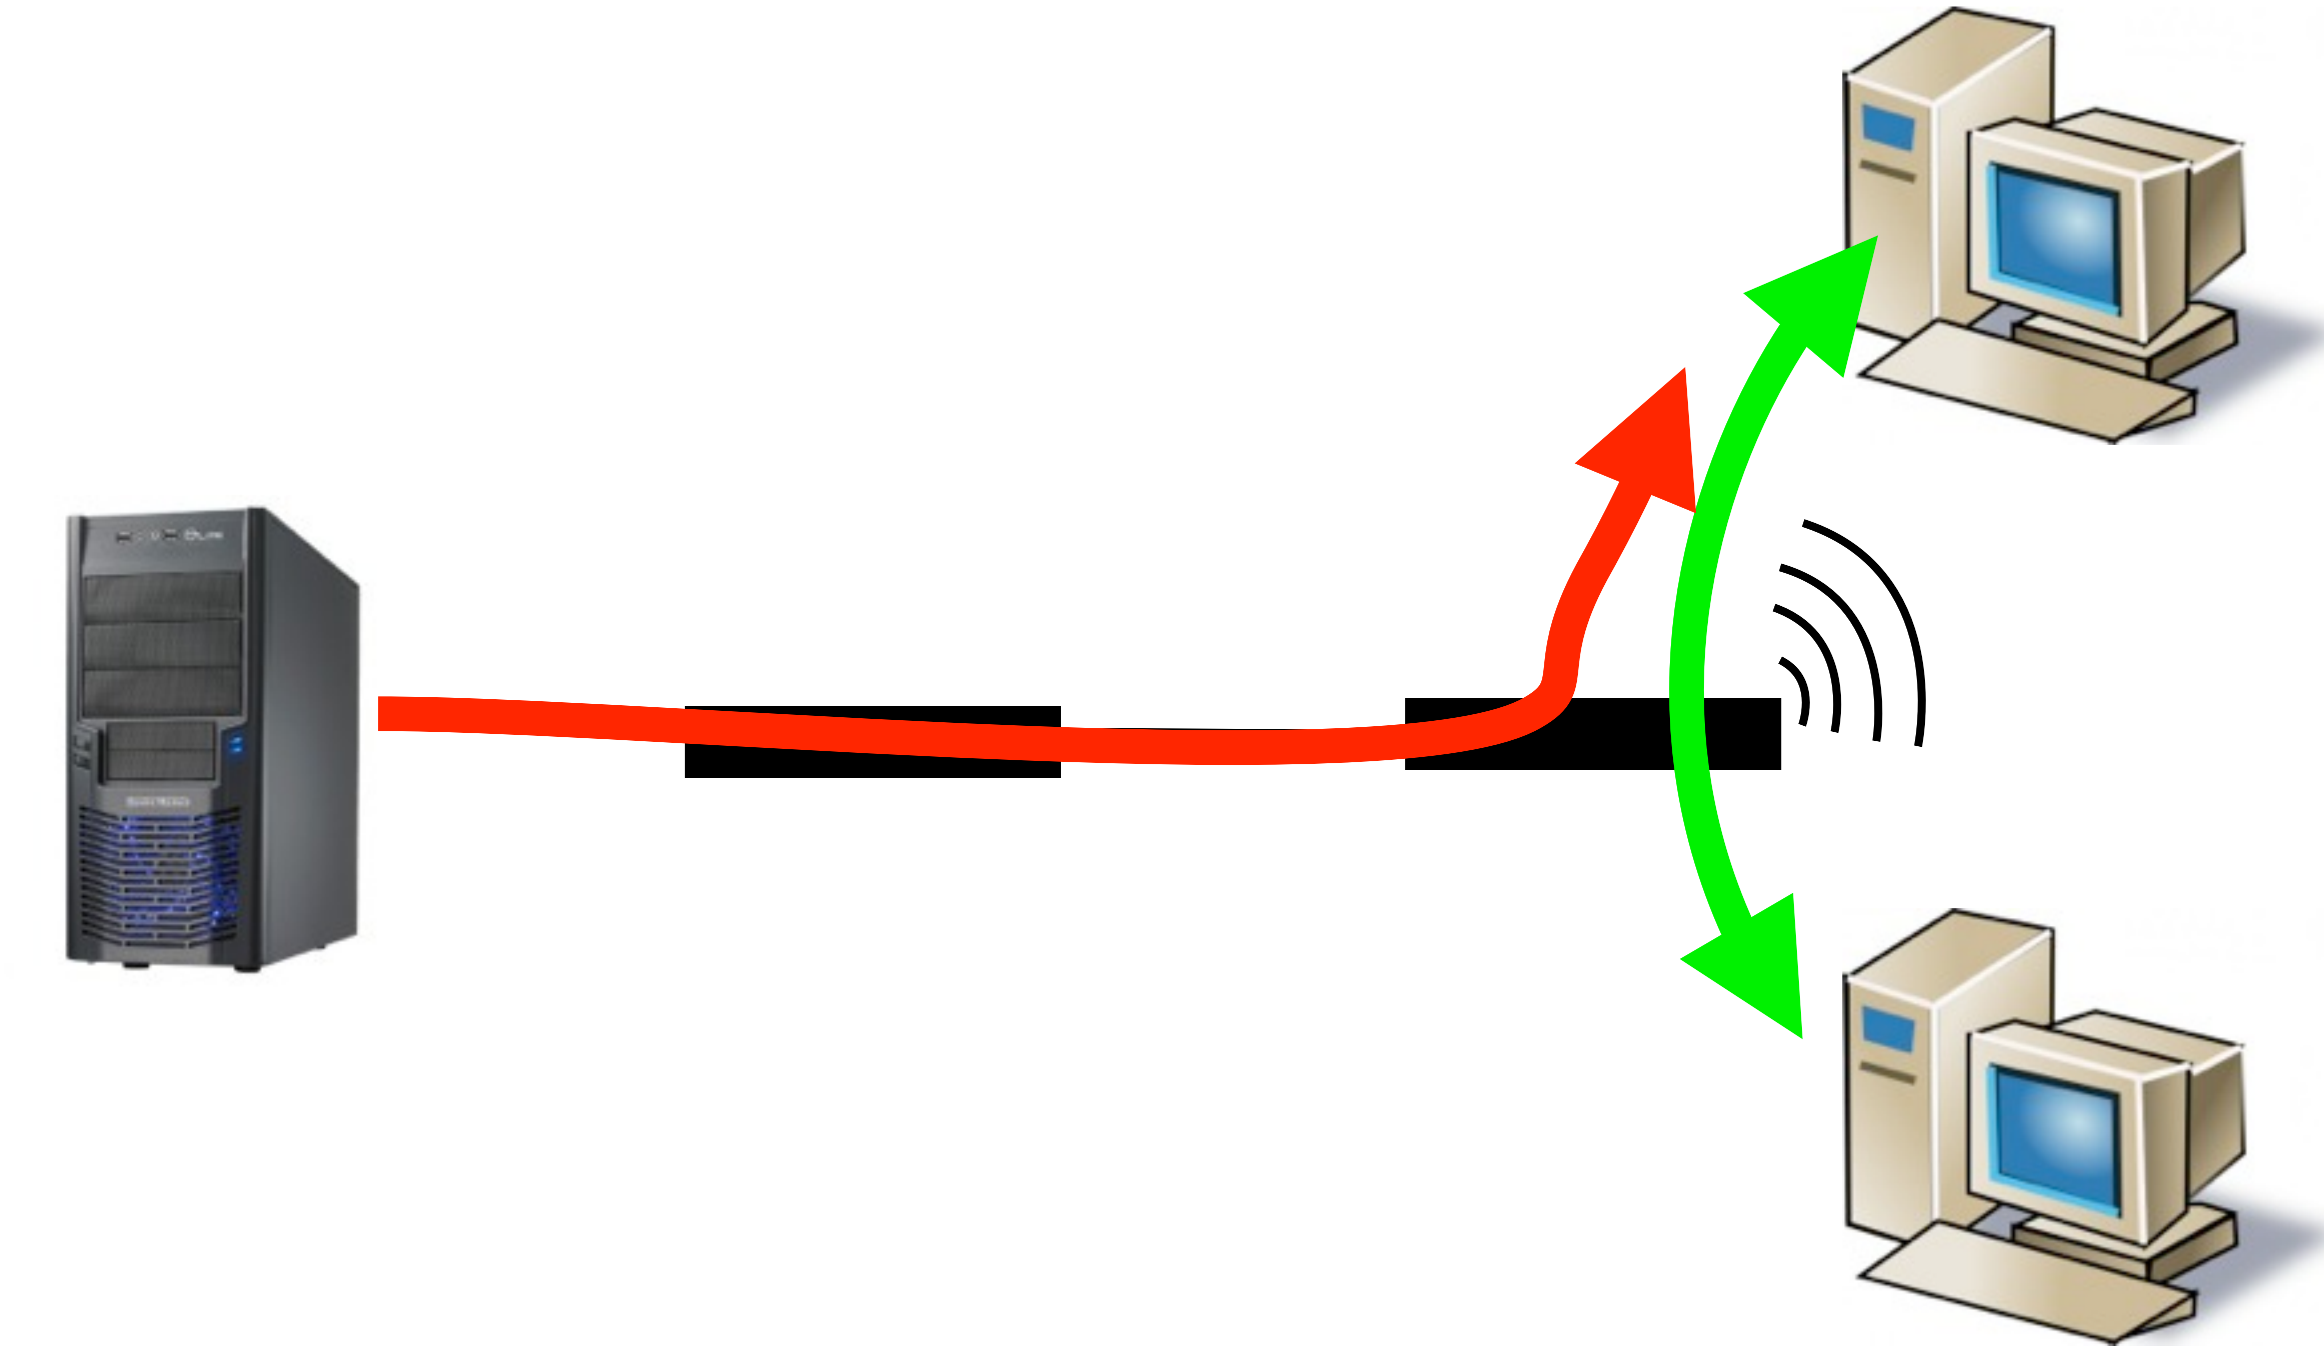

Denial-of-service

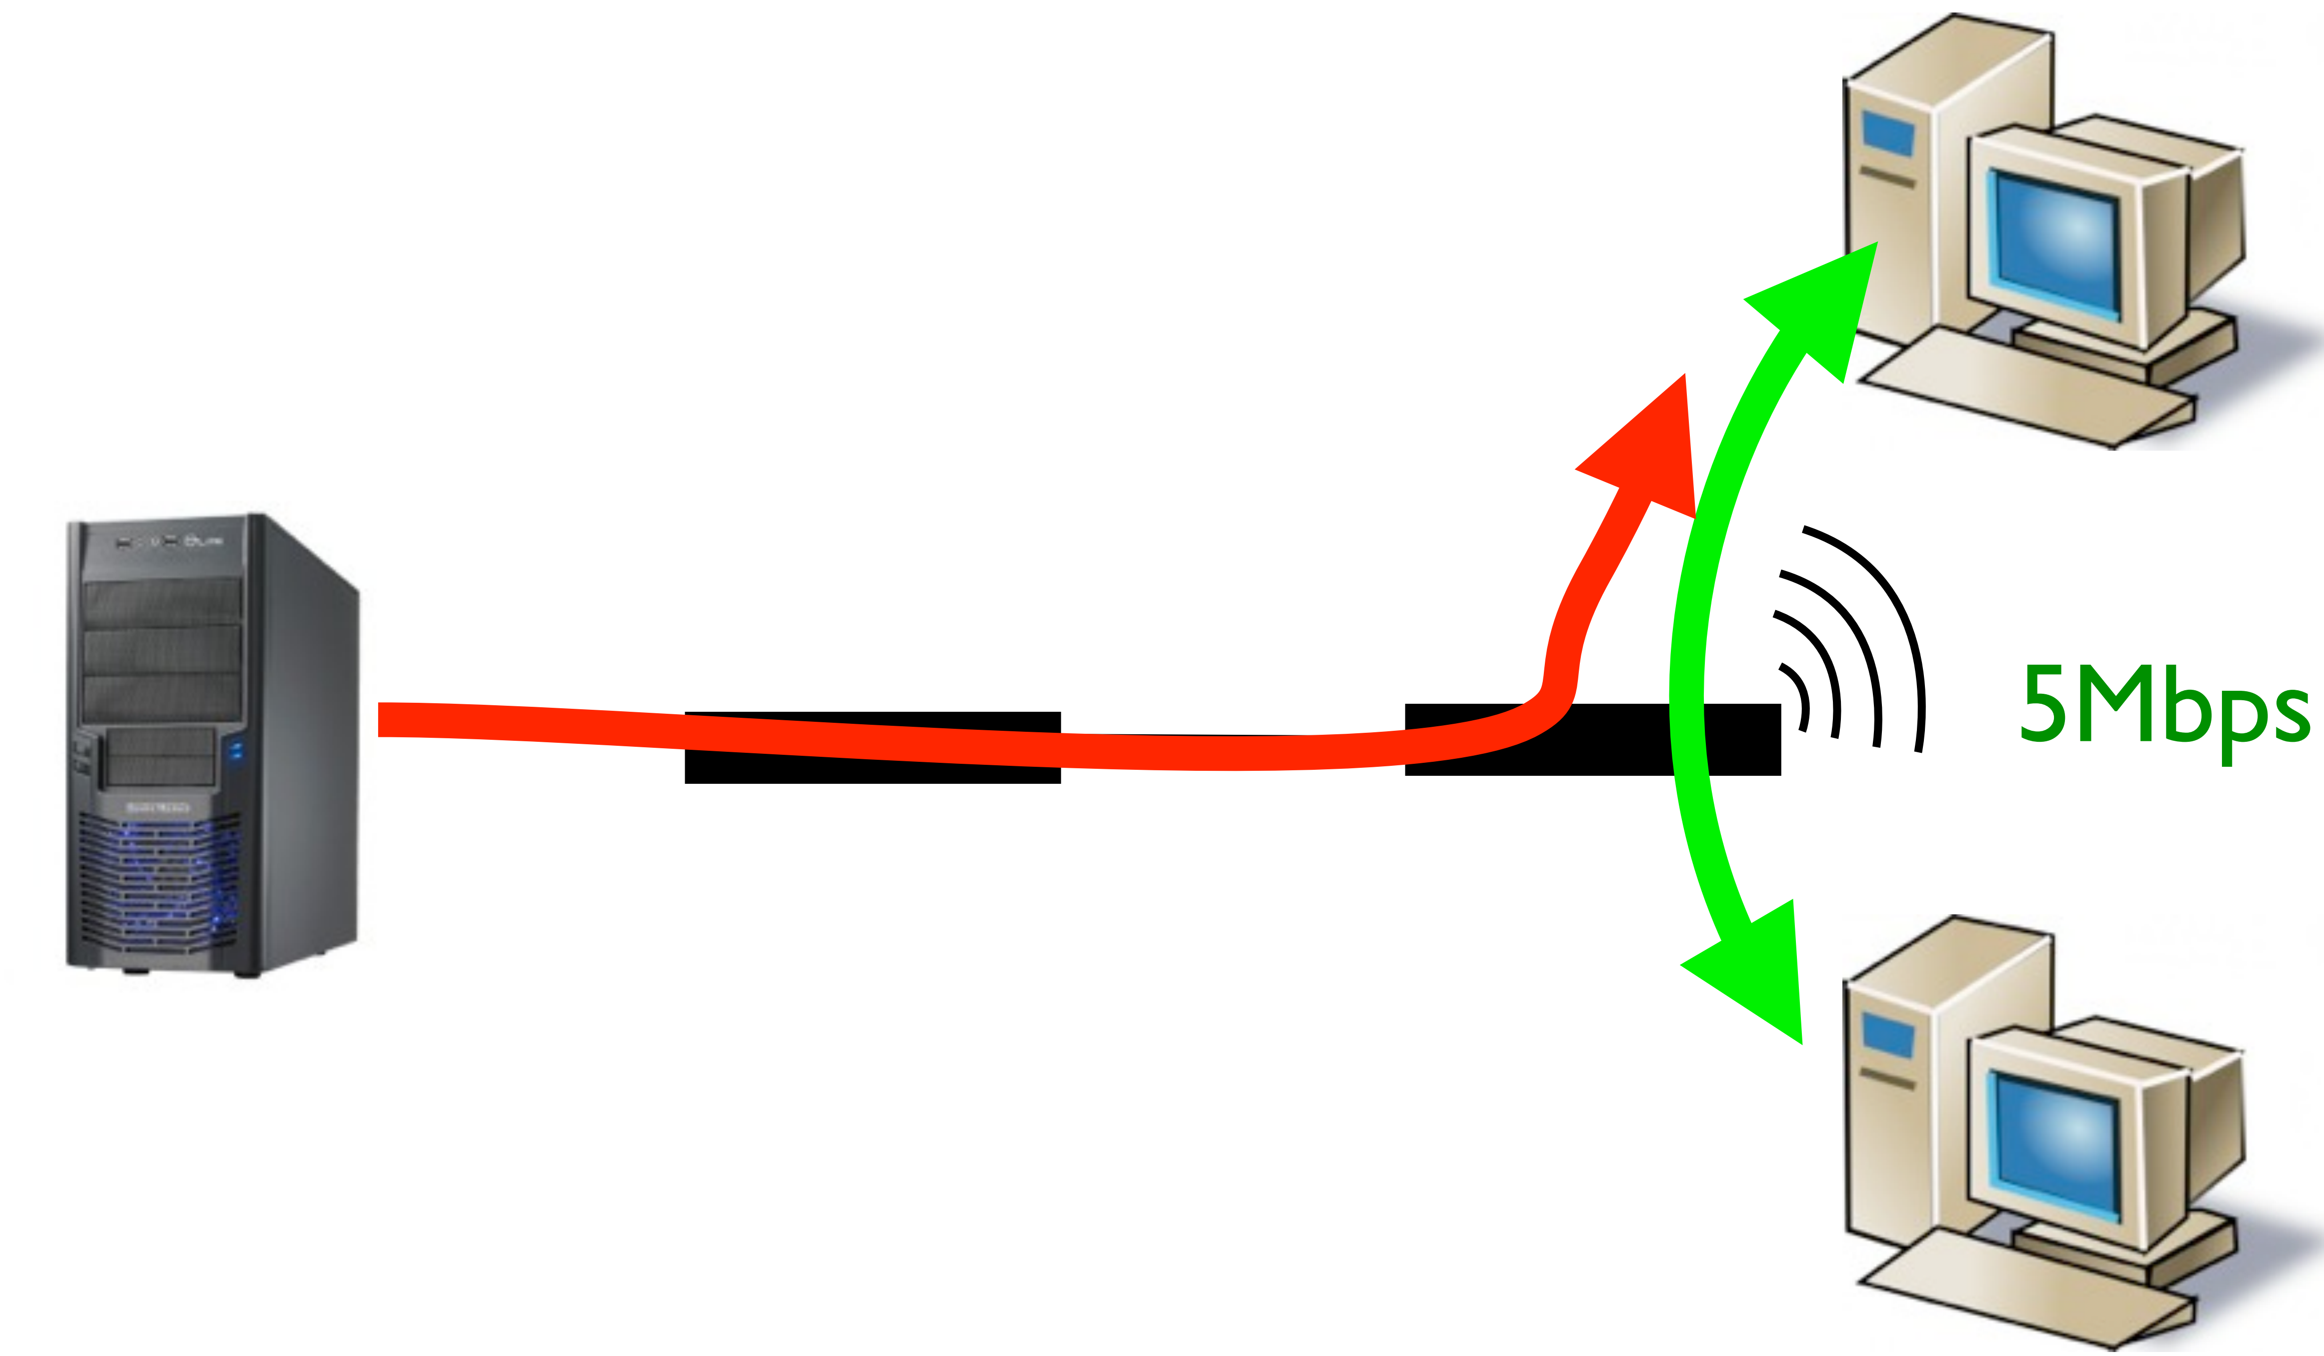

Denial-of-service

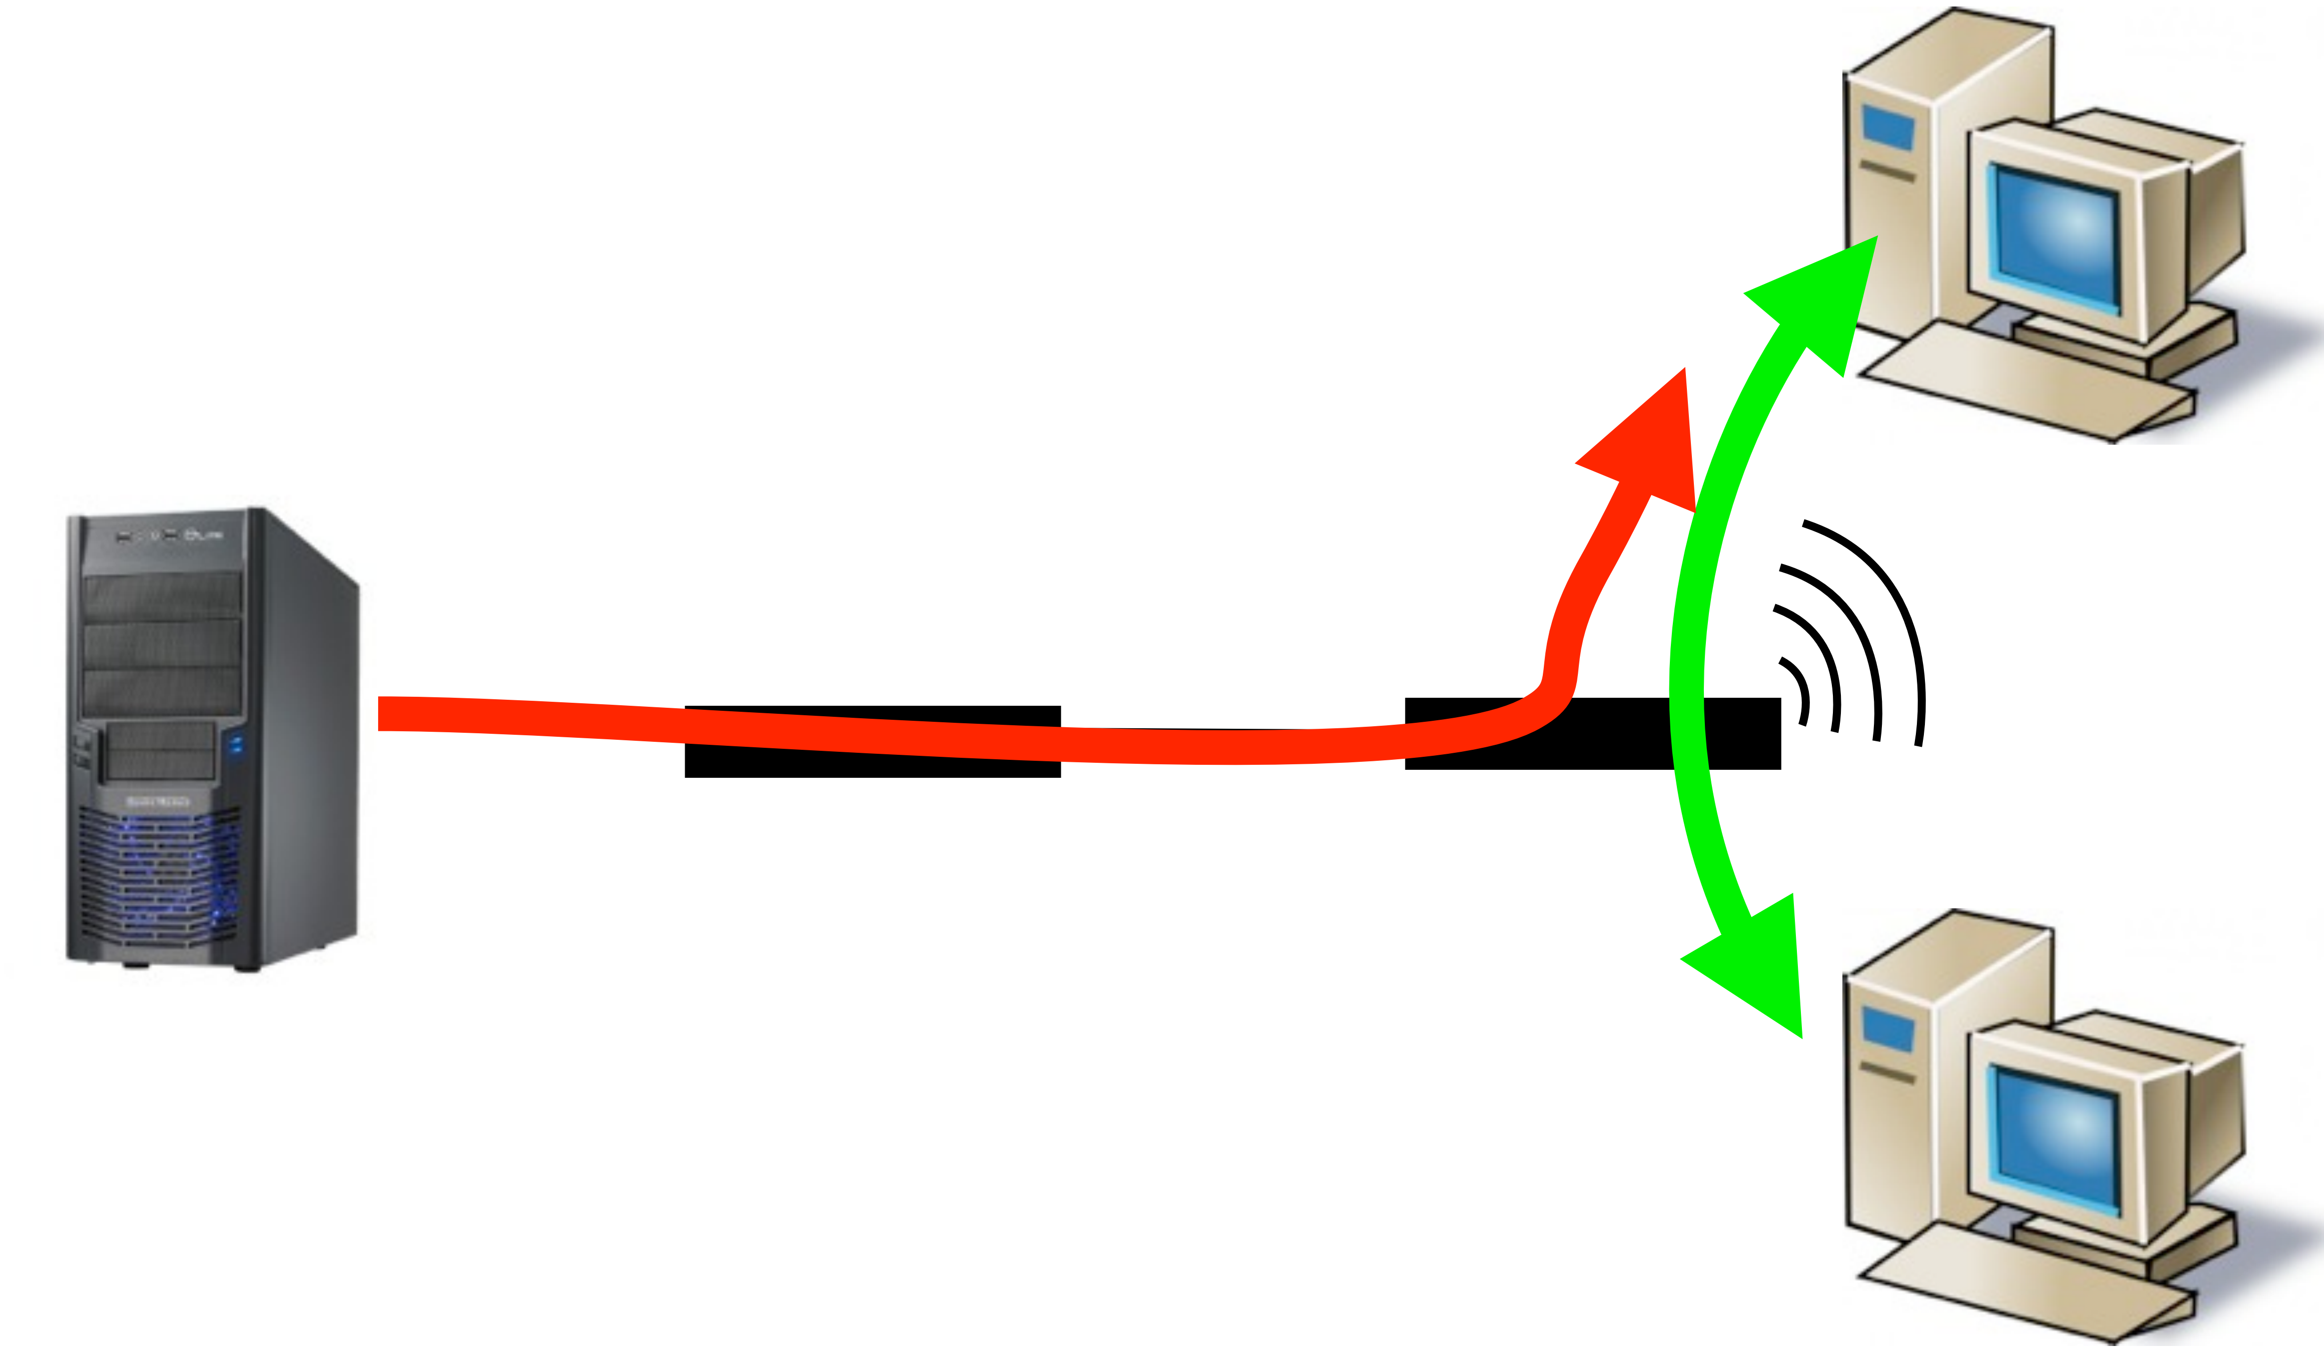

Denial-of-service

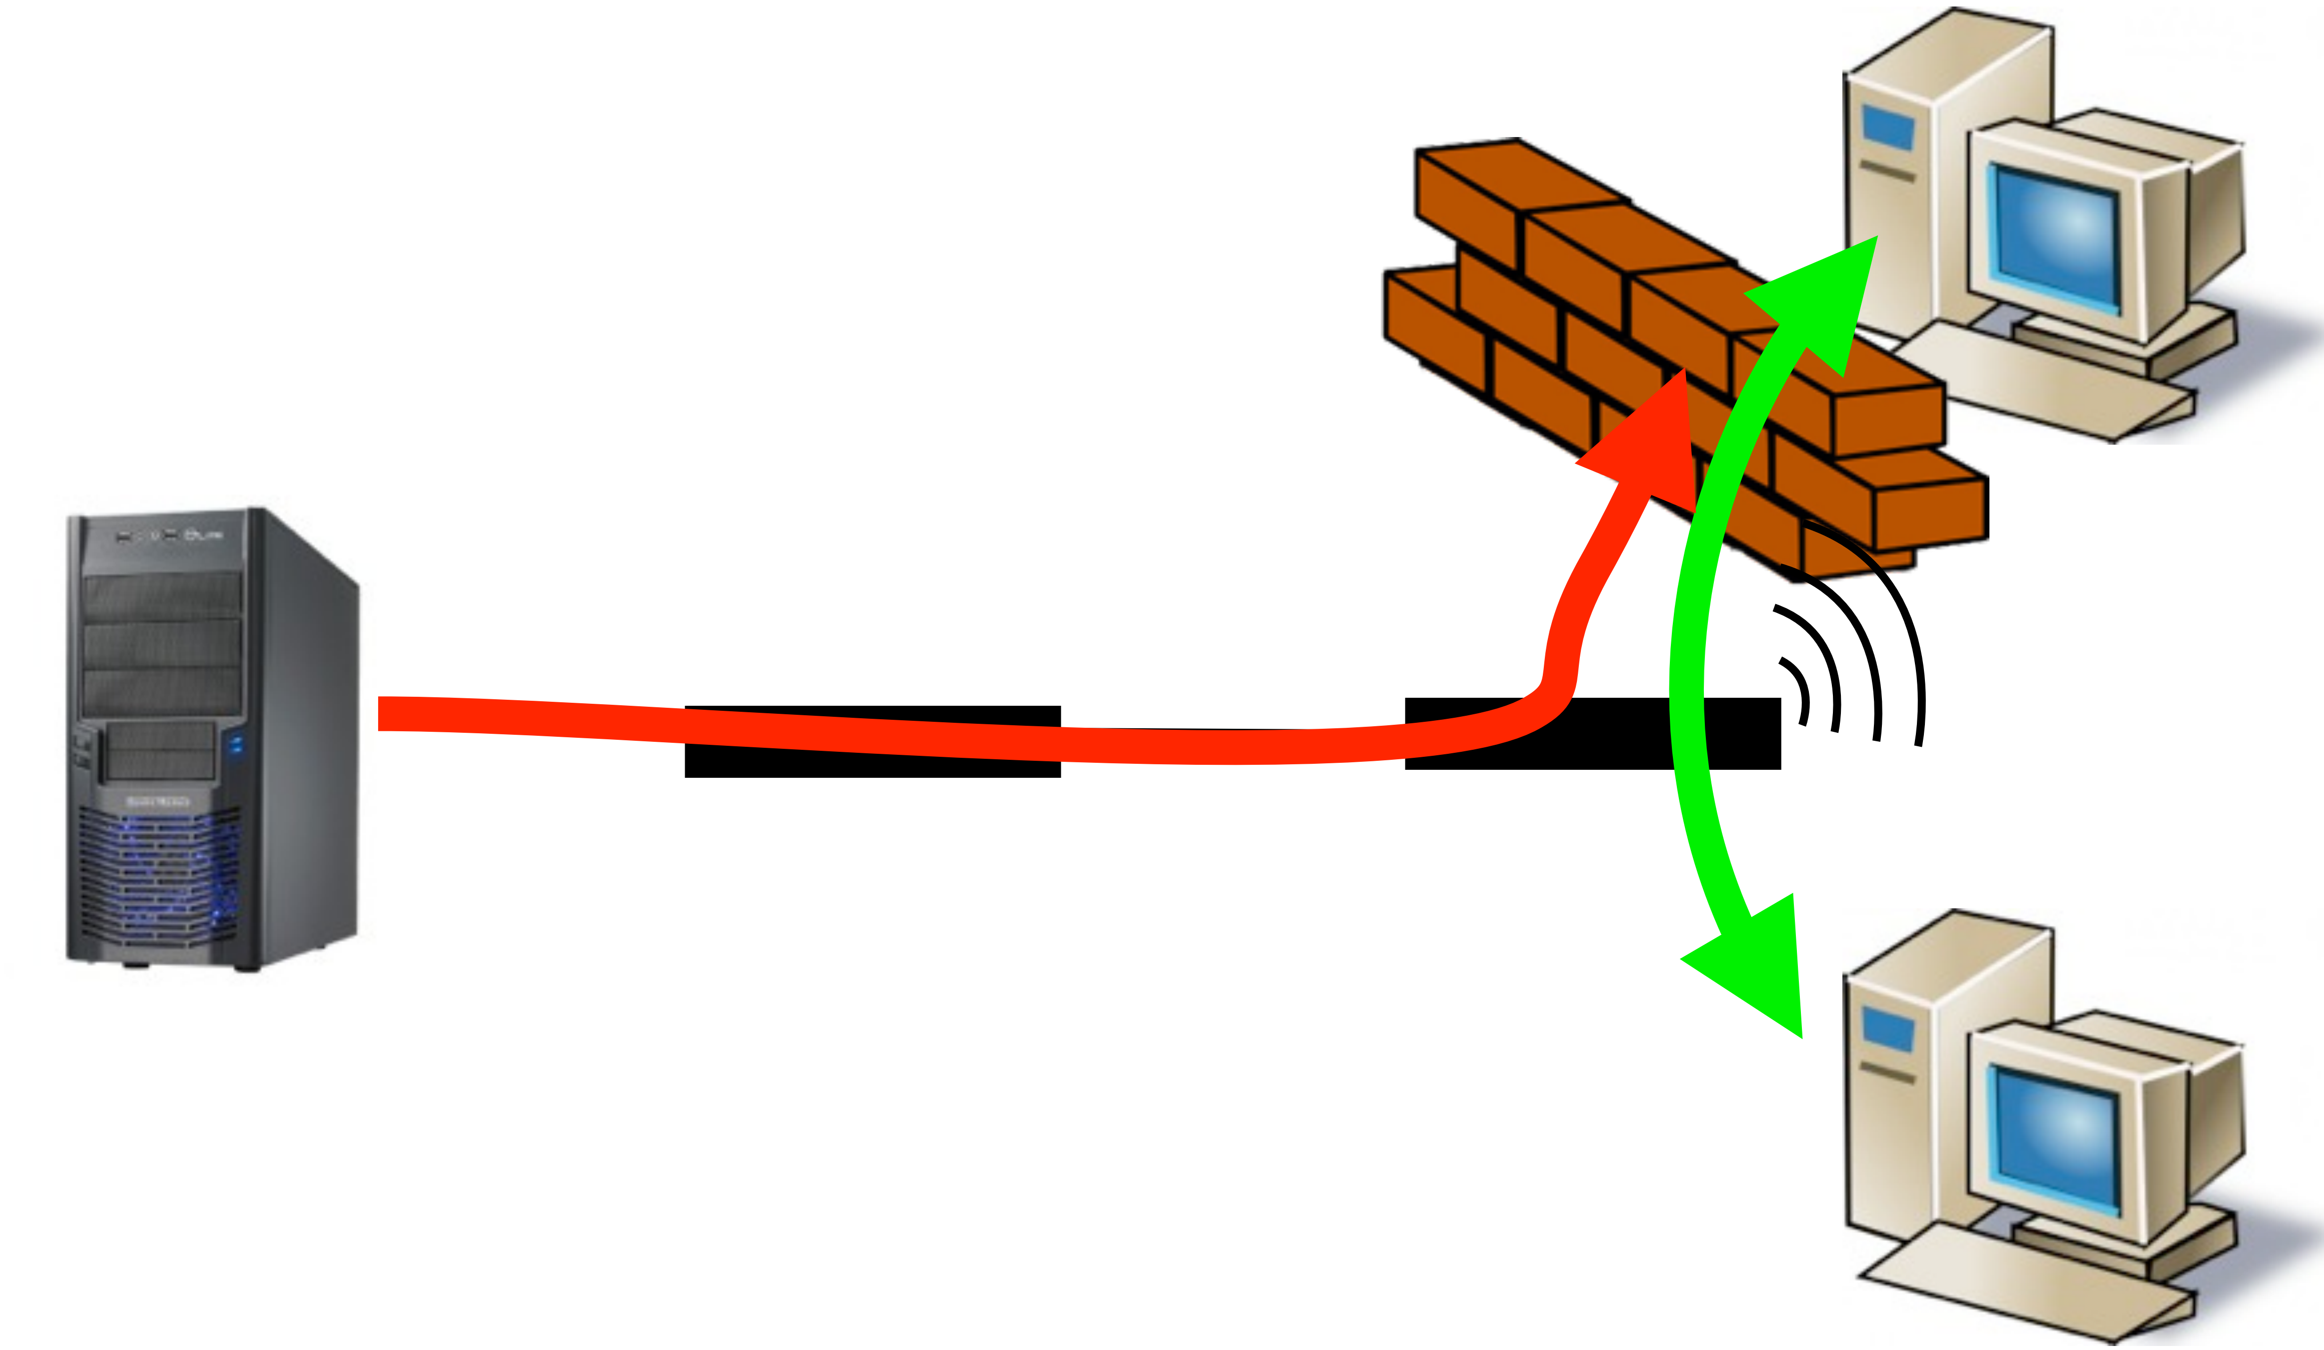

Denial-of-service

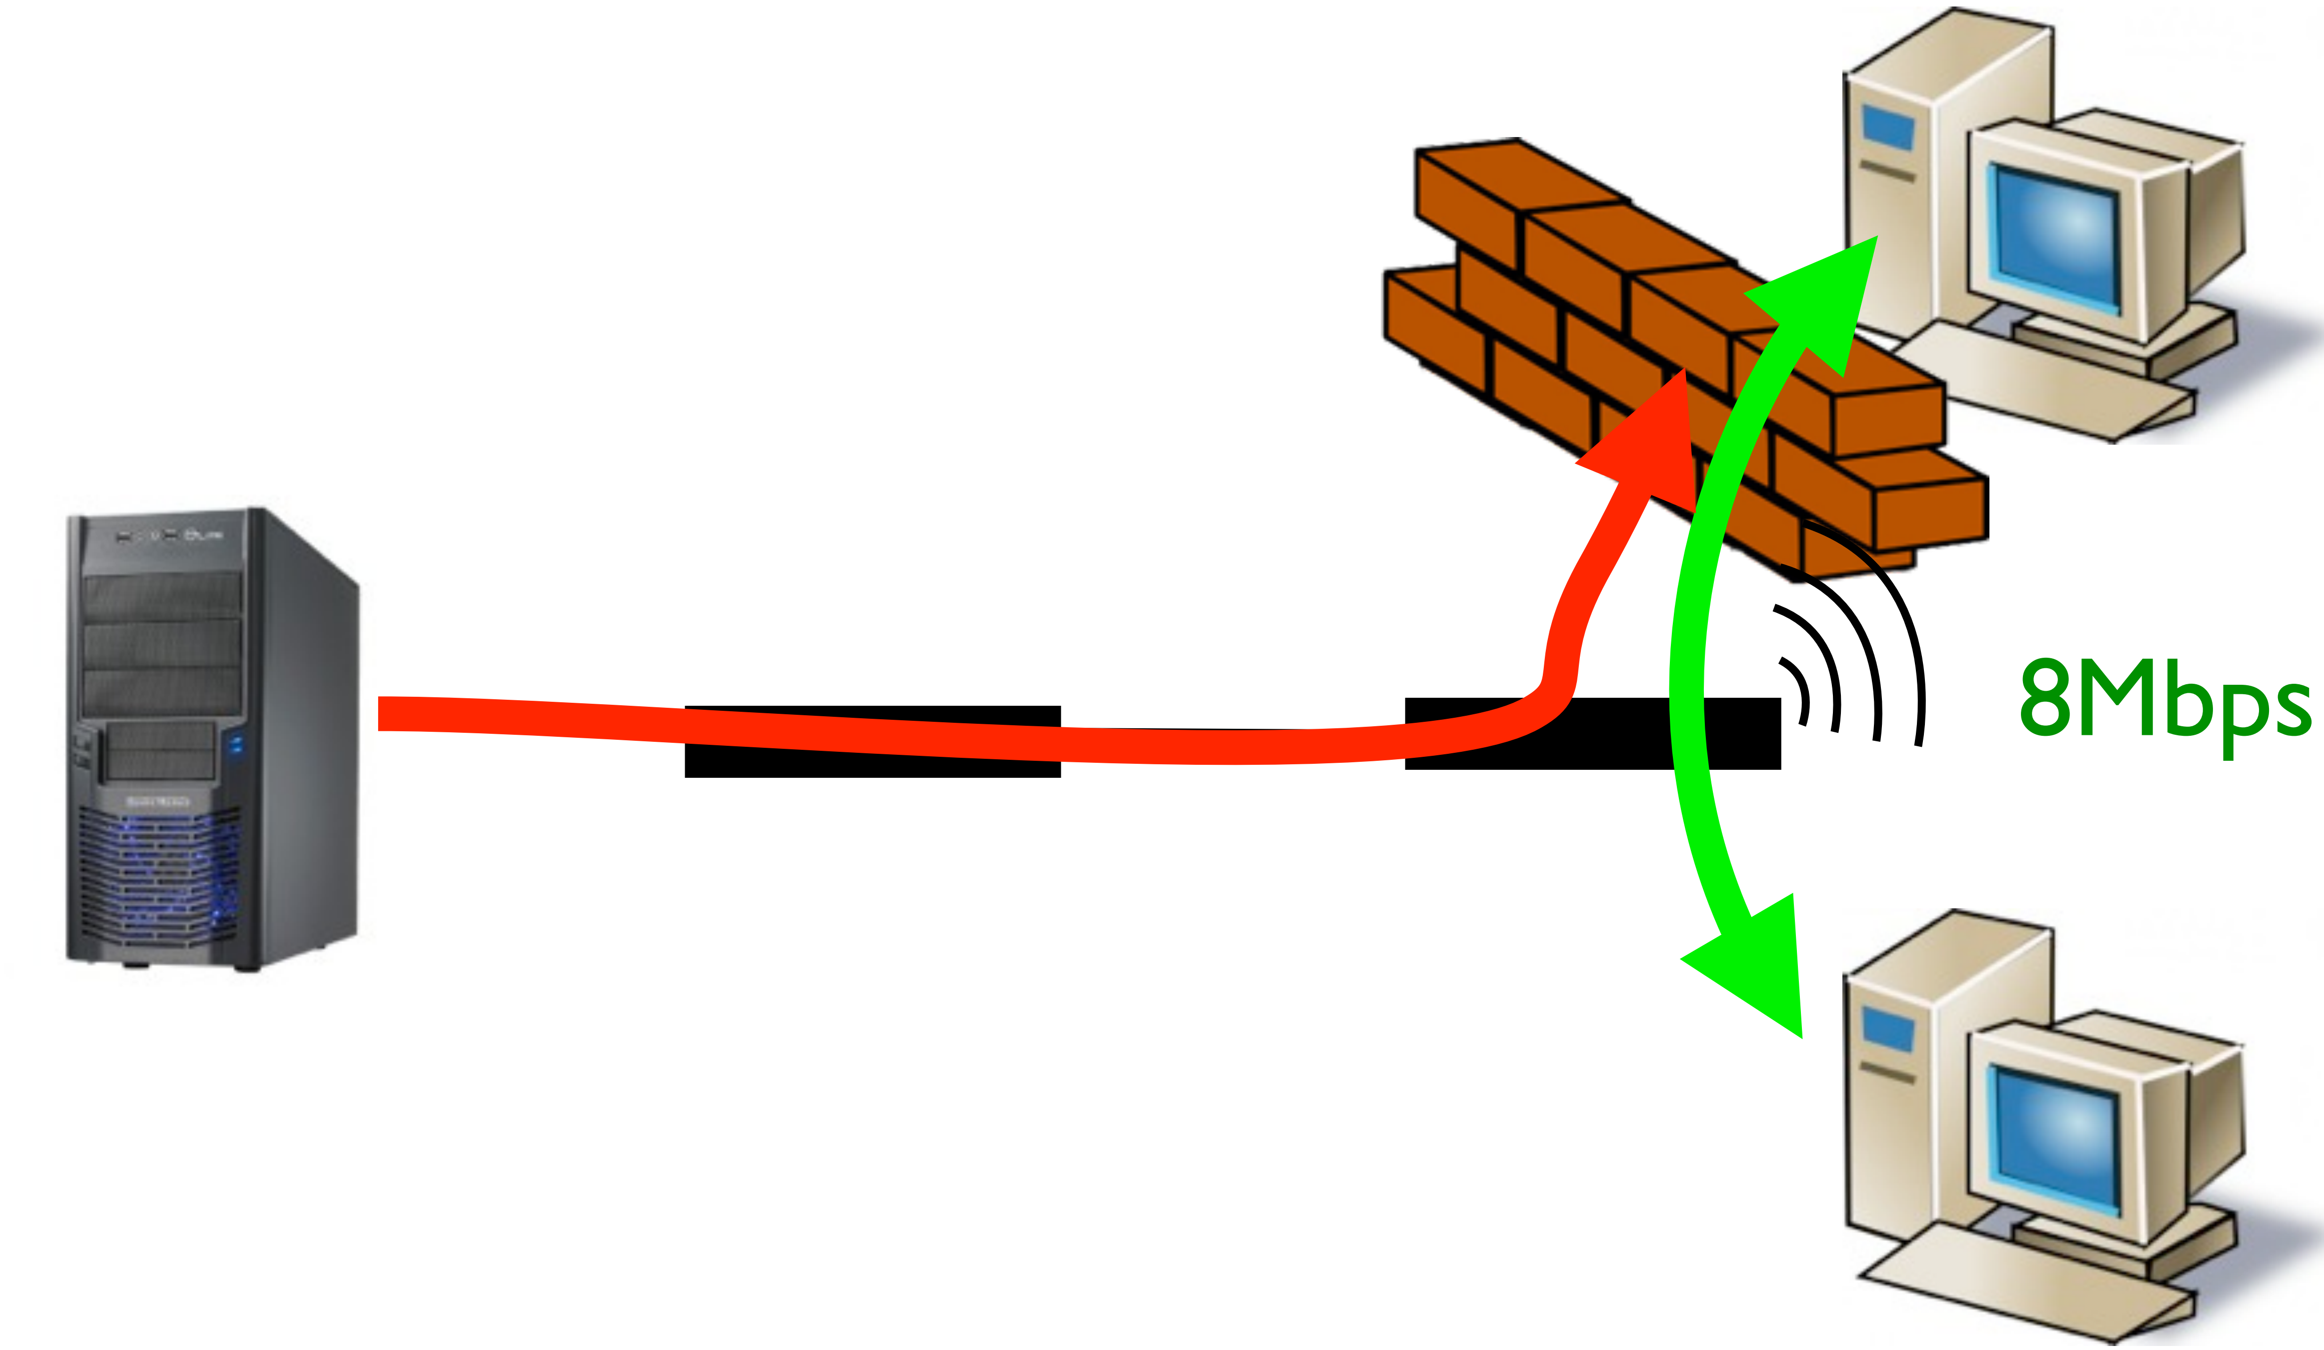

Denial-of-service

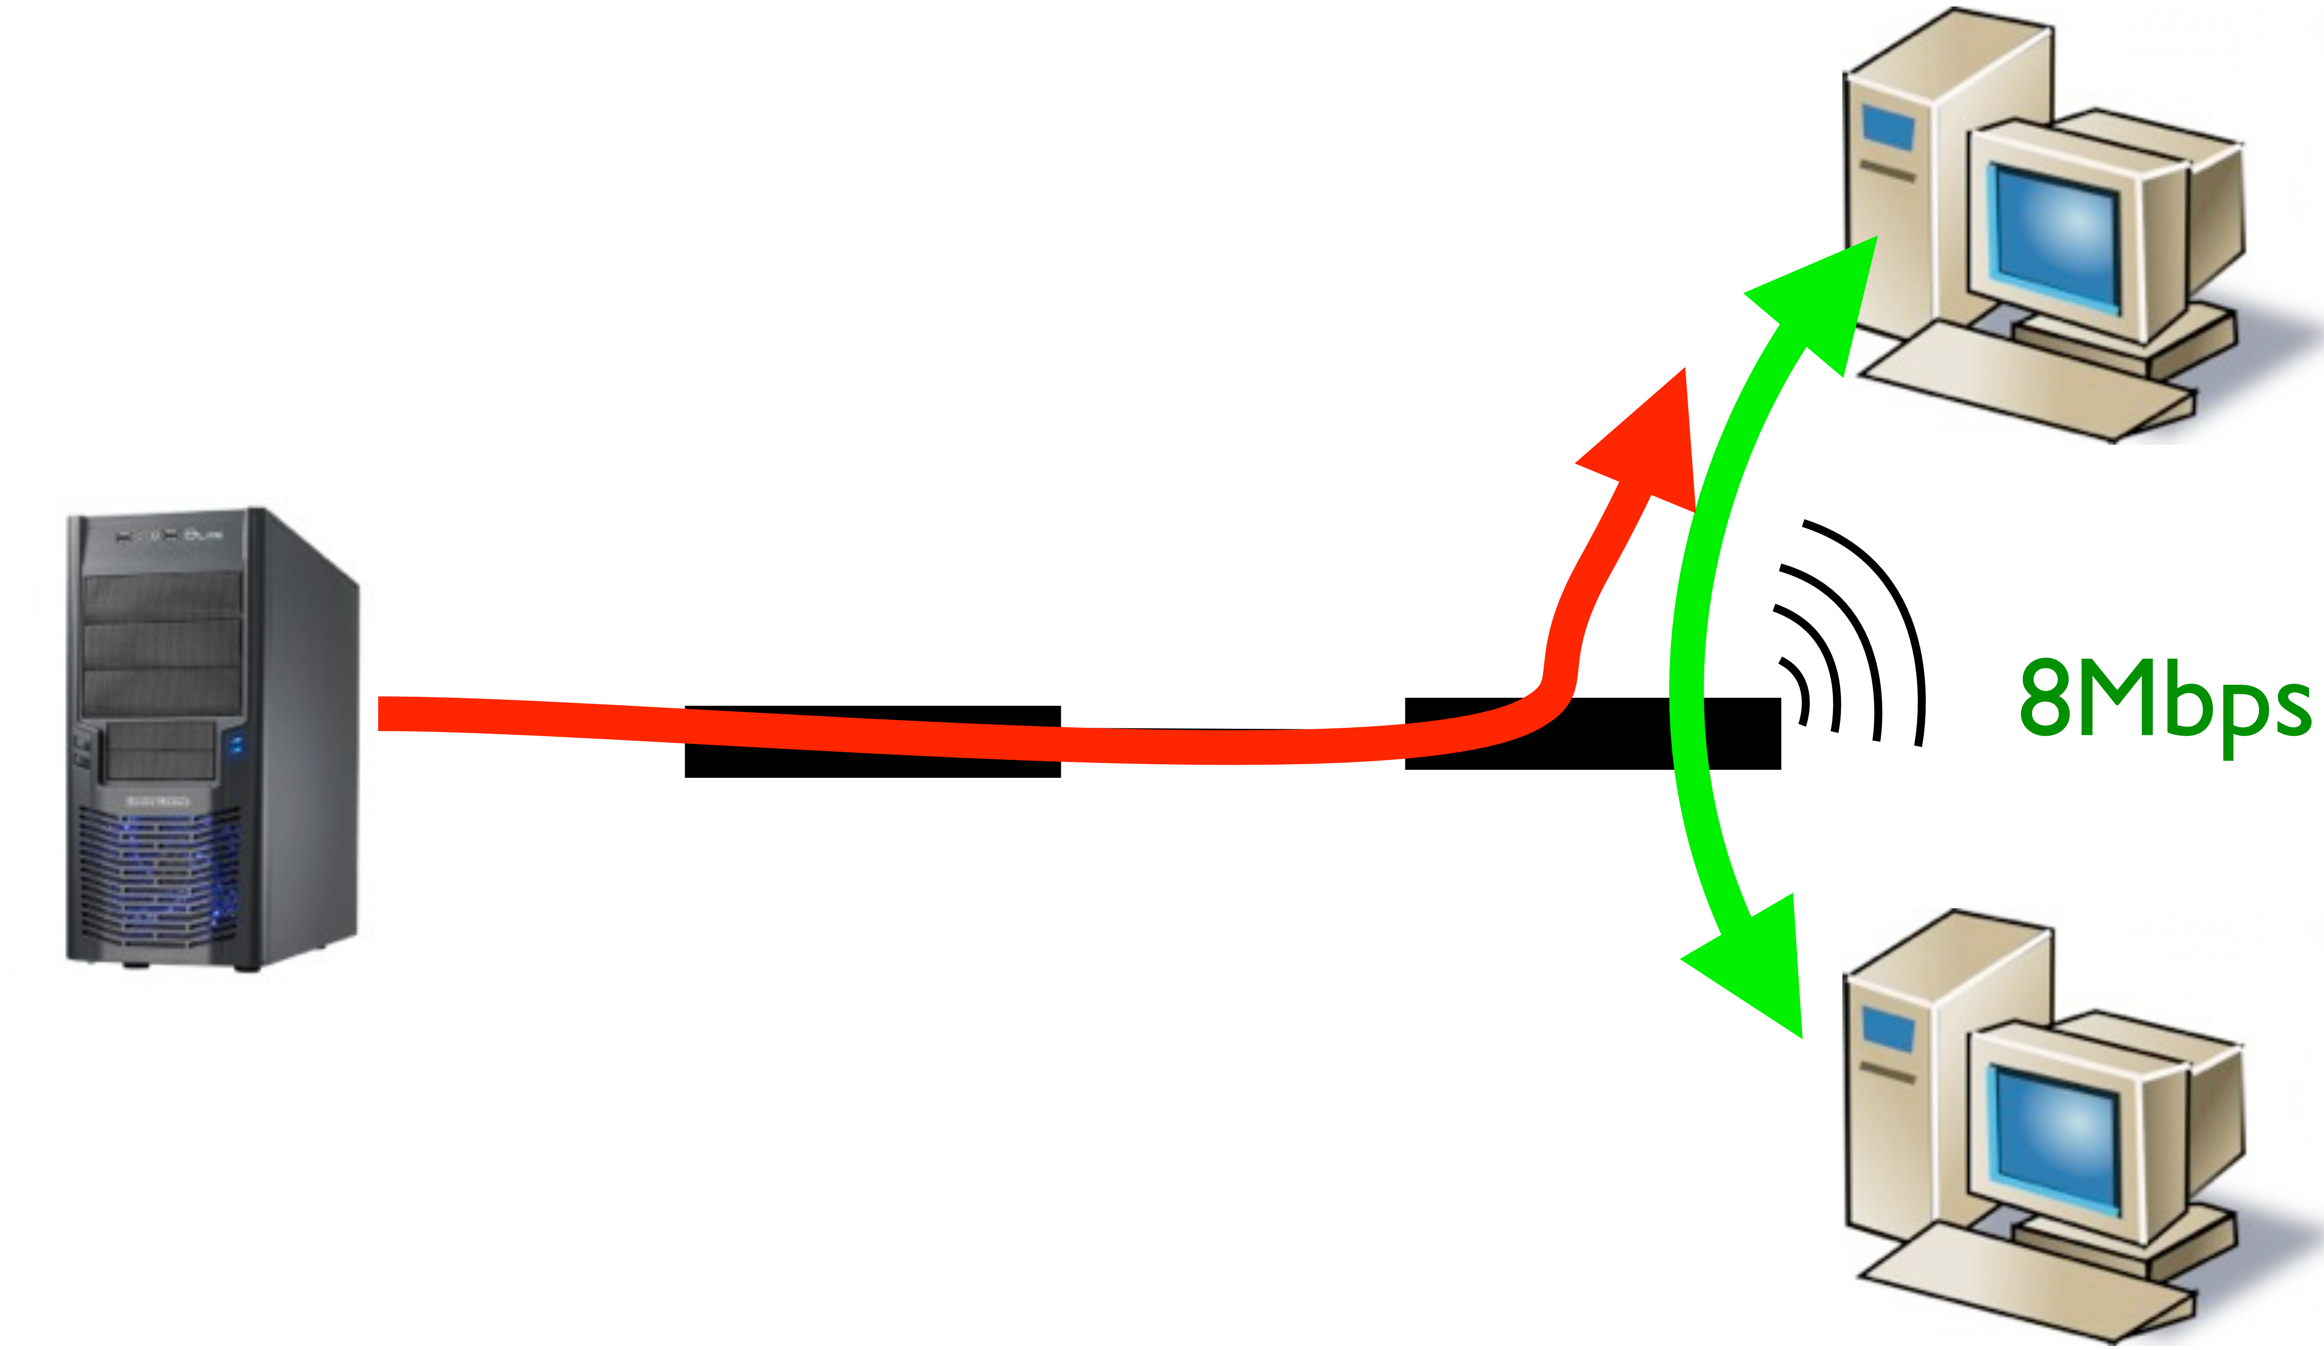

Denial-of-service

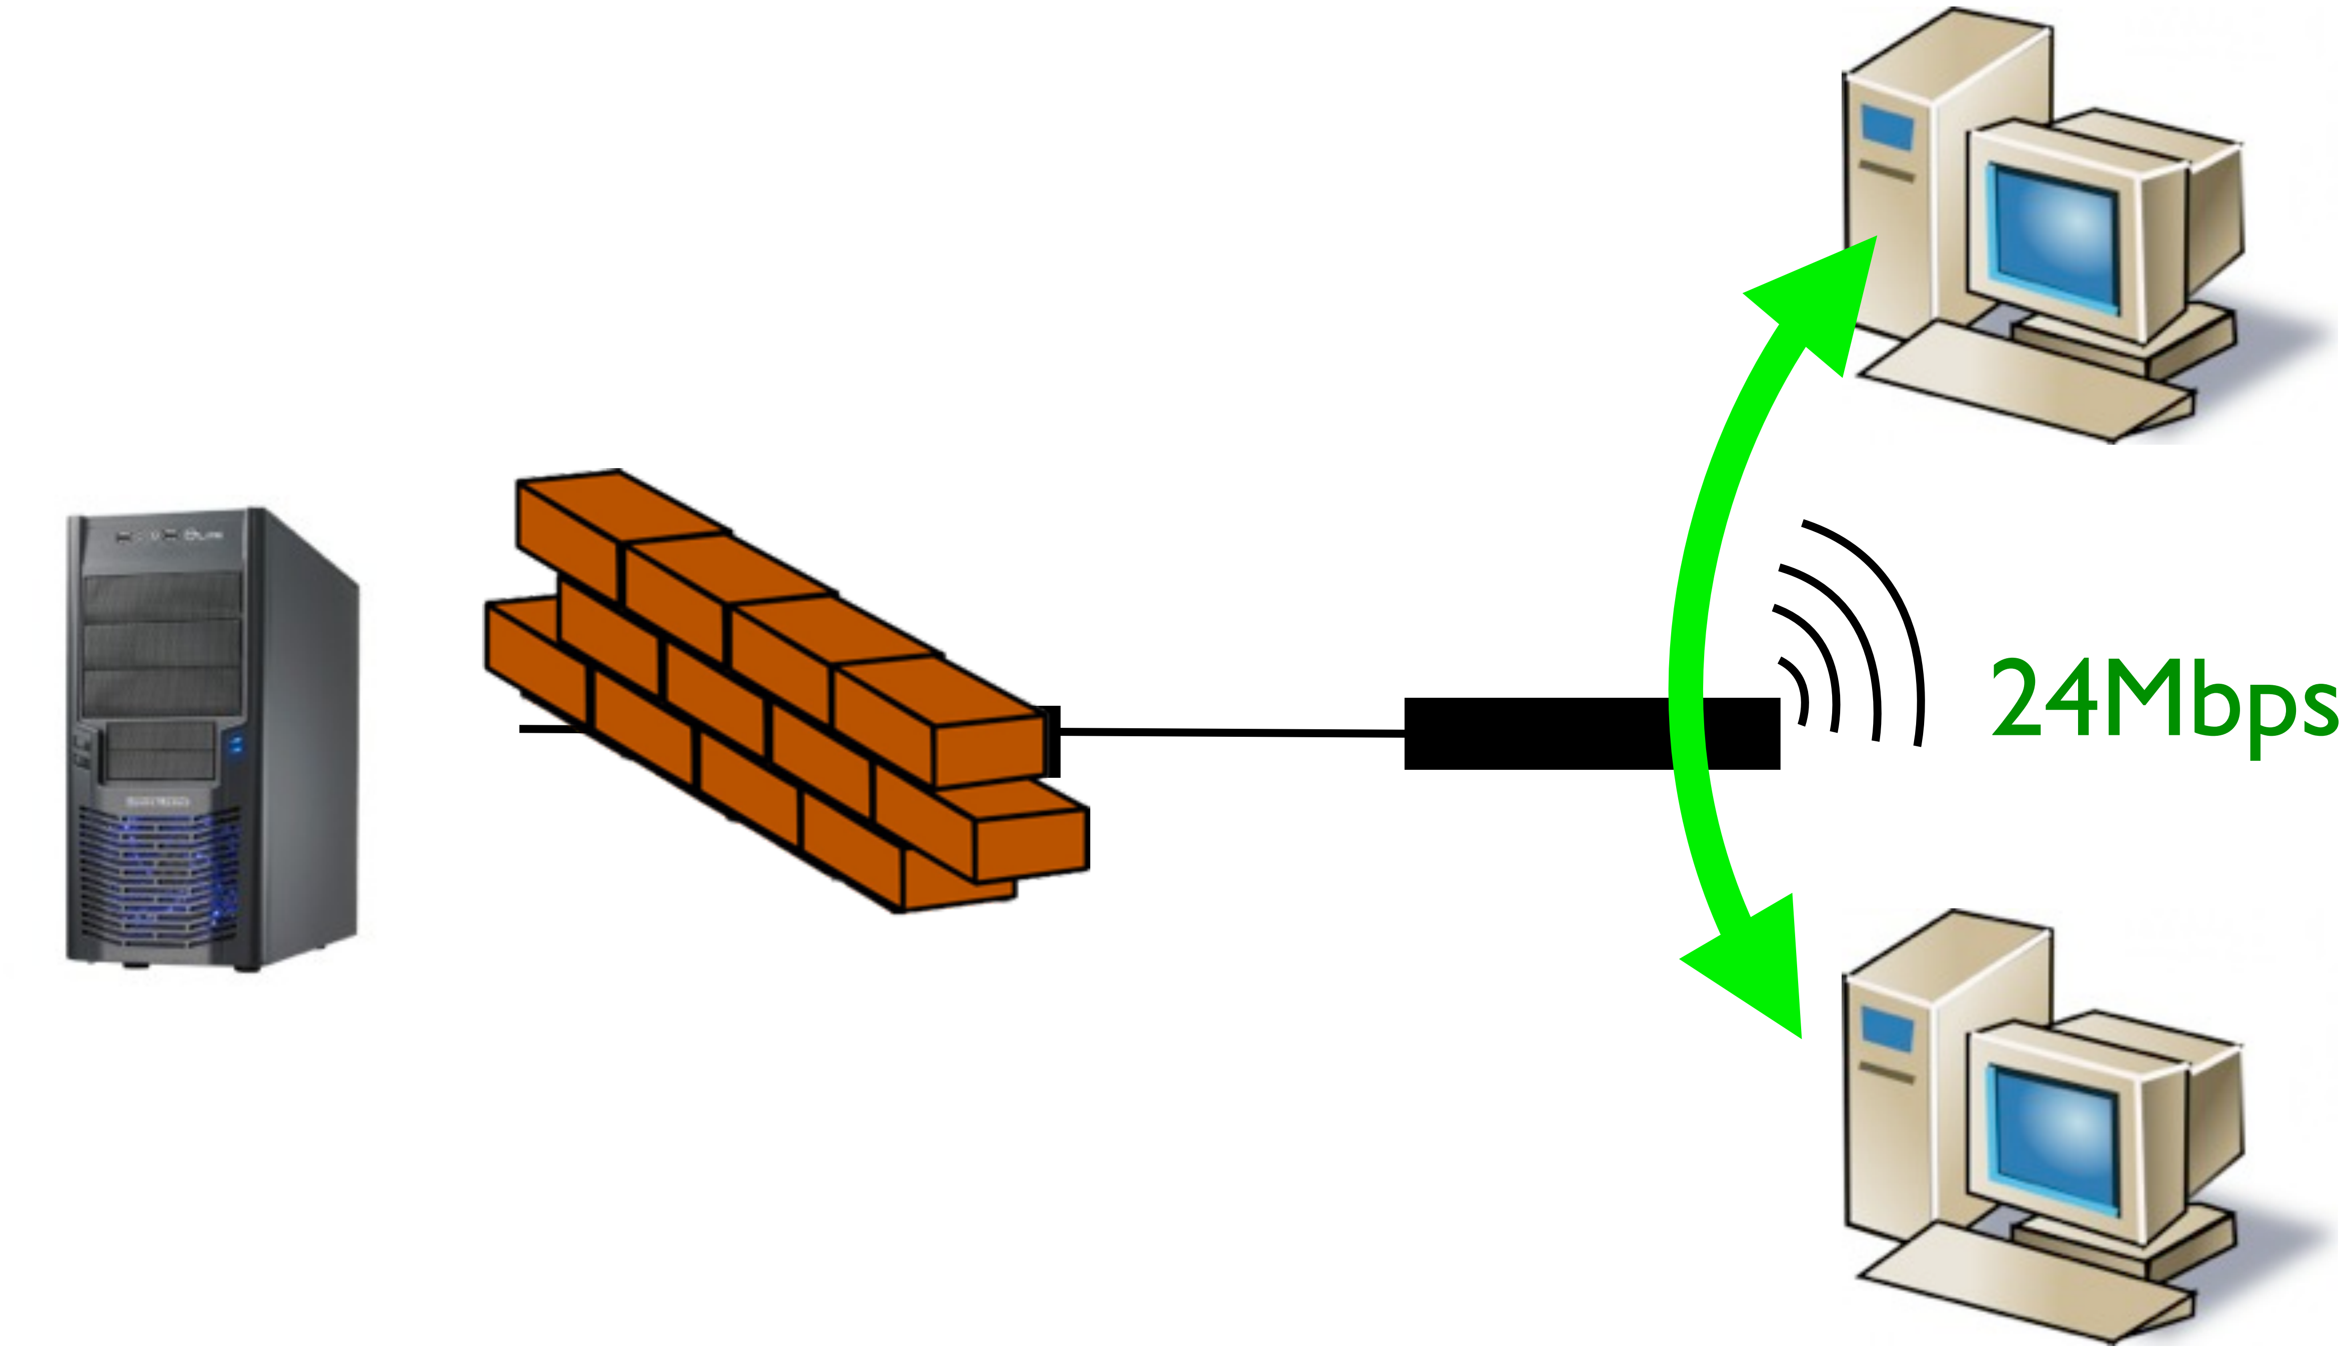

Denial-of-service

Where to go?

Implement more operators!

Where to go?

Implement more operators!

Guaranteed latency

Where to go?

Implement more operators!

Guaranteed latency

Rate limiting

Where to go?

Implement more operators!

Guaranteed latency

Rate limiting

Path properties

Where to go?

Implement more operators!

Guaranteed latency

Rate limiting

Path properties

Hints

Where to go?

Implement more operators!

Guaranteed latency

Rate limiting

Path properties

Hints

Queries

Where to go?

Implement more operators!

Guaranteed latency

Rate limiting

Path properties

Hints

Queries

Your application?

Where to go?

Implement more operators!

Guaranteed latency

Rate limiting

Path properties

Hints

Queries

Your application?

Build a market

Where to go?

Implement more operators!

Guaranteed latency

Rate limiting

Path properties

Hints

Queries

Your application?

Build a market

Scale

Where to go?

Conclusion

Application knowledge can help cloud provider

PANE is our first step in realizing this

Conclusion

```
3640-123#show
running-config
Building
configuration...
Current
configuration :
1432 bytes
version 12.3
service config
service timestamps
debug datetime msec
service timestamps
log datetime msec
service password-
encryption
hostname 3640-123
boot-start-marker
boot-end-marker
enable password 7
02050D4808095E731F
no aaa new-model
resource policy
voice-card 3
ip subnet-zero
ip cef
no ip dhcp use vrf
connected
!--- This is the
Cisco IOS Firewall
configuration.
!--- IN-OUT is the
inspection rule for
traffic that flows
!--- from the
inside interface of
the router to the
outside interface.
ip inspect name IN-
OUT tcp
ip inspect name IN-
OUT udp
ip inspect name IN-
OUT ftp
ip inspect name IN-
OUT http
