CSPE

High Availability Infrastructures for SMBs

Presented by
Alain Casault, Eng. MTCNA, MTCRE, MTCWE
2015 Canadian MUM
Montréal
October 19th 2015

About the trainer

- Electrical Engineering
- Over 20 years of experience with networking and server OS's
- Has experience with many types of clients
- Author of Mikrotik's current **MTCNA** course material

- CSPE (*Centre de Services Professionnels en Éducation*)
  - Educational services
  - My part: MikroTik and Telecommunications training
    - Introduction to the TCP/IP Protocol
    - Introduction to MikroTik Routers

Objectives

We will learn about:

- High availability
- How to configure MikroTik equipment to build robust and flexible networks

Presentation overview

- Introduction
  - Concepts of “HA”
  - Causes of downtime
  - Topology: The anatomy of redundancy
- The technologies involved in “HA”
- Demo

Presentation complexity

Advanced

Intermediate

Introduction

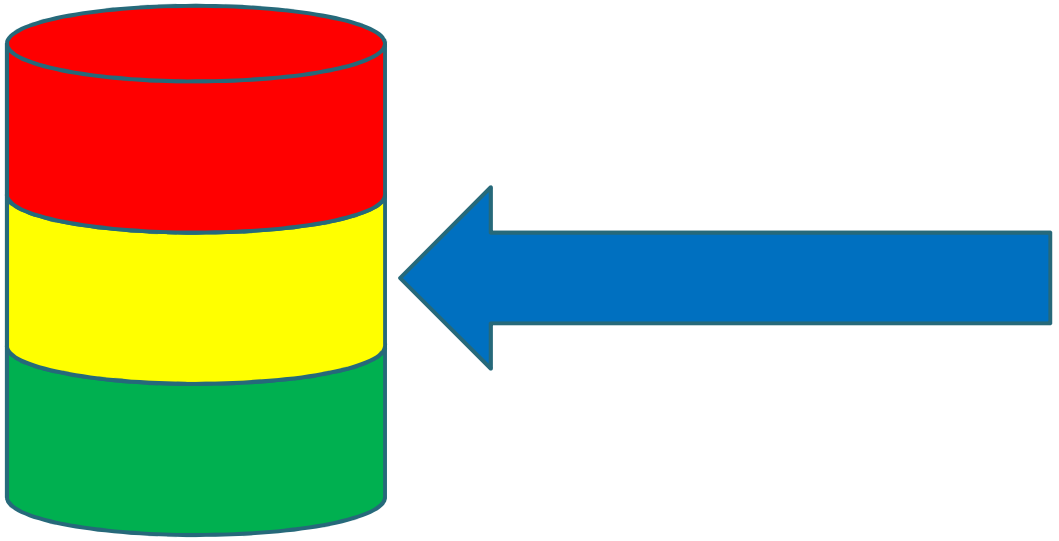

CSPE
Centre de services professionnels en éducation

INTRODUCTION

Concepts

Wikipedia defines "**High Availability**" as

"*… a characteristic of a system, which describes the duration (length of time) for which the system is operational.*"

Concepts

I define "**High Availability**" as

"*Up time*"

"*I-keep-my-job time*"

"*Happy-boss time*"

Concepts

- "**High Availability**" is referred to as a percentage score
- We often hear "nines" of availability

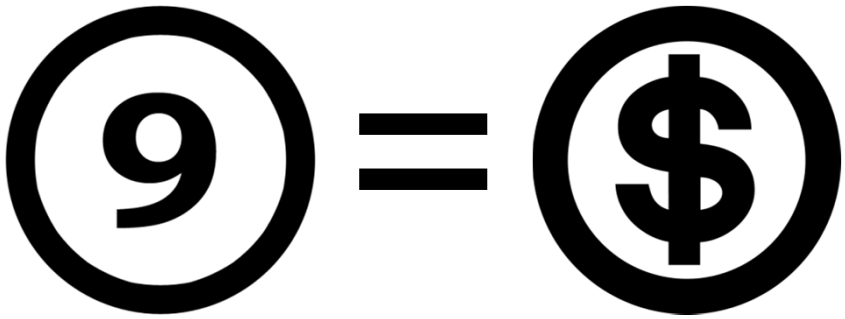

Concepts

- What do the “nines” mean?

Concepts

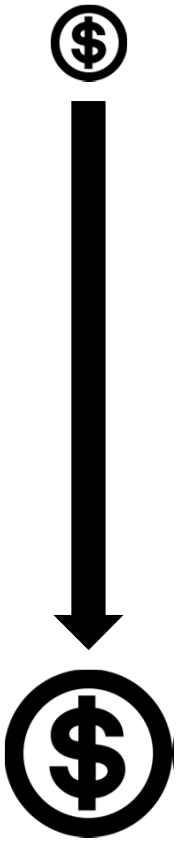

Availability %	Nines	Downtime / year	Downtime / month	Downtime / week
90%	1	36.5 days	72 hours	16.8 hours
95%		18.25 days	36 hours	8.4 hours
97%		10.96 days	21.6 hours	5.04 hours
98%		7.30 days	14.4 hours	3.36 hours
99%	2	3.65 days	7.20 hours	1.68 hours
99.5%		1.83 days	3.60 hours	50.4 Min
99.8%		17.52 hours	86.23 Min	20.16 Min
99.9%	3	8.76 hours	43.8 Min	10.1 Min
99.95%		4.38 hours	21.56 Min	5.04 Min
99.99%	4	52.56 Min	4.38 Min	1.01 Min
99.995%		26.28 Min	2.16 Min	30.24 Sec
99.999%	5	5.26 Min	25.9 Sec	6.05 Sec
99.9999%	6	31.5 Sec	2.59 Sec	604.8 mSec
99.99999%	7	3.15 Sec	262.97 mSec	60.48 mSec
99.999999%	8	315.569 mSec	26.297 mSec	6.048 mSec
99.9999999%	9	31.5569 mSec	2.6297 mSec	0.6048 mSec

Concepts

- For a more realistic “up time” score:
  - Negotiate scheduled downtime
    - Not counted in the “High Availability” score

Causes of downtime

- Some causes are:
  - Hardware failure
  - Network failure
  - **Human error**
  - System overload
  - Electrical supply

Topology

- What would a great “High Availability” scenario look like?

