



modulus

Background

What we do, Technical goals, Infrastructure



What we do

- Greek VoIP services provider
 - Voice Termination/Origination
 - Virtual PBX as a service
- IT services
 - Virtualization
 - Network infrastructure design and installation

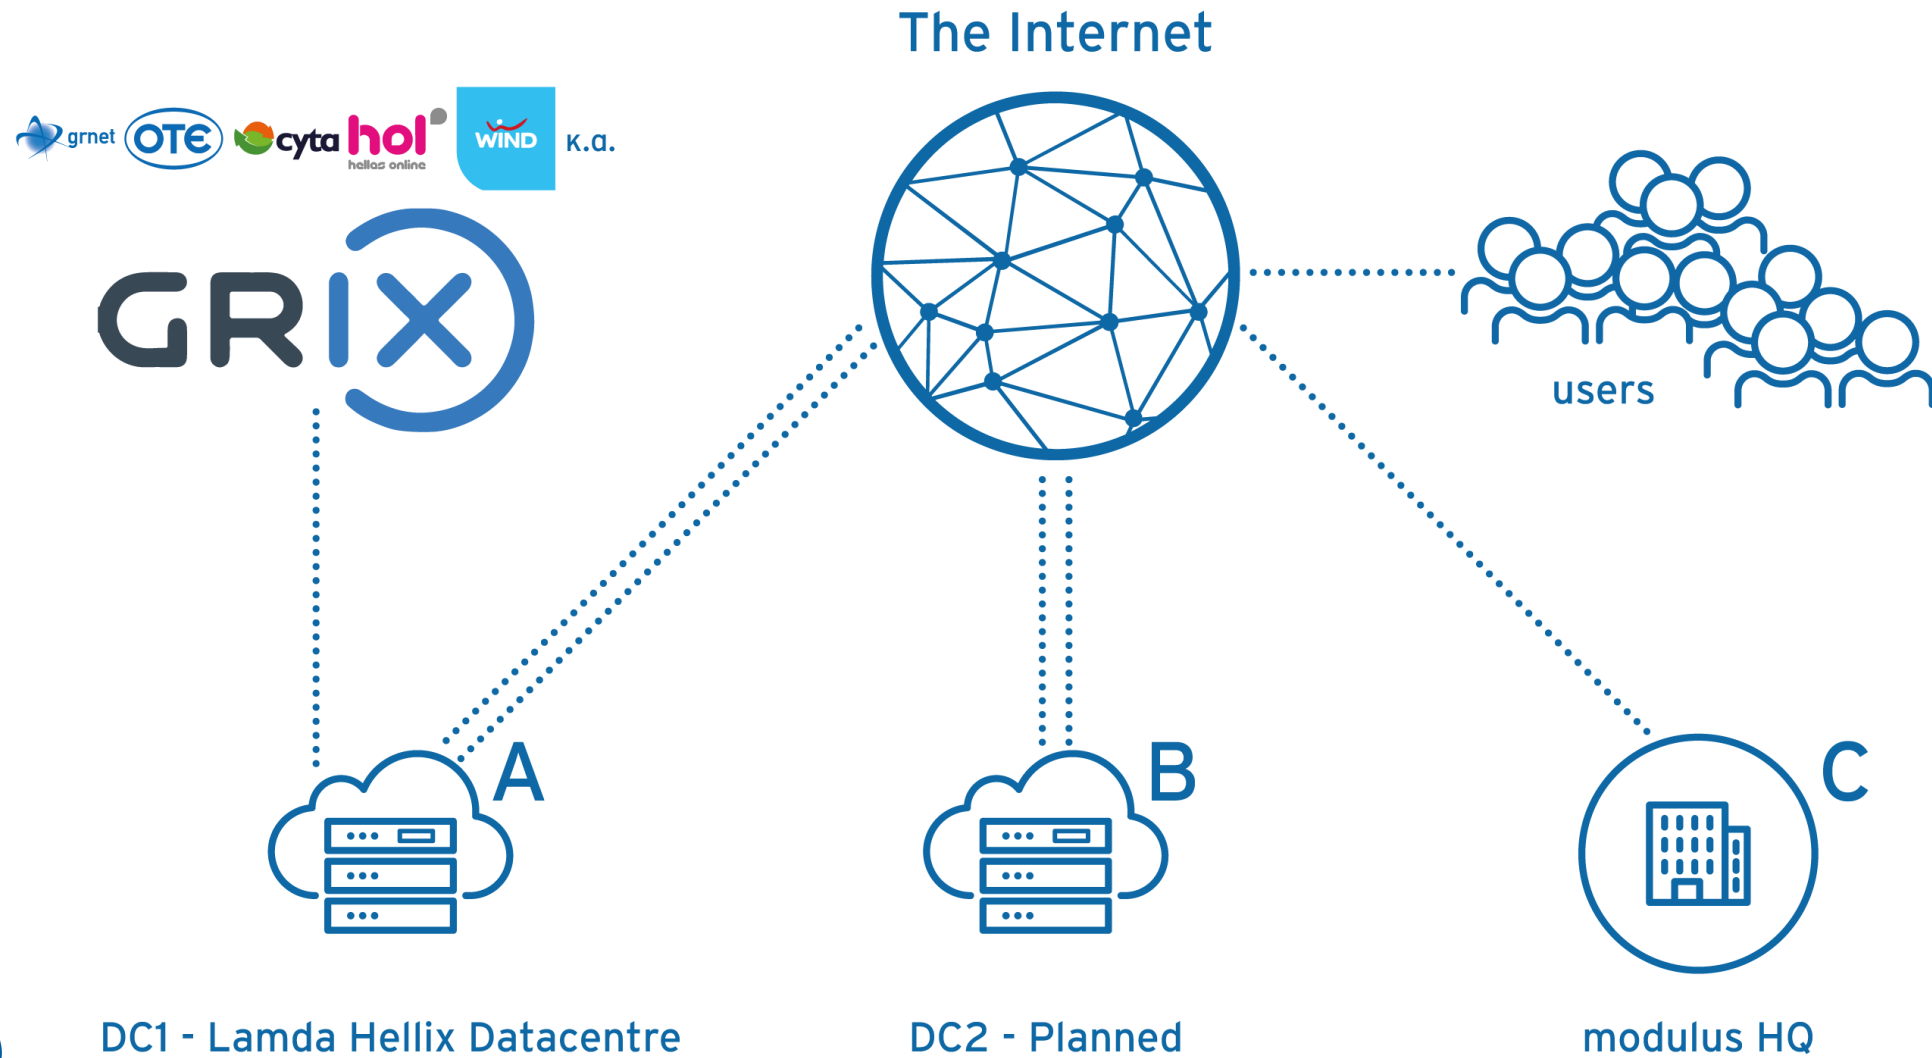


Our Technical Goals

- No single point of failure (SPOF) for services, including
 - Datacenter
 - Power
 - Network
 - Servers / Services
- High Availability > 99.995%
- Low network latency
- Implementation based on the latest technologies available



Our network

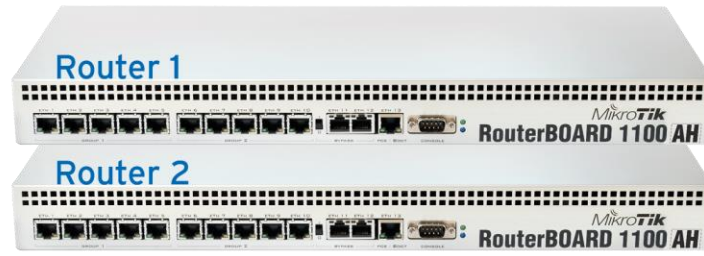


Our Lamda Hellix Infrastructure

- Protected Power Line (A+B)
- 2x Upstream ISP Connections
- 1x GR-IX Connection
- 2x Mikrotik RB1100AH Routers
- 2x Dell 62xx Series Stackable Switches
- 6x Servers