ip inspect name IN-
OUT icmp
!--- OUT-IN is the
inspection rule for
traffic that flows
!--- from the
outside interface
of the router to
the inside
interface.
!--- This rule is
where SMTP/ESMTP
inspection is
specified.
ip inspect name
OUT-IN smtp
no ip ips deny-
action ips-
```

```
no ftp-server
write-enable
controller T1 3/0
 framing sf
 linecode ami
!--- The outside
interface.
interface
Ethernet2/0
 ip address
172.22.1.16
255.255.255.0
!--- Apply the
access list to
permit SMTP/ESMTP
connections
!--- to the mail
server. This also
allows Cisco IOS
Firewall
!--- to inspect
SMTP or ESMTP
commands.
 ip access-group
101 in
 ip nat outside
!--- Apply the
inspection rule
OUT-IN inbound on
this interface.
This is
!--- the rule that
defines SMTP/ESMTP
inspection.
 ip inspect OUT-IN
in
 ip virtual-
reassembly
 half-duplex
interface Serial2/0
 no ip address
 shutdown
!--- The inside
interface.
interface
Ethernet2/1
 ip address
10.10.10.1
255.255.255.0
 ip nat inside
!--- Apply the
inspection rule IN-
OUT inbound on this
interface.
 ip inspect IN-OUT
in
 ip virtual-
reassembly
 half-duplex
ip http server
no ip http secure-
server
ip classless
```

```
!--- The static
translation for the
mail server.
ip nat inside
source static
10.10.10.2
172.22.1.110
ip nat inside
source static
10.10.10.5
172.22.1.111
!--- The access
list to permit SMTP
and ESMTP to the
mail server.
!--- Cisco IOS
Firewall inspects
permitted traffic.
access-list 101
permit tcp any host
172.22.1.110 eq
smtp
control-plane
voice-port 1/0/0
voice-port 1/0/1
voice-port 1/1/0
voice-port 1/1/1
line con 0
line aux 0
line vty 0 4
 password 7
121A0C0411045D5679
 login
end
Current
configuration:
version 12.2
service timestamps
debug uptime
service timestamps
log uptime
no service
password-encryption
hostname sec-3640
aaa new-model
aaa group server
tacacs+ RTP
 server
171.68.120.214
aaa authentication
login default group
RTP none
aaa authorization
exec default group
RTP none
aaa authorization
auth-proxy default
group RTP
enable secret 5
$1$pqRI
$3TDNFT9FdYT8Sd/
q3S0VU1
enable password ww
```

```
ip inspect name
myfw cuseeme
timeout 3600
ip inspect name
myfw ftp timeout
3600
ip inspect name
