

Multi-region Multi-instance MSTP on MikroTik

A deep dive into STP and MSTP on RouterOS

Who am I?

Tomas Kirnak

Team Lead @ Unimus

System & Network Architect
Automation & Monitoring

Ex-MikroTik Trainer, Consultant



About Unimus

Network Automation

(Mass Config Push, network-wide conf. search, etc.)

Configuration Management

(change notifications, historic diffs, etc.)

Disaster recovery

(configuration backup, rapid redeploy)



Unimus

Note for posterity

If you find this presentation online in a .pdf, please watch the video

Proper explanations to every slide and much more information available

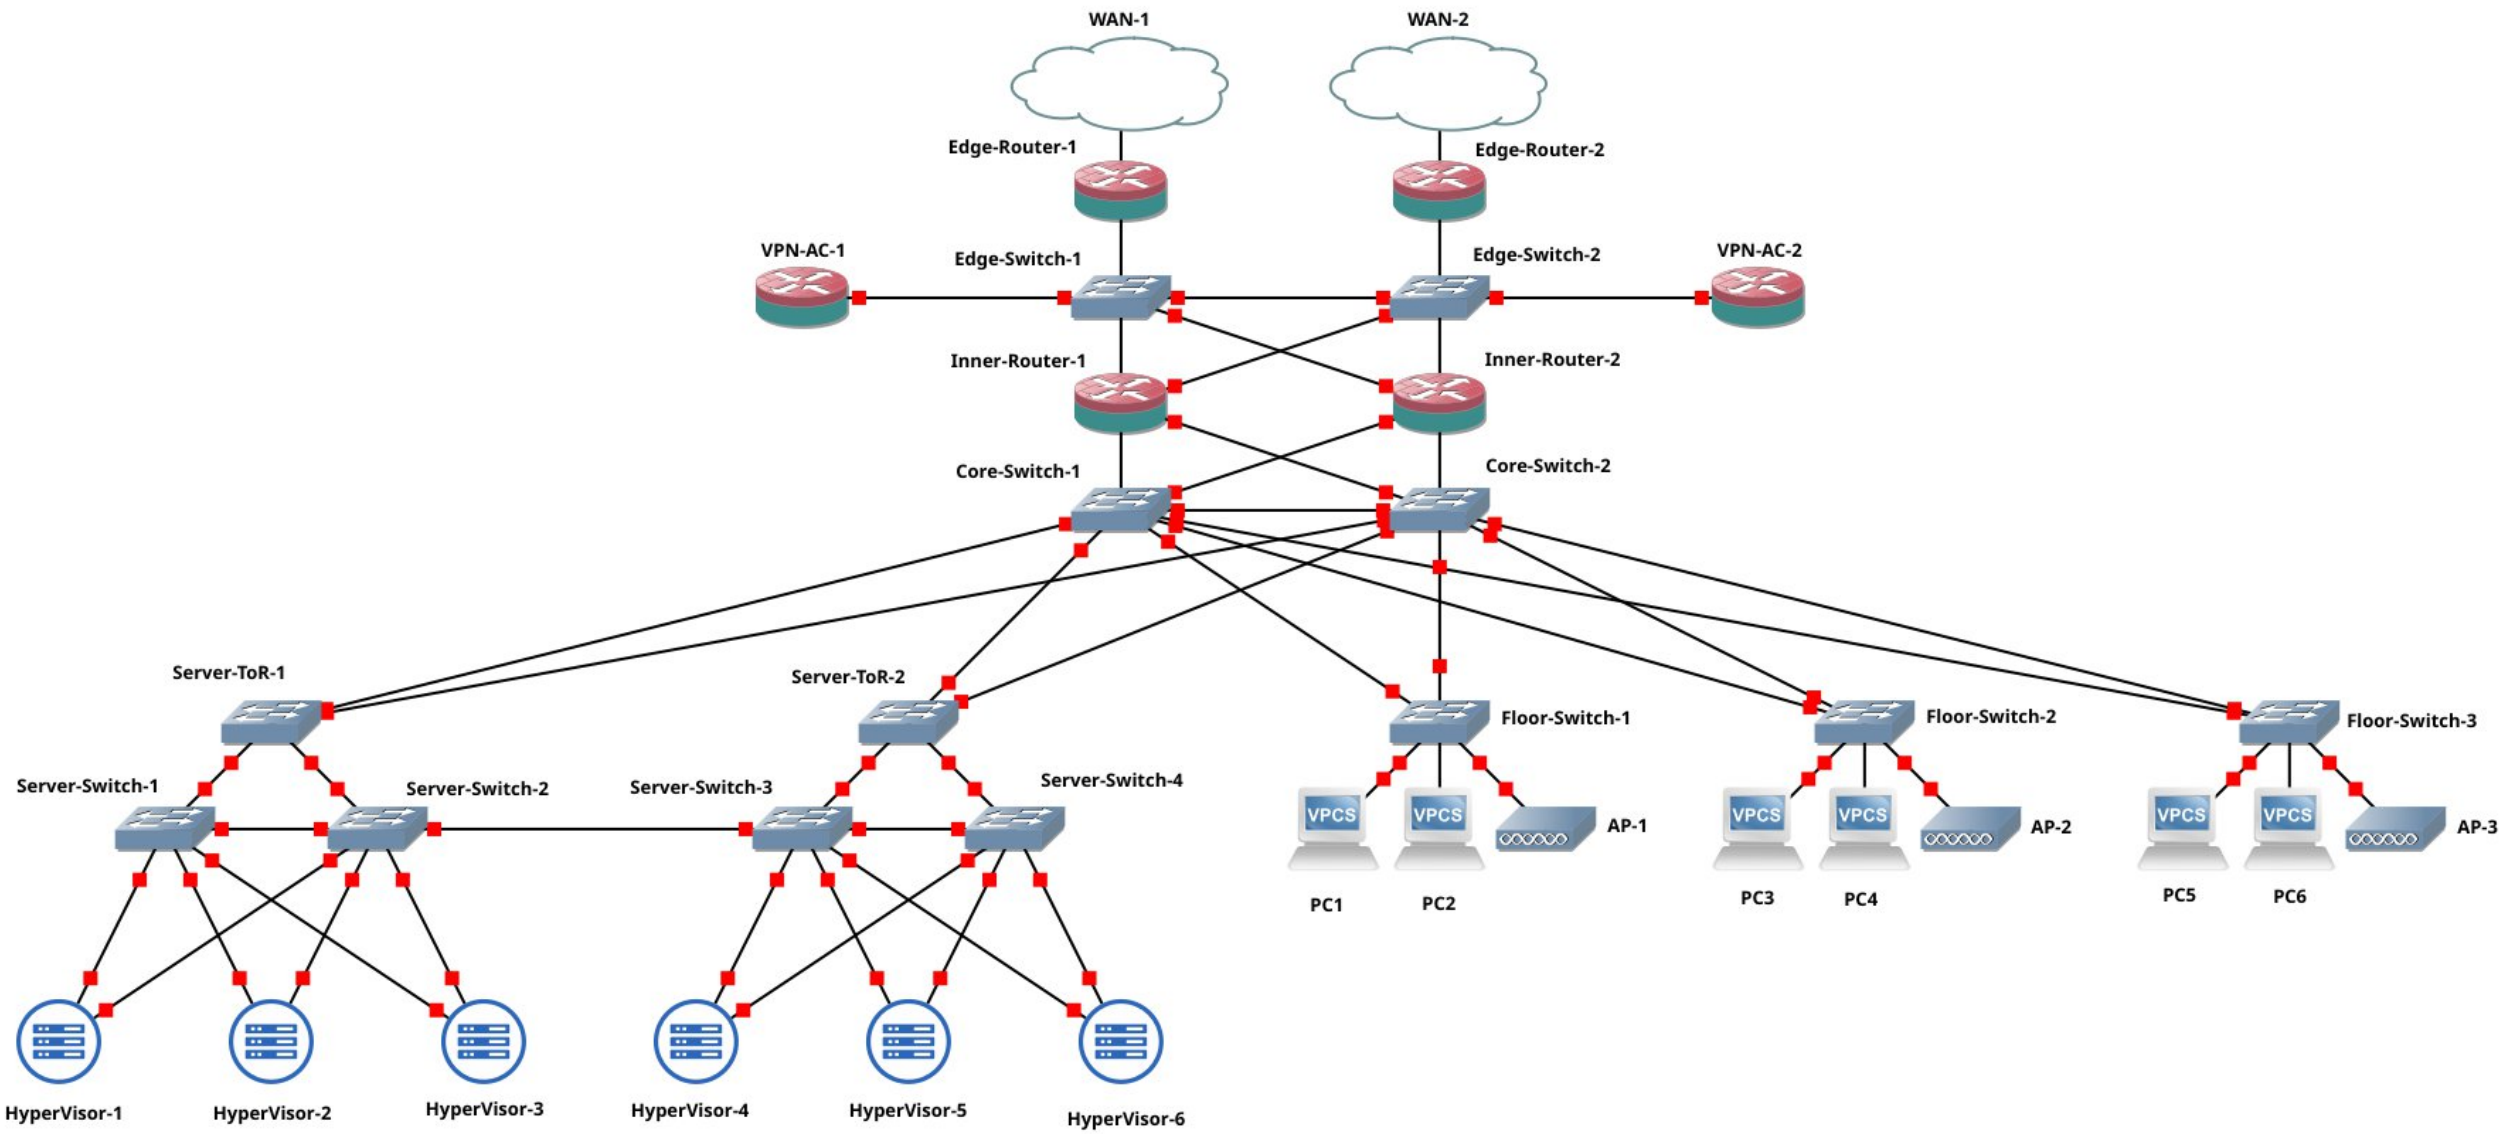
<https://www.youtube.com/c/TomasKirnak/videos>

Why care about MSTP?

MSTP is an often misunderstood, often frowned upon, and not commonly deployed

It's a cool piece of technology, with real use-cases, and solves many architectural challenges in complex networks

It's a very useful tool in your networking "toolkit"



Let's start with an STP intro

Spanning Tree Protocol (STP) is a core protocol in Ethernet networks, ensuring loop-free paths by blocking redundant links.

Due to link blocking, STP decreases available links (and therefore throughput) in the network

STP struggles with large network topologies

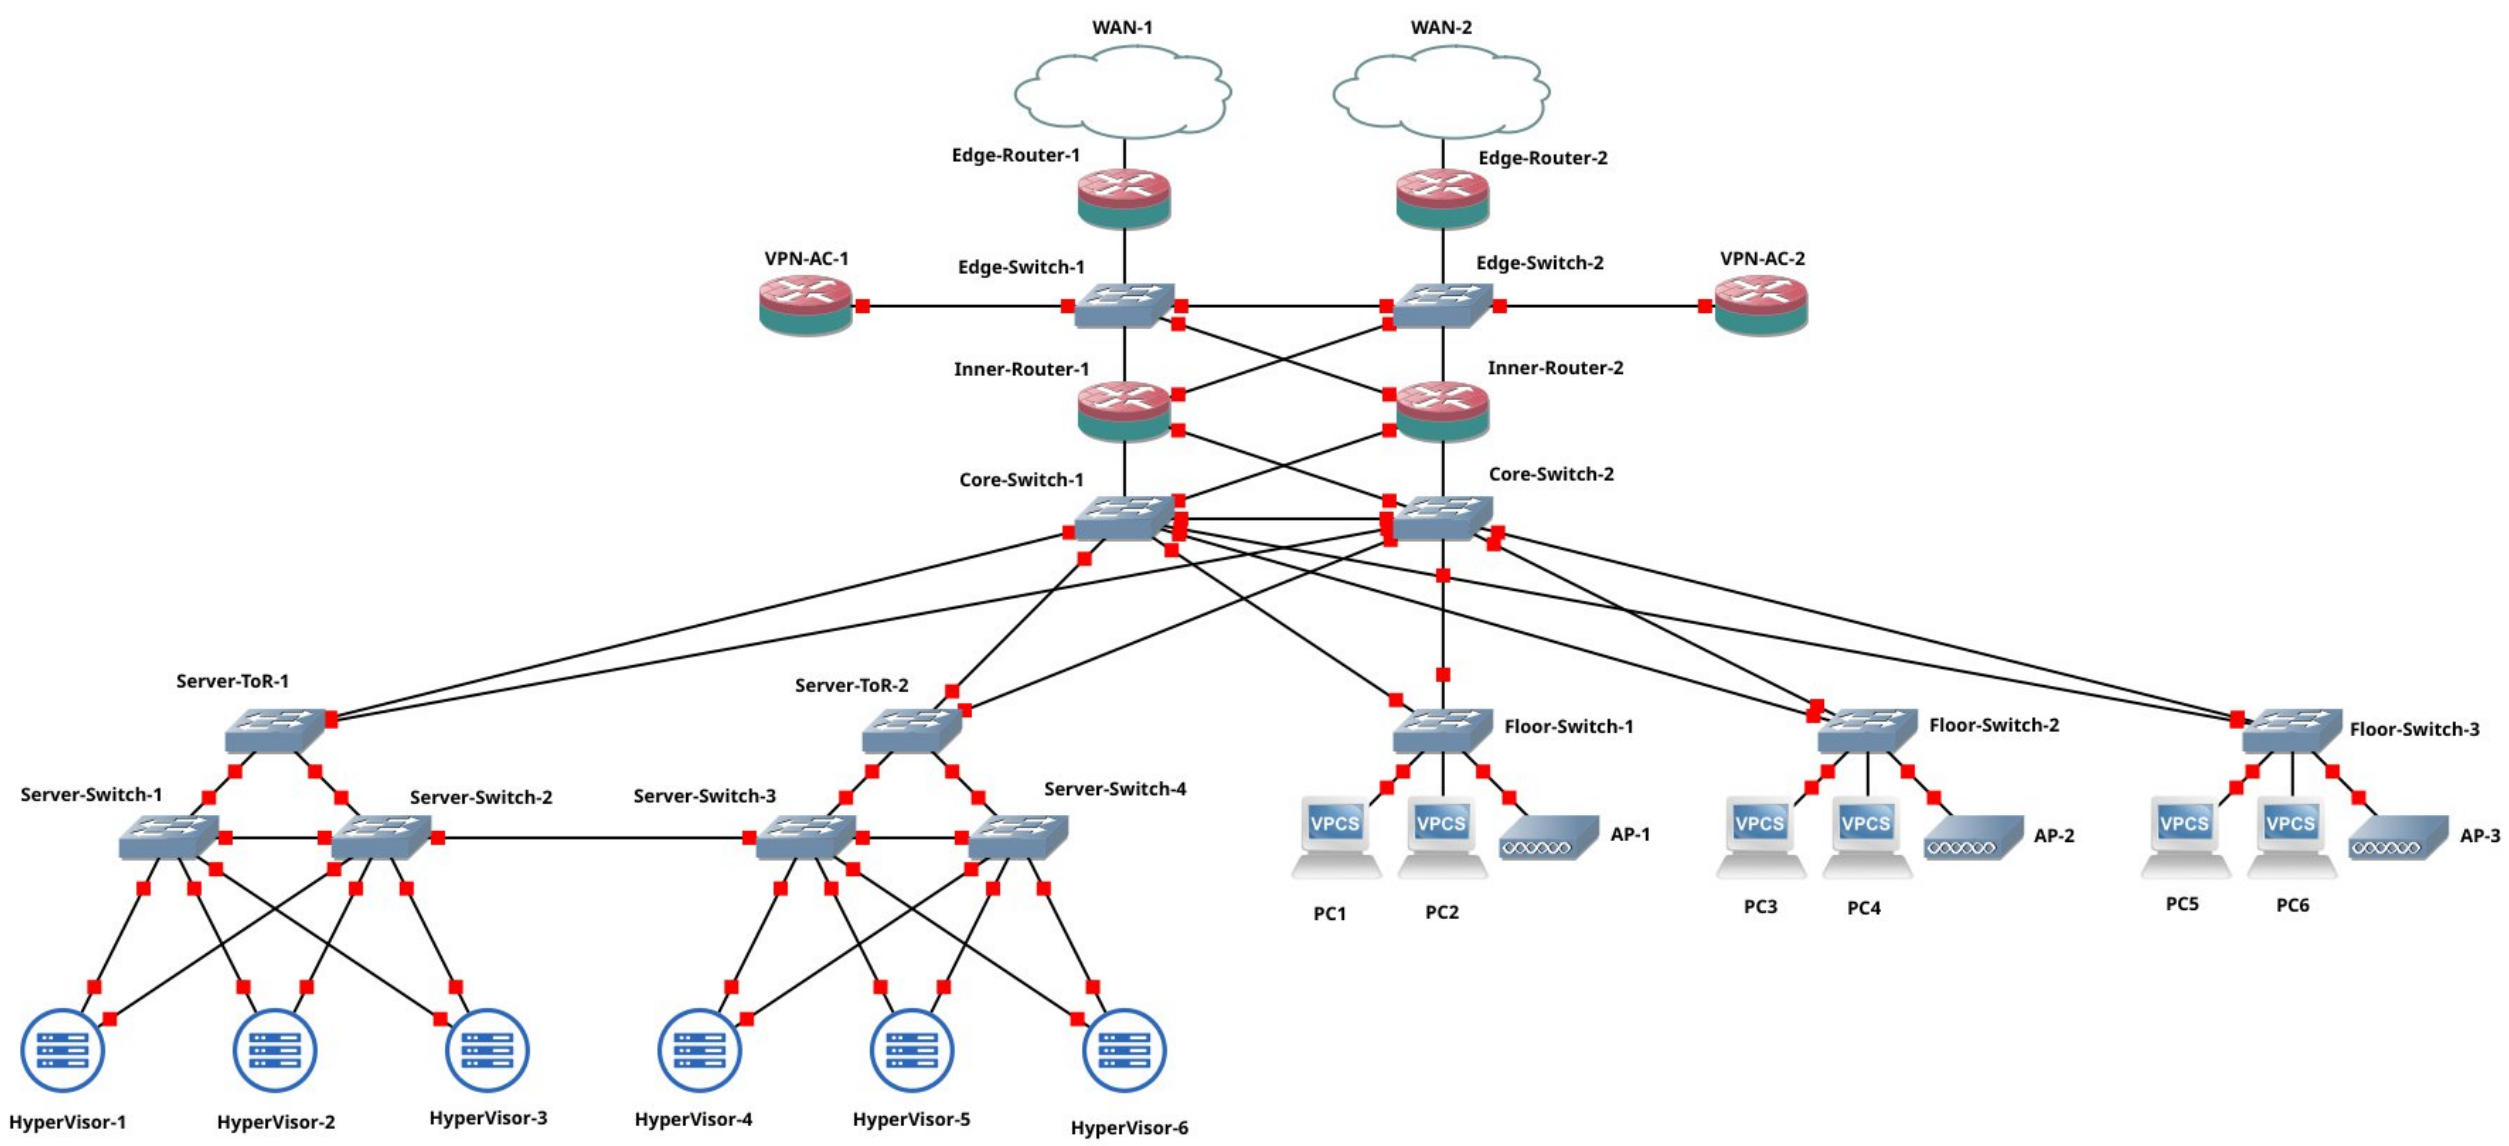
How does STP work?