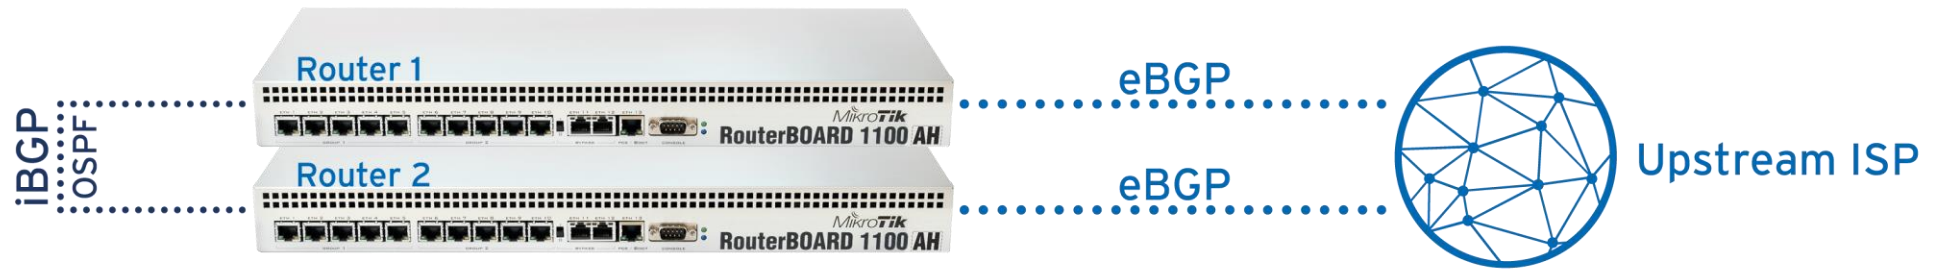
HA Mikrotik-Based Router Infrastructure



HA Mikrotik-Based Router Infrastructure



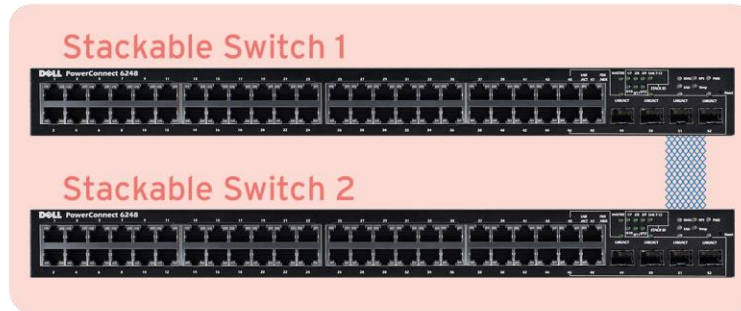
HA Mikrotik-Based Router Infrastructure



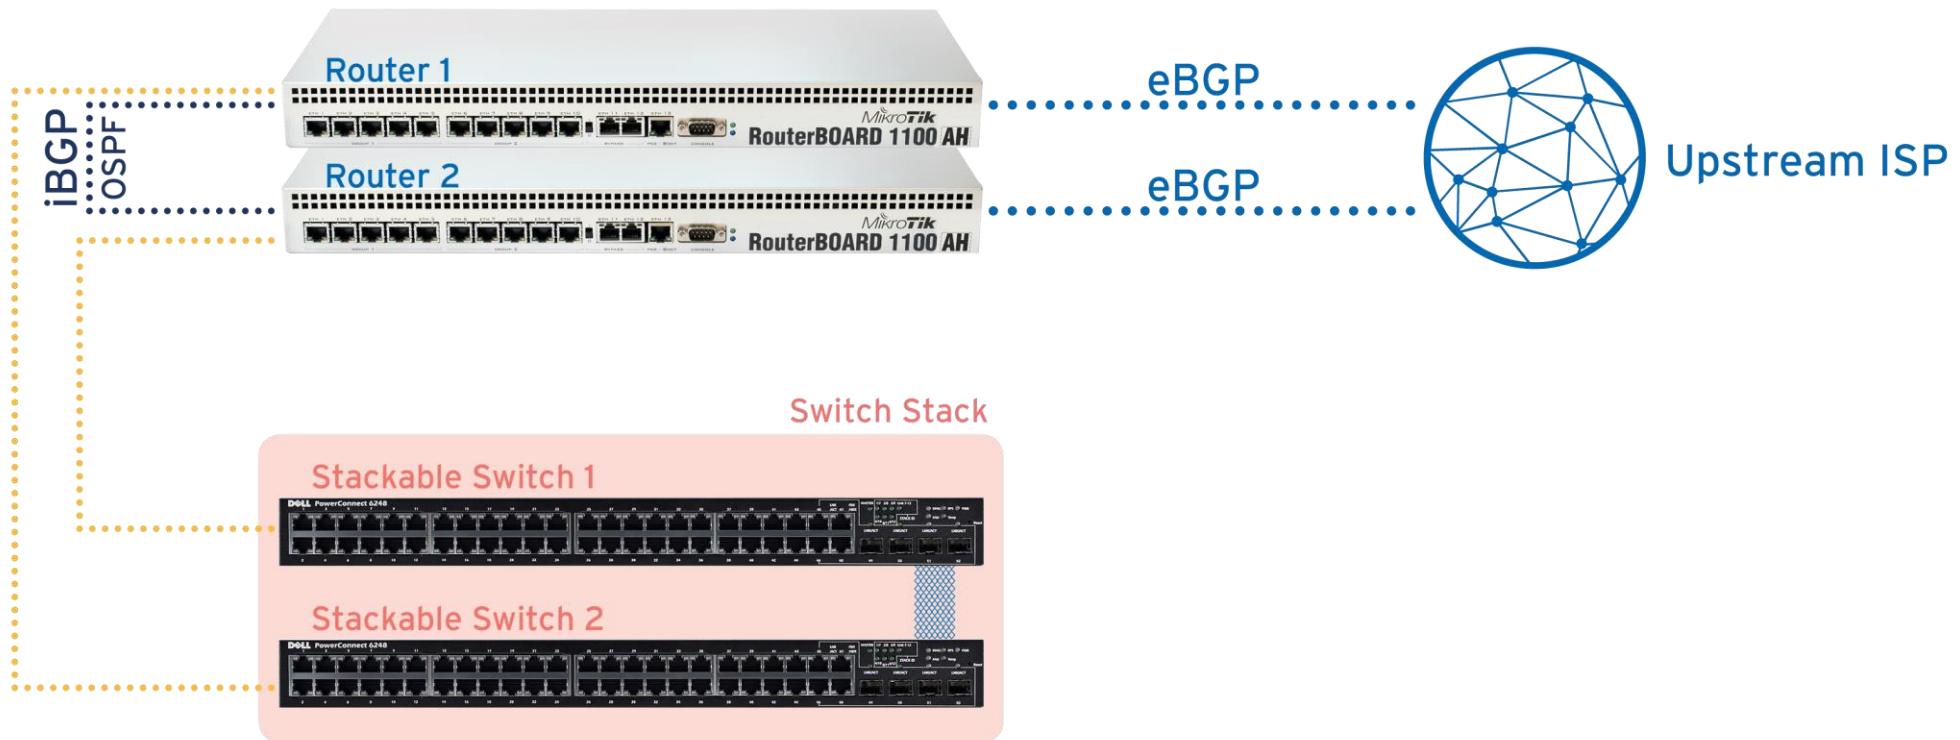
HA Mikrotik-Based Router Infrastructure



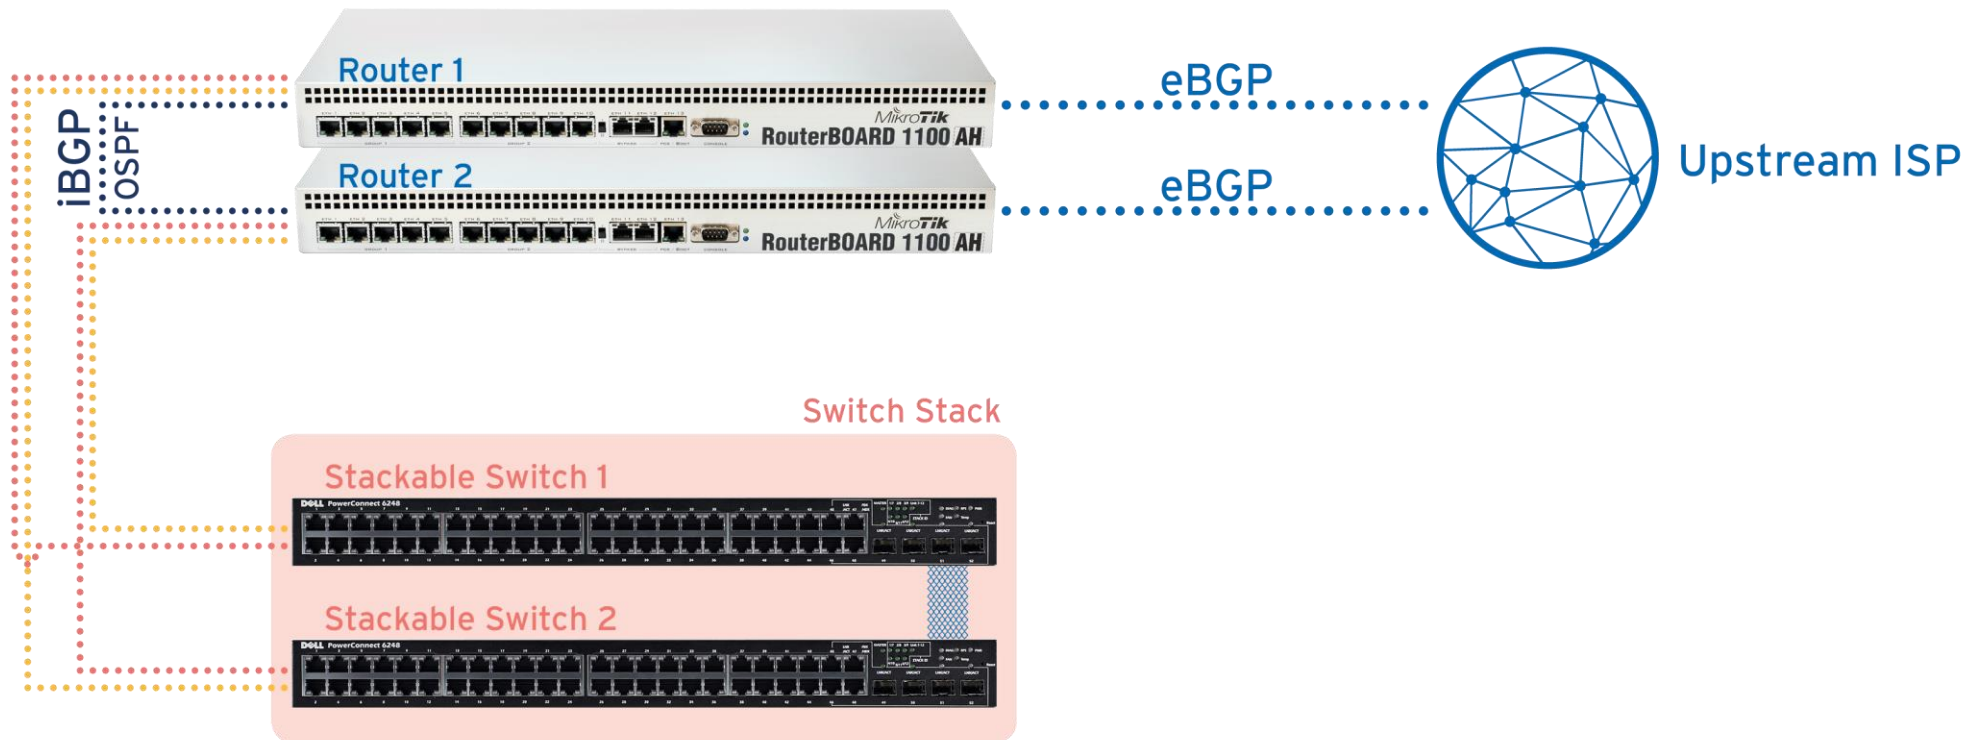
Switch Stack



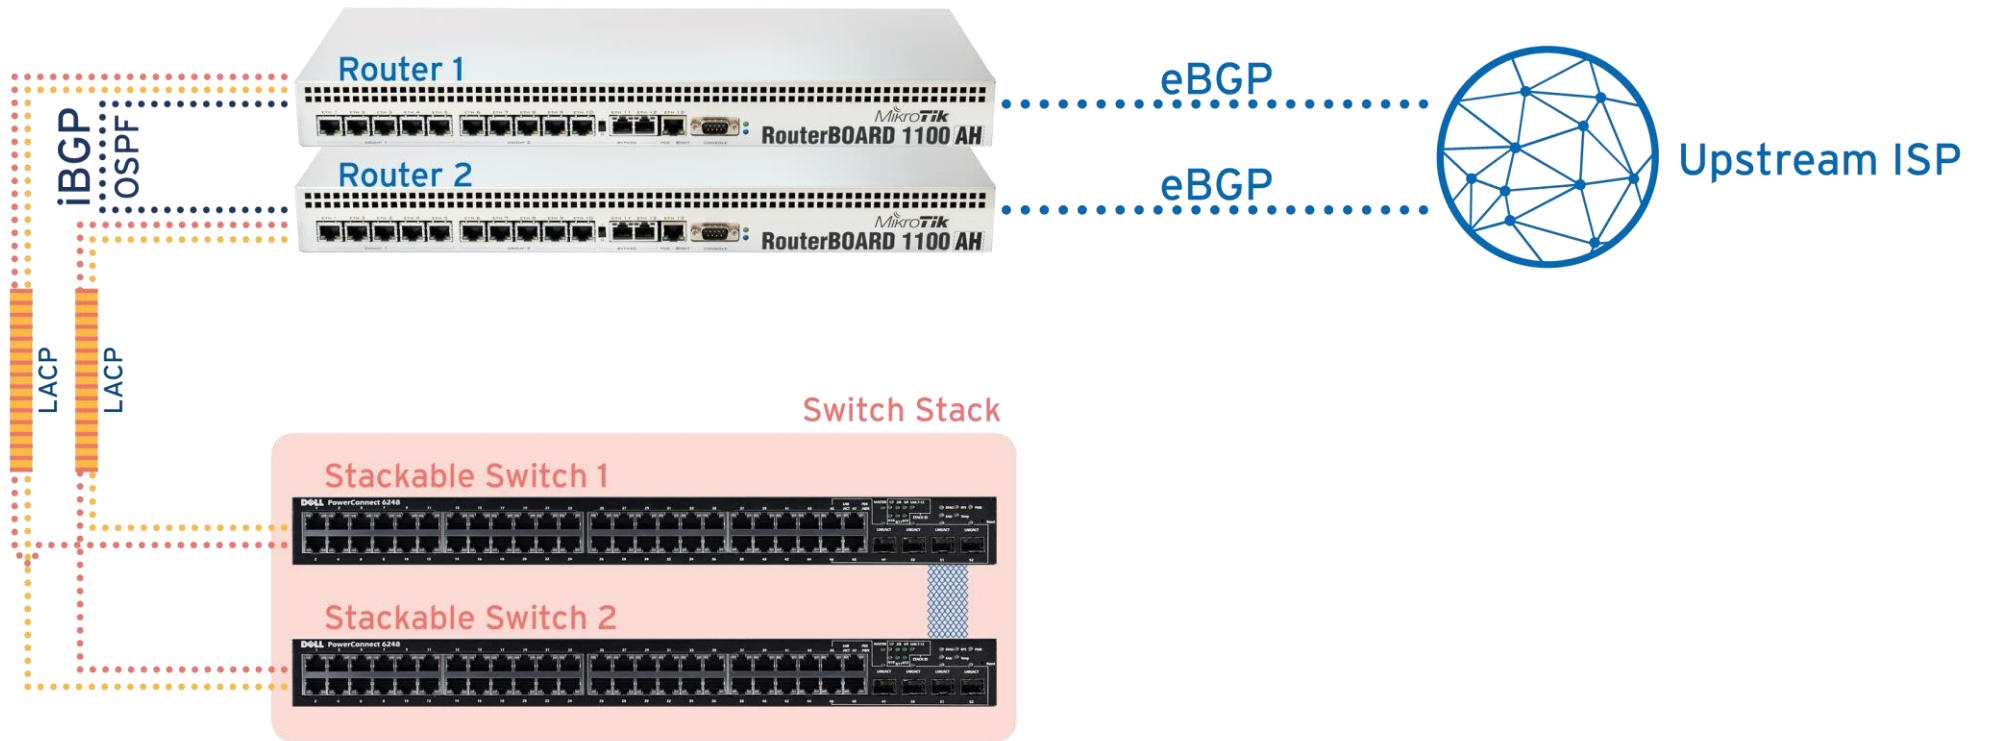
HA Mikrotik-Based Router Infrastructure