myfw http timeout
3600
ip inspect name
myfw rcmd timeout
3600
ip inspect name
myfw realaudio
timeout 3600
ip inspect name
myfw smtp timeout
3600
ip inspect name
myfw sqlnet timeout
3600
ip inspect name
myfw streamworks
timeout 3600
ip inspect name
myfw tftp timeout
30
ip inspect name
myfw udp timeout 15
ip inspect name
myfw tcp timeout
3600
ip inspect name
myfw vdolive
ip auth-proxy auth-
proxy-banner
ip auth-proxy auth-
cache-time 10
ip auth-proxy name
list_a http
ip audit notify log
ip audit po max-
events 100
interface
Ethernet0/0
 ip address
40.31.1.144
255.255.255.0
ip access-group 116
in
 ip nat outside
ip auth-proxy
list_a
 no ip route-cache
 no ip mroute-cache
 speed auto
 half-duplex
 no mop enabled
interface
Ethernet1/0
 ip address
10.14.14.14
255.255.255.0
```

```
   the syslog
server or router.
ip urlfilter audit-
trail
!--- use the ip
urlfilter urlf-
server-log
   command in
global
configuration mode
to enable the
logging of
system messages on
the URL filtering
server.
ip urlfilter urlf-
server-log
!--- use the ip
urlfilter server
vendor command
   in global
configuration mode
to configure a
vendor server for
URL filtering.
   Here we have
configured a
websense server
for URL filtering
ip urlfilter server
vendor websense
192.168.15.15
no ftp-server
write-enable
!--- Below is the
basic interface
configuration
  on the router
interface
FastEthernet0
 ip address
192.168.5.10
255.255.255.0
 ip virtual-
reassembly
!--- use the ip
inspect command in
interface
   configuration
mode
to apply a set of
inspection rules to
an interface.
   Here the
inspection name
TEST is
applied to the
interface
FastEthernet0.
 ip inspect test in
 duplex auto
 speed auto
interface
```

```
 ip virtual-
reassembly
 duplex auto
 speed auto
interface
FastEthernet2
 ip address
10.77.241.109
255.255.255.192
 ip virtual-
reassembly
 duplex auto
 speed auto
interface
FastEthernet2
 no ip address
interface Vlan1
 ip address
10.77.241.111
255.255.255.192
 ip virtual-
reassembly
ip classless
ip route 10.10.10.0
255.255.255.0
172.17.1.2
ip route 10.77.0.0
255.255.0.0
10.77.241.65
!--- Configure the
below commands to
enable
   SDM access to
the cisco routers