STP select a Root Bridge (switch), and calculates a non-looping topology for the network

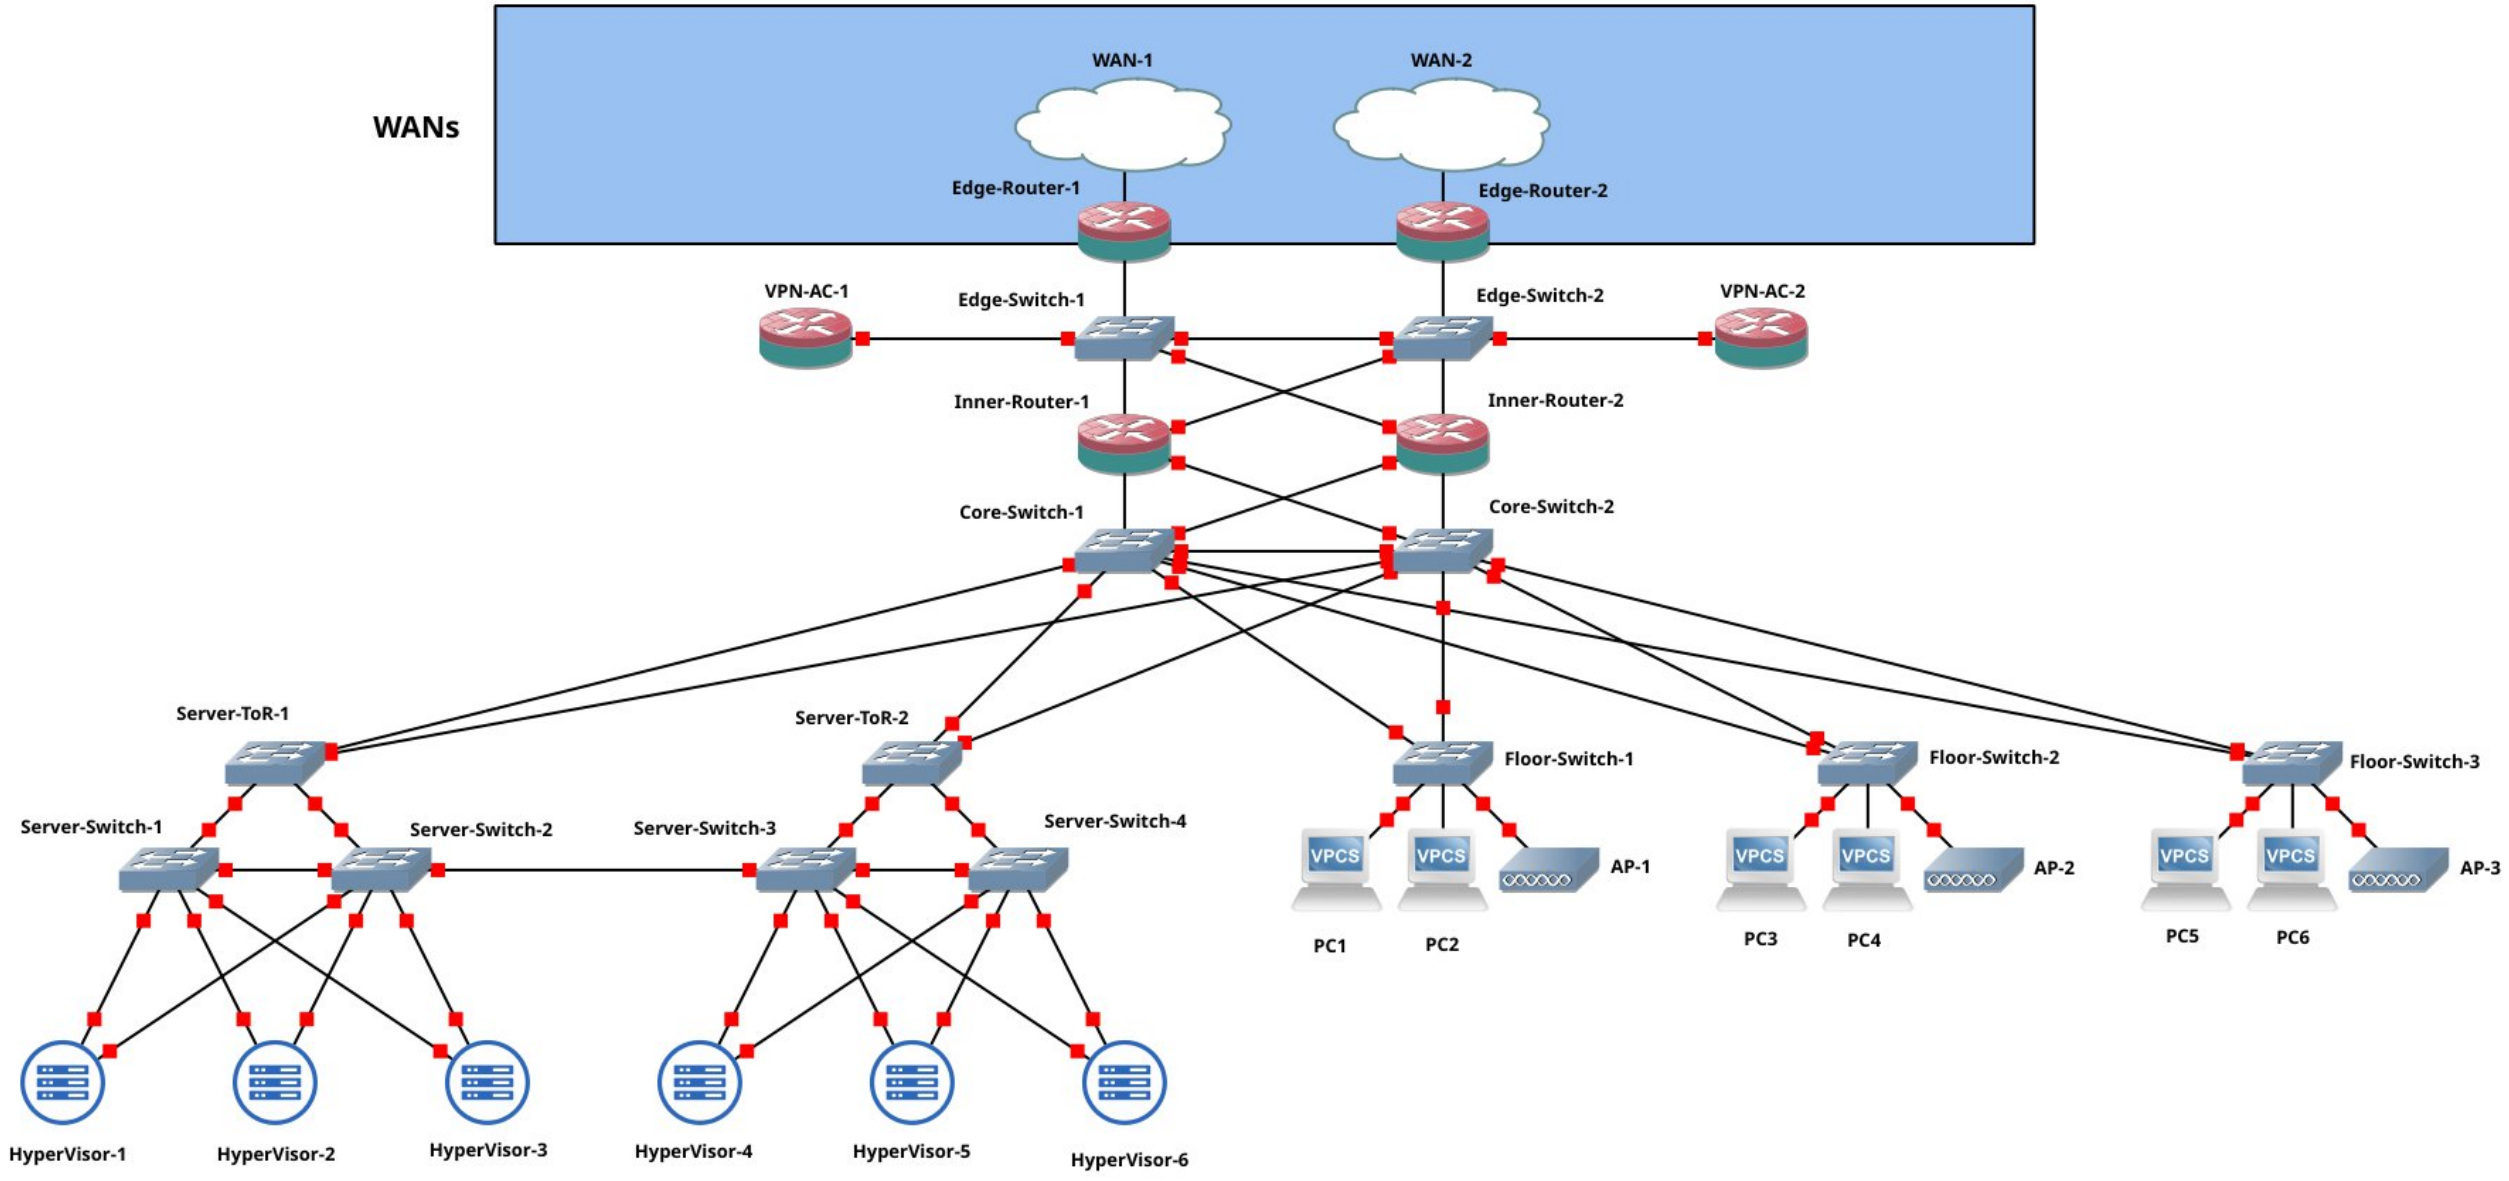
All bridges (switches) calculate the shortest path to the Root Bridge, and all other links between bridges are disabled