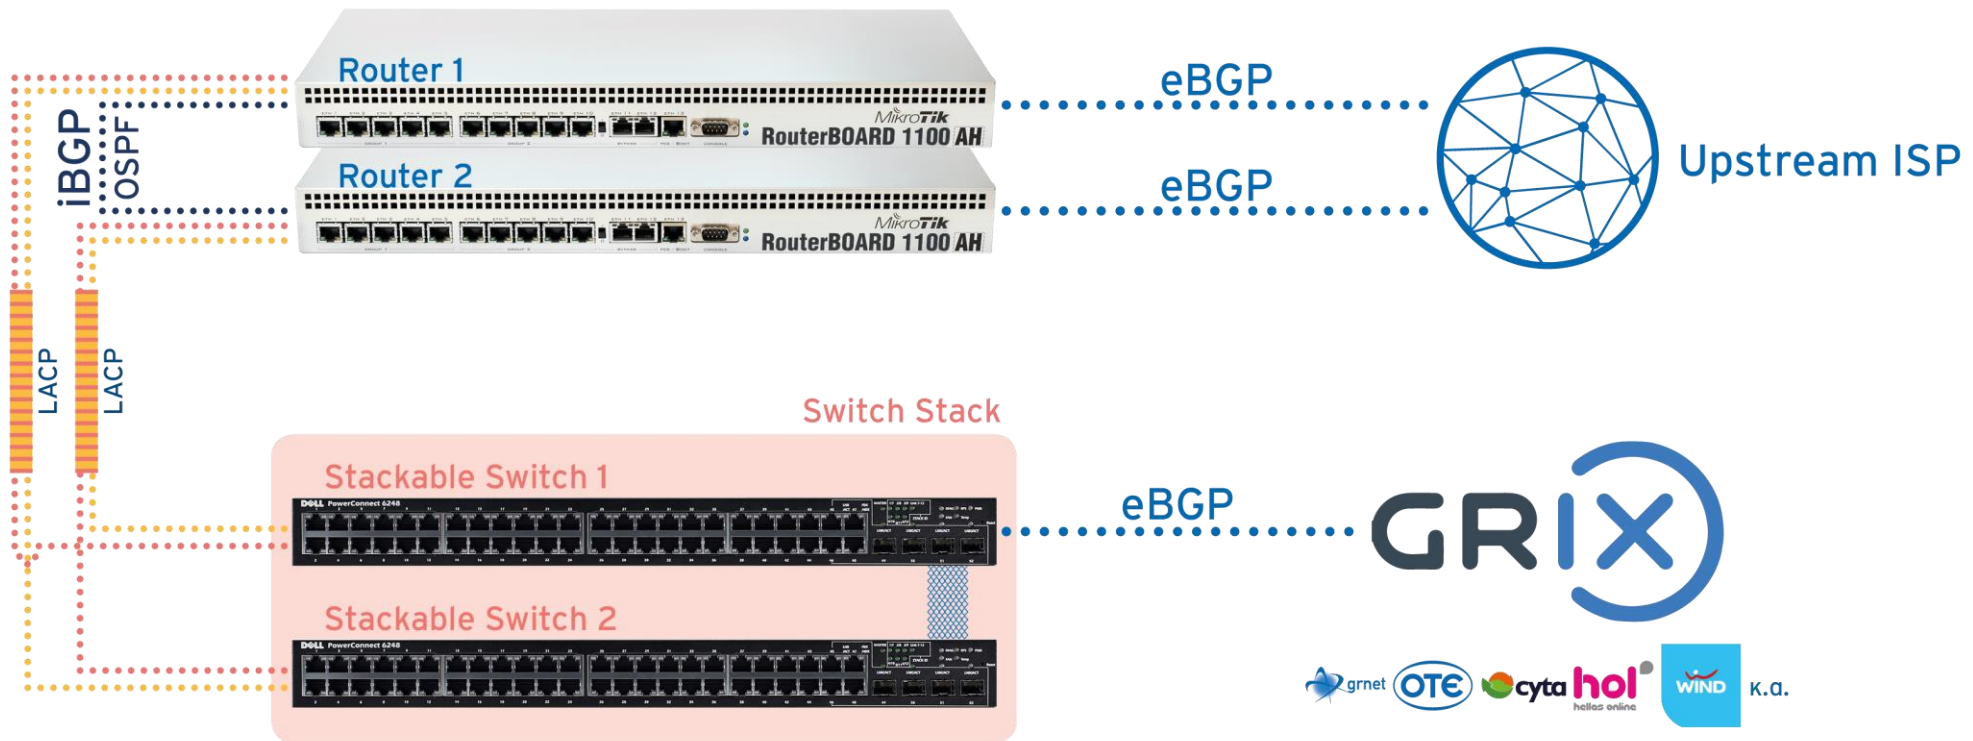
HA Mikrotik-Based Router Infrastructure



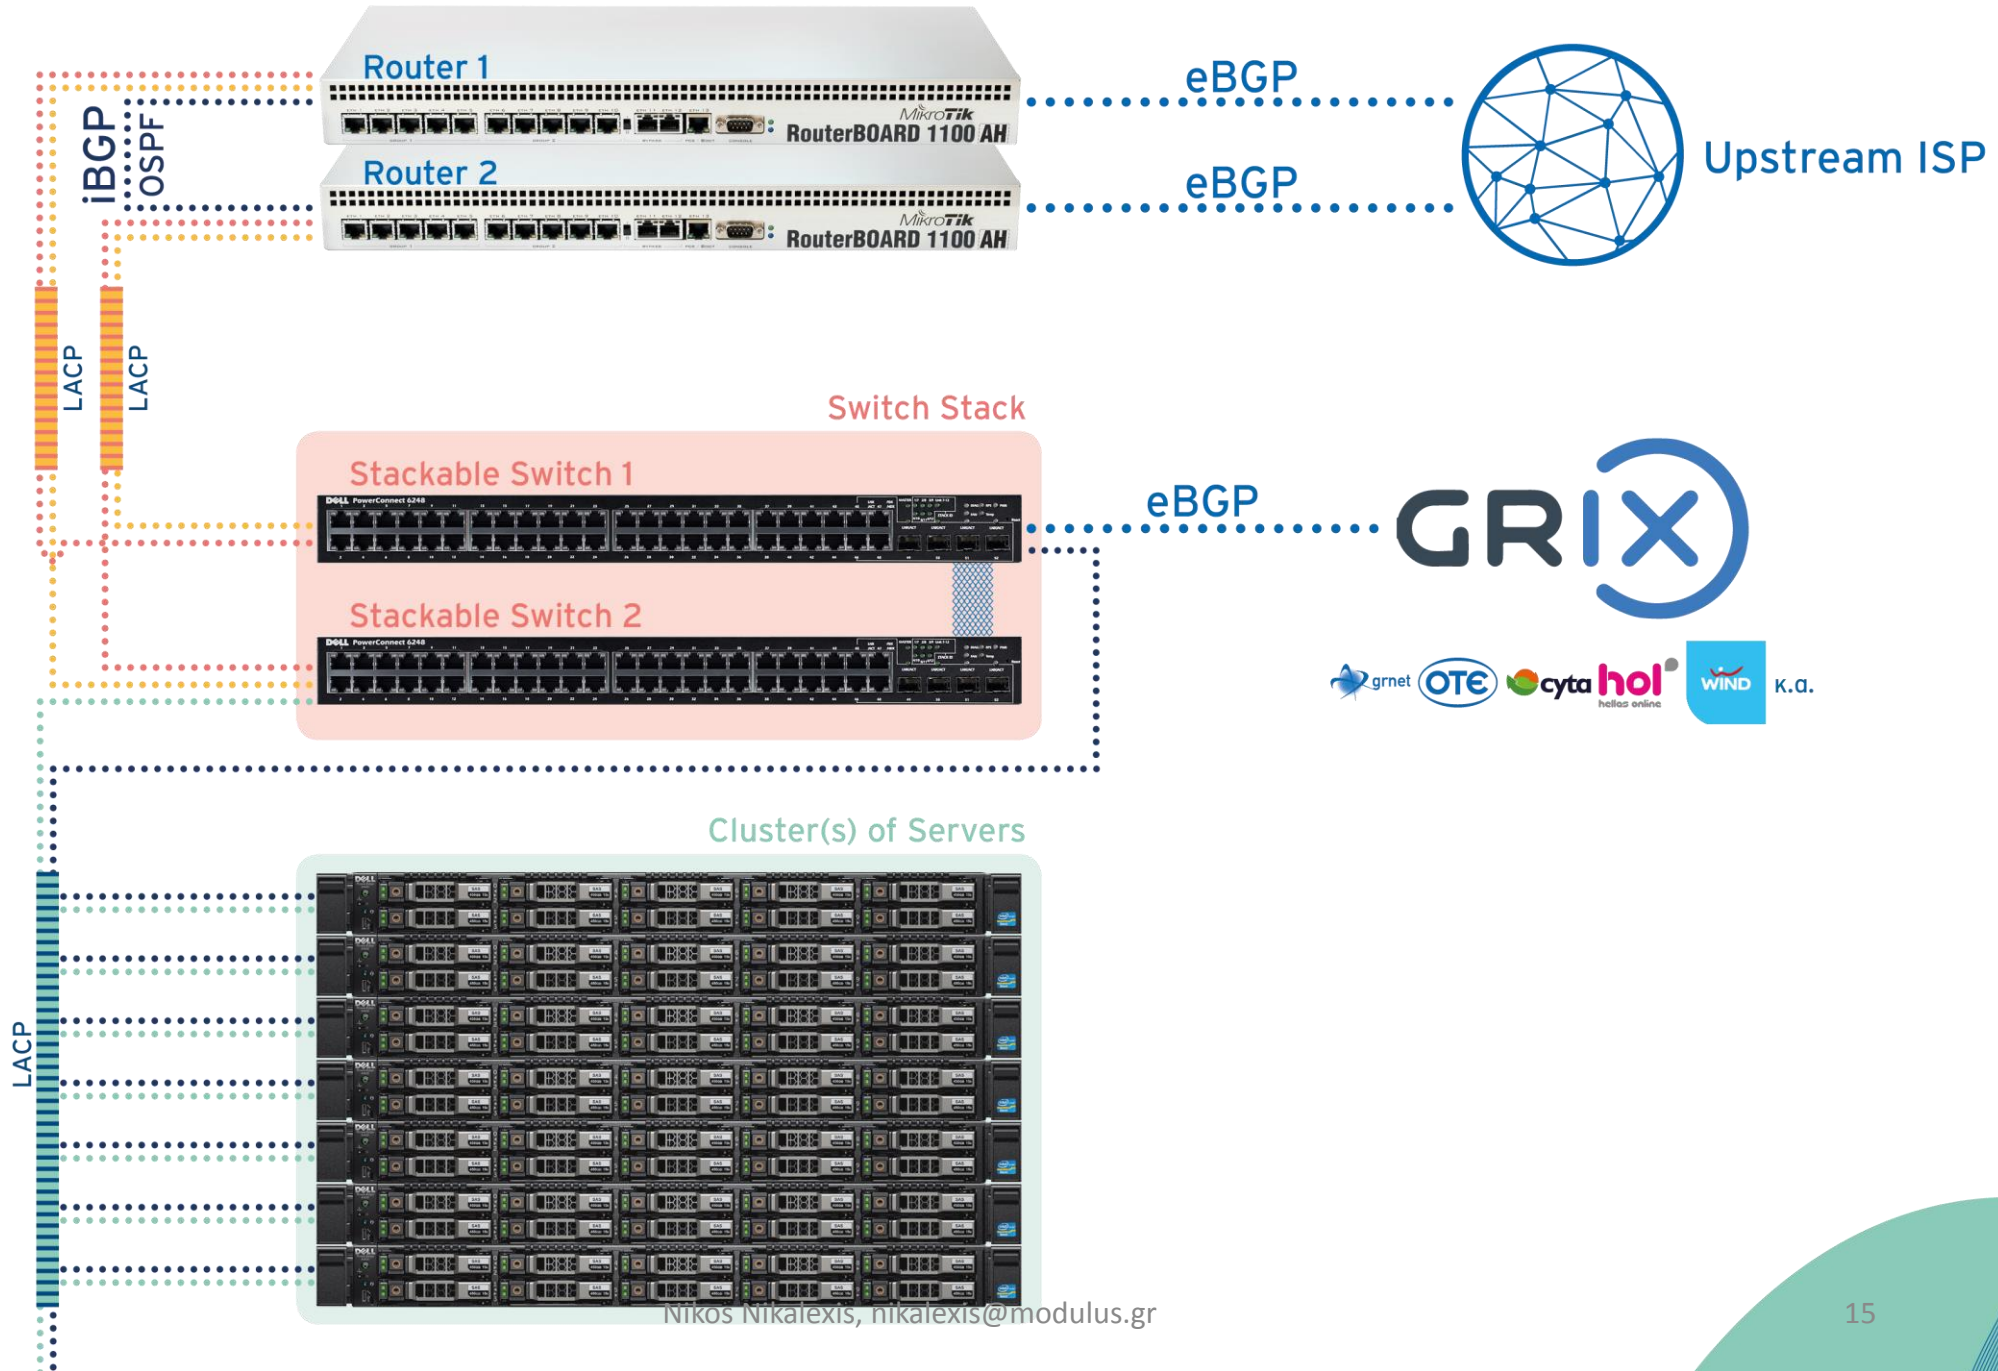
HA Mikrotik-Based Router Infrastructure



HA Mikrotik-Based Router Infrastructure



HA Mikrotik-Based Router Infrastructure



Mikrotik RouterOS setup

Interfaces, Bonding, VRRP, IP Addresses, Dynamic Routing, Traffic flow, Configuration Synchronization, Automatic Backup



Bonding

Interface <bonding1-switch>

General Bonding Traffic