ip http server
ip http
authentication
local
no ip http secure-
server
line con 0
line aux 0
line vty 0 4
 privilege level 15
 transport input
telnet ssh
end
version 12.4
service timestamps
debug datetime msec
service timestamps
log datetime msec
no service
password-encryption
hostname 2851-cme2
logging message-
counter syslog
logging buffered
51200 warnings
no aaa new-model
clock timezone mst
-7
clock summer-time
```

```
no ip dhcp use vrf
connected
ip dhcp pool
pub-112-net
   network
172.17.112.0
255.255.255.0
   default-router
172.17.112.1
   dns-server
172.16.1.22
   option 150 ip
172.16.1.43
   domain-name
bldrtme.com
ip dhcp pool
priv-112-net
   network
192.168.112.0
255.255.255.0
   default-router
192.168.112.1
   dns-server
172.16.1.22
   domain-name
bldrtme.com
   option 150 ip
192.168.112.1
ip domain name
yourdomain.com
no ipv6 cef
multilink bundle-
name authenticated
voice translation-
rule 1
 rule 1 // /1001/
voice translation-
profile default
 translate called 1
voice-card 0
 no dspfarm
interface
GigabitEthernet0/0
 description $ETH-
LAN$$ETH-SW-LAUNCH$
$INTF-INFO-GE 0/0$
 ip address
172.16.112.10
255.255.255.0
 ip nat outside
 ip virtual-
reassembly
 duplex auto
 speed auto
interface
GigabitEthernet0/1
 no ip address
 duplex auto
 speed auto
interface
GigabitEthernet0/1.
132
 encapsulation
```

```
interface
GigabitEthernet0/1.
152
 encapsulation
dot1Q 152
 ip address
192.168.112.1
255.255.255.0
 ip nat inside
 ip virtual-
reassembly
interface
FastEthernet0/2/0
interface
FastEthernet0/2/1
interface
FastEthernet0/2/2
interface
FastEthernet0/2/3
interface Vlan1
 ip address
198.41.9.15
255.255.255.0
router eigrp 1
 network
172.16.112.0
0.0.0.255
 network
172.17.112.0
0.0.0.255
 no auto-summary
ip forward-protocol
nd
ip http server
ip http access-
class 23
ip http
authentication
local
ip http secure-
server
ip http timeout-
policy idle 60 life
86400 requests
10000
ip http path
flash:/gui
ip nat inside
source list 111
interface
GigabitEthernet0/0
overload
access-list 23
permit 10.10.10.0
0.0.0.7
access-list 111
deny   ip
192.168.112.0
0.0.0.255
192.168.0.0
0.0.255.255
access-list 111
permit ip
192.168.112.0
```