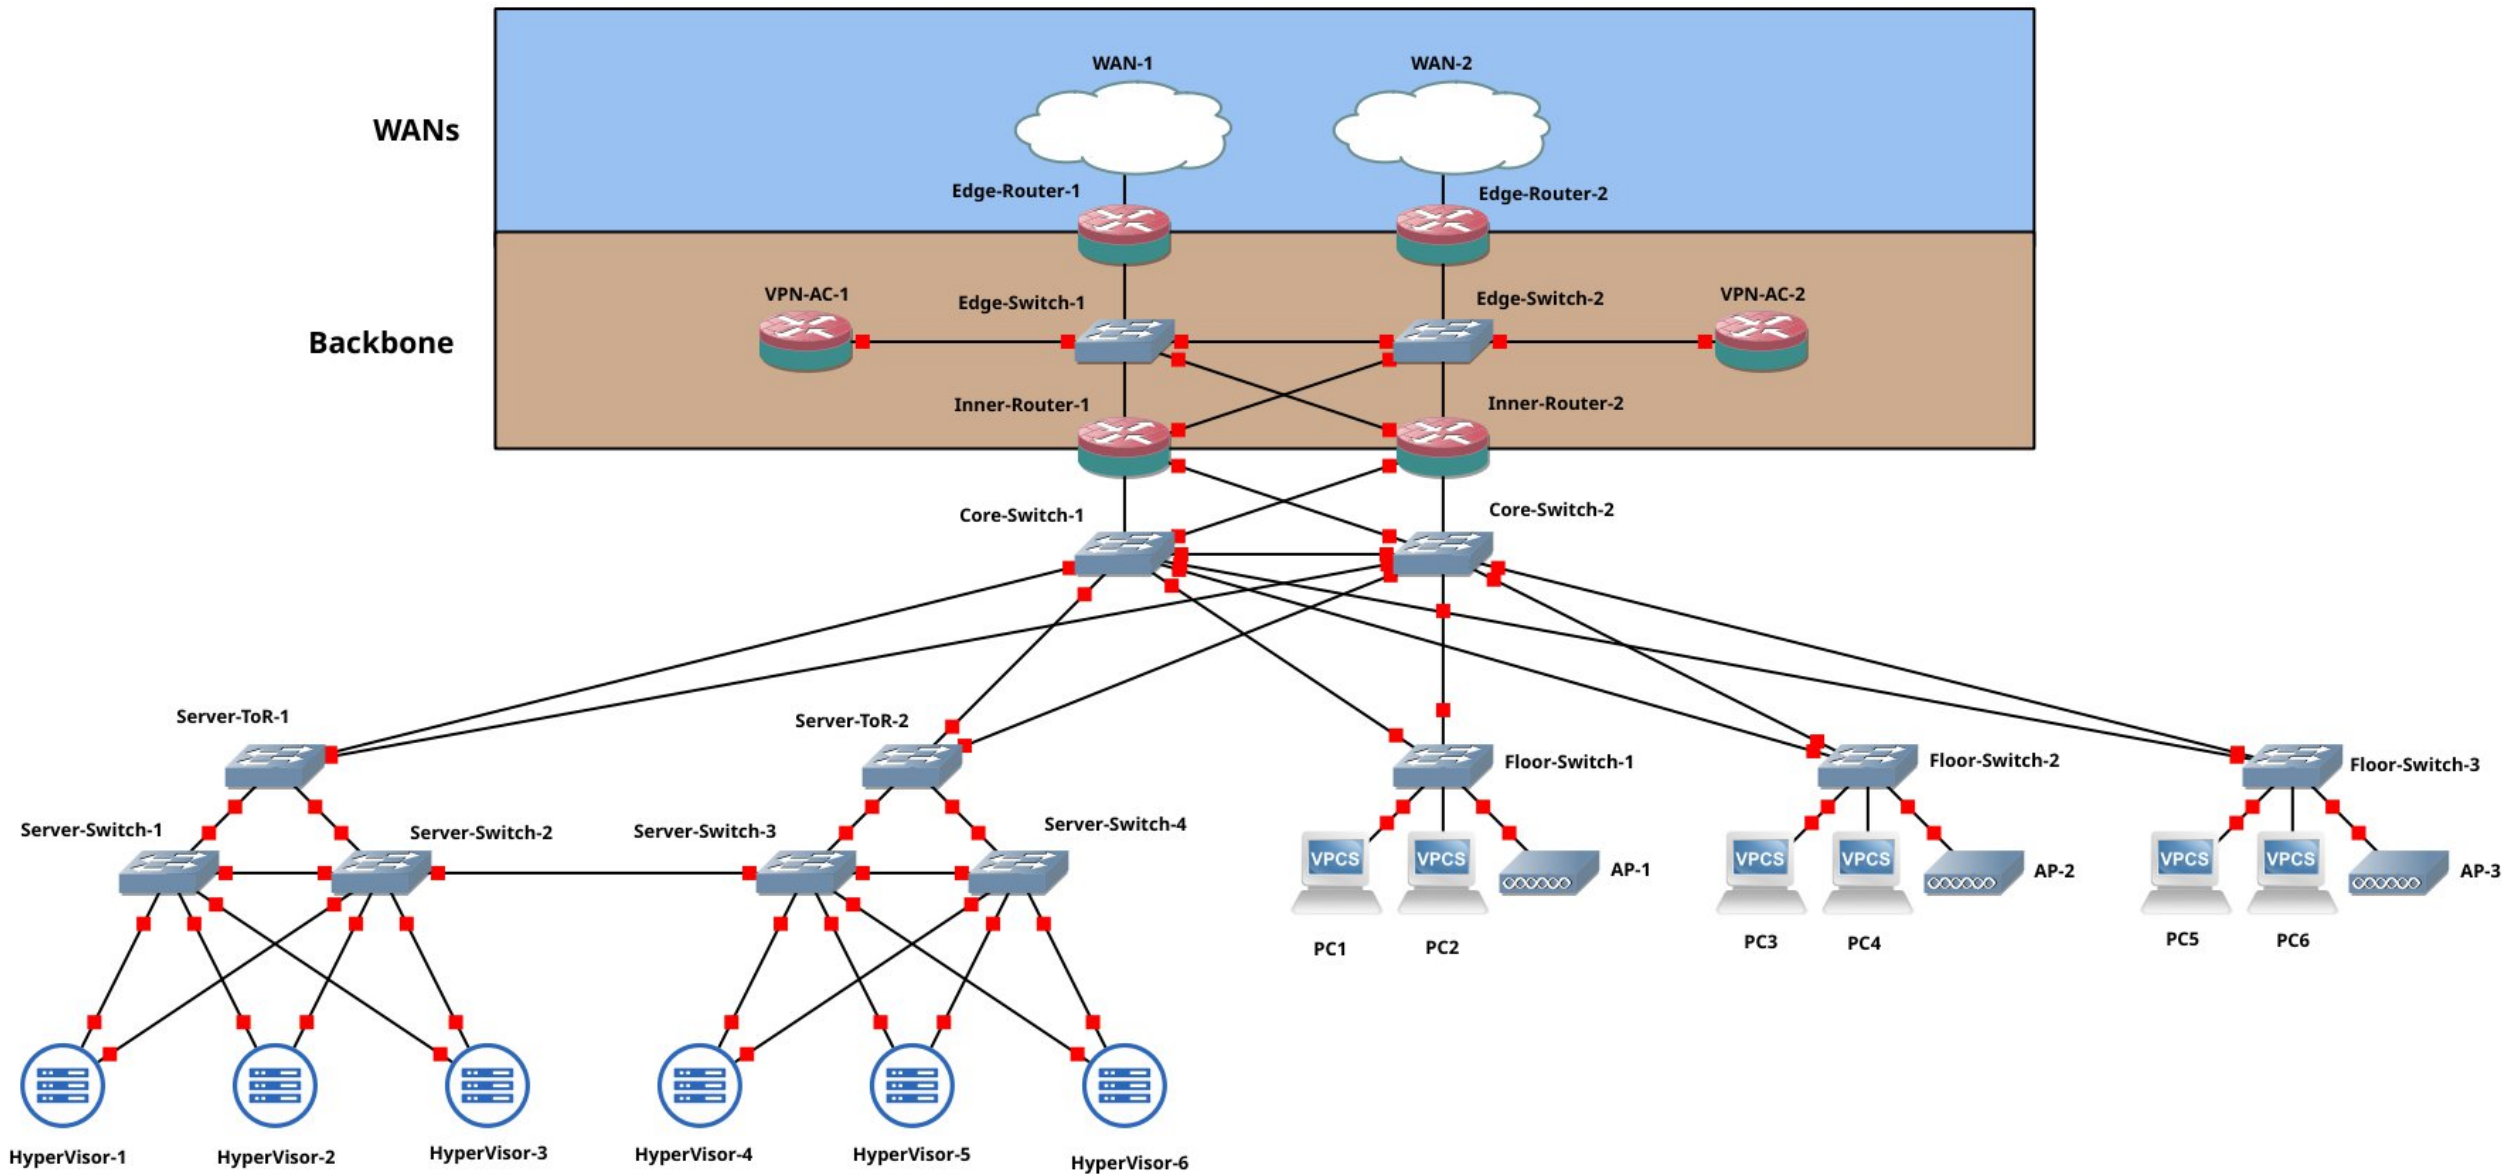
Loop-free network? Yes. All redundant paths un-utilized? Also yes.