Topology

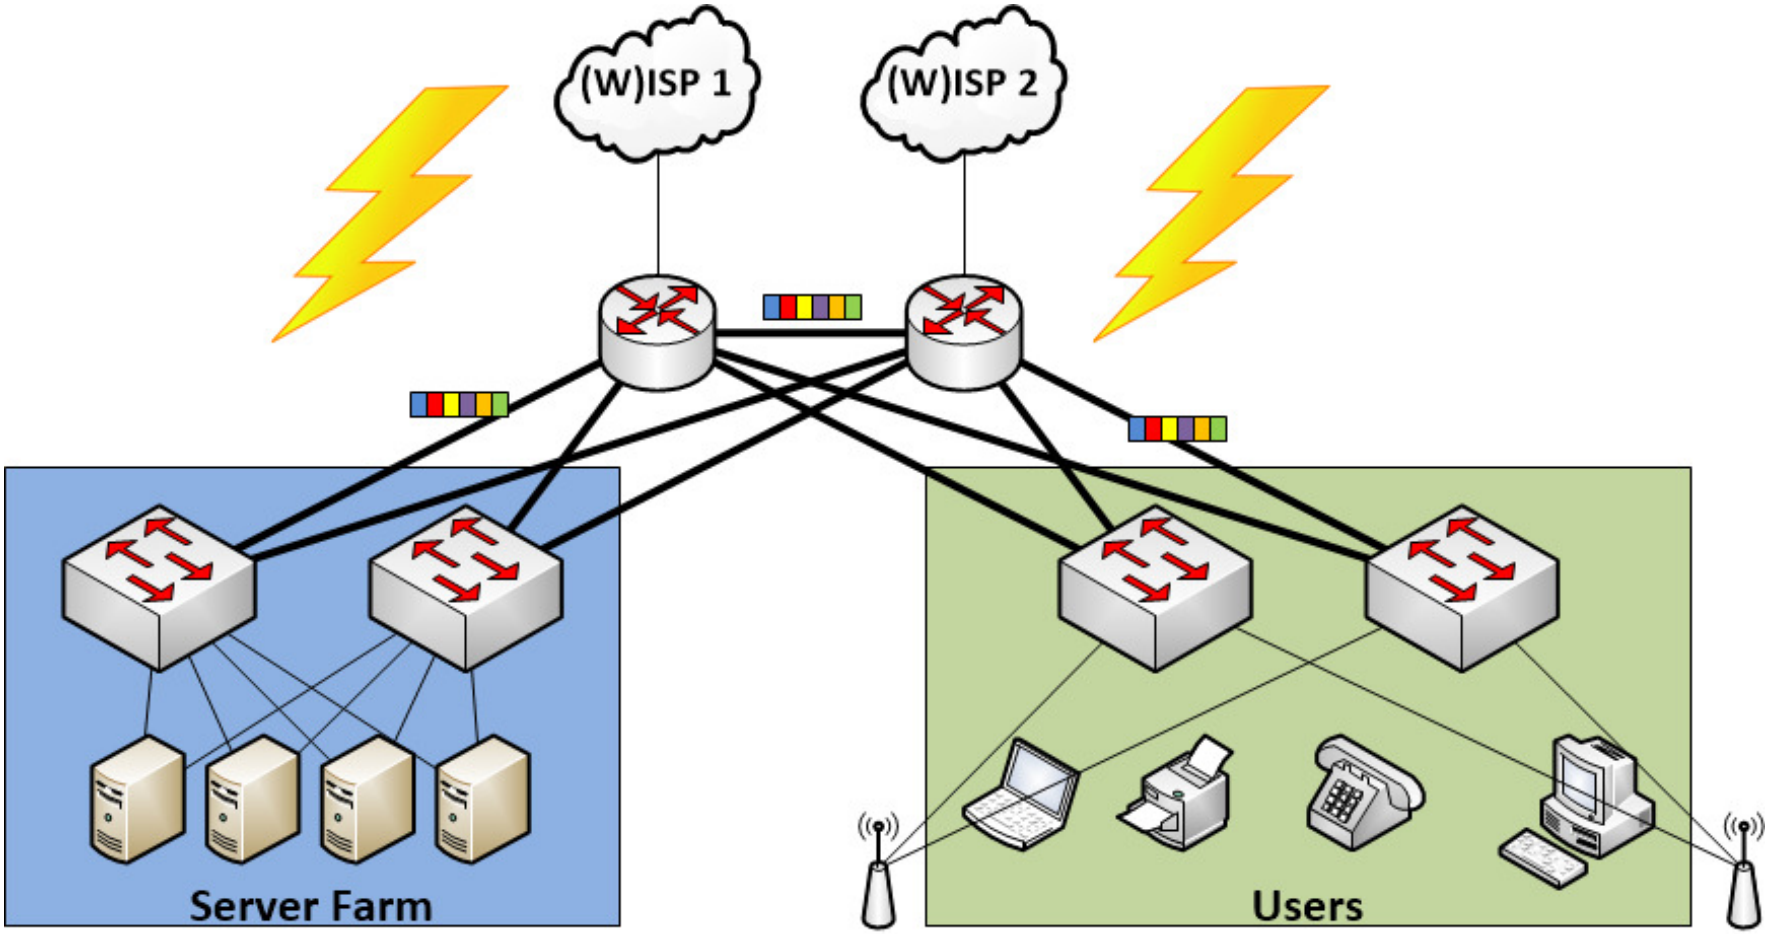

Topology

- How do we get there?
  - Duplicate hardware
  - Duplicate links
  - Duplicate electrical supplies
    - Circuits (120V/240V)
    - Power supplies (in routers)
  - Proper configuration
    - VLANs
    - VRRP
    - Various (optimization)

CSPE
Centre de services professionnels en éducation

TECHNOLOGIES

Legal Disclaimer

What I present in the coming slides is one approach. There are other ways to implement the discussed topics.

CSPE
Centre de services professionnels en éducation

FLEXIBILITY

Definition

Merriam Webster defines "**Flexibility**" as

"*Characterized by a ready capability to adapt to new, different, or changing requirements*"

http://www.merriam-webster.com/dictionary/flexibility

Why design flexible networks?

- Adapt quickly to organizational growth, demands and changes
- Minimize costs associated to those changes

1st design (*newbie*)