99%

```
3640-123#show
running-config
Building
configuration...
Current
configuration :
1432 bytes
version 12.3
service config
service timestamps
debug datetime msec
service timestamps
log datetime msec
service password-
encryption
hostname 3640-123
boot-start-marker
boot-end-marker
enable password 7
02050D4808095E731F
no aaa new-model
resource policy
voice-card 3
ip subnet-zero
ip cef
no ip dhcp use vrf
connected
!--- This is the
Cisco IOS Firewall
configuration.
!--- IN-OUT is the
inspection rule for
traffic that flows
!--- from the
inside interface of
the router to the
outside interface.
ip inspect name IN-
OUT tcp
ip inspect name IN-
OUT udp
ip inspect name IN-
OUT ftp
ip inspect name IN-
OUT http
ip inspect name IN-
OUT icmp
!--- OUT-IN is the
inspection rule for
traffic that flows
!--- from the
outside interface
of the router to
the inside
interface.
!--- This rule is
where SMTP/ESMTP
inspection is
specified.
ip inspect name
OUT-IN smtp
no ip ips deny-
action ips-
```

```
no ftp-server
write-enable
controller T1 3/0
 framing sf
 linecode ami
!--- The outside
interface.
interface
Ethernet2/0
 ip address
172.22.1.16
255.255.255.0
!--- Apply the
access list to
permit SMTP/ESMTP
connections
!--- to the mail
server. This also
allows Cisco IOS
Firewall
!--- to i
SMTP or ES
commands.
 ip access
101 in
 ip nat out
!--- Apply
inspection
OUT-IN inbou
this interfa
This is
!--- the rule
defines SMTP/
inspection.
 ip inspect OU
in
 ip virtual-
reassembly
 half-duplex
interface Seria
 no ip address
 shutdown
!--- The inside
interface.
interface
Ethernet2/1
 ip address
10.10.10.1
255.255.255.0
 ip nat inside
!--- Apply the
inspection rule IN-
OUT inbound on this
interface.
 ip inspect IN-OUT
in
 ip virtual-
reassembly
 half-duplex
ip http server
no ip http secure-
server
ip classless
```

```
!--- The static
translation for the
mail server.
ip nat inside
source static
10.10.10.2
172.22.1.110
ip nat inside
source static
10.10.10.5
172.22.1.111
!--- The access
list to permit SMTP
and ESMTP to the
mail server.
!--- Cisco IOS
Firewall inspects
permitted traffi
```

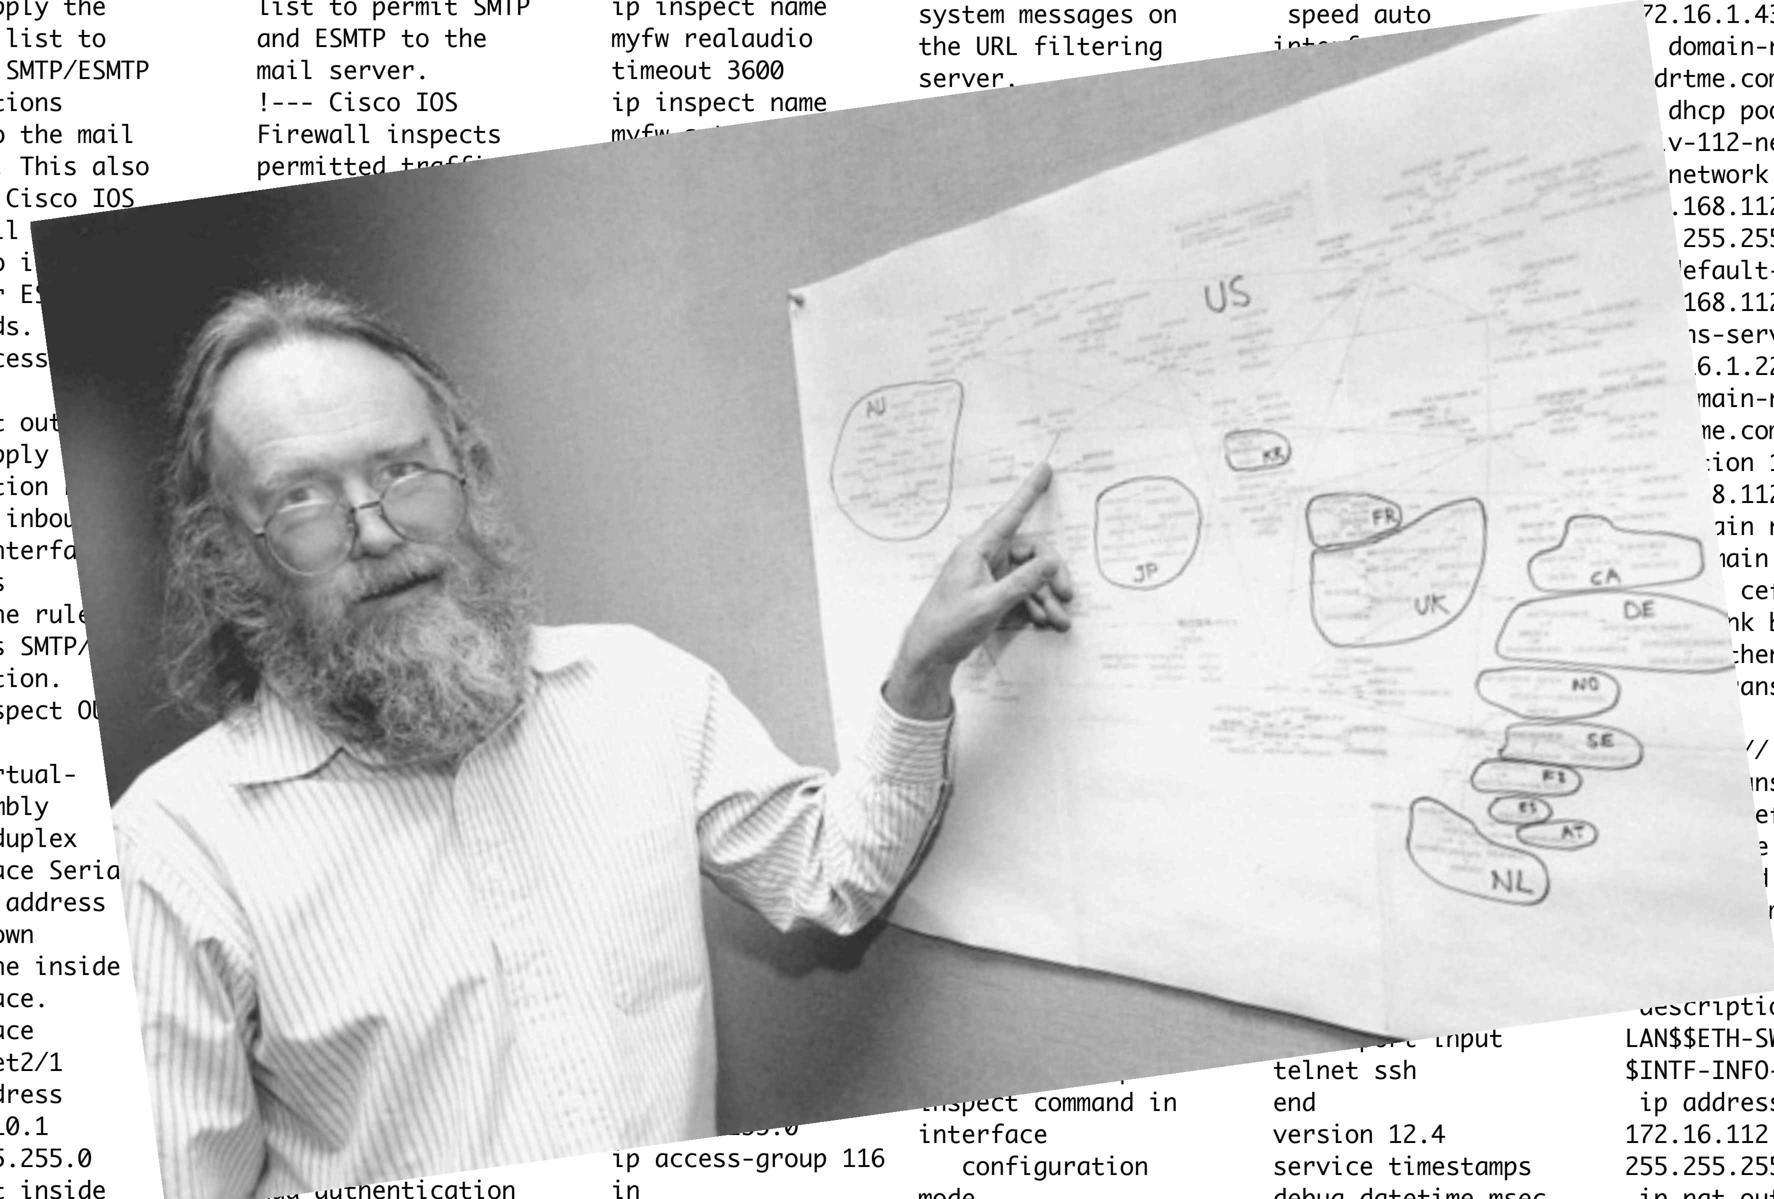


```
aaa authentication
login default group
RTP none
aaa authorization
exec default group
RTP none
aaa authorization
auth-proxy default
group RTP
enable secret 5
$1$pqRI
$3TDNFT9FdYT8Sd/
q3S0VU1
enable password ww
```

```
ip inspect name
myfw cuseeme
timeout 3600
ip inspect name
myfw ftp timeout
3600
ip inspect name
myfw http timeout
3600
ip inspect name
myfw rcmd timeout
3600
ip inspect name
myfw realaudio
timeout 3600
ip inspect name
```

```
ip access-group 116
in
 ip nat outside
ip auth-proxy
list_a
 no ip route-cache
 no ip mroute-cache
 speed auto
 half-duplex
 no mop enabled
interface
Ethernet1/0
 ip address
10.14.14.14
255.255.255.0
```

```
   the syslog
server or router.
ip urlfilter audit-
trail
!--- use the ip
urlfilter urlf-
server-log
   command in
global
configuration mode
to enable the
logging of
system messages on
the URL filtering
server.
```

```
inspect command in
interface
   configuration
mode
to apply a set of
inspection rules to
an interface.
   Here the
inspection name
TEST is
applied to the
interface
FastEthernet0.
 ip inspect test in
 duplex auto
 speed auto
interface
```

```
 ip virtual-
reassembly
 duplex auto
 speed auto
interface
FastEthernet2
 ip address
10.77.241.109
255.255.255.192
 ip virtual-
reassembly
 duplex auto
 speed auto
```

```
telnet ssh
end
version 12.4
service timestamps
debug datetime msec
service timestamps
log datetime msec
no service
password-encryption
hostname 2851-cme2
logging message-
counter syslog
logging buffered
51200 warnings
no aaa new-model
clock timezone mst
-7
clock summer-time
```

```
no ip dhcp use vrf
connected
ip dhcp pool
pub-112-net
   network
172.17.112.0
255.255.255.0
   default-router
172.17.112.1
   dns-server
172.16.1.22
   option 150 ip
```

```
description $ETH-
LAN$$ETH-SW-LAUNCH$
$INTF-INFO-GE 0/0$
 ip address
172.16.112.10
255.255.255.0
 ip nat outside
 ip virt
reassembly
 duplex aut
 speed auto
interface
 no ip
 duplex auto
 speed auto
interface
GigabitEthernet0/1.
132
 encapsulation
```

```
interface
GigabitEthernet0/1.
152
 encapsulation
dot1Q 152
 ip address
192.168.112.1
255.255.255.0
 ip nat inside
 ip virtual-
reassembly
interface
FastEthernet0/2/0
interface
FastEthernet0/2/1
interface
FastEthernet0/2/2
interface
FastEthernet0/2/3
interface Vlan1
 ip address
198.41.9.15
255.255.255.0
router eigrp 1
 network
172.16.112.0
0.0.0.255
 network
172.17.112.0
0.0.0.255
 no auto-summary
ip forward-protocol
nd
ip http server
ip http access-
class 23
ip http
authentication
local
ip http secure-
server
ip http timeout-
policy idle 60 life
86400 requests
10000
ip http path
flash:/gui
ip nat inside
source list 111
interface
GigabitEthernet0/0
overload
access-list 23
permit 10.10.10.0
access-list 111
deny ip
0.0.0.255
192.168.0.0
0.0.255.255
access-list 111
permit ip
```

<1%


CENTRE STREET
SOUTH
OCCUPY
EVERYTHING
#OCCUPYWALLST
WE ALREADY KNOW THAT WE OWN EVERYTHING--THE TASK IS TO EXCLUDE THE INTRUSIONS OF CAPITAL AND POWER
Beware
BANKS
BAILED

Participatory Networking

Participatory Networking

1. management API

Participatory Networking

1. management API
2. network controller

Participatory Networking

1. management API
2. network controller

Safe

Participatory Networking

1. management API
2. network controller

Safe
Secure

Participatory Networking

1. management API
2. network controller

Safe
Secure
Fair

Questions?

Rodrigo Fonseca
rfonseca@cs.brown.edu

Andrew Ferguson
adf@cs.brown.edu

Arjun Guha
arjun@cs.brown.edu

Shriram Krishnamurthi
sk@cs.brown.edu


2012 USENIX Federated Conferences Week · June 12–15, 2012
4th USENIX Workshop on Hot Topics in Cloud Computing
HotCloud '12
JUNE 12—13, 2012
BOSTON, MA
usenix
THE ADVANCED COMPUTING SYSTEMS ASSOCIATION