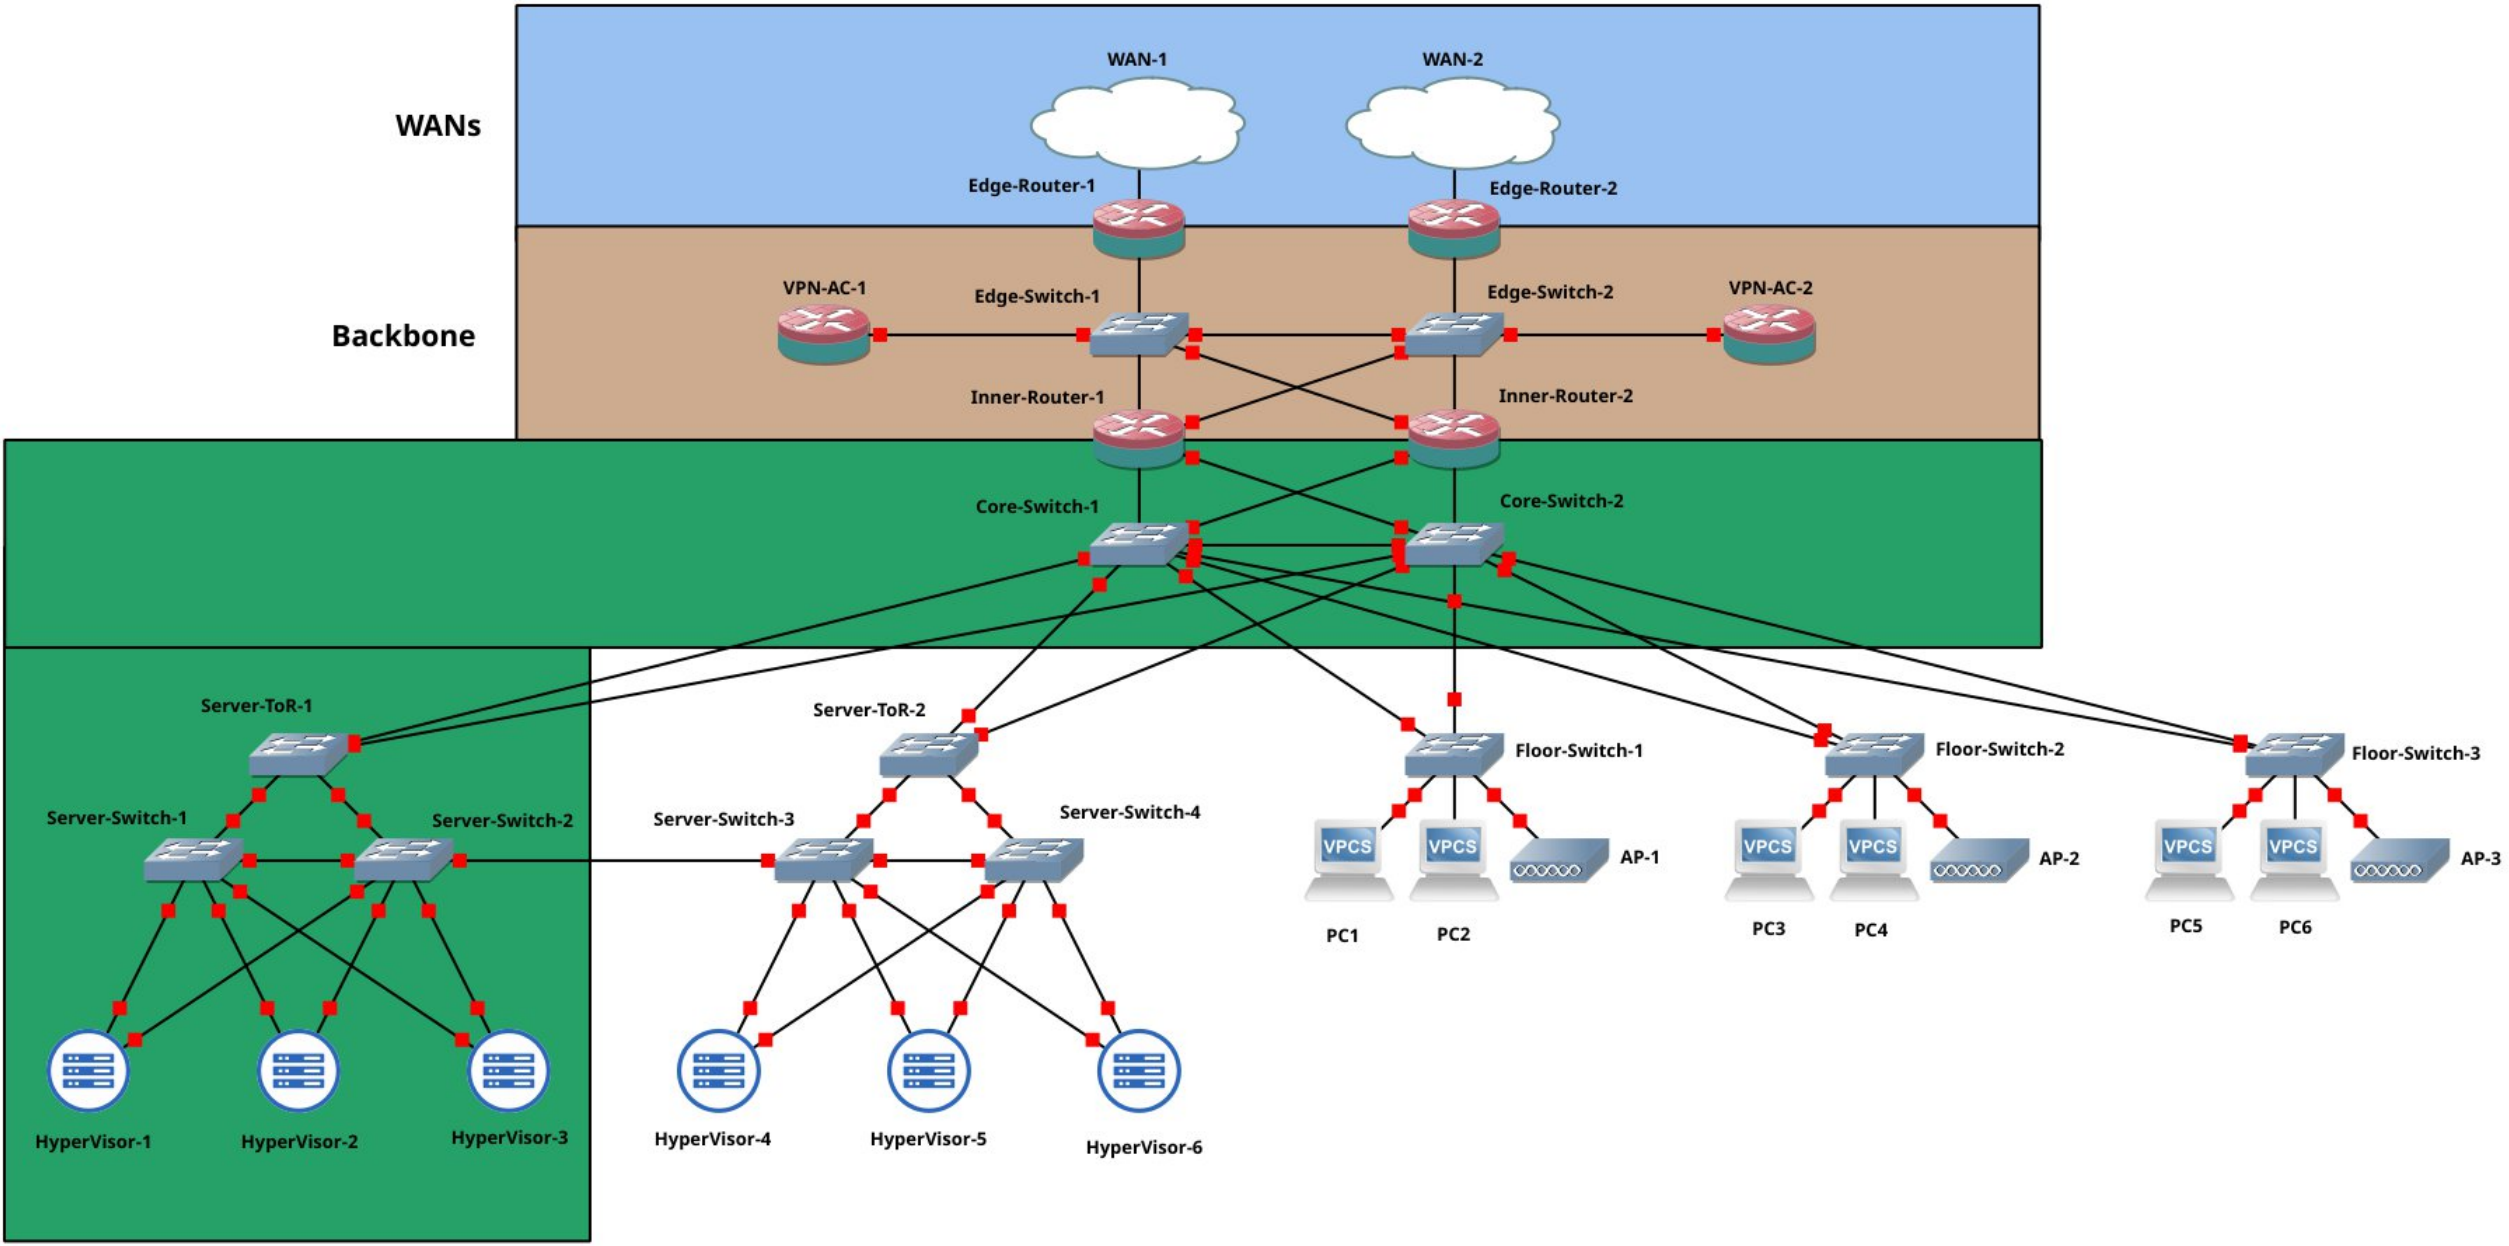
Let's see STP at work



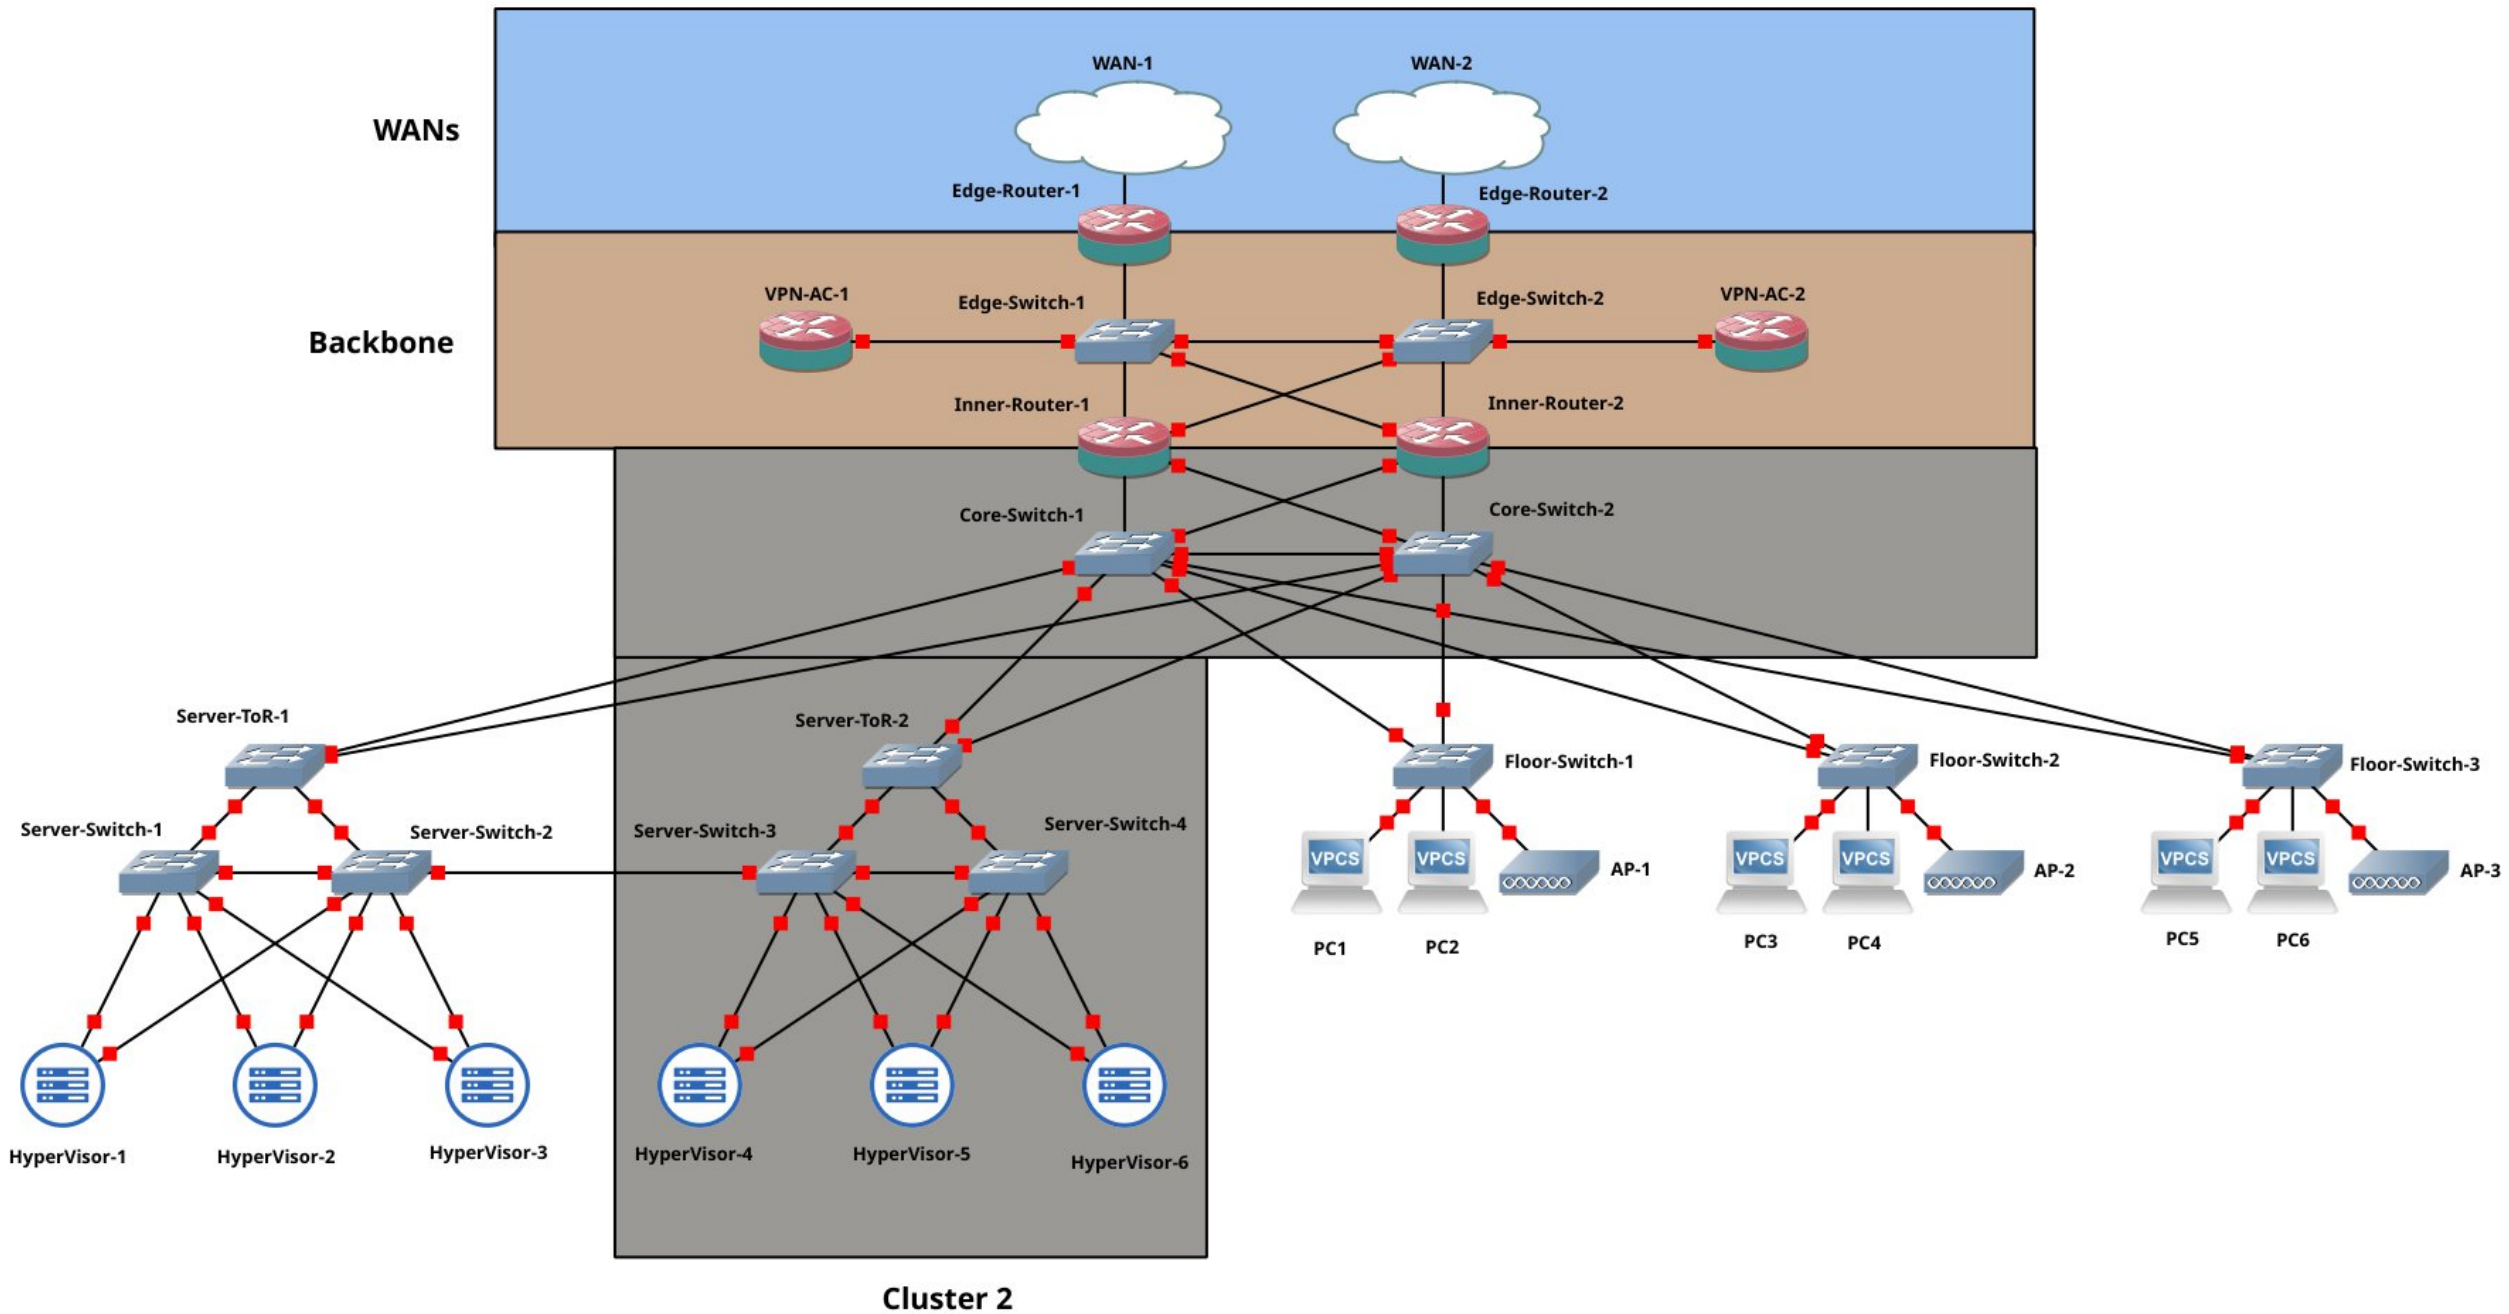
WANs

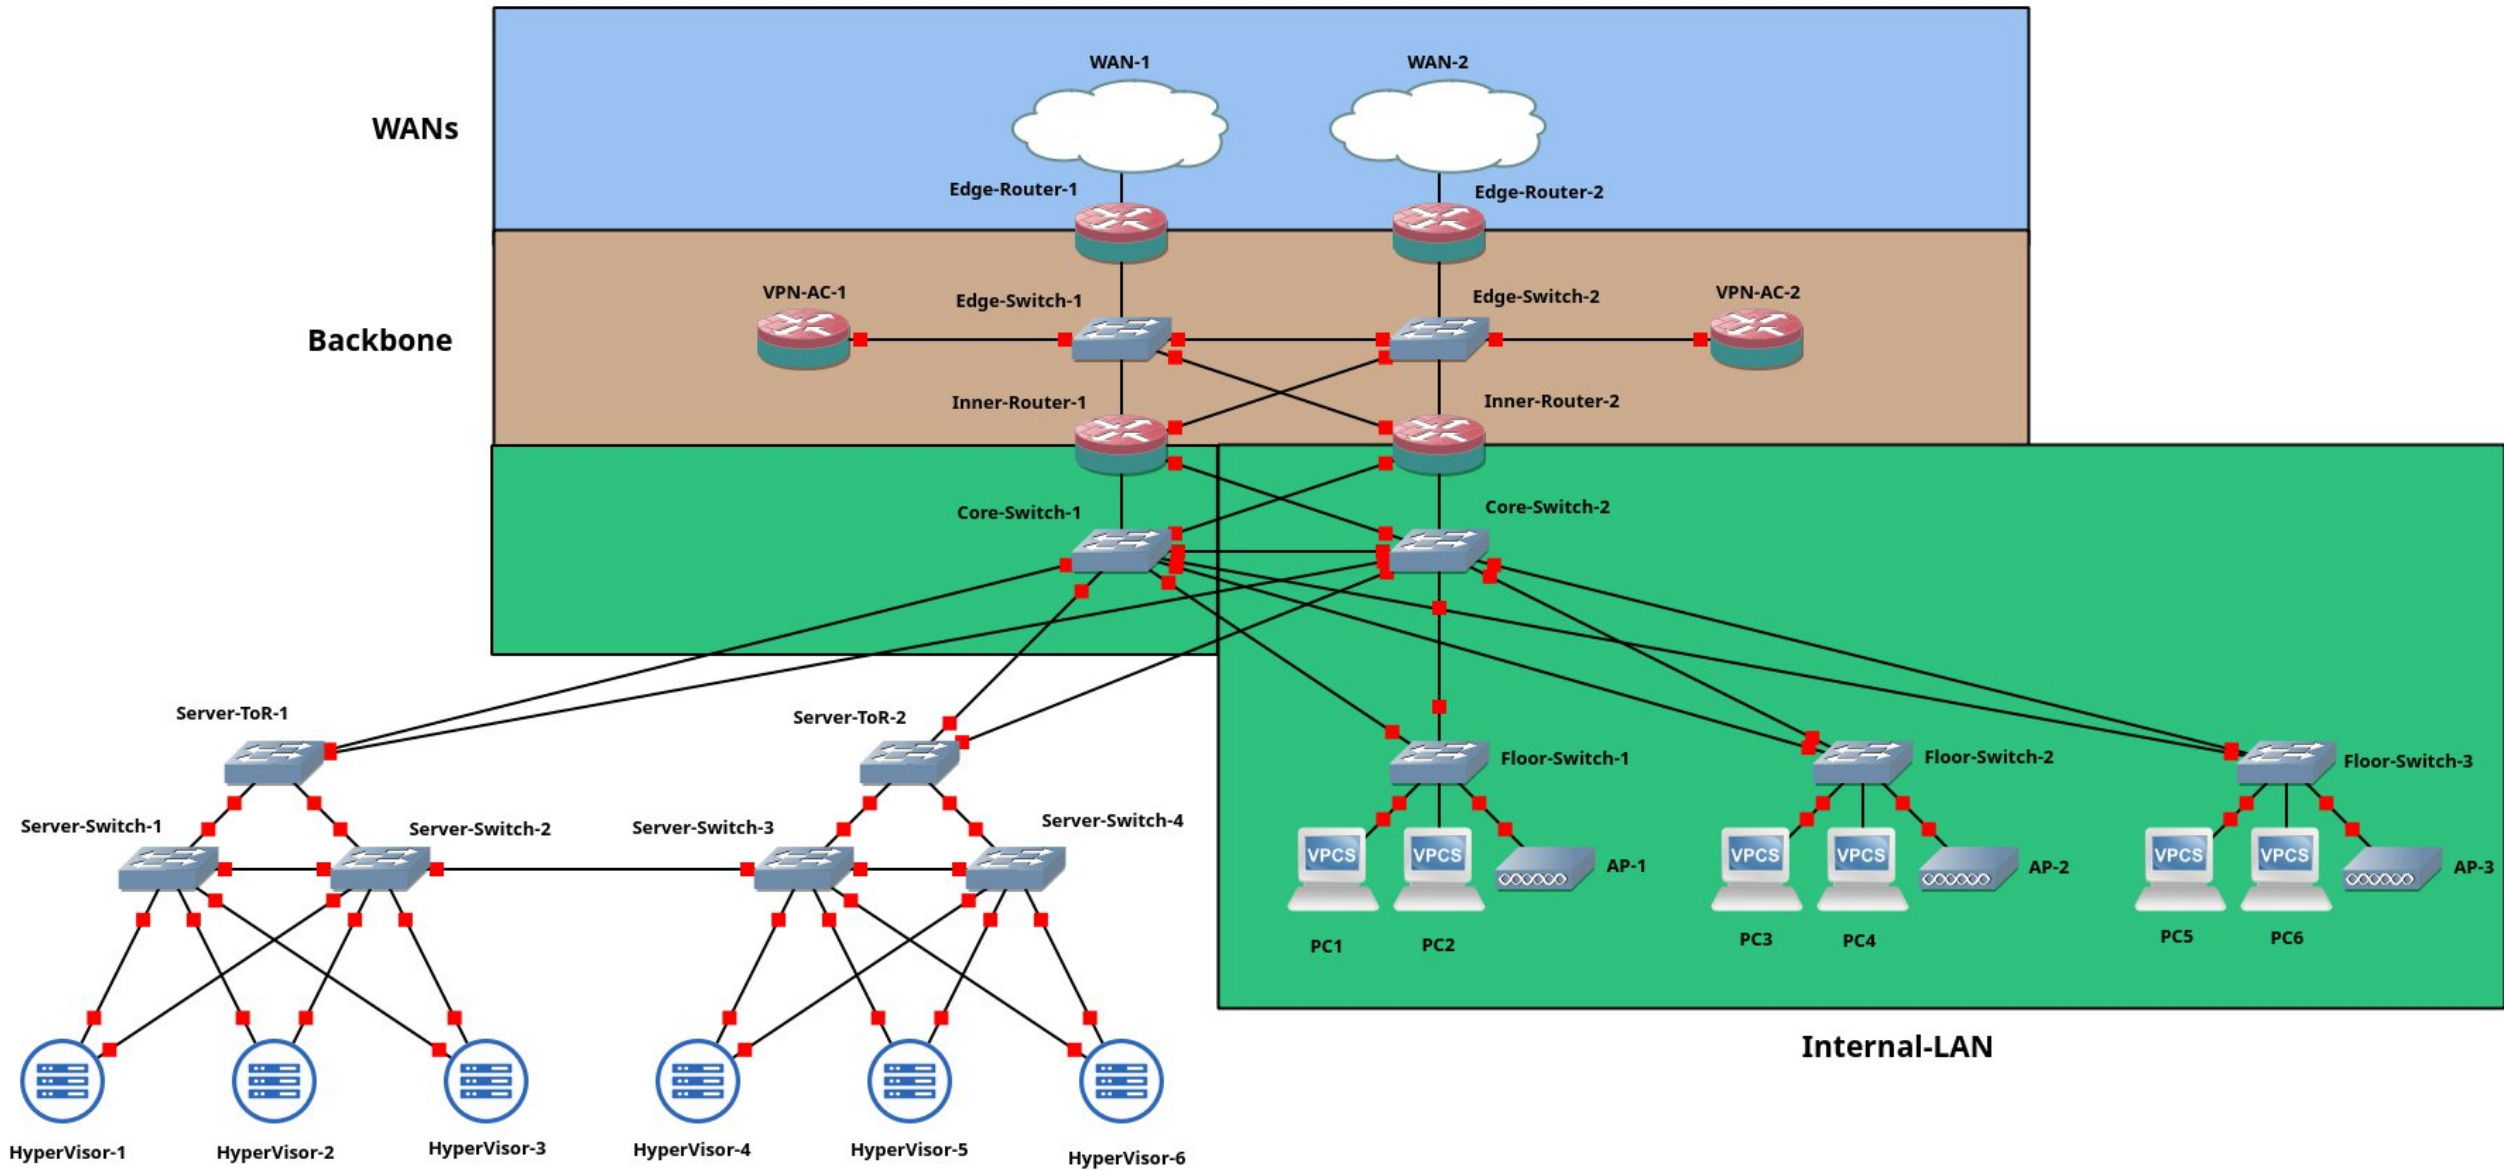


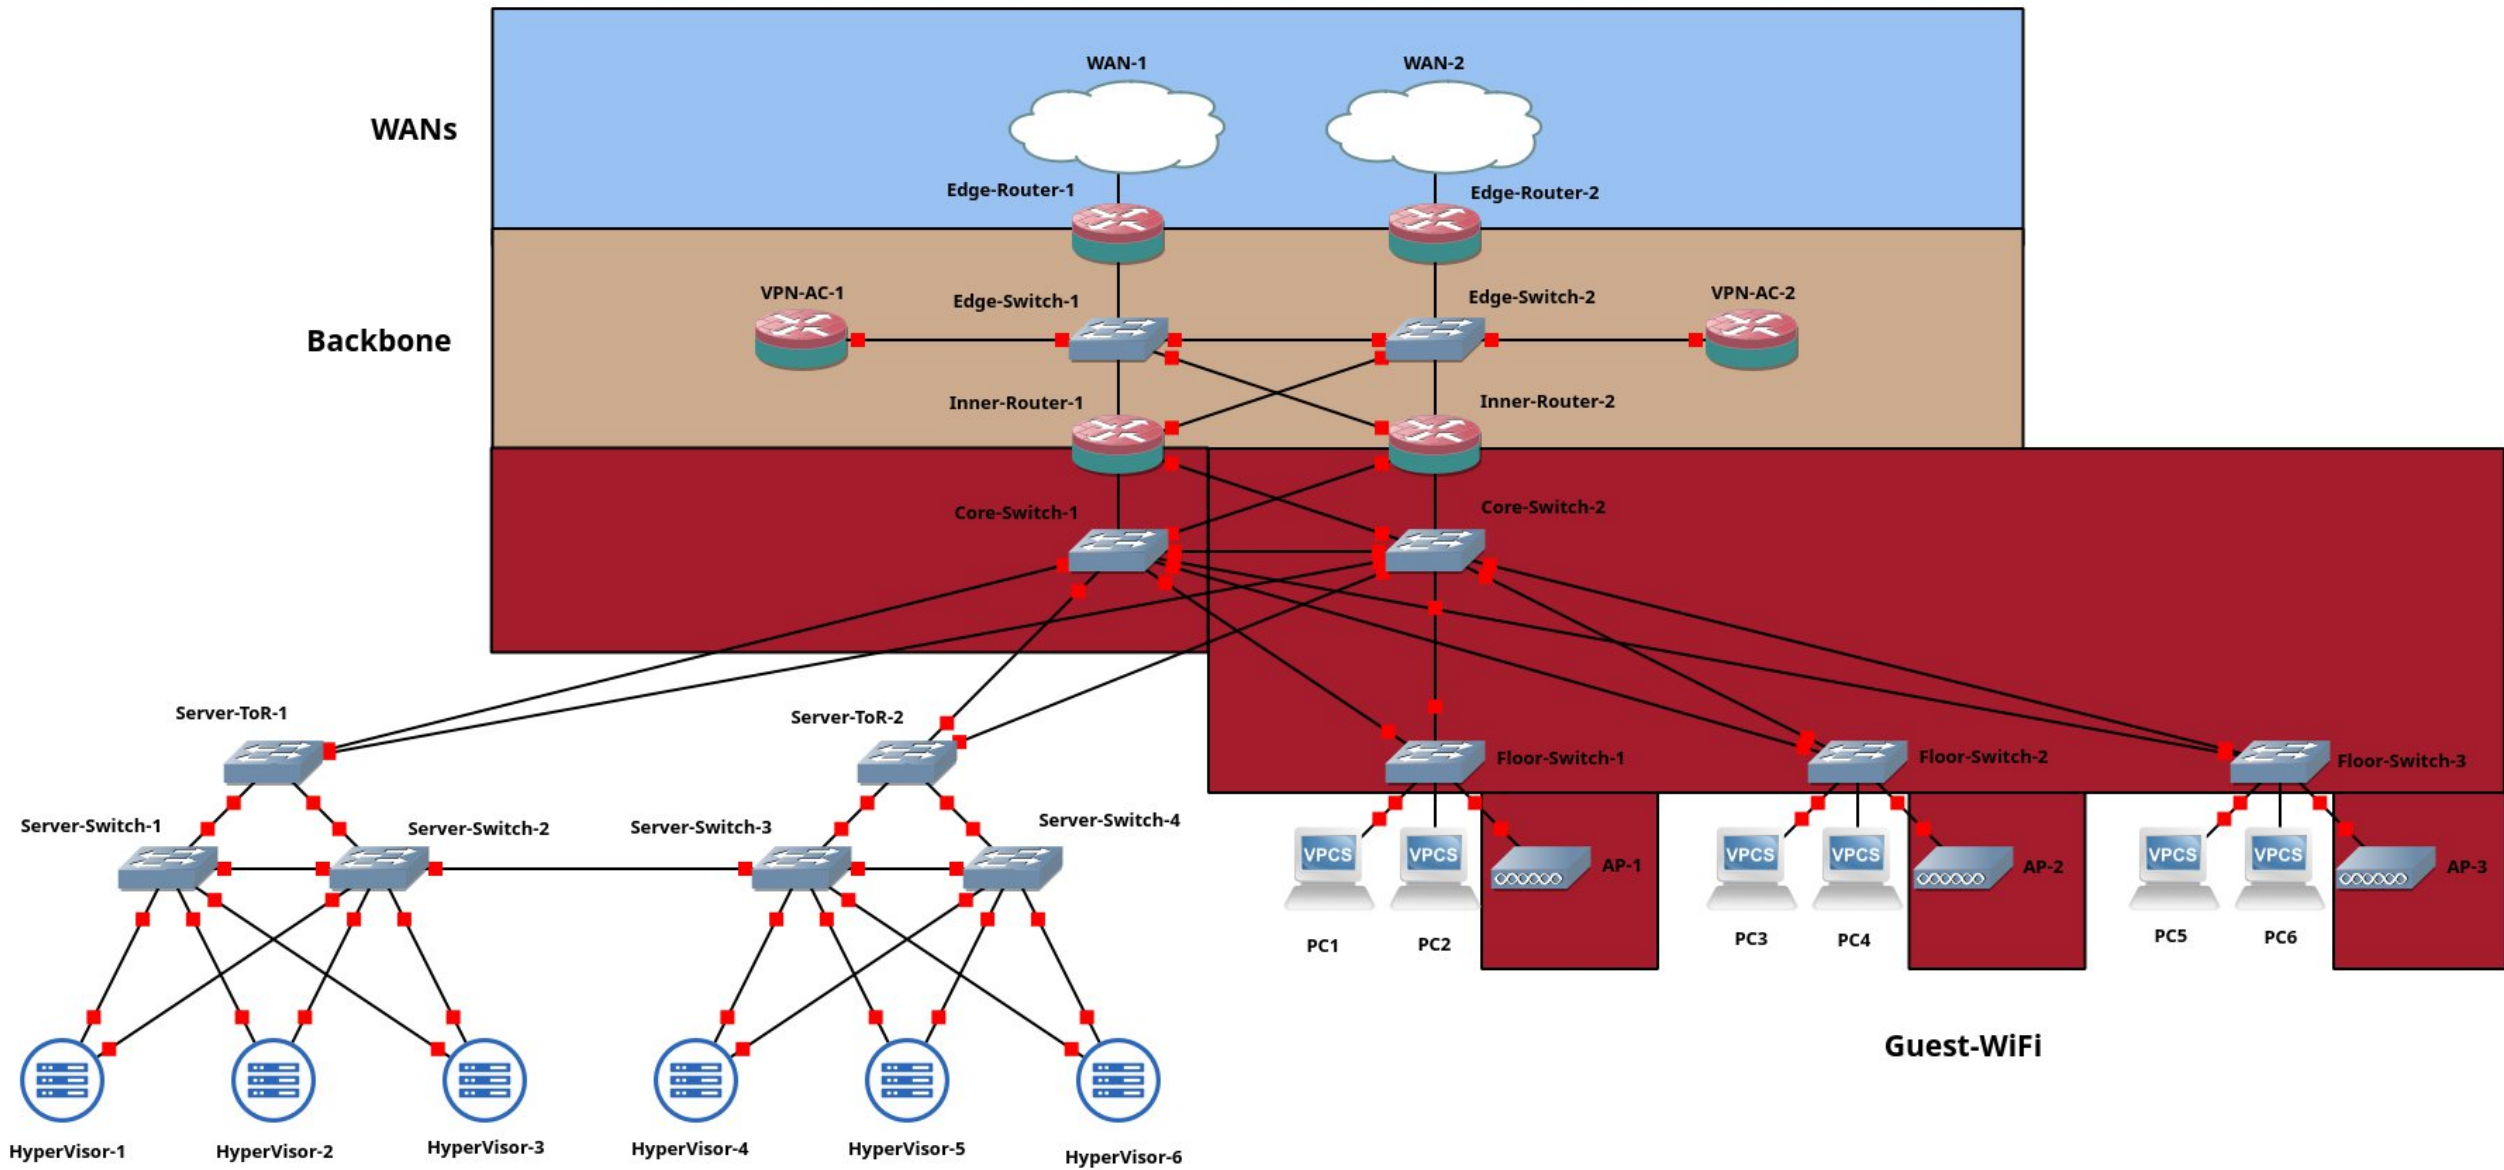


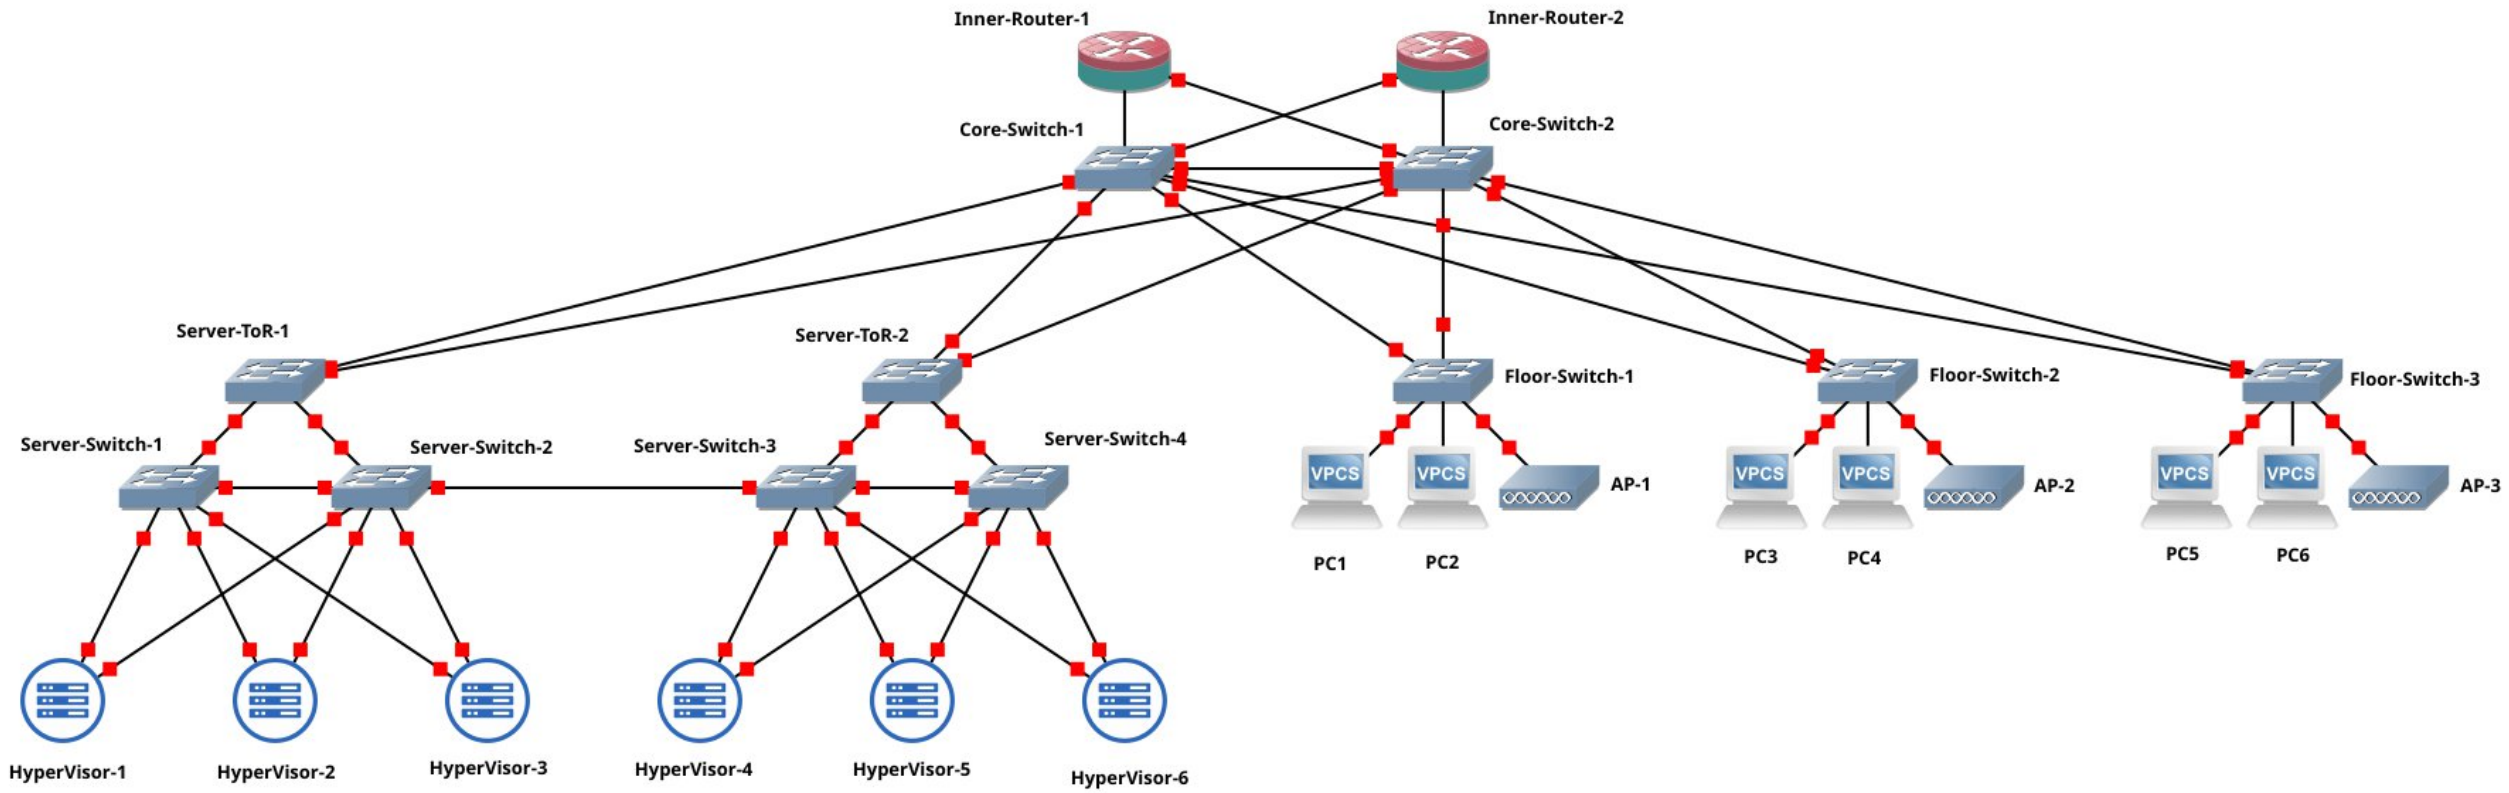


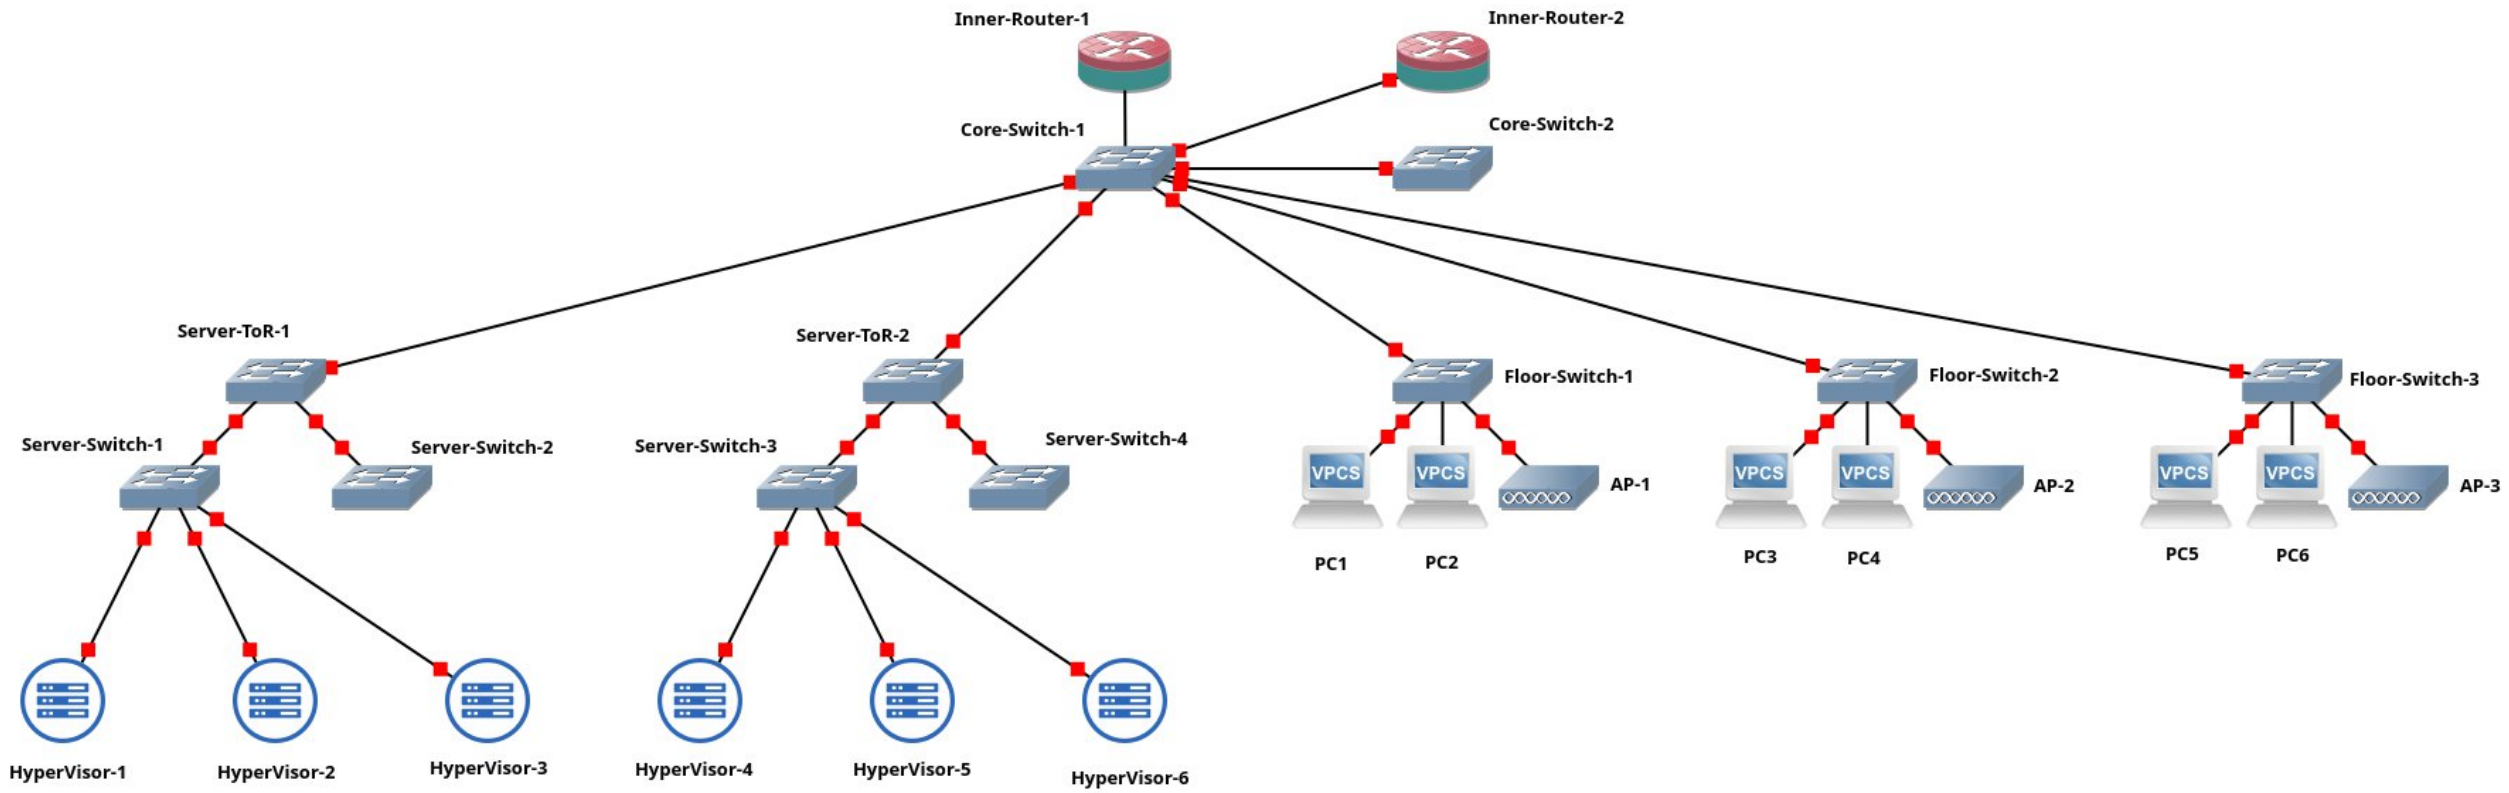
Cluster 1











Downsides of STP

Original STP was slow. One of the things that gave it a bad name.

It was also not ready for multi-gig links, so it struggles with proper path selection in modern networks.

Please DO NOT use STP in today's networks.