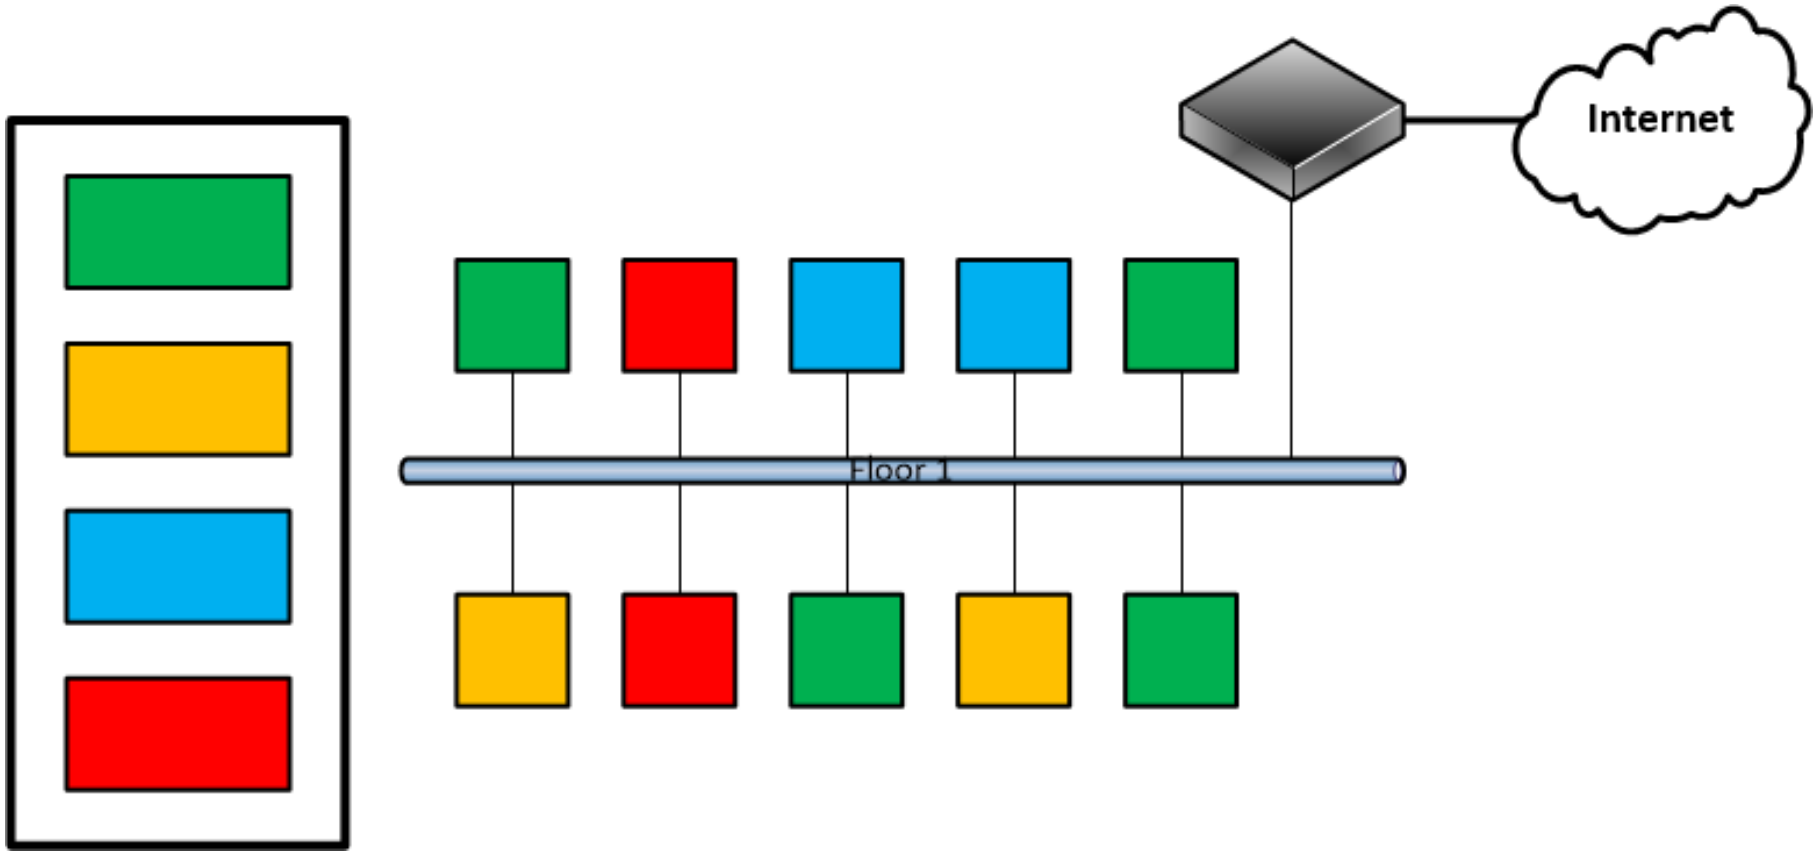

Design confronts reality

- Move Orange group to 2nd floor
- 1 Red and 1 Green on 2nd floor but connected to their LAN (*1st floor*)
- 1 Orange on 1st floor but connected to his LAN (*2nd floor*)
- Isolate Red group that plays Halo (*and slows down the network*)

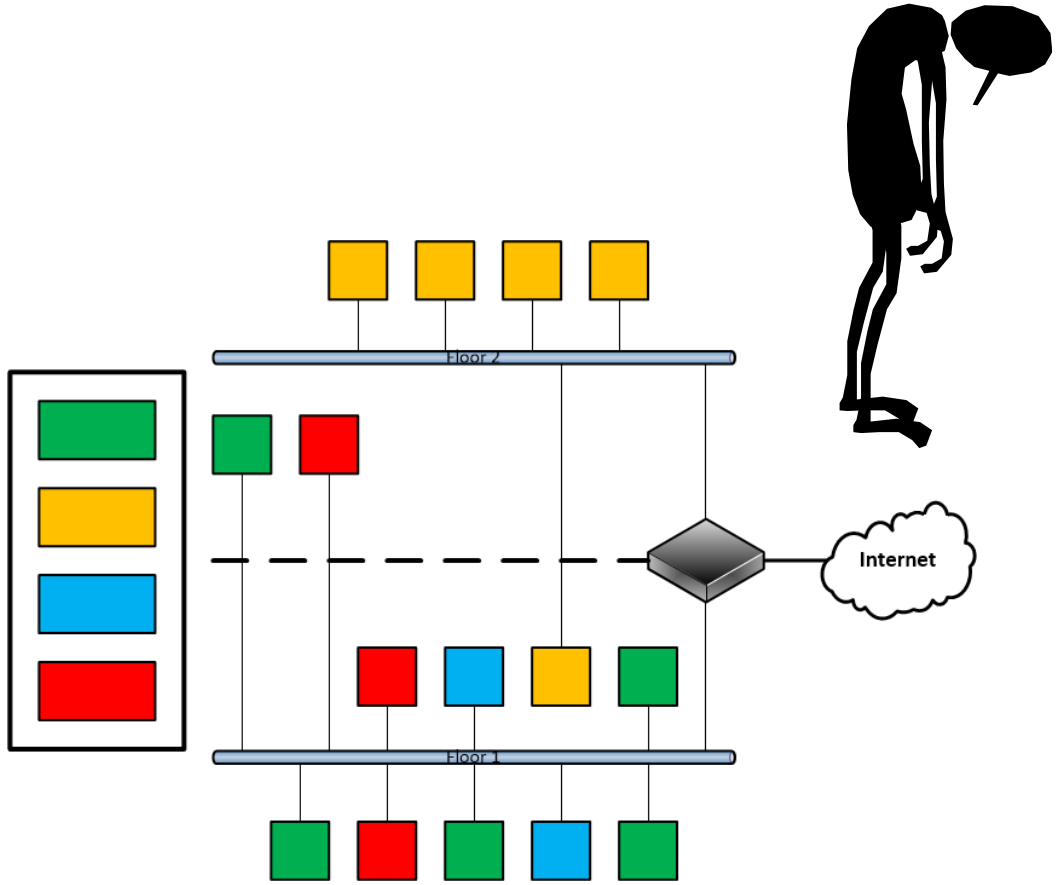

What to do?

- VLANs

What are VLANs?

- Layer 2 technology used to partition networks into separate “virtual” broadcast domains