A problem in the home

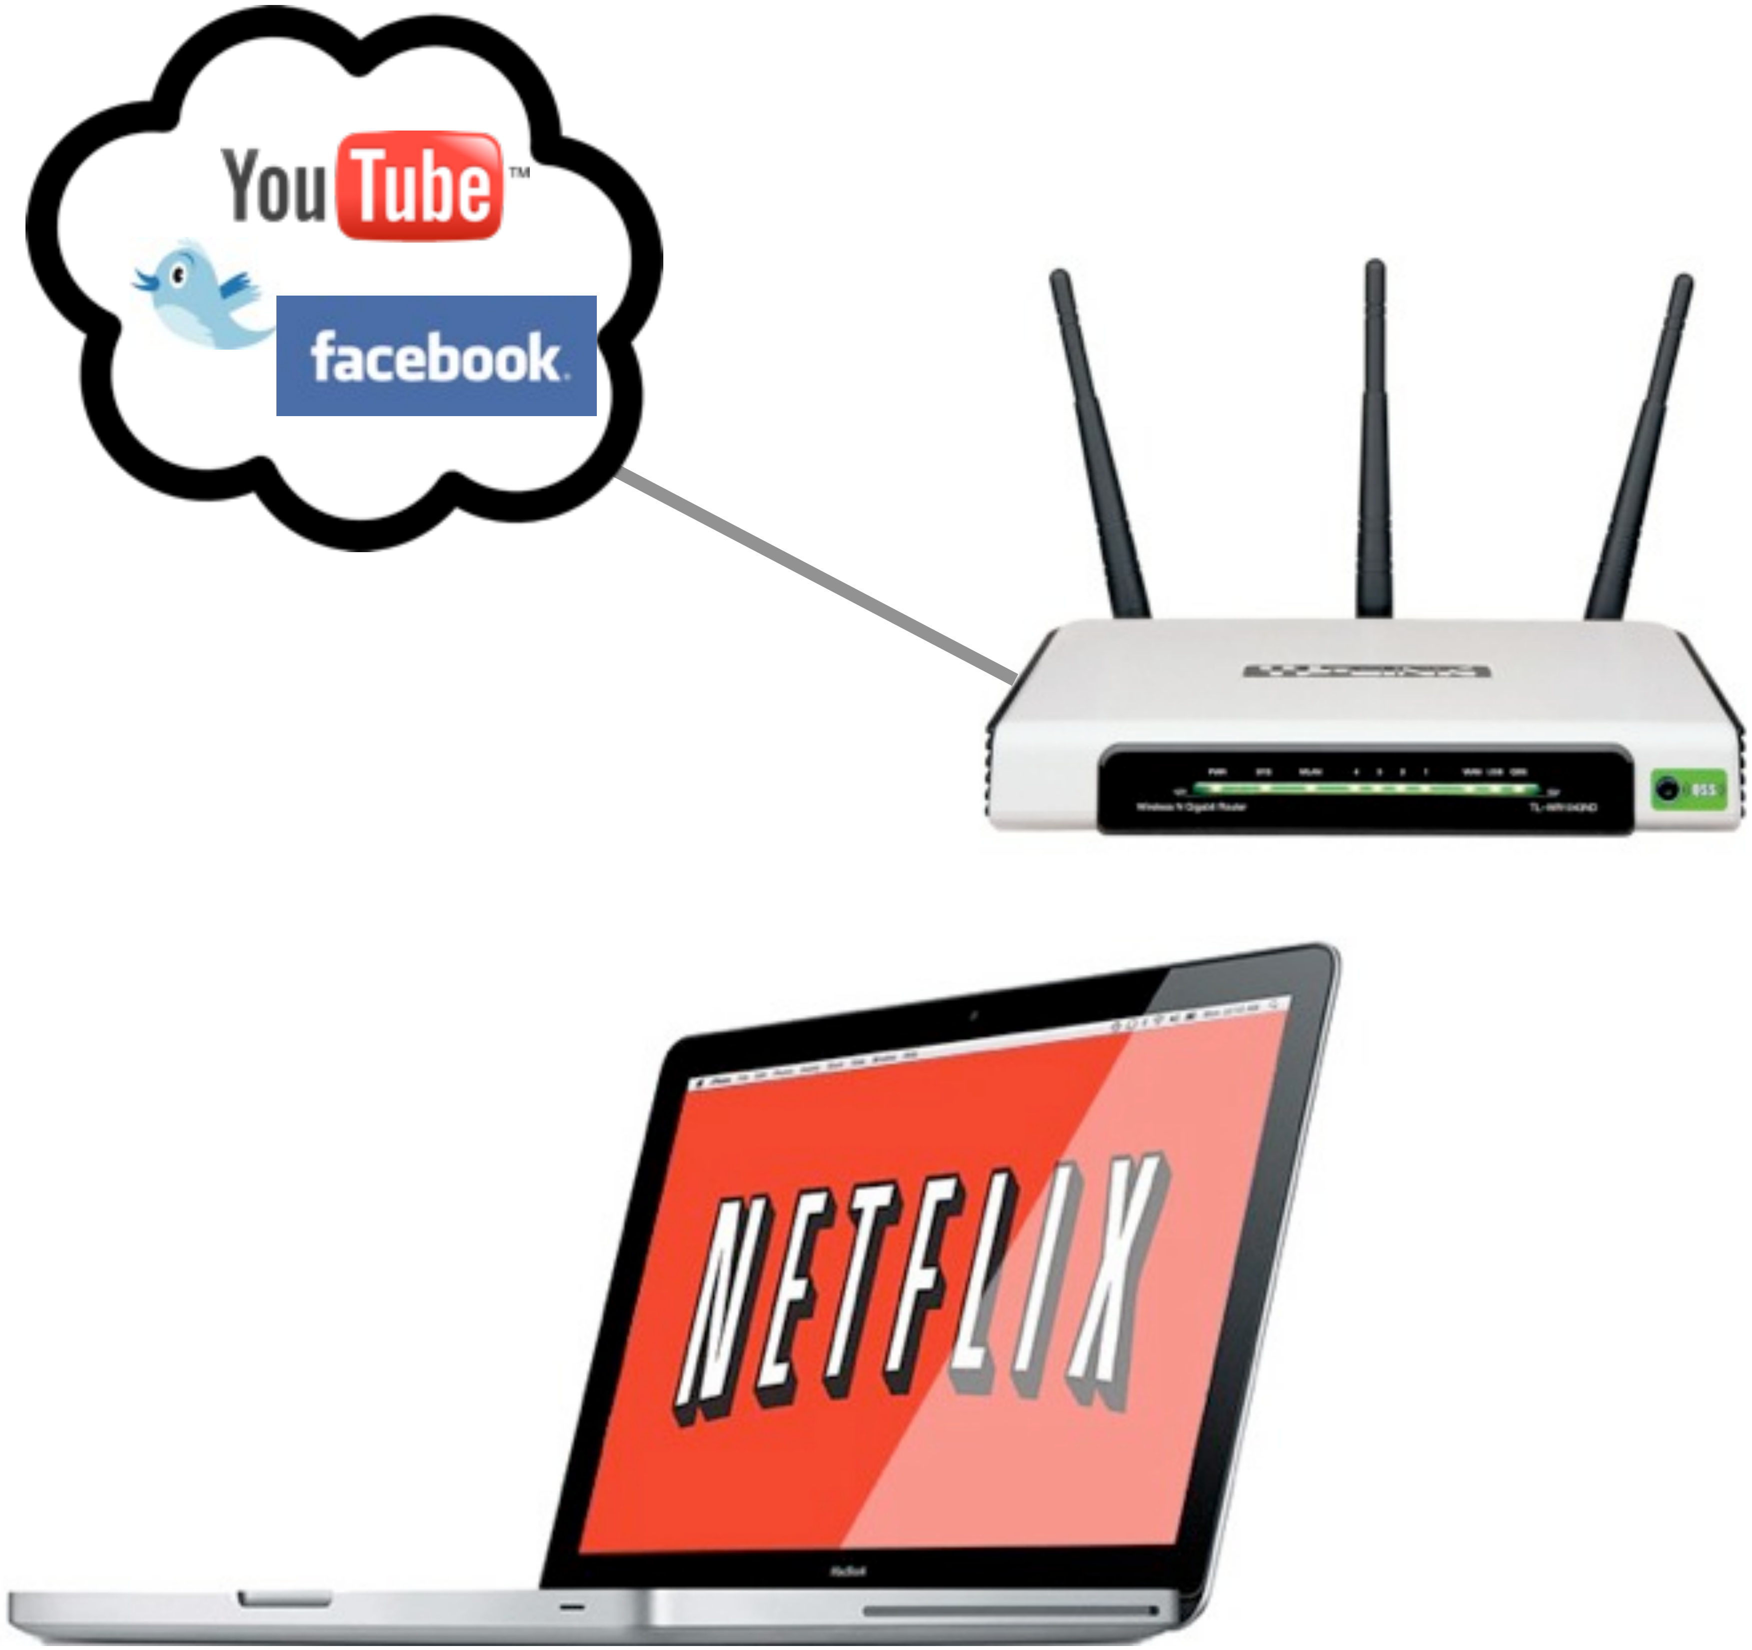
YouTube
facebook
NETFLIX


NETFLIX
89%
Buffering
Full Screen
More Episodes
Back to Browsing

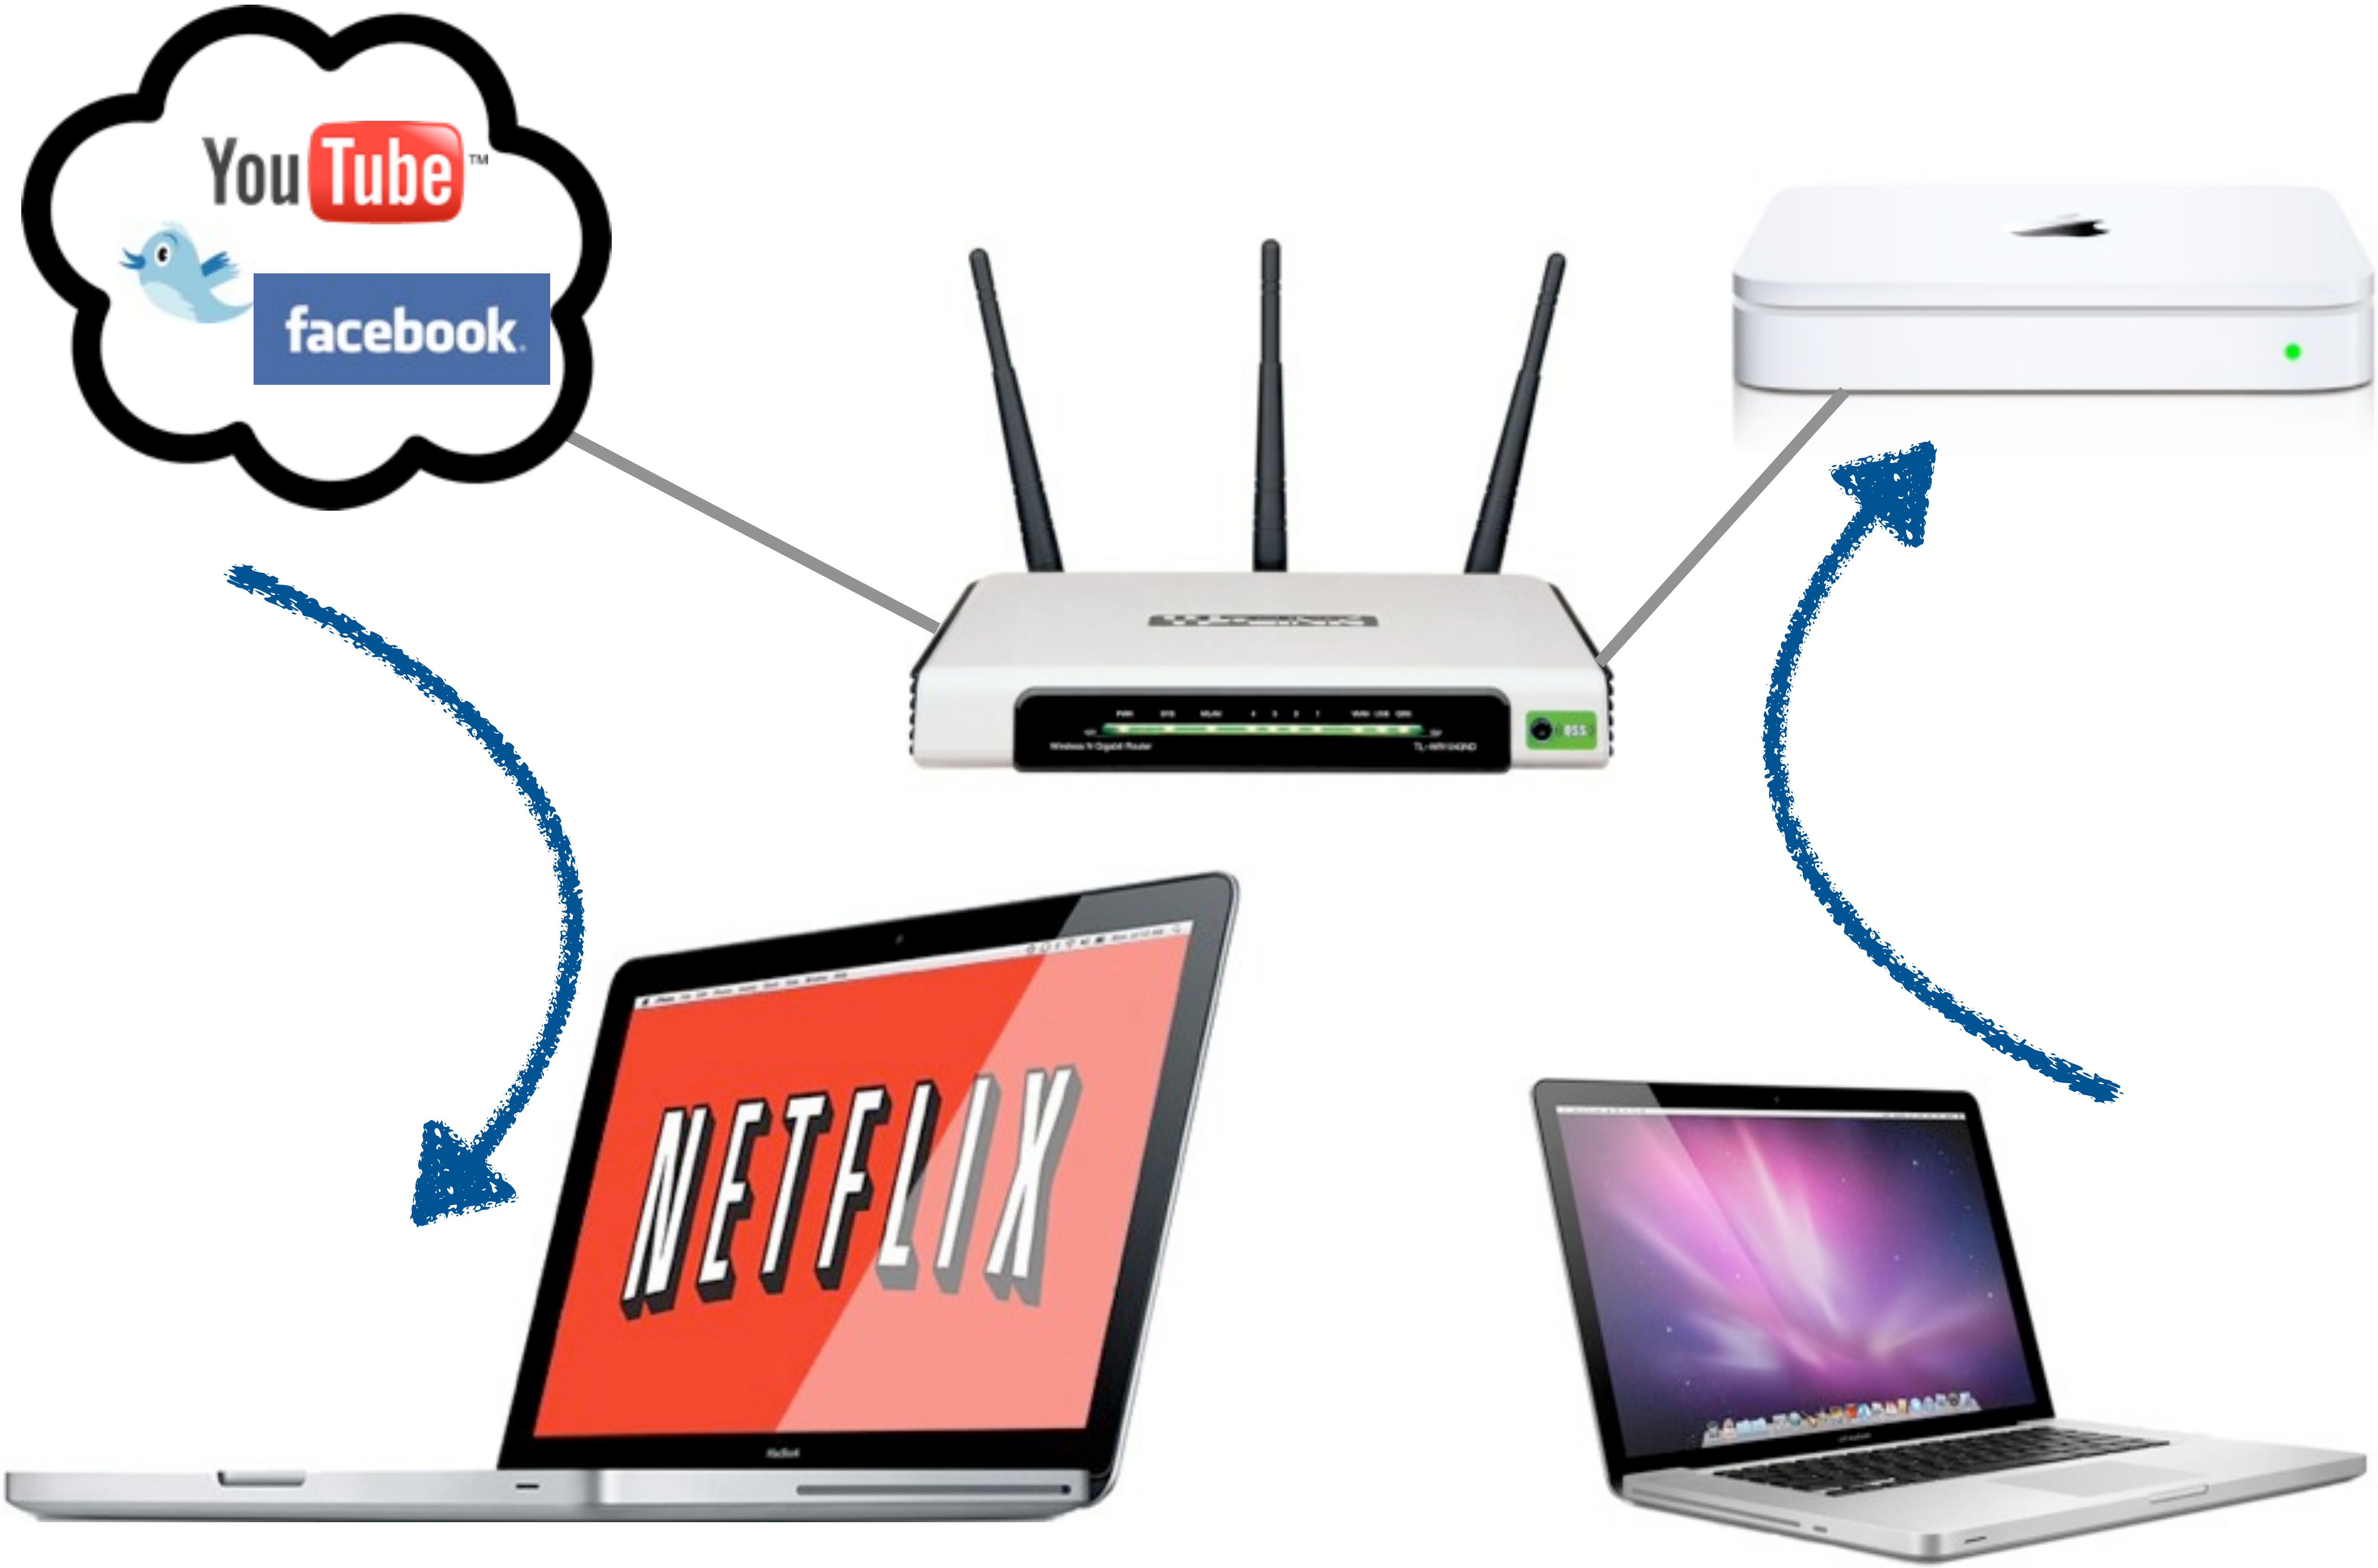
YouTube
facebook
NETFLIX

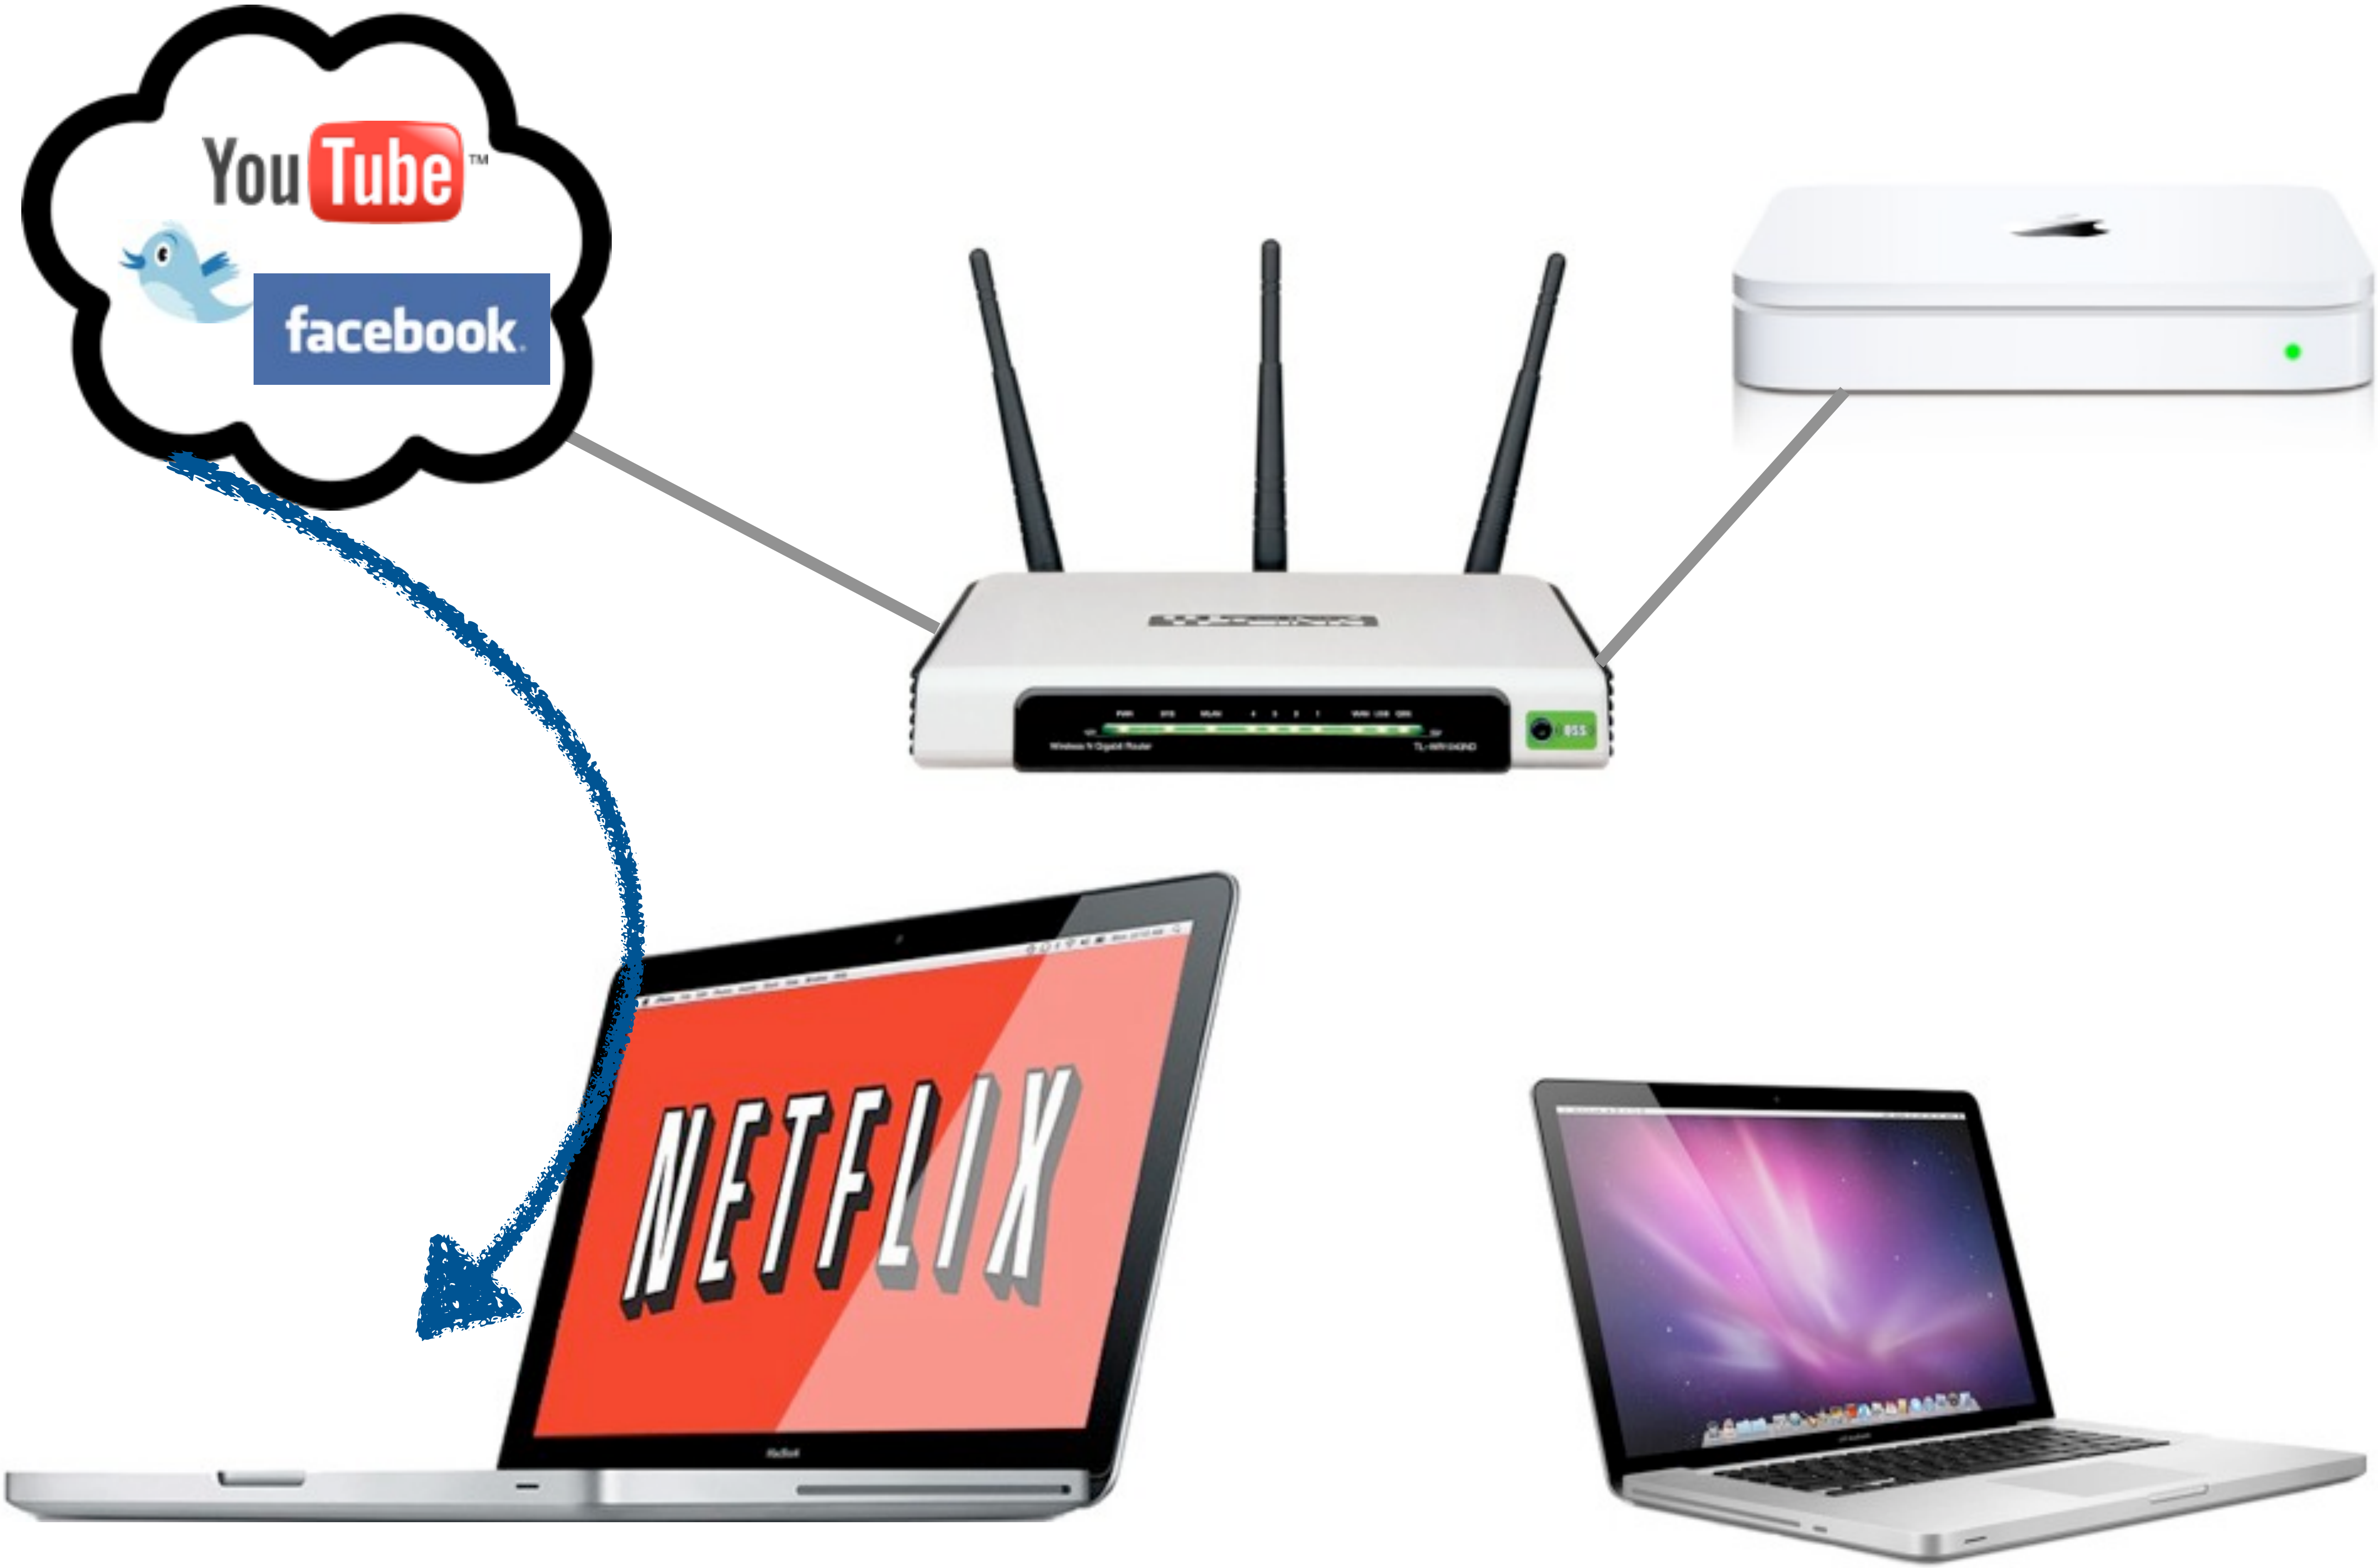
YouTube
facebook
NETFLIX

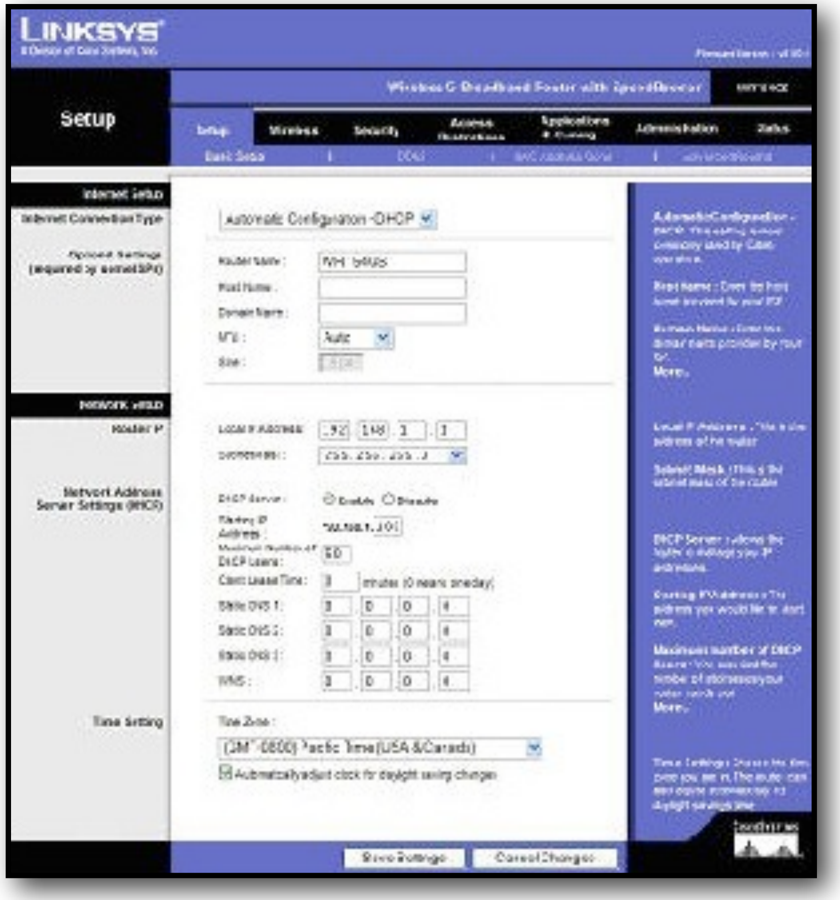
LINKSYS
Setup
Wireless
Security
Administration
Status
Internet Setup
Internet Connection Type
Automatic Configuration - DHCP
Network Setup
Router IP
Network Address Server Settings (DHCP)
Time Setting
Save Settings
Cancel Changes

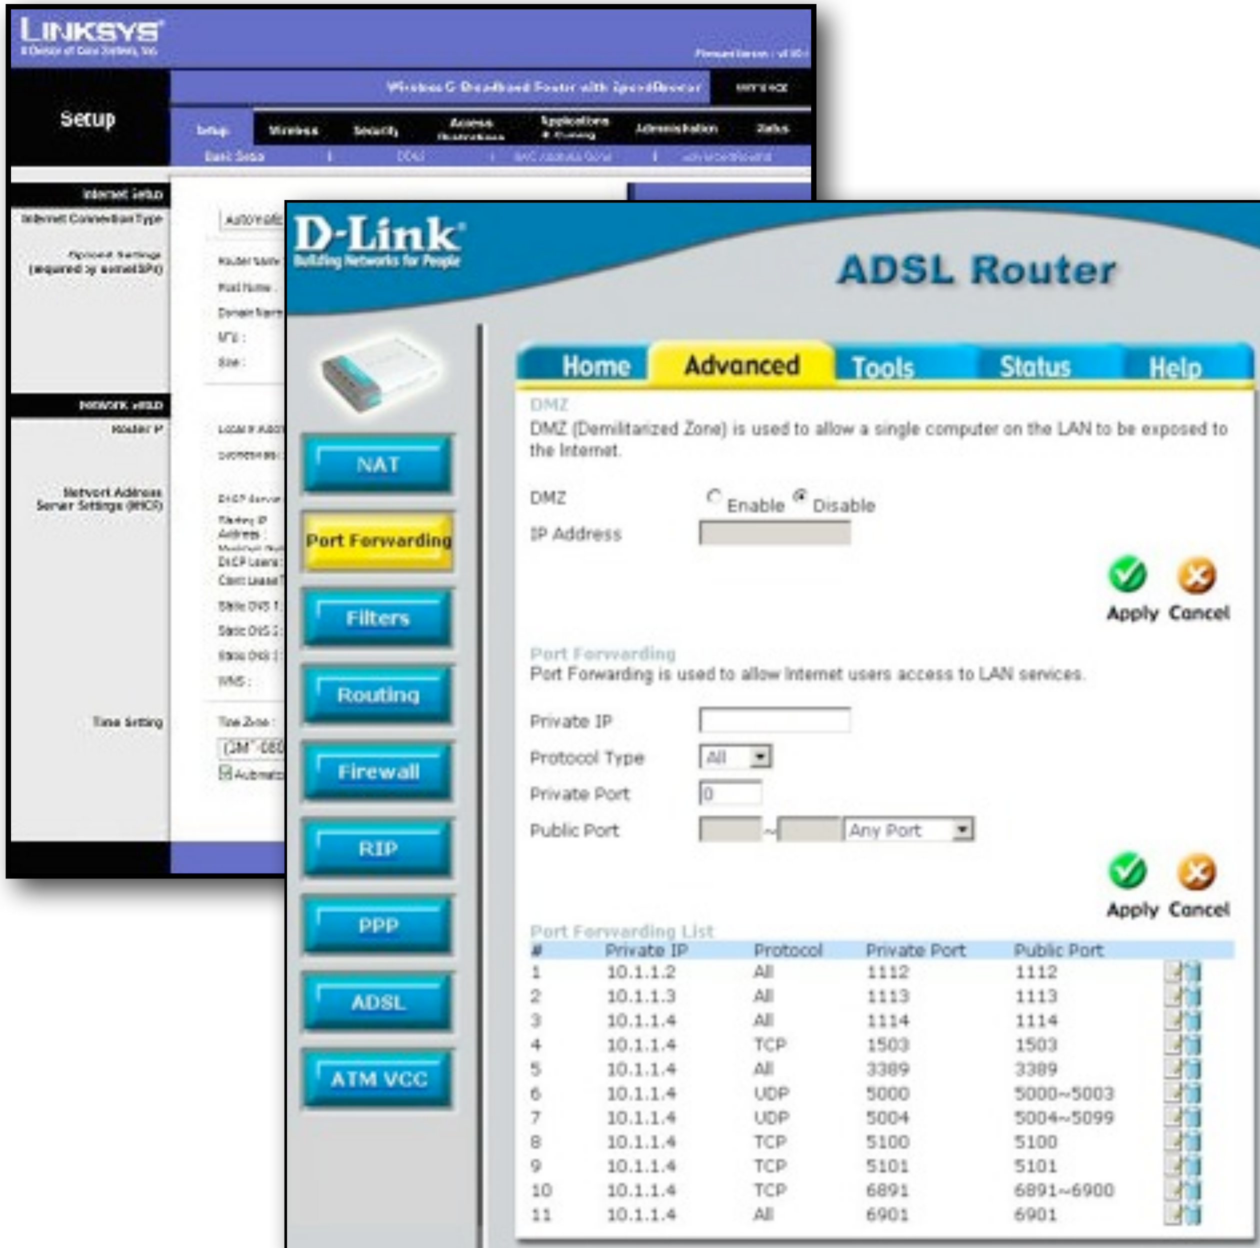
LINKSYS
Setup
D-Link
Building Networks for People
ADSL Router
Home
Advanced
Tools
Status
Help
NAT
Port Forwarding
Filters
Routing
Firewall
RIP
PPP
ADSL
ATM VCC
DMZ
DMZ (Demilitarized Zone) is used to allow a single computer on the LAN to be exposed to the Internet.
DMZ
Enable
Disable
IP Address
Apply
Cancel
Port Forwarding
Port Forwarding is used to allow Internet users access to LAN services.
Private IP
Protocol Type
All
Private Port
0
Public Port
Any Port
Apply
Cancel
Port Forwarding List
Private IP Protocol Private Port Public Port
1 10.1.1.2 All 1112 1112
2 10.1.1.3 All 1113 1113
3 10.1.1.4 All 1114 1114
4 10.1.1.4 TCP 1503 1503
5 10.1.1.4 All 3389 3389
6 10.1.1.4 UDP 5000 5000~5003
7 10.1.1.4 UDP 5004 5004~5099
8 10.1.1.4 TCP 5100 5100
9 10.1.1.4 TCP 5101 5101
10 10.1.1.4 TCP 6891 6891~6900
11 10.1.1.4 All 6901 6901