RSTP comes into play

RSTP improves STP by reducing convergence time (quite massively). This is done by introducing 2 new port states, and 4 new port roles.

In a converged RSTP network, port transition time is typically measured in milliseconds.

Edge ports (end-user devices) transition to forwarding immediately, allowing for no penalty to end users even when using RSTP in the network.

RSTP is great, please use it

RSTP is a great protocol, you should use it in every single of your networks.

HOWEVER, don't just turn it on, configure it properly! Most issues arise because administrators just turn on STP without configuring it. This leads to issues and then gives STP a bad reputation.

More on proper STP configuration later.

Downsides of RSTP

RSTP is fast to converge, doesn't penalize end-users.

RSTP can deal with fast modern links.

RSTP will however still turn off redundant links, leading to inefficient link / traffic utilization in the network.

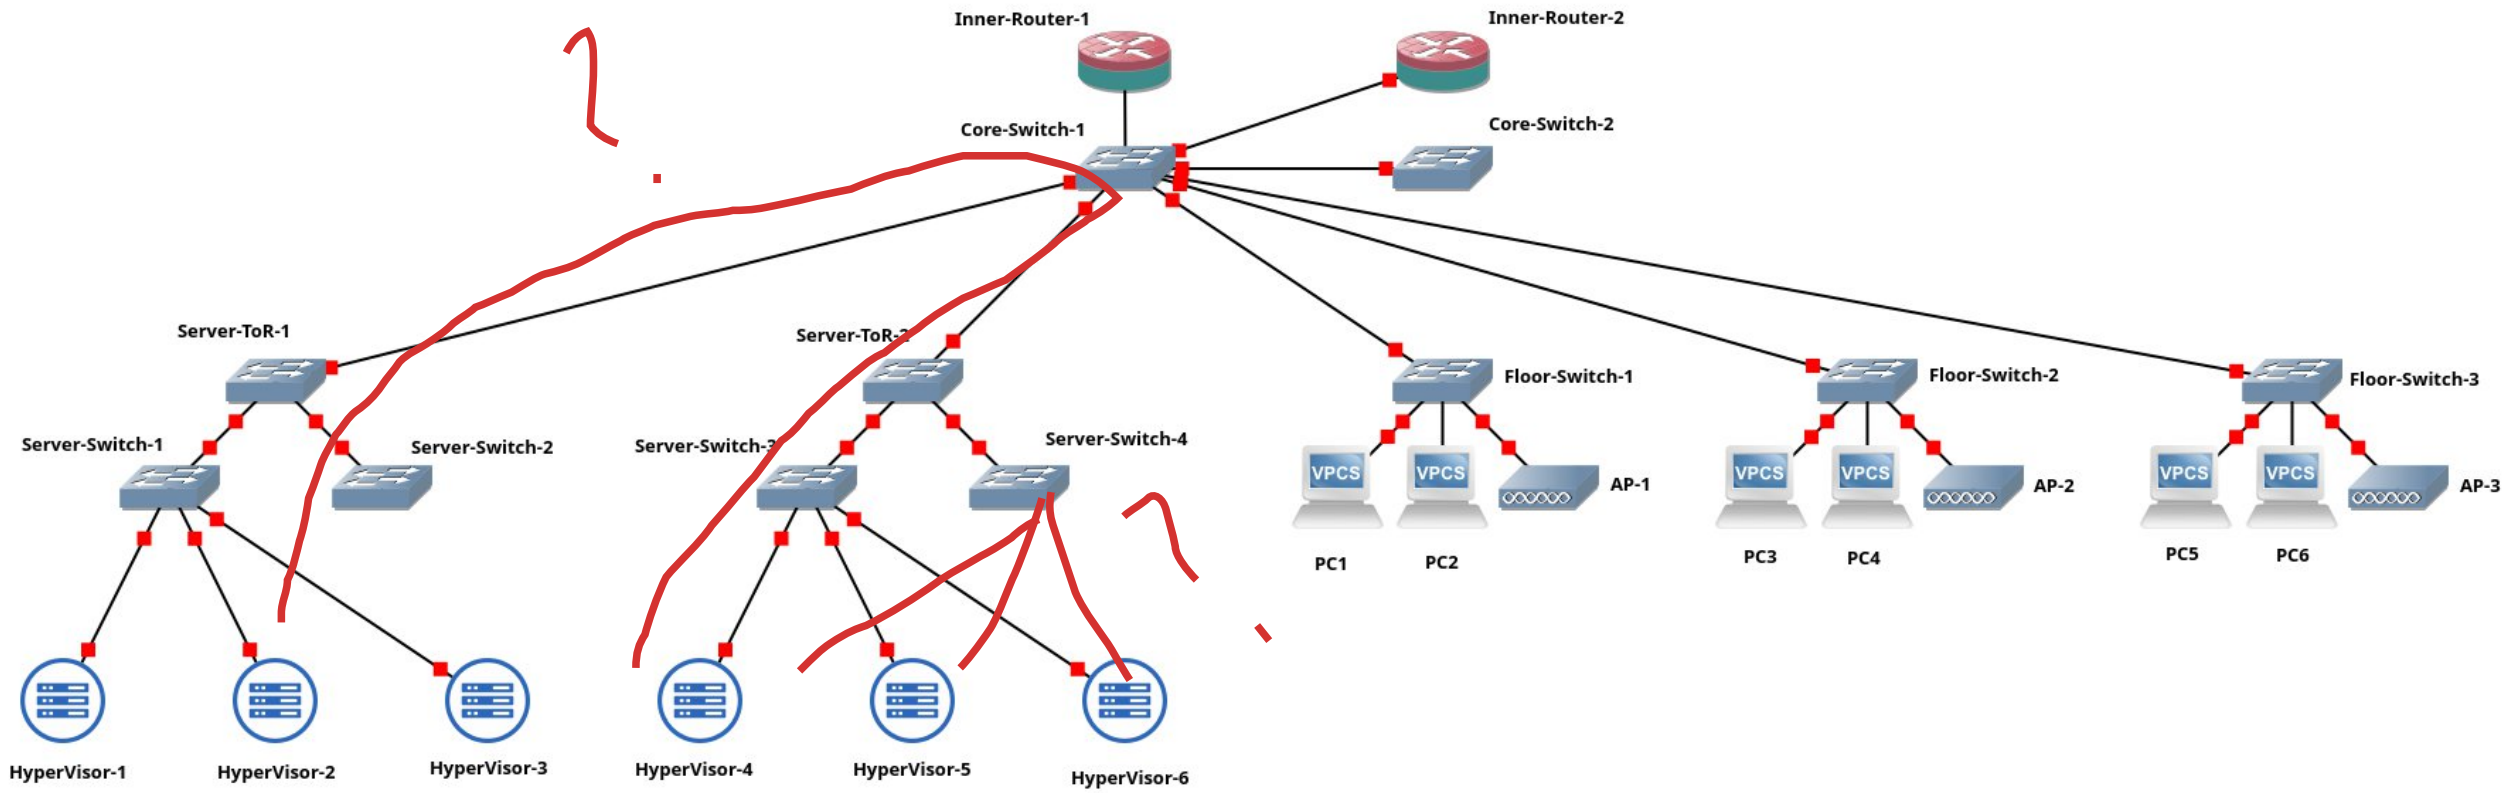
"I have all of these links, why should they be idle?"