Slaves: ether12-switch-6248
ether13-switch-6224

Mode: 802.3ad

Primary: none

Link Monitoring: mii

Transmit Hash Policy: layer 3 and 4

Down Delay: 0 ms

Up Delay: 0 ms

LACP Rate: 1 s

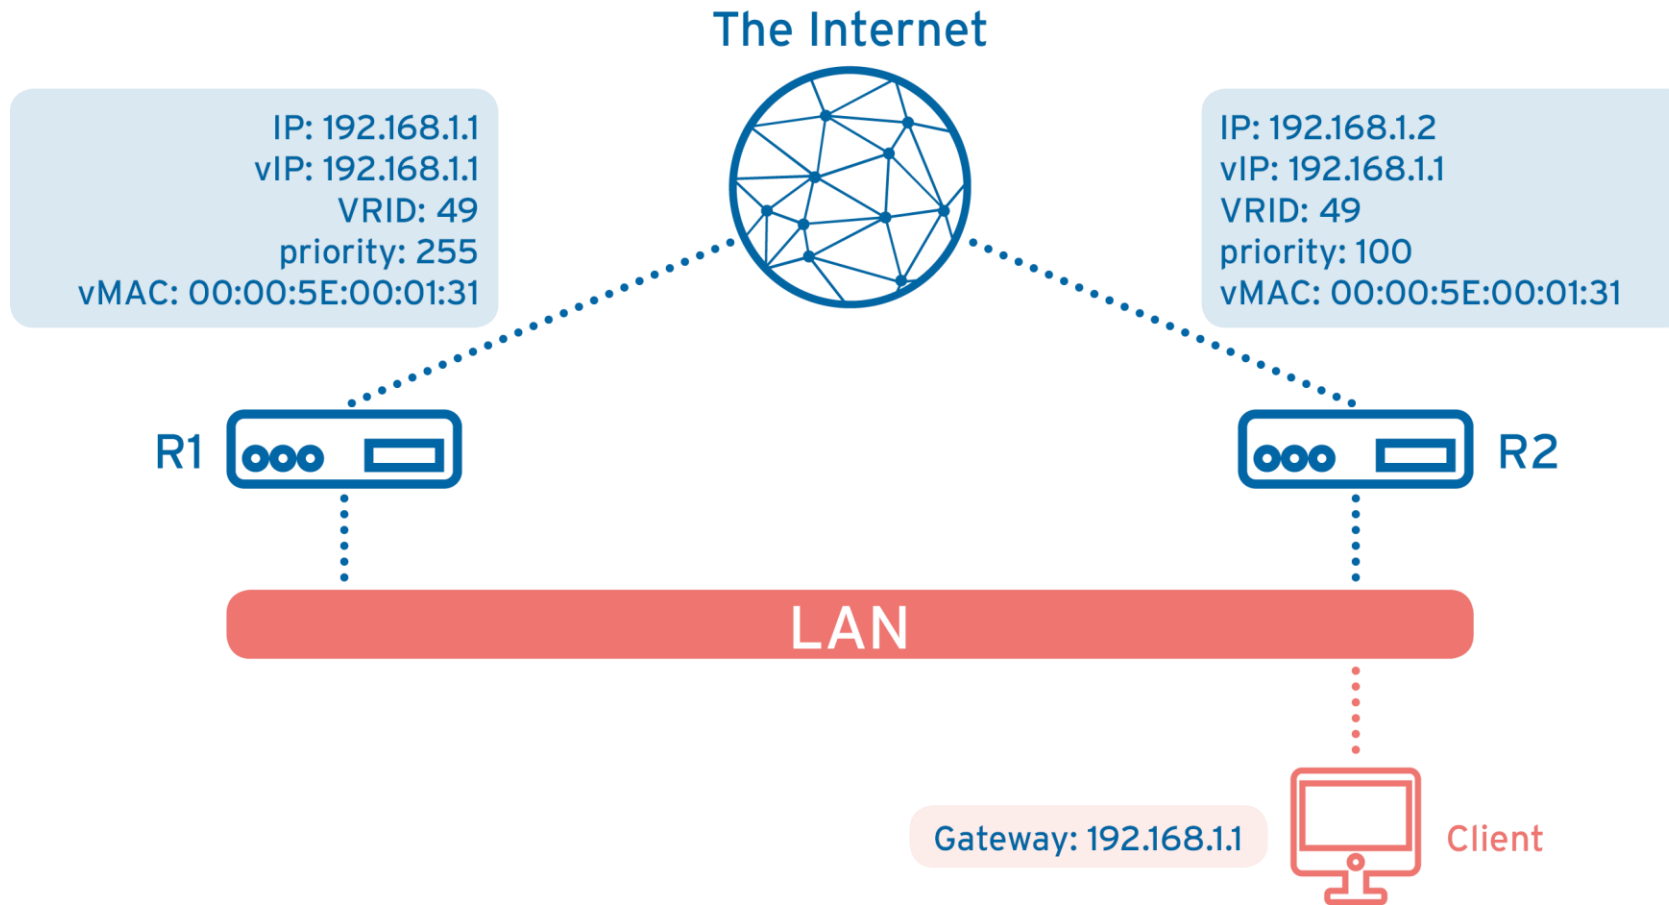
MII Interval: 100 ms

OK
Cancel
Apply
Disable
Comment
Copy
Remove
Torch

enabled running slave



VRRP (1/4)



VRRP (2/4)

- Features
 - Automatic Master / Backup mode
 - Optional preemption mode
- Pros
 - Easy configuration
 - Small transition time (a few seconds)
- Cons
 - Needs a separate IP for each router
 - Plus one for the Virtual IP (gateway)
- Summary
 - For every LAN with two redundant routers, 5 IPs are wasted:
 - Network, Broadcast, Virtual IP, 2x Router IPs
 - For large subnets ($> /26$), this is not a big problem
 - Considering recent IPv4 space exhaustion, we had to seek a smarter solution



VRRP (3/4)

A solution hidden in RouterOS!!!