2nd design

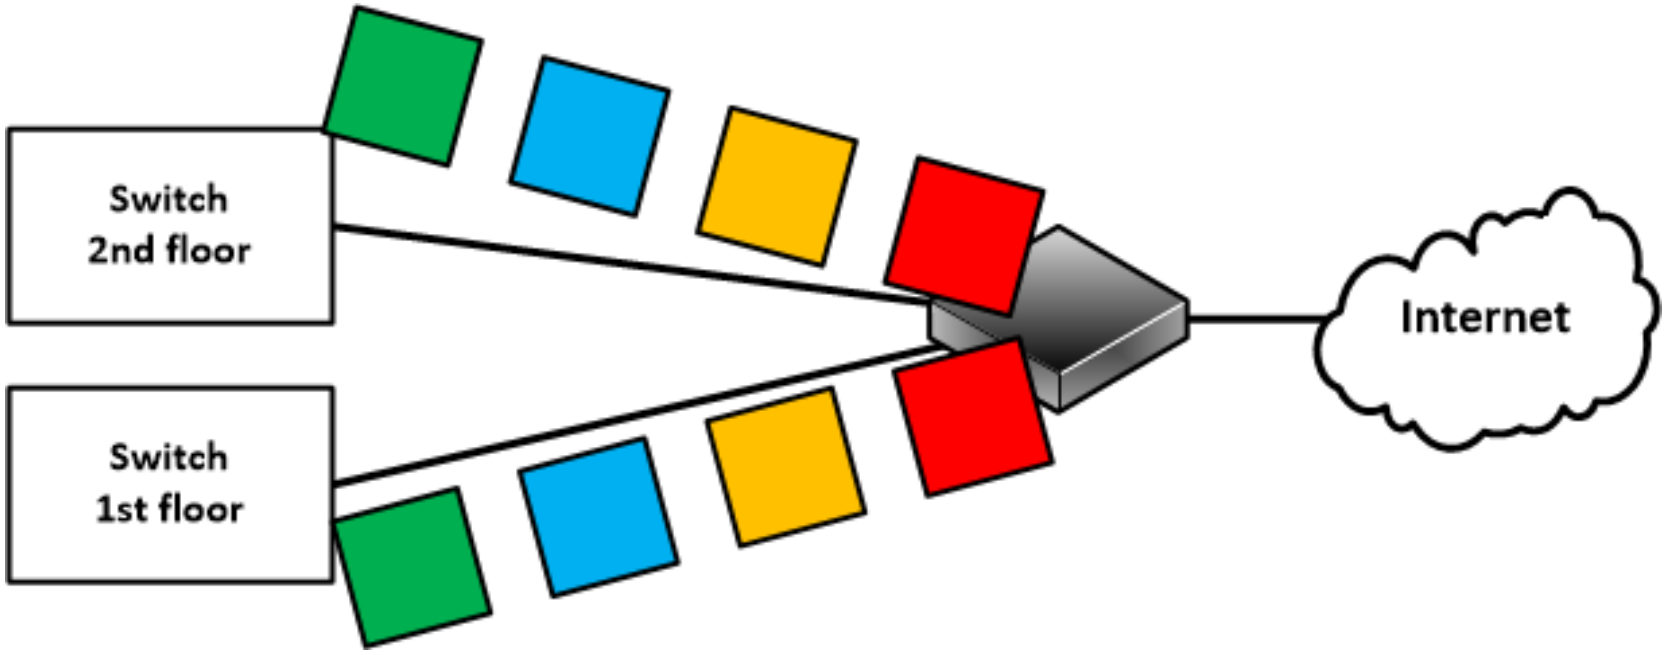

How to visualise VLAN configs

- Trunk ports

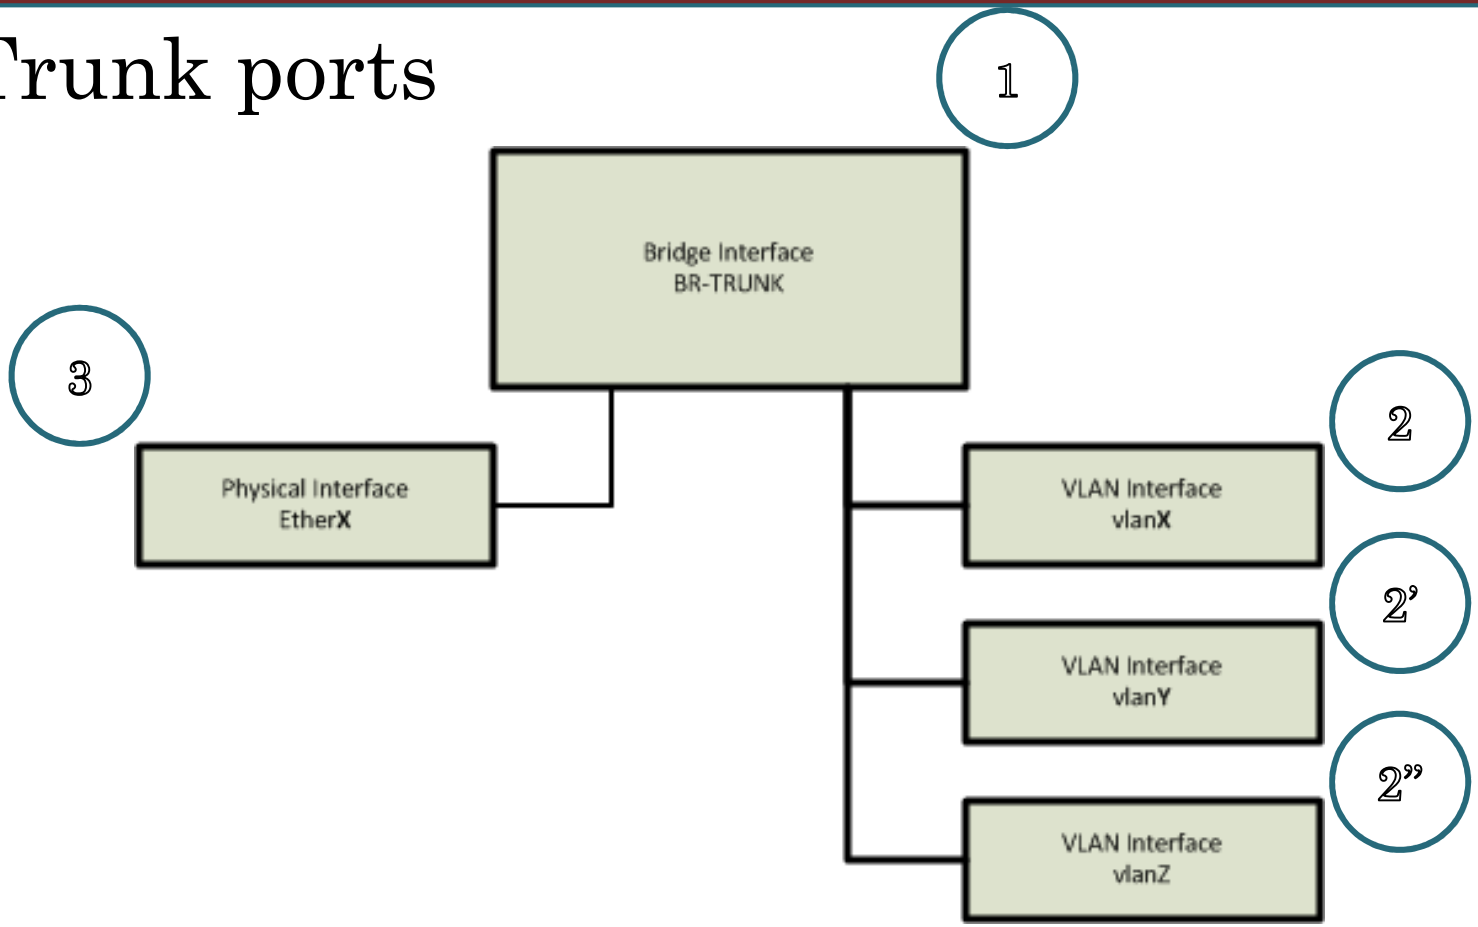

How to visualise VLAN configs

- Regular VLANs

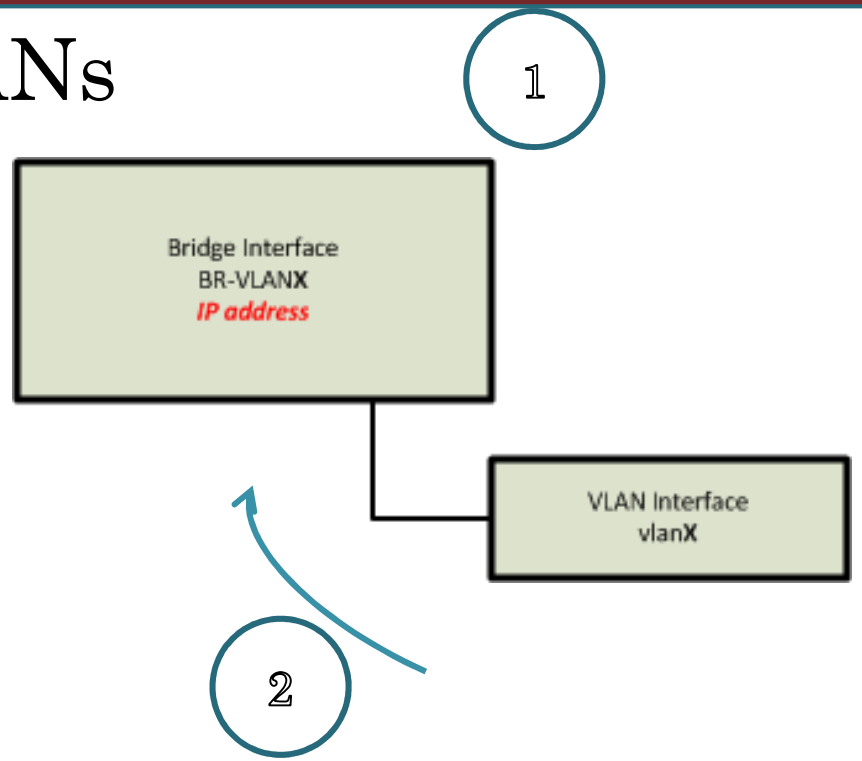

How to visualise VLAN configs

- Access port(s)