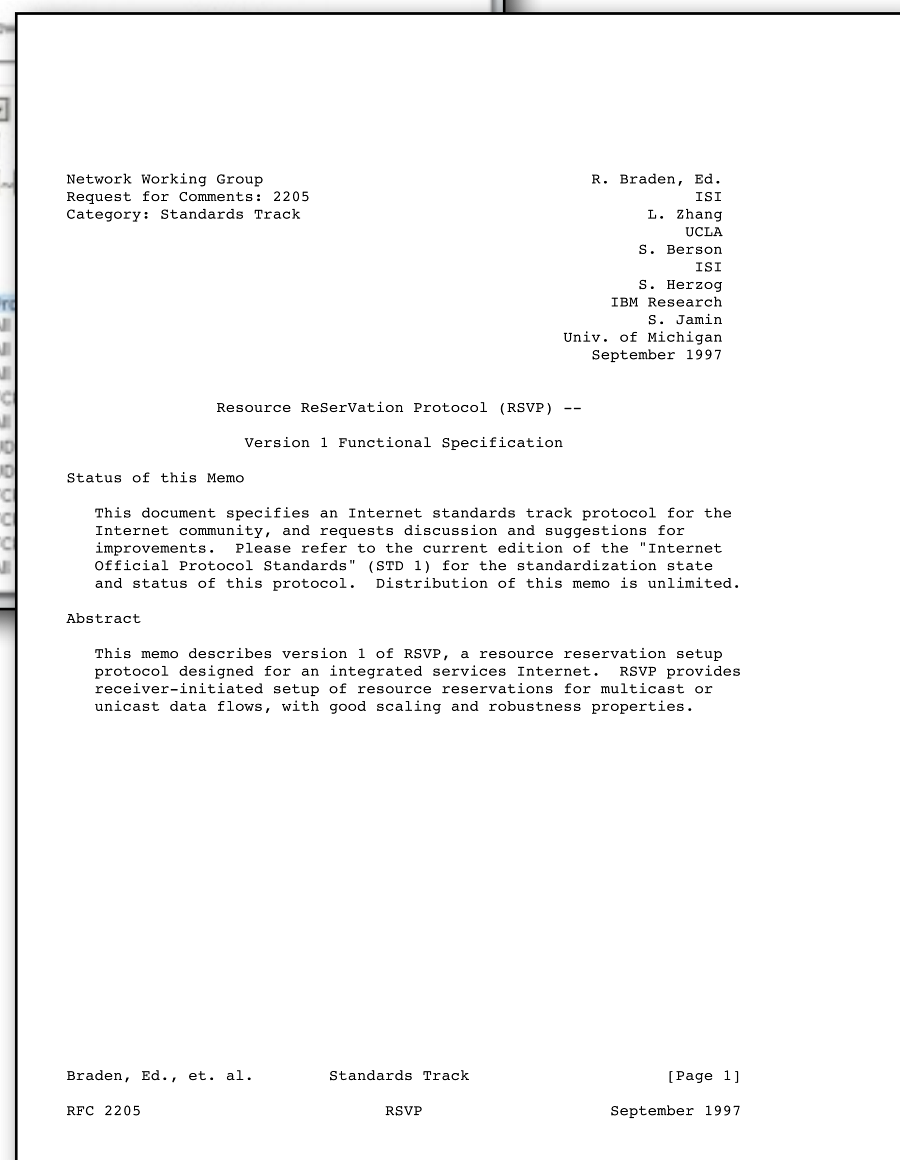

```
Network Working Group                                   R. Braden, Ed.
Request for Comments: 2205                                         ISI
Category: Standards Track                                     L. Zhang
                                                                  UCLA
                                                             S. Berson
                                                                   ISI
                                                             S. Herzog
                                                          IBM Research
                                                              S. Jamin
                                                     Univ. of Michigan
                                                        September 1997


                Resource ReSerVation Protocol (RSVP) --

                   Version 1 Functional Specification

Status of this Memo

   This document specifies an Internet standards track protocol for the
   Internet community, and requests discussion and suggestions for
   improvements.  Please refer to the current edition of the "Internet
   Official Protocol Standards" (STD 1) for the standardization state
   and status of this protocol.  Distribution of this memo is unlimited.

Abstract

   This memo describes version 1 of RSVP, a resource reservation setup
   protocol designed for an integrated services Internet.  RSVP provides
   receiver-initiated setup of resource reservations for multicast or
   unicast data flows, with good scaling and robustness properties.



Braden, Ed., et. al.        Standards Track                    [Page 1]

RFC 2205                          RSVP                   September 1997
```

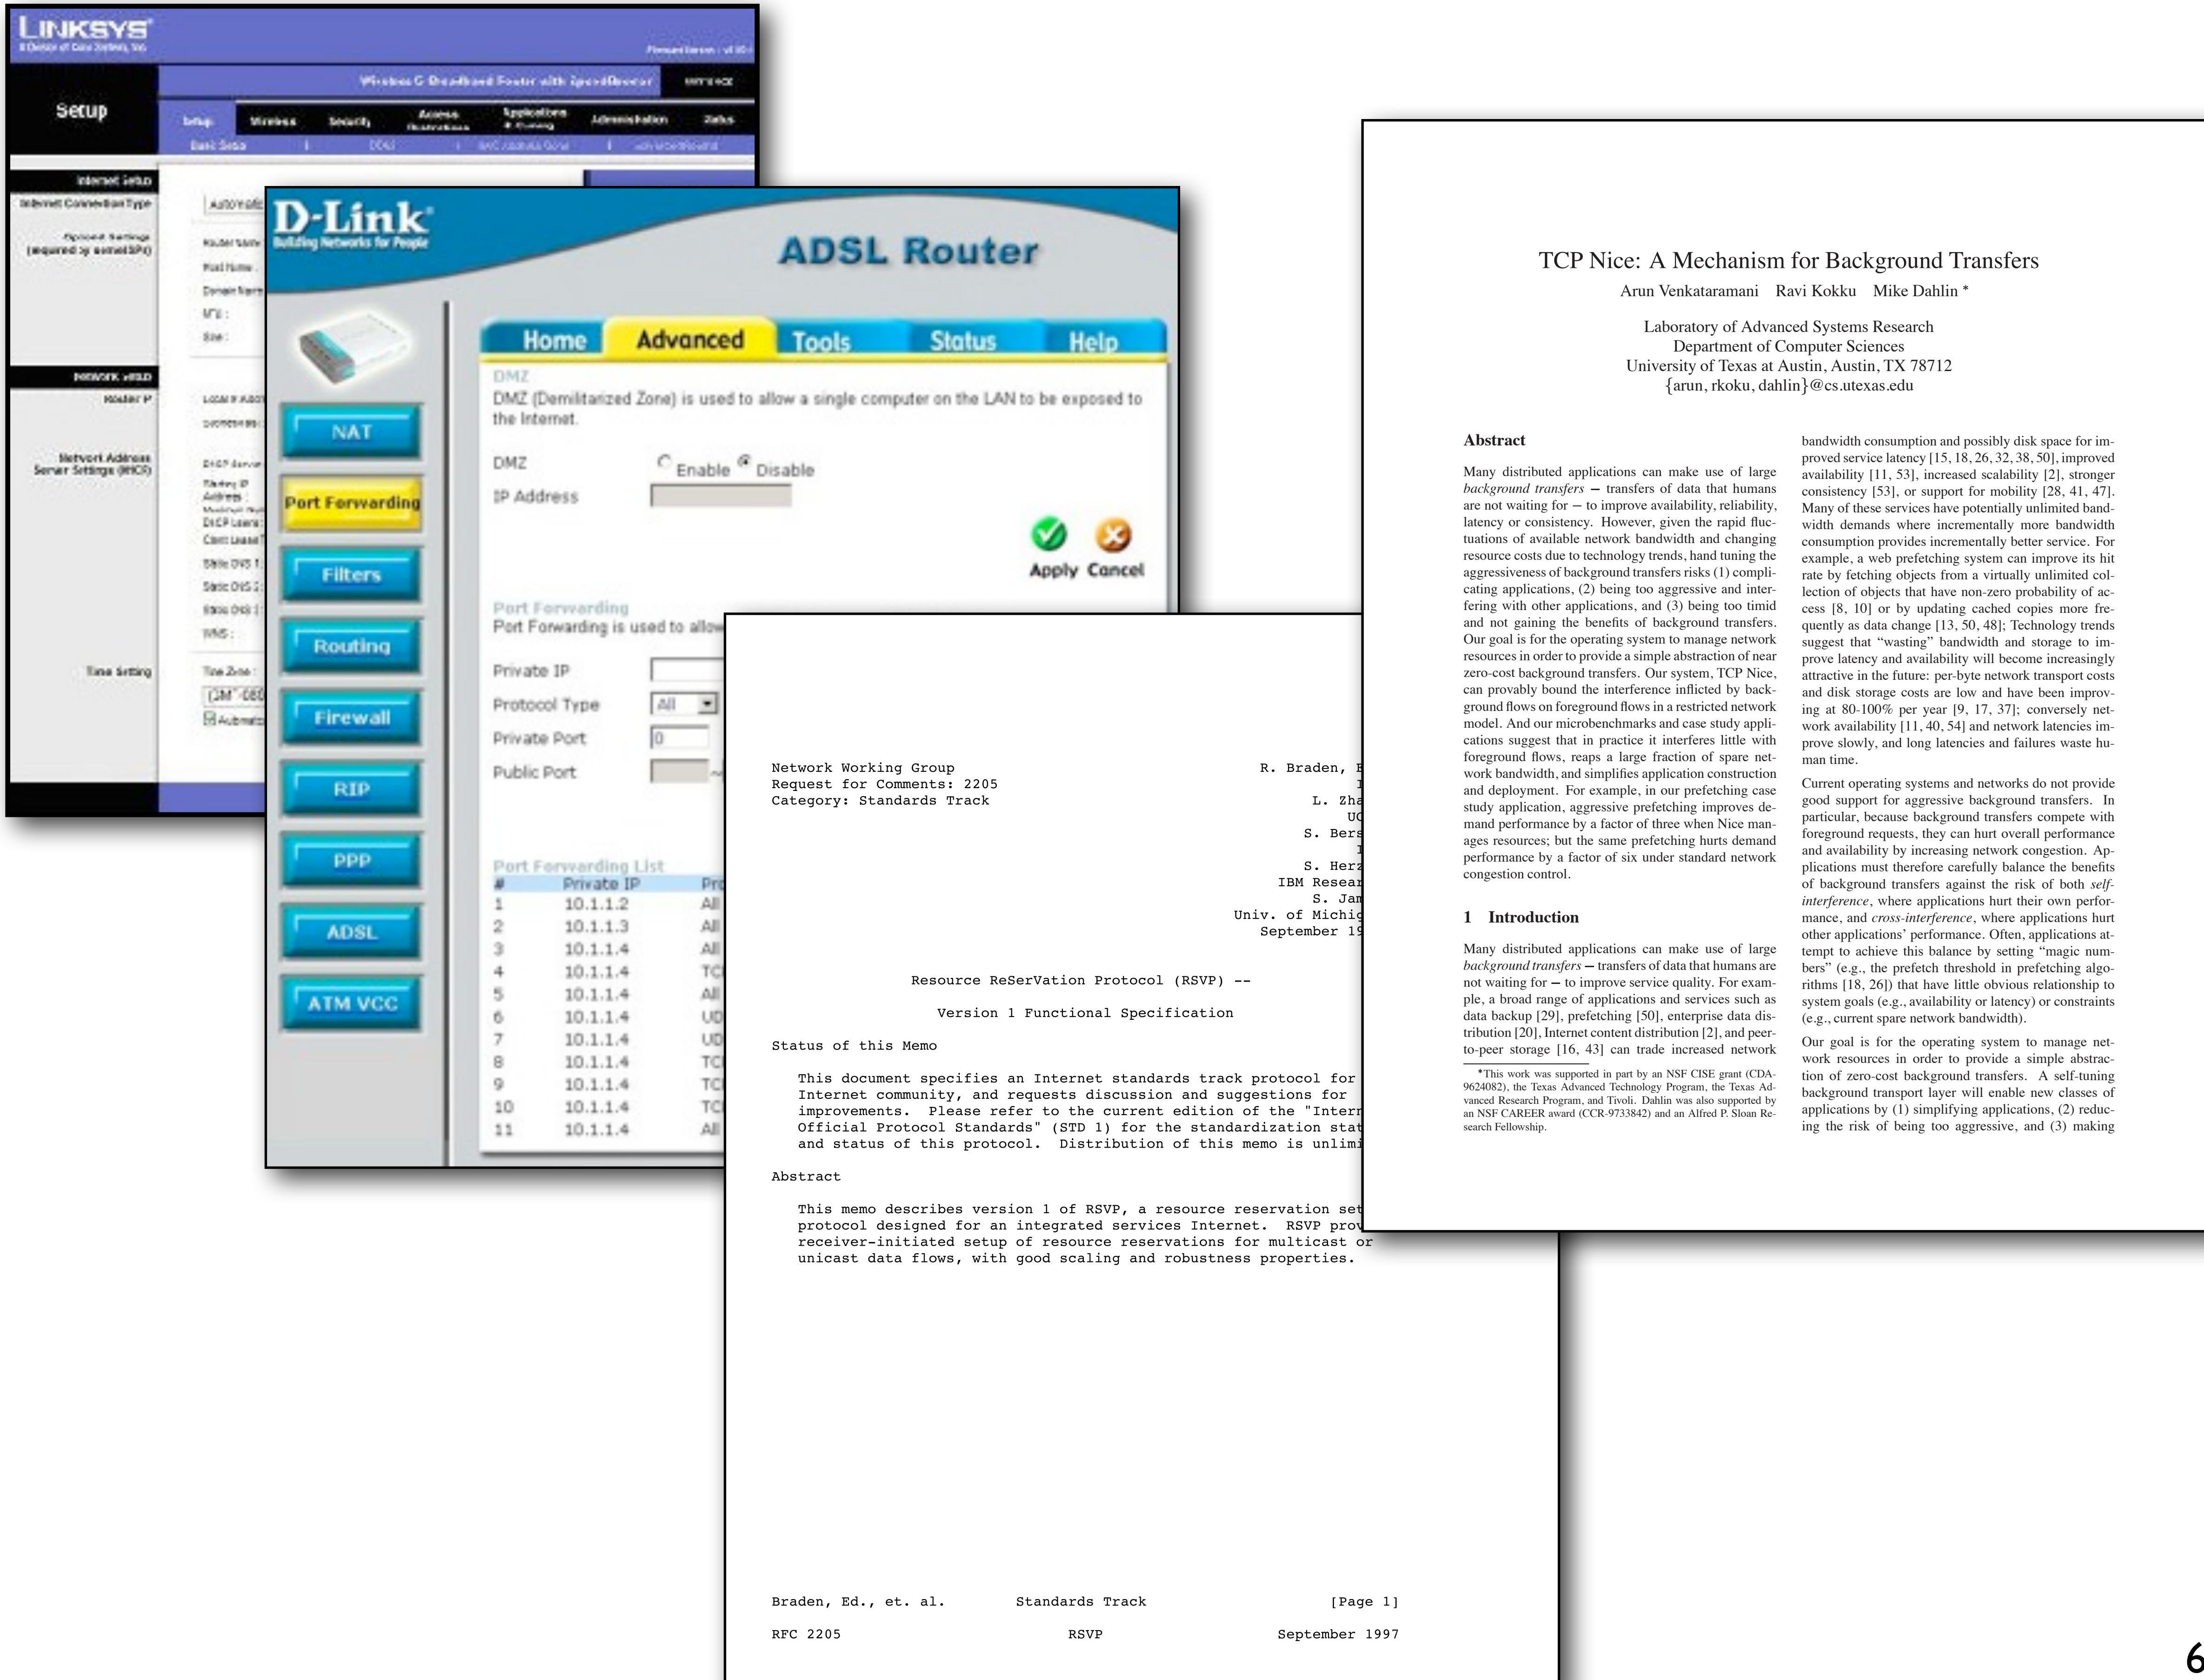


```
Network Working Group                                        R. Braden,
Request for Comments: 2205                                         
Category: Standards Track                                      L. Zha
                                                                   UC
                                                             S. Bers

                                                             S. Herz
                                                          IBM Resear
                                                              S. Jam
                                                       Univ. of Michig
                                                          September 19

               Resource ReSerVation Protocol (RSVP) --

                  Version 1 Functional Specification

Status of this Memo

   This document specifies an Internet standards track protocol for
   Internet community, and requests discussion and suggestions for
   improvements.  Please refer to the current edition of the "Intern
   Official Protocol Standards" (STD 1) for the standardization stat
   and status of this protocol.  Distribution of this memo is unlimi

Abstract

   This memo describes version 1 of RSVP, a resource reservation set
   protocol designed for an integrated services Internet.  RSVP prov
   receiver-initiated setup of resource reservations for multicast or
   unicast data flows, with good scaling and robustness properties.

Braden, Ed., et. al.        Standards Track                  [Page 1]

RFC 2205                         RSVP                   September 1997
```

TCP Nice: A Mechanism for Background Transfers

Arun Venkataramani   Ravi Kokku   Mike Dahlin *

Laboratory of Advanced Systems Research
Department of Computer Sciences
University of Texas at Austin, Austin, TX 78712
{arun, rkoku, dahlin}@cs.utexas.edu

Abstract

Many distributed applications can make use of large *background transfers* — transfers of data that humans are not waiting for — to improve availability, reliability, latency or consistency. However, given the rapid fluctuations of available network bandwidth and changing resource costs due to technology trends, hand tuning the aggressiveness of background transfers risks (1) complicating applications, (2) being too aggressive and interfering with other applications, and (3) being too timid and not gaining the benefits of background transfers. Our goal is for the operating system to manage network resources in order to provide a simple abstraction of near zero-cost background transfers. Our system, TCP Nice, can provably bound the interference inflicted by background flows on foreground flows in a restricted network model. And our microbenchmarks and case study applications suggest that in practice it interferes little with foreground flows, reaps a large fraction of spare network bandwidth, and simplifies application construction and deployment. For example, in our prefetching case study application, aggressive prefetching improves demand performance by a factor of three when Nice manages resources; but the same prefetching hurts demand performance by a factor of six under standard network congestion control.

1 Introduction

Many distributed applications can make use of large *background transfers* — transfers of data that humans are not waiting for — to improve service quality. For example, a broad range of applications and services such as data backup [29], prefetching [50], enterprise data distribution [20], Internet content distribution [2], and peer-to-peer storage [16, 43] can trade increased network bandwidth consumption and possibly disk space for improved service latency [15, 18, 26, 32, 38, 50], improved availability [11, 53], increased scalability [2], stronger consistency [53], or support for mobility [28, 41, 47]. Many of these services have potentially unlimited bandwidth demands where incrementally more bandwidth consumption provides incrementally better service. For example, a web prefetching system can improve its hit rate by fetching objects from a virtually unlimited collection of objects that have non-zero probability of access [8, 10] or by updating cached copies more frequently as data change [13, 50, 48]; Technology trends suggest that "wasting" bandwidth and storage to improve latency and availability will become increasingly attractive in the future: per-byte network transport costs and disk storage costs are low and have been improving at 80-100% per year [9, 17, 37]; conversely network availability [11, 40, 54] and network latencies improve slowly, and long latencies and failures waste human time.

Current operating systems and networks do not provide good support for aggressive background transfers. In particular, because background transfers compete with foreground requests, they can hurt overall performance and availability by increasing network congestion. Applications must therefore carefully balance the benefits of background transfers against the risk of both *self-interference*, where applications hurt their own performance, and *cross-interference*, where applications hurt other applications' performance. Often, applications attempt to achieve this balance by setting "magic numbers" (e.g., the prefetch threshold in prefetching algorithms [18, 26]) that have little obvious relationship to system goals (e.g., availability or latency) or constraints (e.g., current spare network bandwidth).

Our goal is for the operating system to manage network resources in order to provide a simple abstraction of zero-cost background transfers. A self-tuning background transport layer will enable new classes of applications by (1) simplifying applications, (2) reducing the risk of being too aggressive, and (3) making

*This work was supported in part by an NSF CISE grant (CDA-9624082), the Texas Advanced Technology Program, the Texas Advanced Research Program, and Tivoli. Dahlin was also supported by an NSF CAREER award (CCR-9733842) and an Alfred P. Sloan Research Fellowship.