Undocumented but working

- Setup only one VRRP interface (in private space?)
- Set this interface as a child for your VLANs
- When VRRP is in MASTER mode:
 - Every child VLAN is RUNNING
 - IP addresses on that VLAN interface are ACTIVE
- When VRRP is in BACKUP mode:
 - Every child VLAN is DOWN
 - IP addresses on that VLAN interface are INVALID



VRRP (4/4)

Interface <vrrp1-switch>

General VRRP Scripts Traffic

Interface: bonding1-switch

VRID: 22

Priority: 100

Interval: 1.00 s

Preemption Mode

Authentication

none simple ah

Password:

Version: 2

V3 Protocol: IPv4

OK Cancel Apply Disable Comment Copy Remove Torch

enabled running slave master

Interface <vrrp1-switch>

General VRRP Scripts Traffic

Interface: bonding1-switch

VRID: 22

Priority: 50

Interval: 1.00 s

Preemption Mode

Authentication

none simple ah

Password:

Version: 2

V3 Protocol: IPv4

OK Cancel Apply Enable Comment Copy Remove Torch

disabled running slave



Interfaces overview

Interface List

Interface Ethernet EoIP Tunnel IP Tunnel GRE Tunnel VLAN VRRP Bonding LTE

+ - ✓ ✗ 📄 🔍

	Name	Type	L2 MTU	Tx	Rx	Tx Packet (p/s)	Rx Packet (p/s)
R	↔ bonding1-switch	Bonding		19.6 Mbps	15.3 Mbps	3,670	
RM	↔ vrrp1-switch	VRRP					
R	↔ vlan64	VLAN					
R	↔ vlan72	VLAN					
R	↔ vlan80	VLAN					
R	↔ vlan132	VLAN					
R	↔ vlan133	VLAN					
R	↔ vlan1126	VLAN					
R	↔ vlan1331	VLAN					
R	↔ vlan1332	VLAN					
R	↔ vlan3000-grix	VLAN					
R	↔ ether1-routers	Ethernet	1598				
	↔ ether2	Ethernet	1598				
	↔ ether3	Ethernet	1598				
	↔ ether4	Ethernet	1598				
	↔ ether5	Ethernet	1598				
	↔ ether6	Ethernet	1598				
	↔ ether7	Ethernet	1598				
	↔ ether8	Ethernet	1598				
	↔ ether9	Ethernet	1598				
	↔ ether10	Ethernet	1598				
R	↔ ether11-lamda-hellix	Ethernet	1600				
RS	↔ ether12-switch-6248	Ethernet	1600				
RS	↔ ether13-switch-6224	Ethernet	1600				



IP addresses overview

MASTER Router

Address	Network	Interface
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]

BACKUP Router

Address	Network	Interface
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]
[blurred]	[blurred]	[blurred]

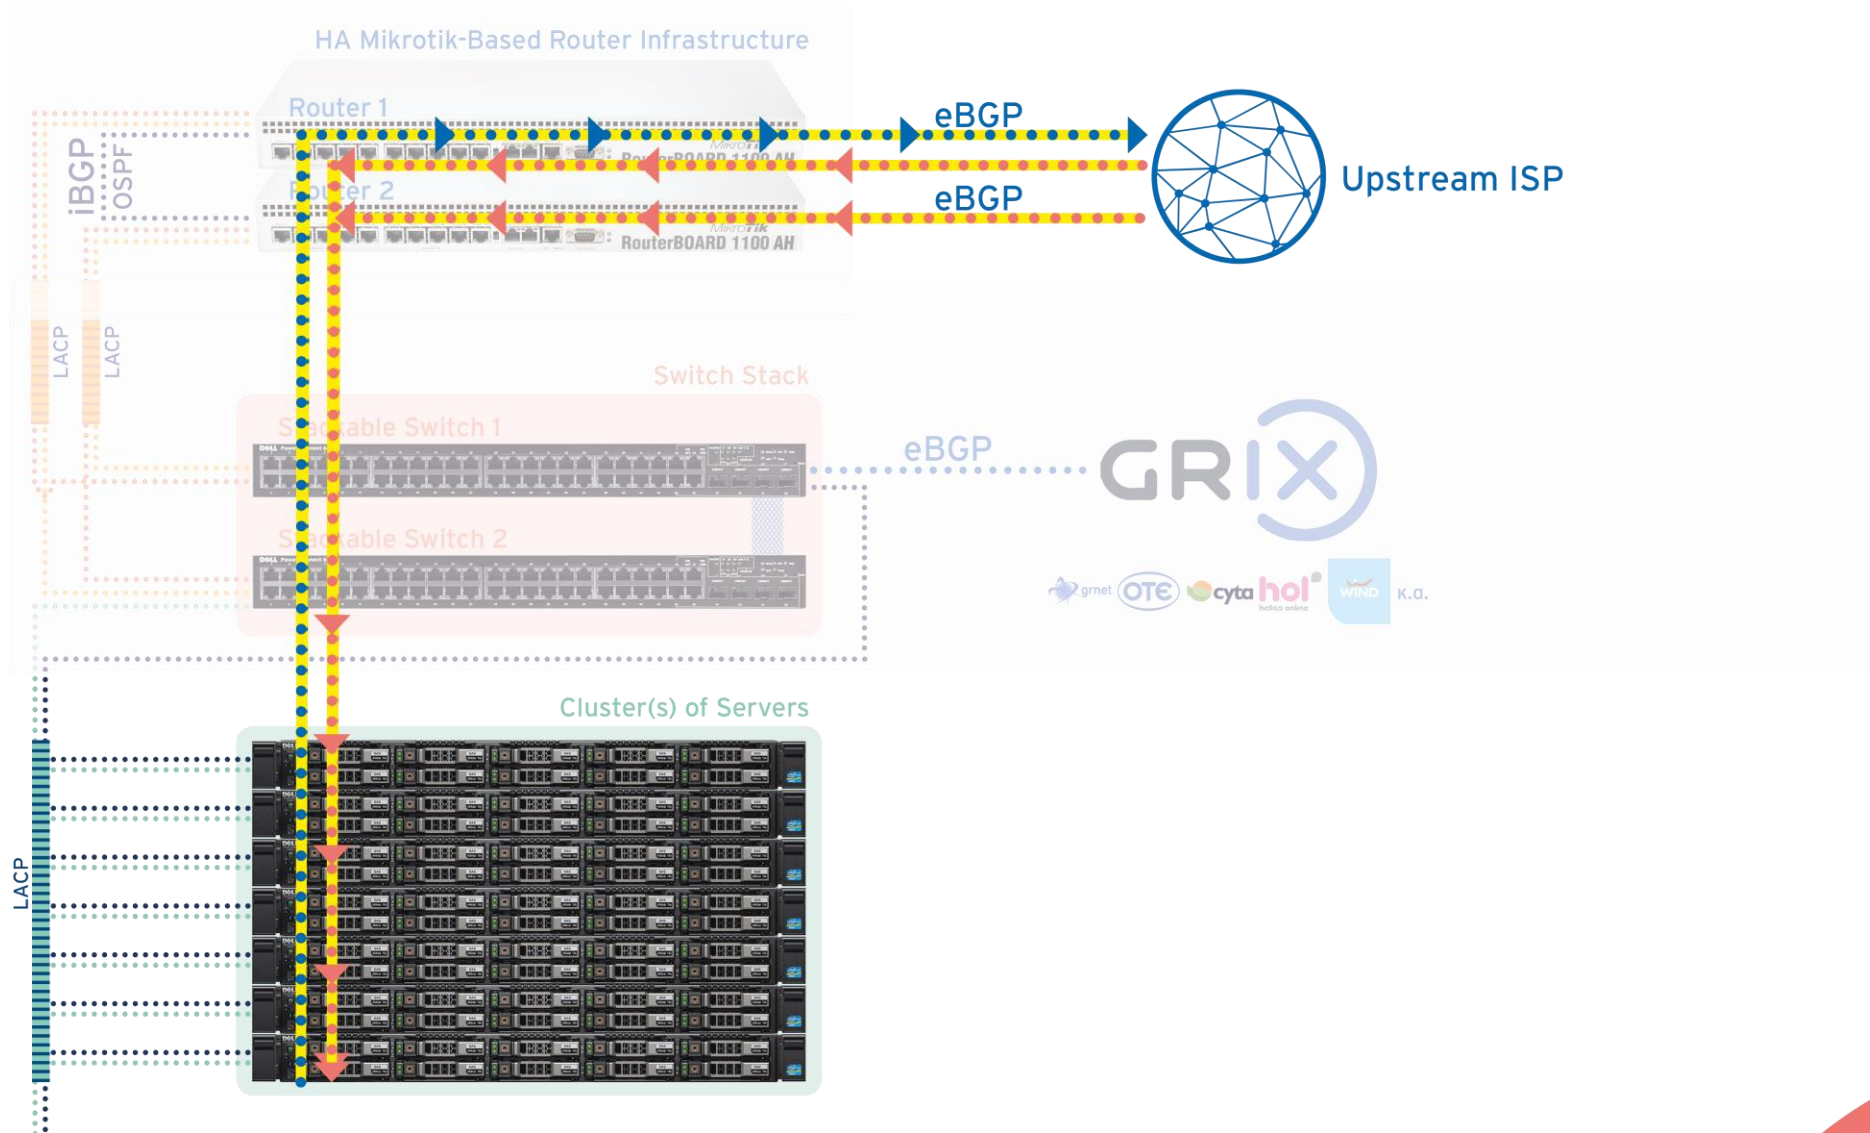


Dynamic Routing

- eBGP
 - Upstream eBGP for each Router
 - Connect each upstream link directly with each router
 - You don't lose access to your routers in a case of a hardware/software failure
 - This way, we avoid using a switch device for upstream connectivity
 - GR-IX eBGP through a VLAN configured on the Switch Stack
 - This is not an upstream interconnection, we can afford losing it
- iBGP / OSPF
 - Activated on both routers
- BFD with each peer (RFC 5880)
 - Rapid fault detection (< 1 second)