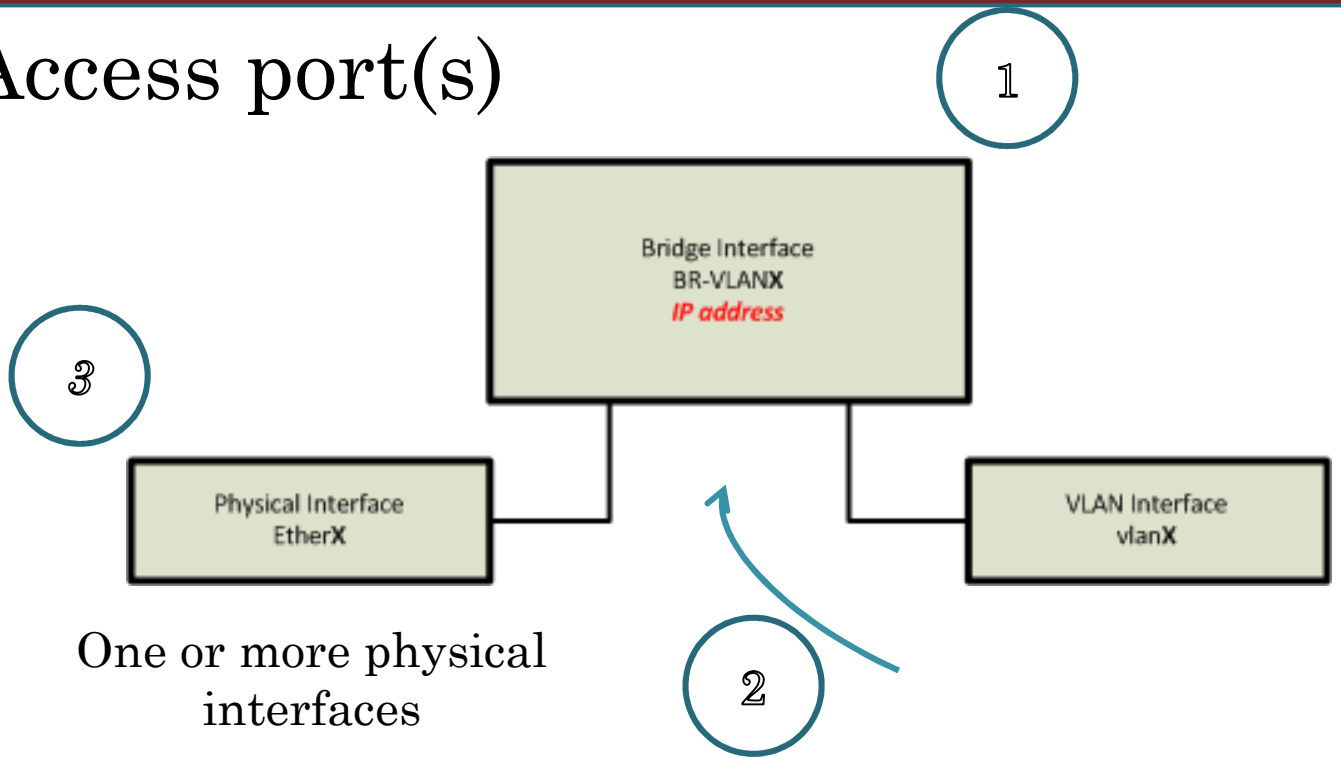

One or more physical interfaces

VLANs

- How to do it

Creating VLANs (step by step)

- **Bridge interfaces**
  - 1 to host all VLANs

    and
  - 1 per VLAN

```
/interface bridge
add name=BR-TRUNK
add name=BR-VLAN1
add name=BR-VLAN2
```

Creating VLANs (step by step)

- **VLAN interfaces**
  - 1 per VLAN

```
/interface vlan
add interface=BR-TRUNK name=vlan1 vlan-id=1
add interface=BR-TRUNK name=vlan2 vlan-id=2
```

Creating VLANs (step by step)

Bind interfaces

- Trunk ports
  - 1 or more physical interfaces (*ex.* →*BR-TRUNK*)
- VLANs and access ports
  - 1 VLAN interface (*ex.* → *BR-VLANX*) (*must*)
  - 1 or more physical interfaces (*ex.* → *BR-VLANX*) (*for access ports if present*)

```
/interface bridge port
add bridge=BR-TRUNK interface=ether5
add bridge=BR-VLAN1 interface=vlan1
add bridge=BR-VLAN1 interface=ether1
```

Finishing touches

- Create:
  - IP addresses on VLANs
    - *BR-VLANx*
    - ***Trunk bridge does not get IP address***
  - DHCP server on VLANs
    - *BR-VLANx*
    - ***Trunk bridge does not get DHCP server***
  - Others parameters
    - DNS, (S)NTP, Identity, etc.

Optimization

- Bridge STP Priorities
  - Cores have higher priority
    - **Master** has higher (*ex. 1000 hex*)
    - **Backup** has lower (*ex. 1001 hex*)
  - Others routers keep default (*8000 hex*)

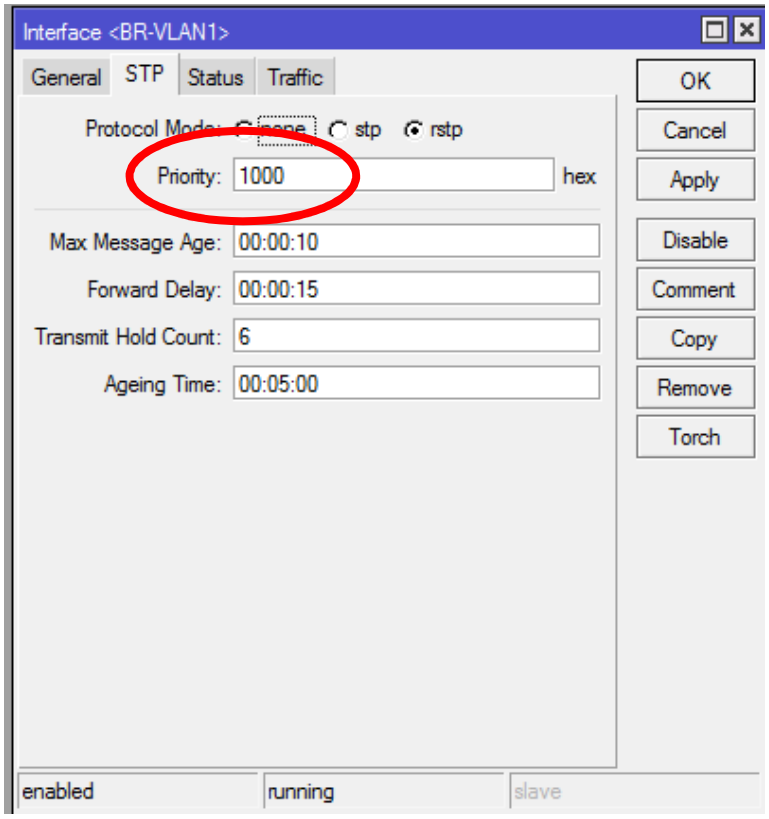

Optimization

- Why change STP priority?
  - Anecdote: ***The forgotten bridge!!***

Optimization (anecdote)