Evolution to MSTP

MSTP was developed to address (R)STP limitations. It introduces the concepts of Regions and Instances.

MSTP can divide the network into multiple Regions, each with its own spanning tree

MSTP allows use of redundant paths within each region for improved link / traffic utilization.







Let's look at how regions work




Regions allow us to “carve out” a part of the network, and that part will have its own spanning tree

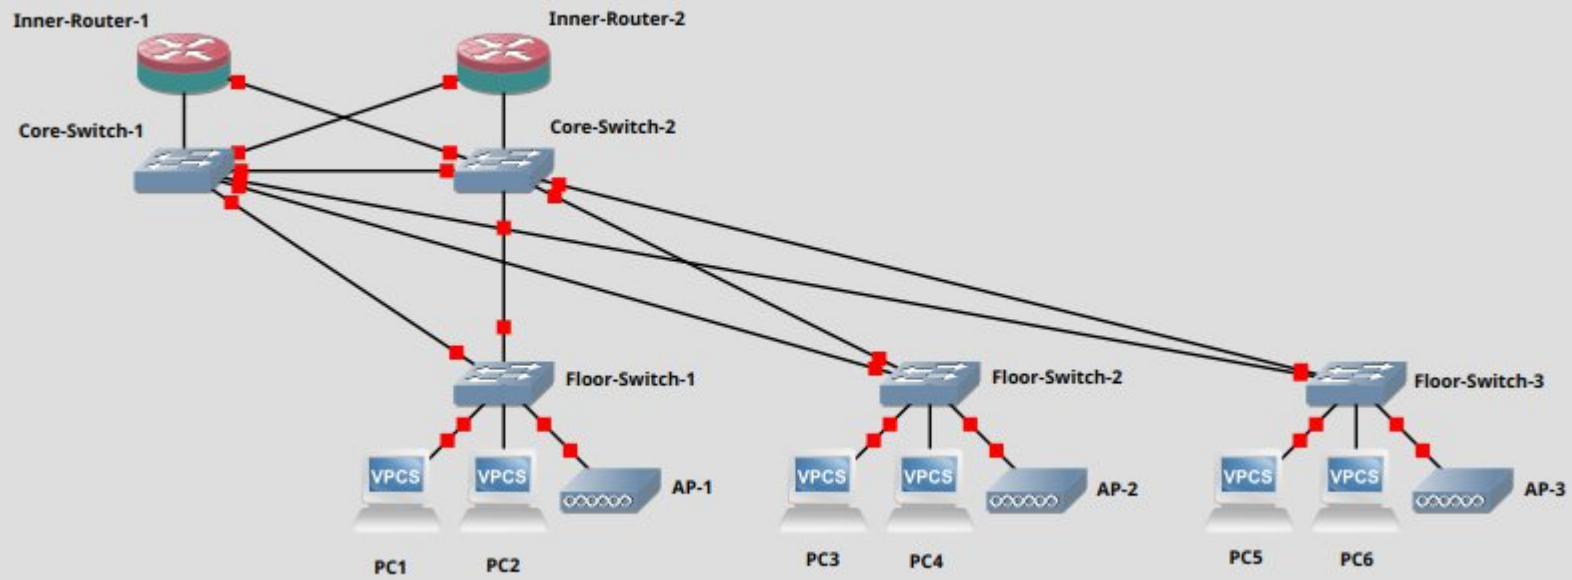
Each region has its own Root Bridge, and therefore its own internal topology

Let's carve out our big network into multiple regions







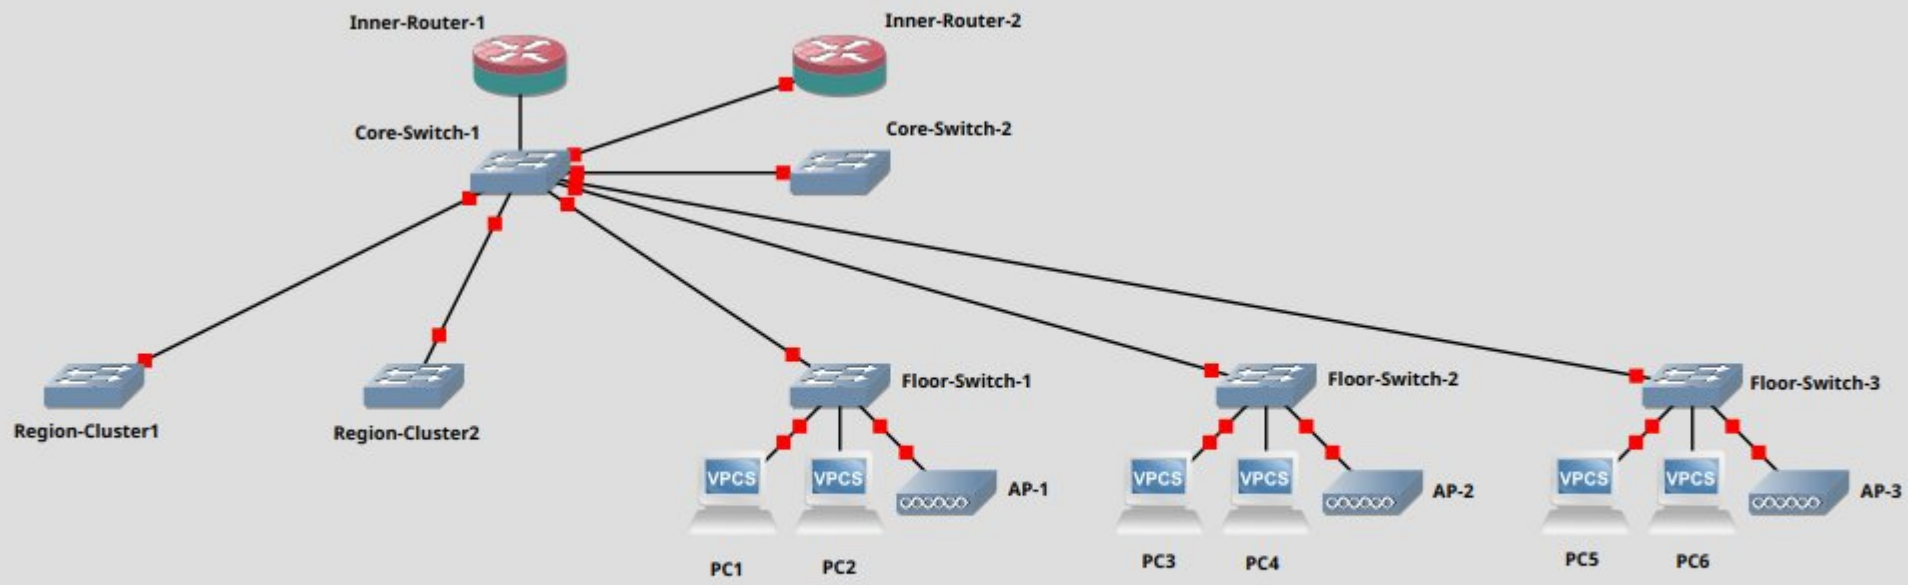


How do regions interact?

Regions appear to “edge” bridges (switches) in other regions just as a single bridge (switch)

Internal topology of each region is hidden from other regions

We can visualize it like this:

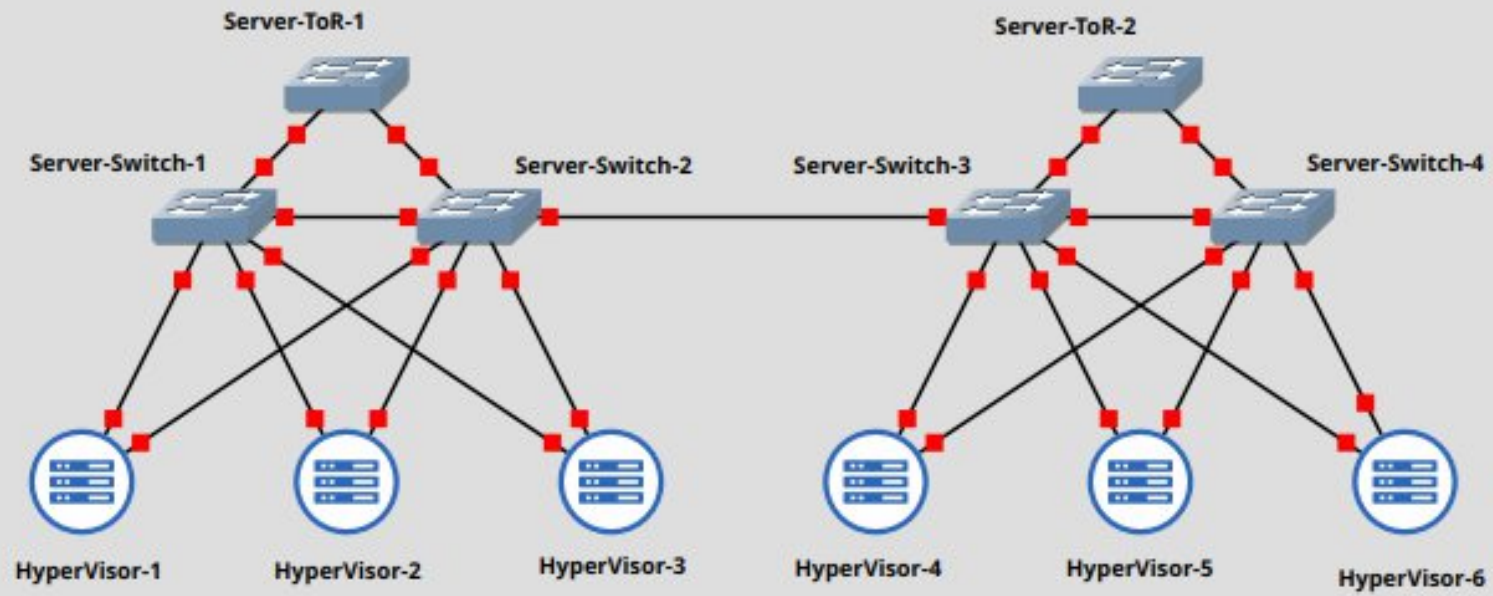


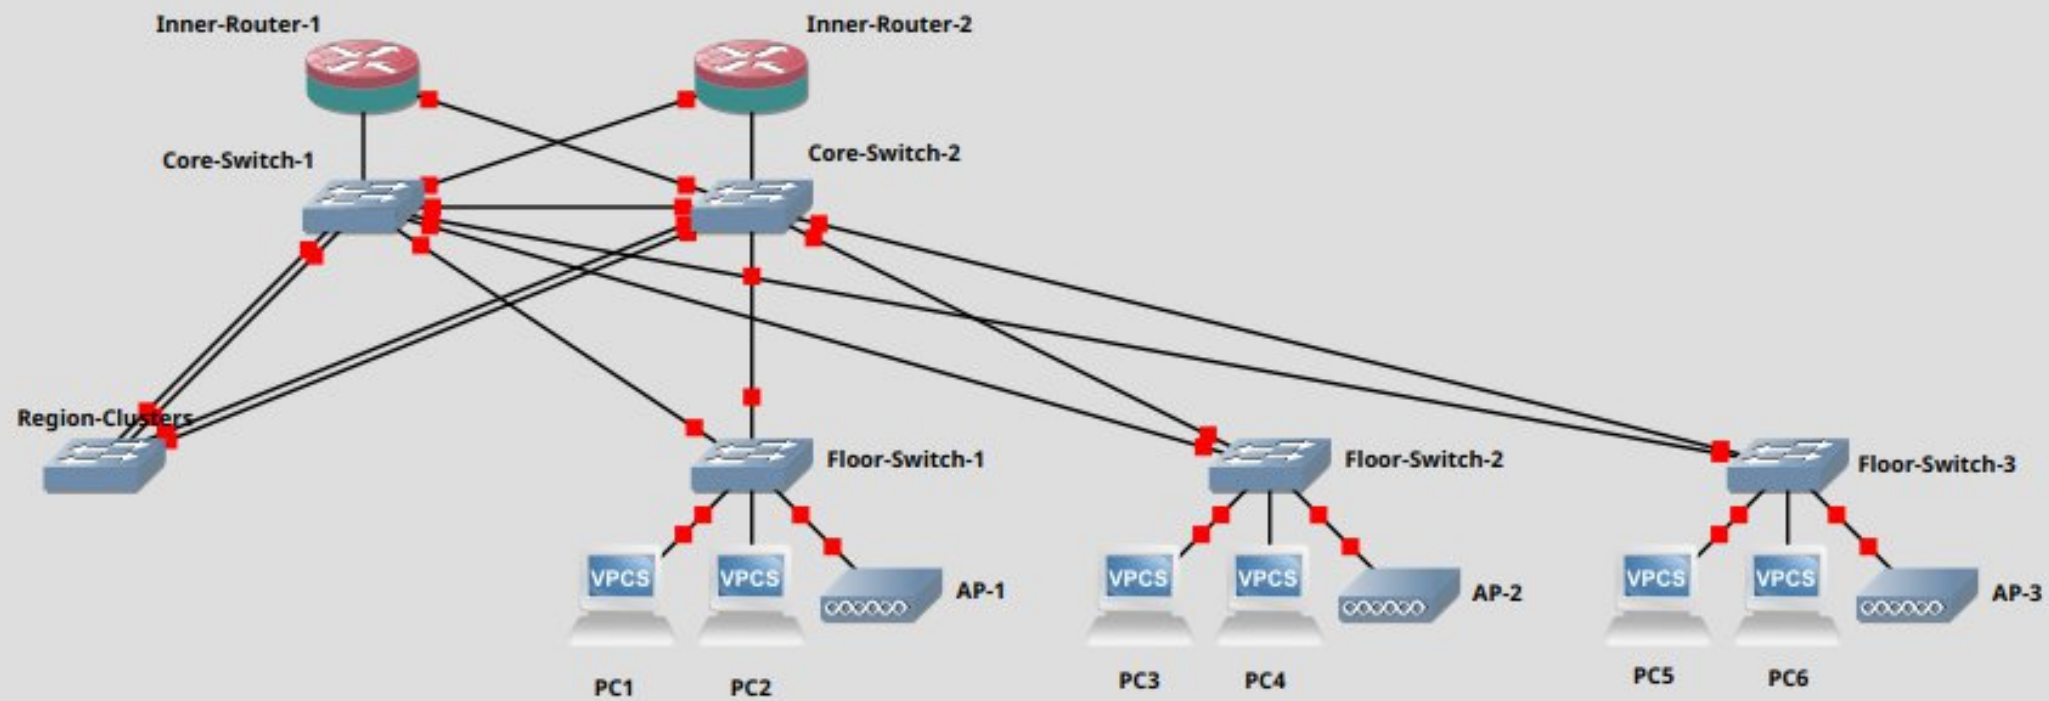
Not great, not terrible

This particular region topology doesn't solve our problem

Traffic between cluster regions still must transit the "main" region


Lets fix that:

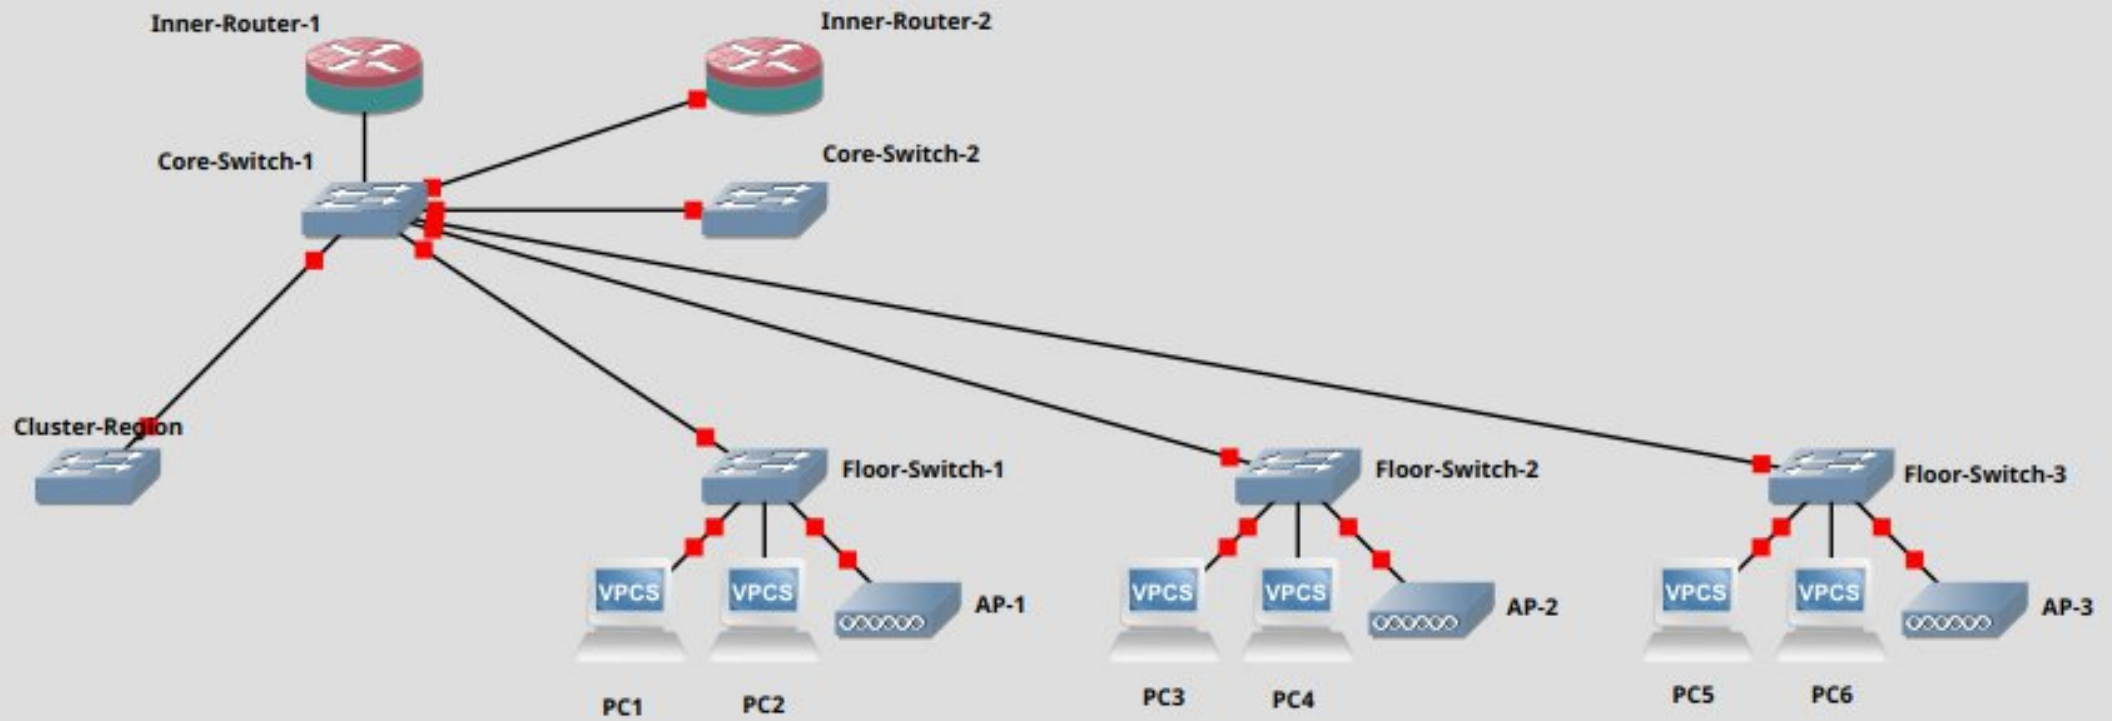




Topology after convergence

Let's see how the internal topology of our regions looks now





Great success!

Communication between servers now takes the shortest path (between racks), rather than going through the upstream network.

Regions allow you to design topologies with flexibility, and give you the power to enforce traffic routing according to your needs even when using STP.

Just one small issue

With traffic routing now solved, we still have a bunch of links that are not being utilized.

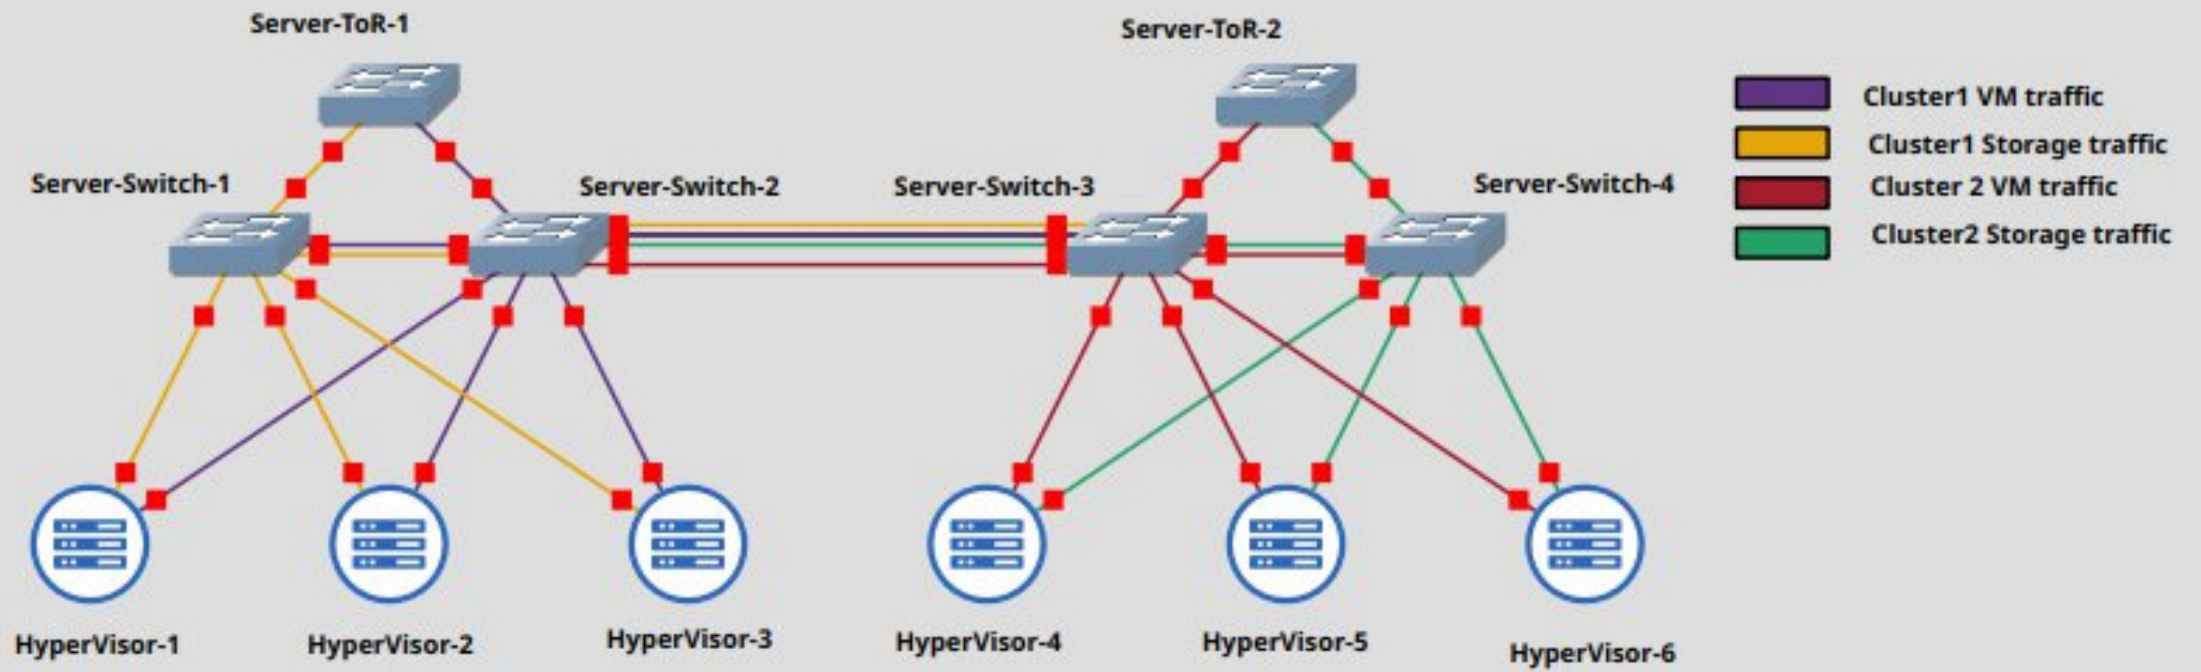
It would be nice if we could utilize our spare throughput.


MSTP Instances

Each region can encompass one or more VLANs, allowing for customized spanning tree topologies for each VLAN.

You can freely assign VLANs into Multiple Spanning Tree Instances (MSTIs).

This means when you have multiple VLANs, you can use your spare links by “assigning” different links to different MSTIs.





Acrony-whaaat?

Now that we understand advanced MSTP concepts, let's explain a few dreaded acronyms:

MSTI - Multiple Spanning Tree Instance (region, vlan-based instance, etc.).

IST - Internal Spanning Tree (each region maintains a single IST). A region that doesn't have a vlan-based MSTI only has the IST.

CST - Common Spanning Tree - Tree containing only regions and their interconnects. If an (R)STP bridge connects to the topology, it's actually served the CST. CST is just another MSTI.

How to configure MSTP

Bridge setup

New Interface

General STP VLAN Status Traffic

Protocol Mode: none STP RSTP MSTP

Priority: 8000 hex

Port Cost Mode: short long

Region Name:

Region Revision: 0

Max Message Age: 00:00:20

Forward Delay: 00:00:15

Transmit Hold Count: 6

Max Hops: 20

OK
Cancel
Apply
Disable
Comment
Copy
Remove
Torch
Reset Traffic Counters

enabled running slave passthrough

Set Protocol Mode

Set Priority (IST priority)

Set Path Cost mode

Set Region Name & Revision

MSTP needs VLAN Filtering

Make sure to set it up properly

Path Cost Mode

Data rate	Long	Short
10 Mbps	2,000,000	100
100 Mbps	200,000	19
1 Gbps	20,000	4
10 Gbps	2,000	2
25 Gbps	800	1
40 Gbps	500	1
50 Gbps	400	1
100 Gbps	200	1

Path Cost Mode selects what values are used for link costs

“Short” is the original STP path costs

User Long, make sure all devices in topology use it

Port setup

New Bridge Port

General STP VLAN Status

Priority: 80 hex

Path Cost: ▼

Internal Path Cost: ▼

Edge: auto ▼

Point To Point: auto ▼

Auto Isolate

Restricted Role

Restricted TCN

BPDU Guard

OK

Cancel

Apply

Disable

Comment

Copy

Remove

enabled inactive Hw. Offload

Path cost can be overridden here. Useful for LAGs (bonds)

Set security settings properly:

- Edge
- Isolation
- Role / TCN restrictions
- BPDU guard

Proper STP security

For ports to other STP bridges:

- Edge: no
- Auto-isolate: yes
- Restricted role: no
- Restricted TCN: no
- BPDU guard: no

For end-device ports:

- Edge: yes
- Auto-isolate: no
- Restricted role: yes
- Restricted TCN: yes
- BPDU guard: yes

MSTI setup

New Bridge MSTI

General Status

Bridge: br1

Identifier: 1

Priority: 8000 hex

VLAN Mapping: 1-4094

OK

Cancel

Apply

Disable

Comment

Copy

Remove

enabled

Select bridge, set MSTI identifier

Set root priority within this MSTI

Set which VLANs are a part of this MSTI

MST Overrides setup


The screenshot shows a window titled "New Bridge MST Override" with two tabs: "General" and "Status". The "General" tab is selected and contains the following fields:

- Interface:** A dropdown menu with "ether1" selected.
- Identifier:** A text input field containing the number "1".
- Priority:** A text input field containing "80", with a "hex" label to its right.
- Internal Path Cost:** A dropdown menu.

On the right side of the dialog, there is a vertical stack of buttons: "OK", "Cancel", "Apply", "Disable", "Comment", "Copy", and "Remove". At the bottom left of the dialog, the text "enabled" is displayed.

Normally, ports inherit all their properties set in “/interface bridge port” in all MSTIs they are present

You can override some of port properties for a particular MSTI here



Bridge configuration

```
/interface bridge
```

```
add name=bridge1 frame-types=admit-only-vlan-tagged priority=0x6000 protocol-mode=mstp region-name=cluster-region region-revision=1 vlan-filtering=yes
```

```
/interface bridge port
```

```
add interface=ether1 bridge=bridge1 auto-isolate=yes edge=no frame-types=admit-only-vlan-tagged
```

```
add interface=ether2 bridge=bridge1 auto-isolate=yes edge=no frame-types=admit-only-vlan-tagged
```

```
/interface bridge msti
```

```
add bridge=bridge1 identifier=1 priority=0x6000 vlan-mapping=201
```

```
add bridge=bridge1 identifier=2 priority=0x8000 vlan-mapping=202
```

Tips and tricks

- Make sure your whole network is using the same path cost model. Especially in multi-vendor networks.
- Set path cost properly on LAGs / bonds, on tunnels, etc.
- Monitor who is root-bridge in topology, have warning when it's not a bridge you expect.
- Validate if config works as expected (beware of bugs – not MikroTik specific, even “big boy” vendors have tons of STP bugs).

That's it, thank you!

Q&A session