Traffic Flow



VRRP Scripting

- On MASTER:

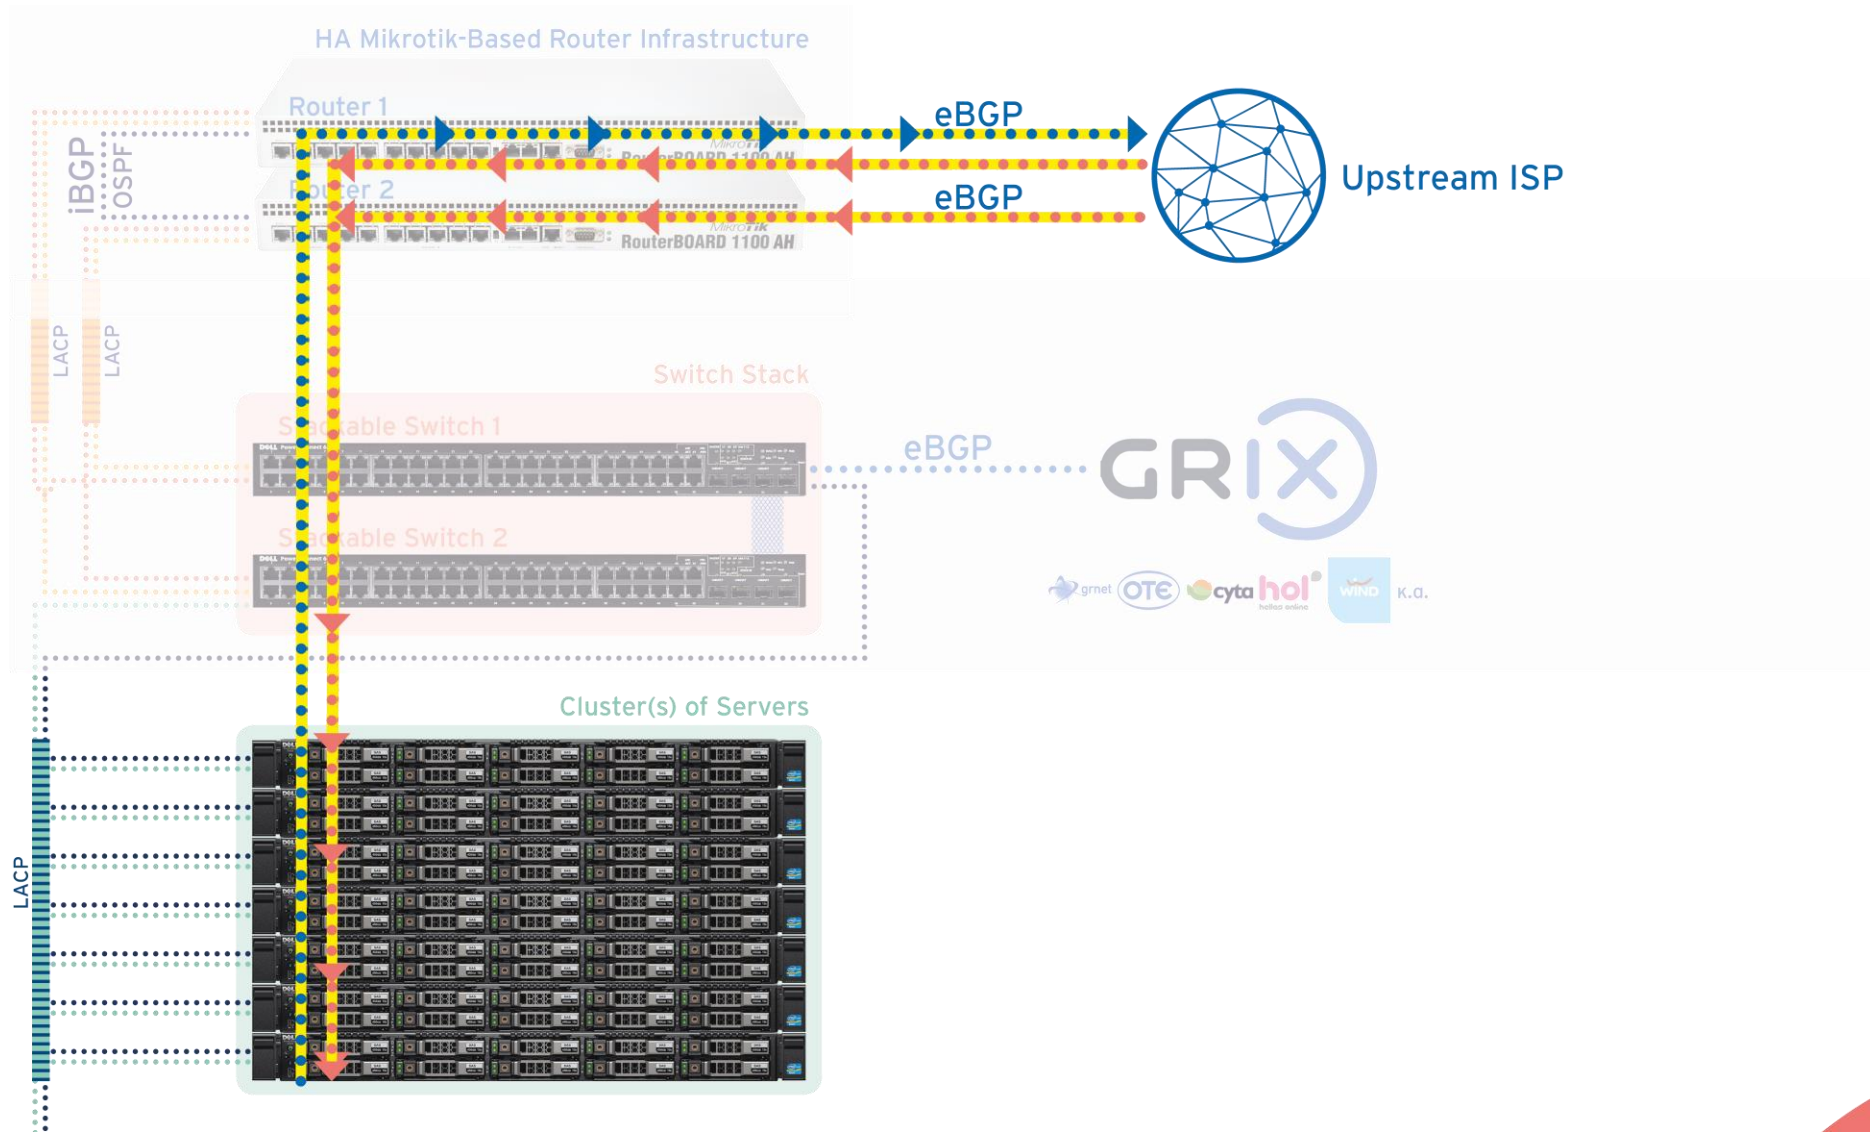
```
{  
  /routing filter set [find chain="providers-out" action="passthrough" set-bgp-med=200] set-  
  bgp-med=100;  
}
```

- On BACKUP:

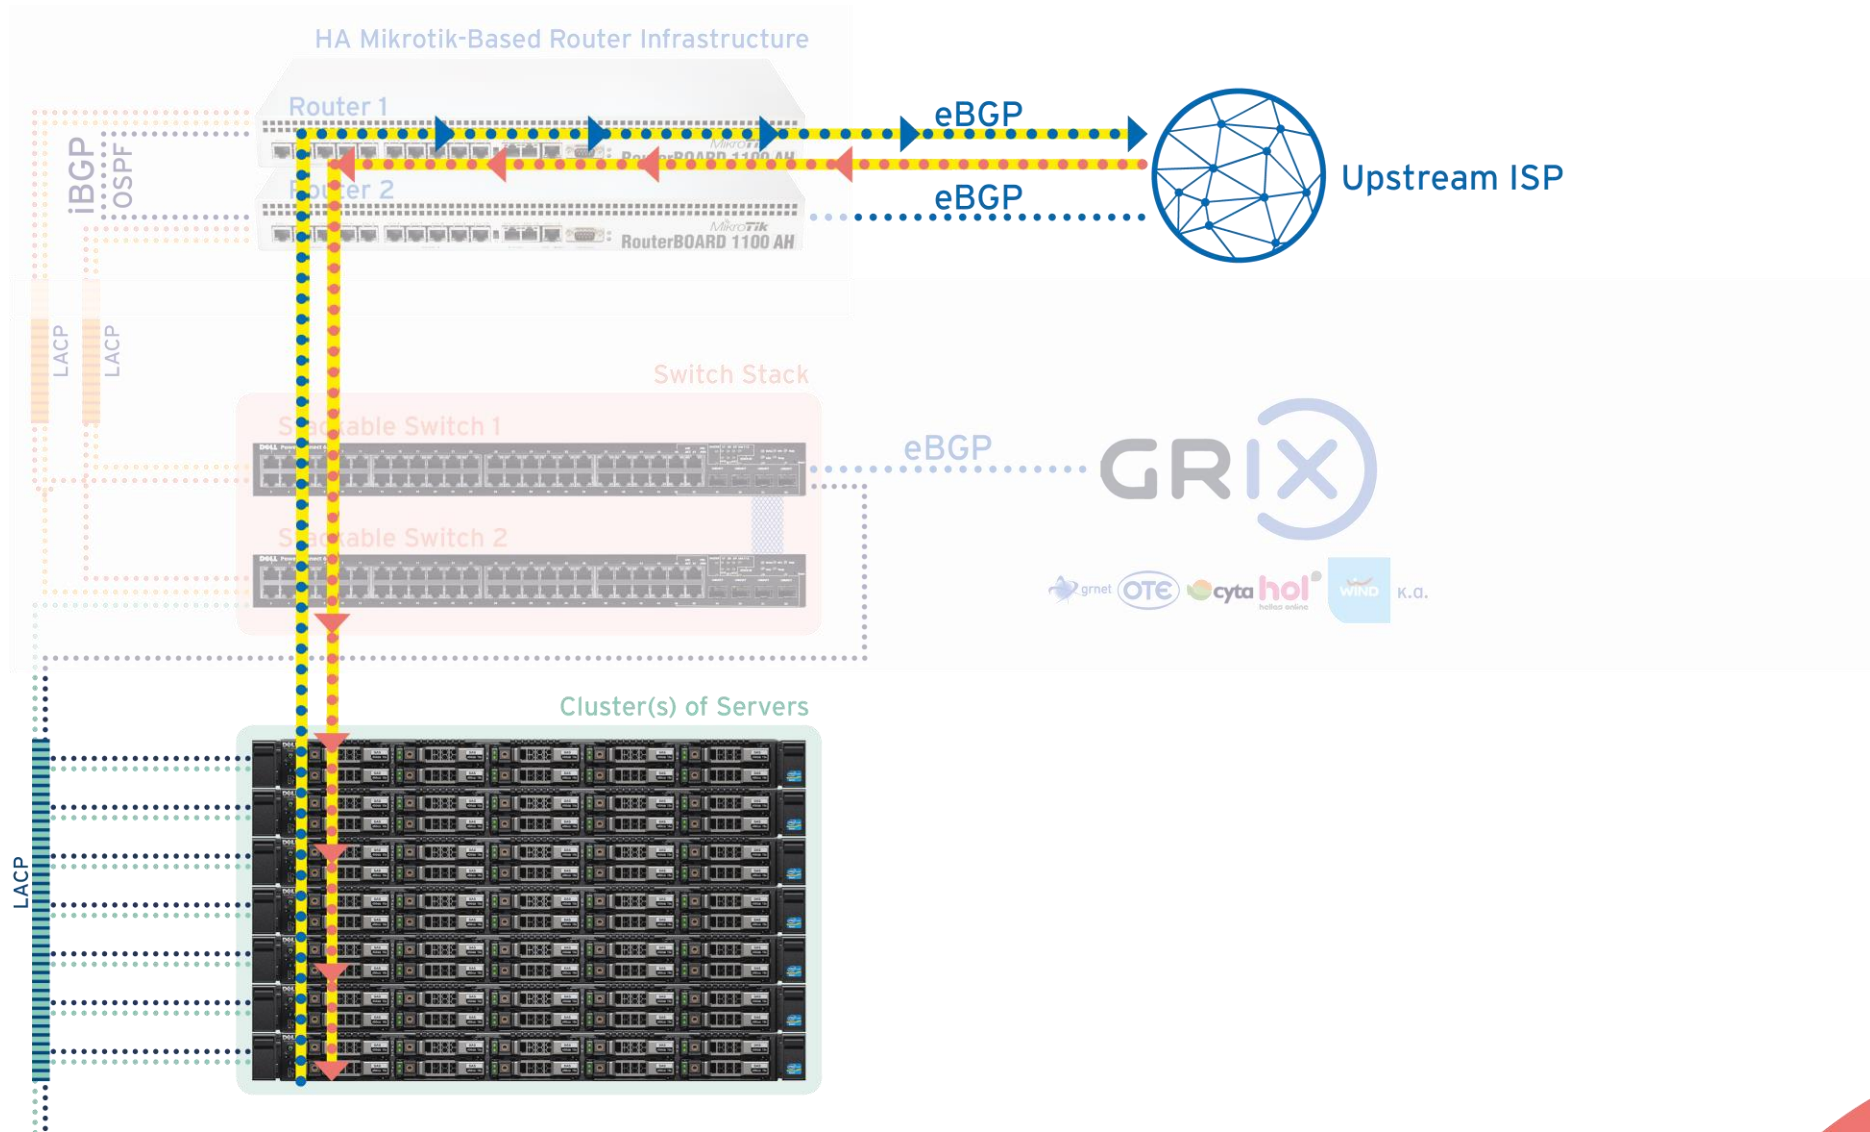
```
{  
  /routing filter set [find chain="providers-out" action="passthrough" set-bgp-med=100] set-  
  bgp-med=200;  
}
```



Traffic Flow (before scripting)



Traffic Flow Fixed (after scripting)



Configuration Synchronization (1/3)

- Our 2 Routers have:
 - Shared config
 - Interfaces
 - VLAN IP addresses
 - Firewall rules
 - QoS rules
 - Routing filters
 - Discreet config
 - VRRP Priority option
 - Non-VLAN IP addresses
 - Upstream eBGP configs



Configuration Synchronization (2/3)

- Develop a python script that:
 - Connects to each router through SSH
 - Exports the full config
 - Calculates diffs between configs and...
 - sends it in an e-mail to the admin team
- Run this script
 - Periodically to be up to date
 - Manually to check your setup on demand



Configuration Synchronization (3/3)

```
13 /interface vrrp
14 add interface=bonding1-switch name=vrrp1-switch on-backup="{\r\
15 \n/routing filter set [find chain=\"providers-out\" action=\"passthrough\"
16 \_set-bgp-med=100] set-bgp-med=200;\r\
17 \n/interface disable vpn-modulus;\r\
18 \n}\r\
19 \n" on-master="{\r\
20 \n/routing filter set [find chain=\"providers-out\" action=\"passthrough\"
21 \_set-bgp-med=200] set-bgp-med=100;\r\
22 \n/interface enable vpn-modulus;\r\
23 \n}\r\
n 24 \n" preemption-mode=no version=2 vrid=22
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
n 44
n 45
n 46
47
48
49
50
51
n 52 add limit-at=10M/10M max-limit=10M/10M name=
53 target= /32
54 add limit-at=10M/10M max-limit=10M/10M name= target=\
55 /32
```

```
13 /interface vrrp
14 add disabled=yes interface=bonding1-switch name=vrrp1-switch on-backup="{\r\
15 \n/routing filter set [find chain=\"providers-out\" action=\"passthrough\"
16 \_set-bgp-med=100] set-bgp-med=200;\r\
17 \n/interface disable vpn-modulus;\r\
18 \n}\r\
19 \n" on-master="{\r\
20 \n/routing filter set [find chain=\"providers-out\" action=\"passthrough\"
21 \_set-bgp-med=200] set-bgp-med=100;\r\
22 \n/interface enable vpn-modulus;\r\
23 \n}\r\
n 24 \n" preemption-mode=no priority=50 version=2 vrid=22
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
n 44
n 45
n 46
47
48
49
50
51
n
```



Automatic Backup

- Develop a python script that:
 - Connects to each router through SSH
 - Exports the full config
 - Creates a backup file
 - Transfers the backup file to a safe location via FTP
- Run this script
 - Periodically, e.g. every 2 days
 - We schedule different days for each router
 - Avoid bugs in export and backup



Conclusion

Testing, Goals achieved, ToDo, Feature Requests



Testing

- MASTER router failure
 - <3 seconds downtime until BACKUP router takeover
 - <1 second downtime until BFD marks our peer as down
- BACKUP router failure
 - No downtime
- MASTER switch failure
 - <10 seconds downtime on some sessions until LACP recovers on backplane
- BACKUP switch failure
 - No downtime



Goals achieved

- No SPOF Network
- Network High Availability
- Configuration Synchronization
- Configuration Backup with Easy Restoration
- Low cost, commodity hardware



ToDo

- Use Ansible
 - Centrally manage all HA routers & more...
 - Store all configuration data in the Ansible inventory
 - Use group variables for common config
 - Use host variables for discreet config
 - Use GIT for keeping track of changes
 - Write a module talking to RouterOS API
 - Write roles for master / backup configurations
 - Write playbook for deploying HA router infrastructure
- Upgrade to CCR
 - More powerful
 - Redundant Power Supply
 - Supports SFP interfaces



Feature Requests

- Hardware
 - No SPOF / Single Unit Fully Redundant Router
 - 2xPSU, 2xBackplanes, 2xLinecards
 - Stackable switches
- Software
 - Configuration Synchronization
 - Single interface point (winbox, console, api etc)
 - Connection tracking synchronization (like linux conntrackd) to achieve:
 - Connection-based firewall rules
 - NATed connections



Thank you!

Any Questions?