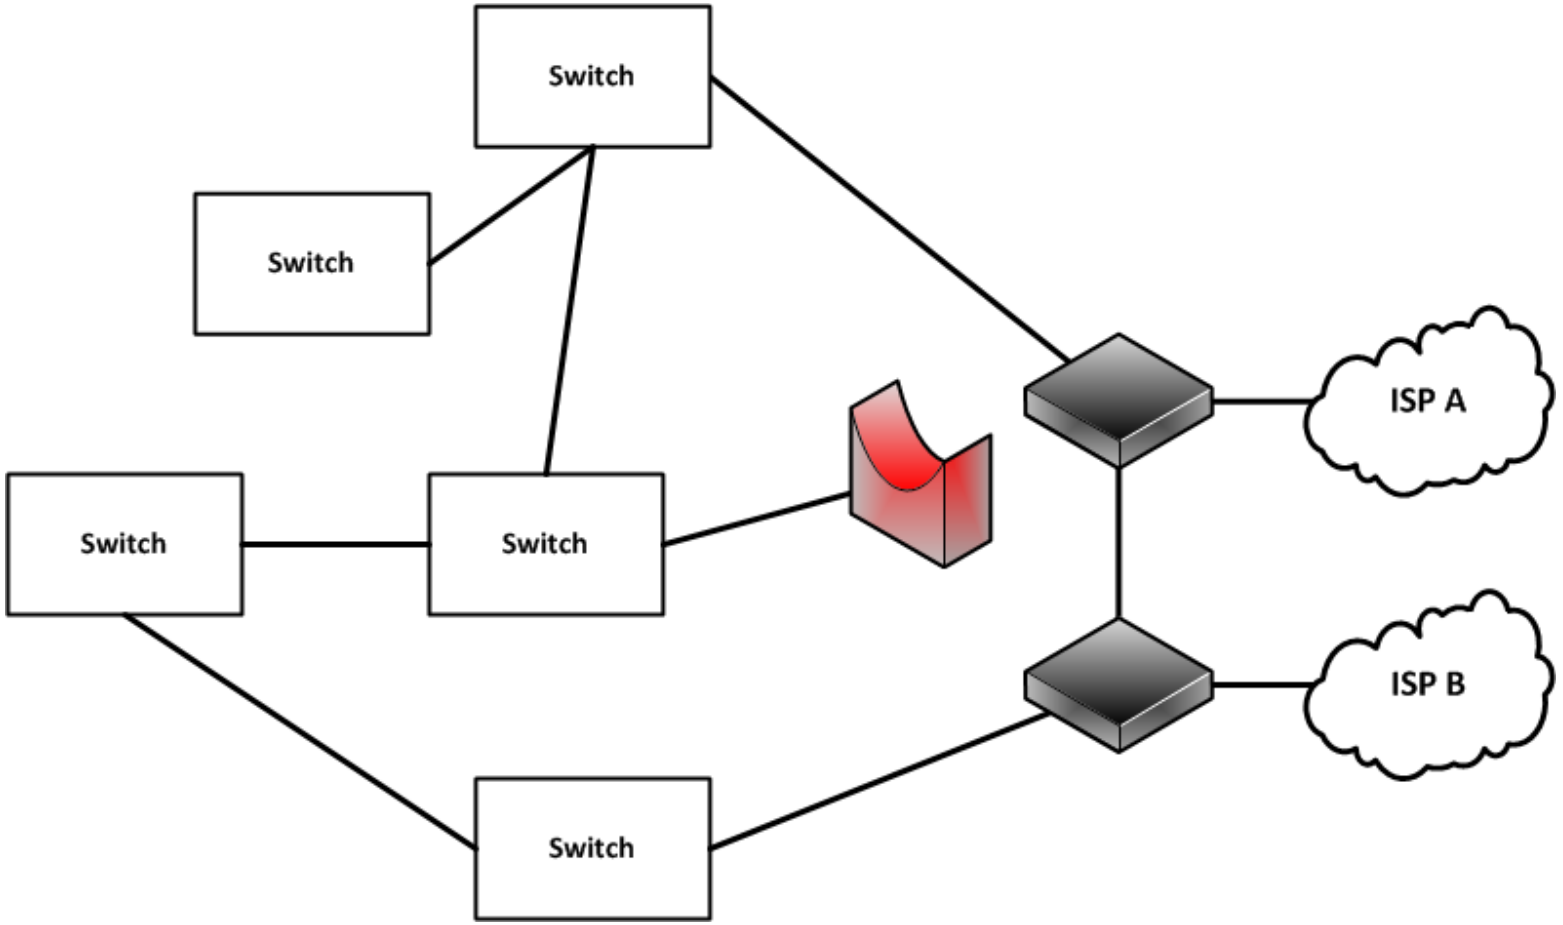

Optimization

- Bridge ports
  - Edge (Trunk=no, Access=yes)
  - Point-to-point (Full duplex=yes, Half duplex=no)

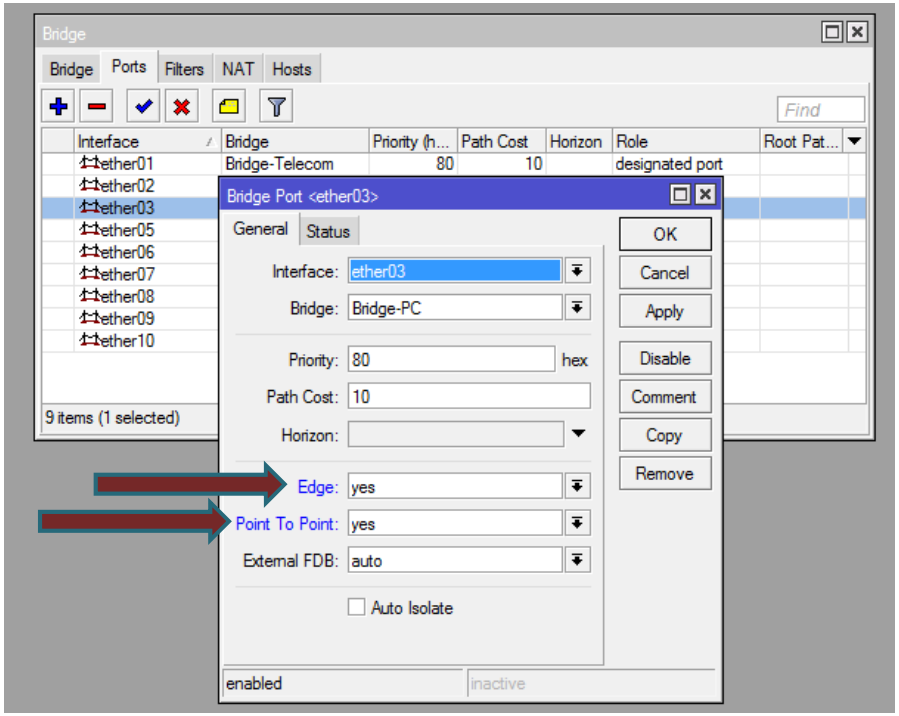

Optimization

- Goals of all these steps:
  - Speed up convergence
  - Be able to predict network behavior and operations

CSPE
Centre de services professionnels en éducation

REDUNDANCY

Definition

Merriam Webster defines "**Redundancy**" as

"*A part in a machine, system, etc., that has the same function as another part and that exists so that the entire machine, system, etc., will not fail if the main part fails.*"

http://www.merriam-webster.com/dictionary/redundancy

Why should we have a redundant network?

- Avoid downtime
  - Loss of productivity
  - Loss of revenue
- SLAs (*Service Level Agreements*) can carry cash penalties if ISPs fail to meet contractual goals

2nd design

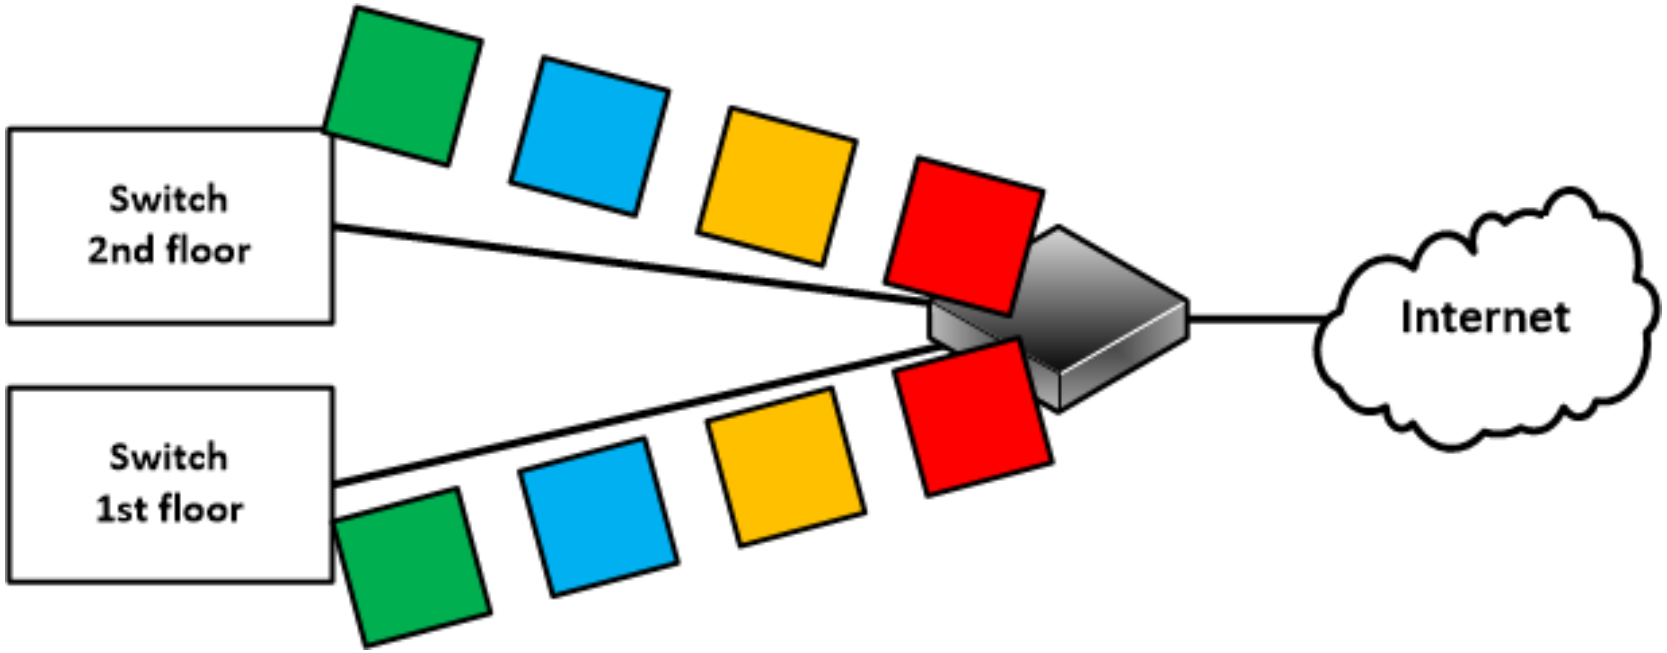

Design confronts reality

What to do?

- Hardware redundancy
- Link redundancy
- VRRP

3rd design

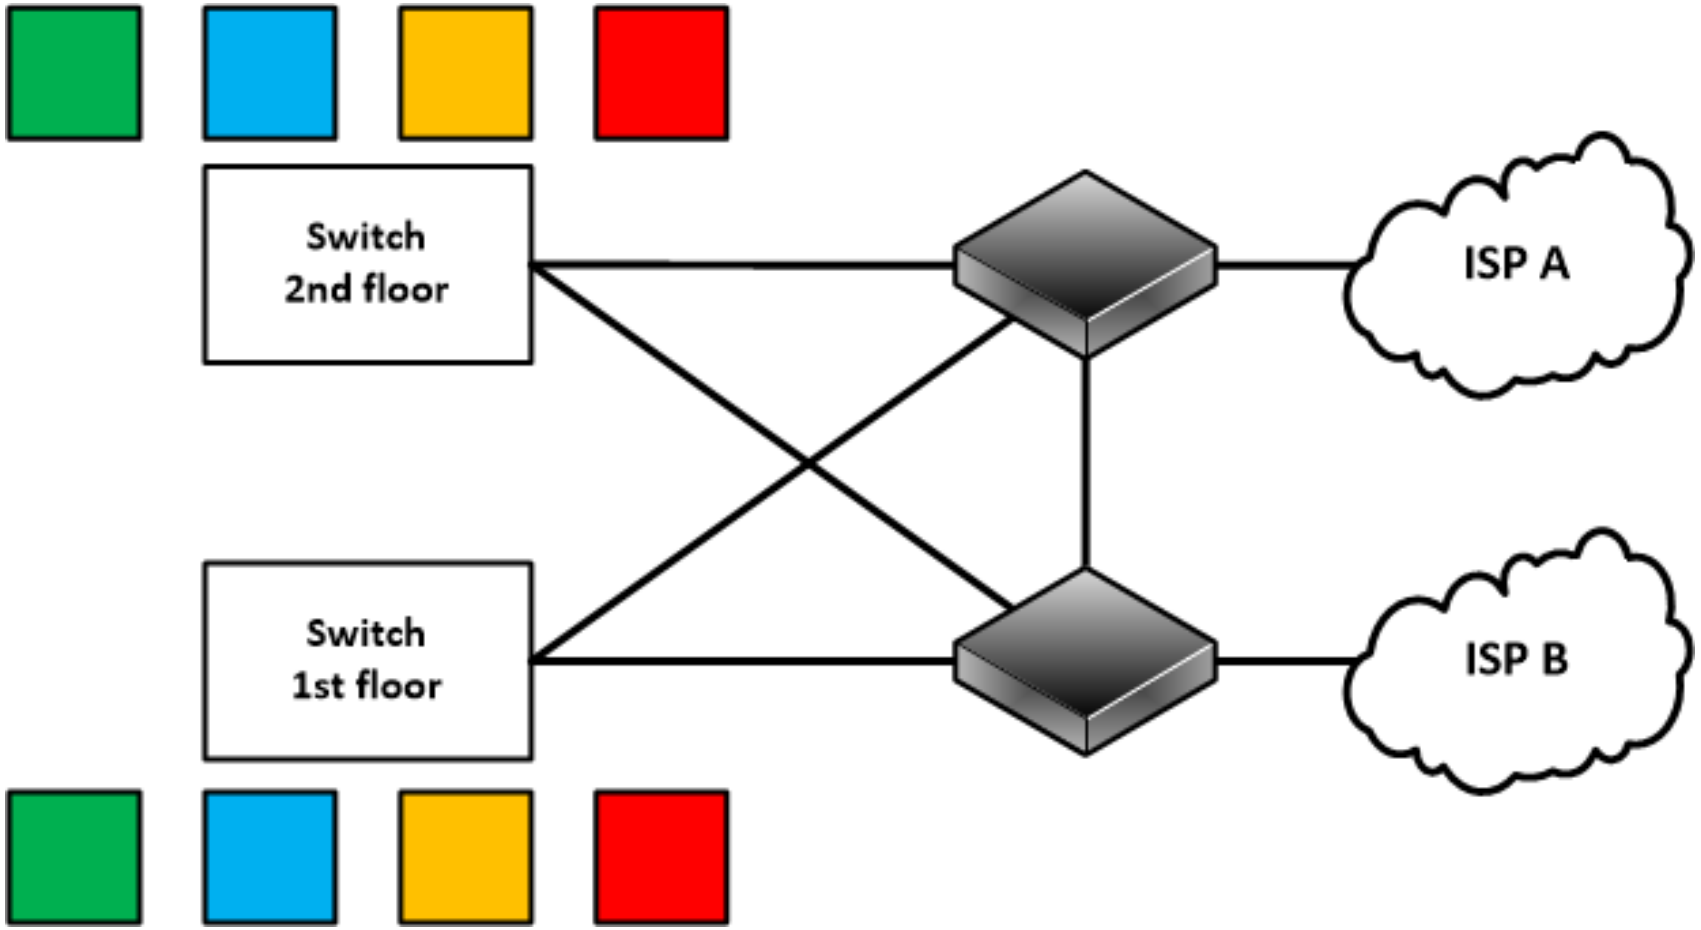

3rd design

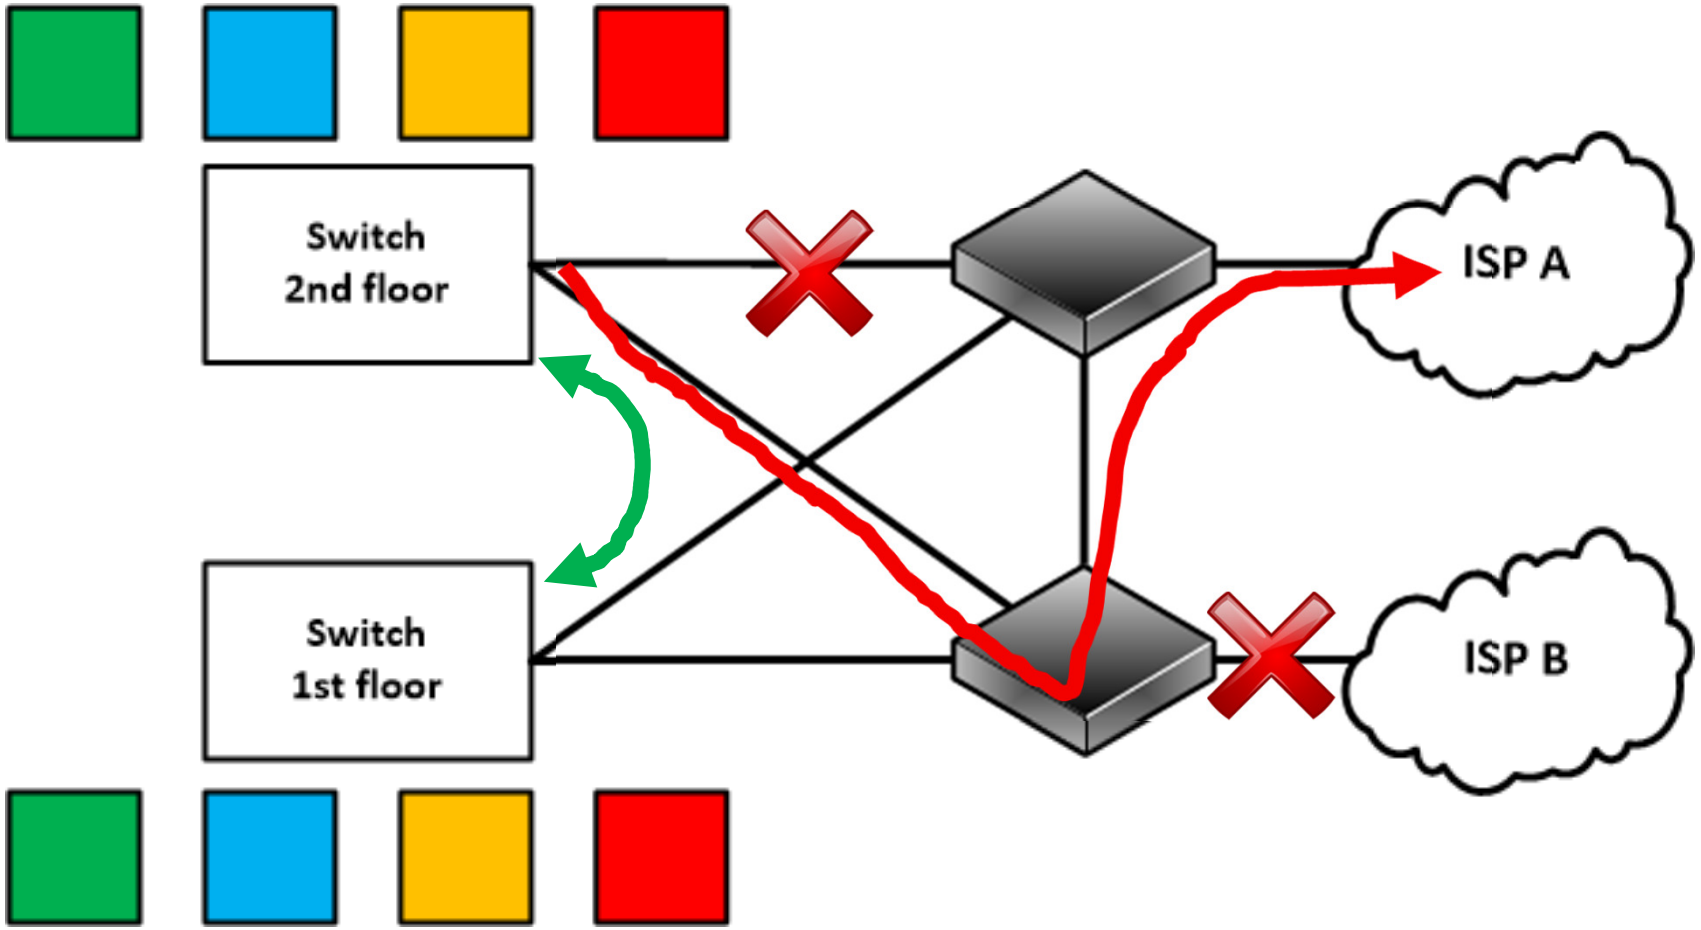

3rd design

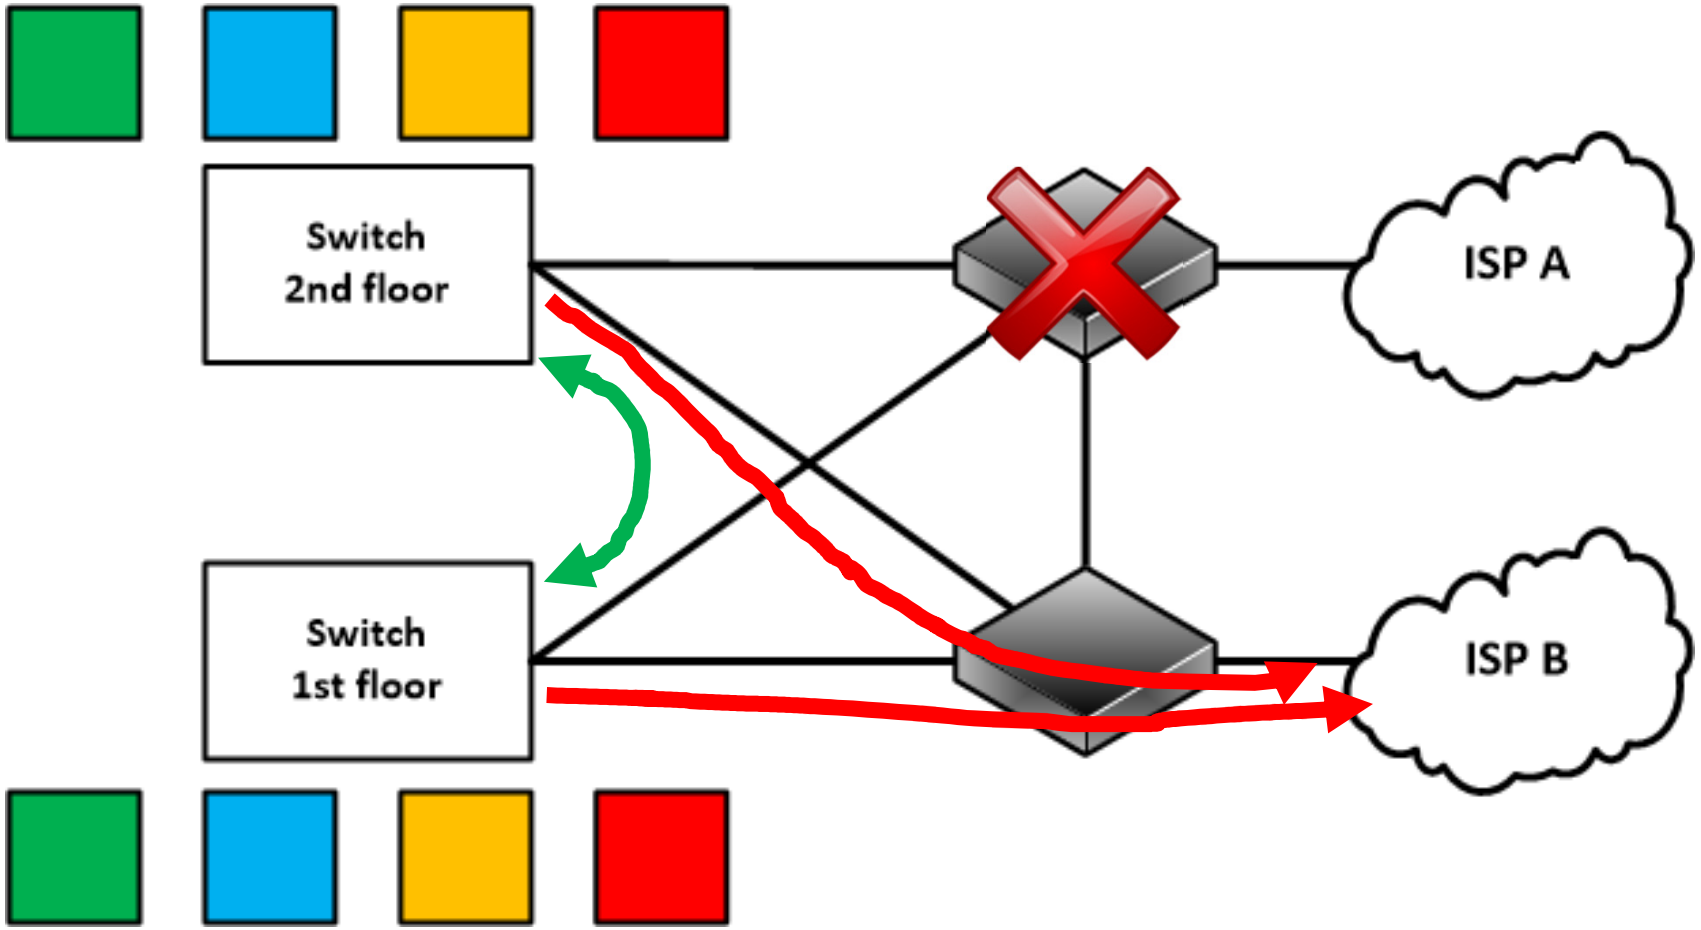

Is hardware redundancy enough?

- NO!
  - Answer: VRRP

What is VRRP?

VRRP: *virtual router redundancy protocol*

- A protocol that assigns responsibility for a virtual router (***VR***) to **one** physical router inside a group (*two or more routers*)
- Shares the control of a virtual IP address between those members

Why use VRRP?

- Insures the availability of the default gateway as long as one member remains active

How to visualise VRRP configs

- VLAN α

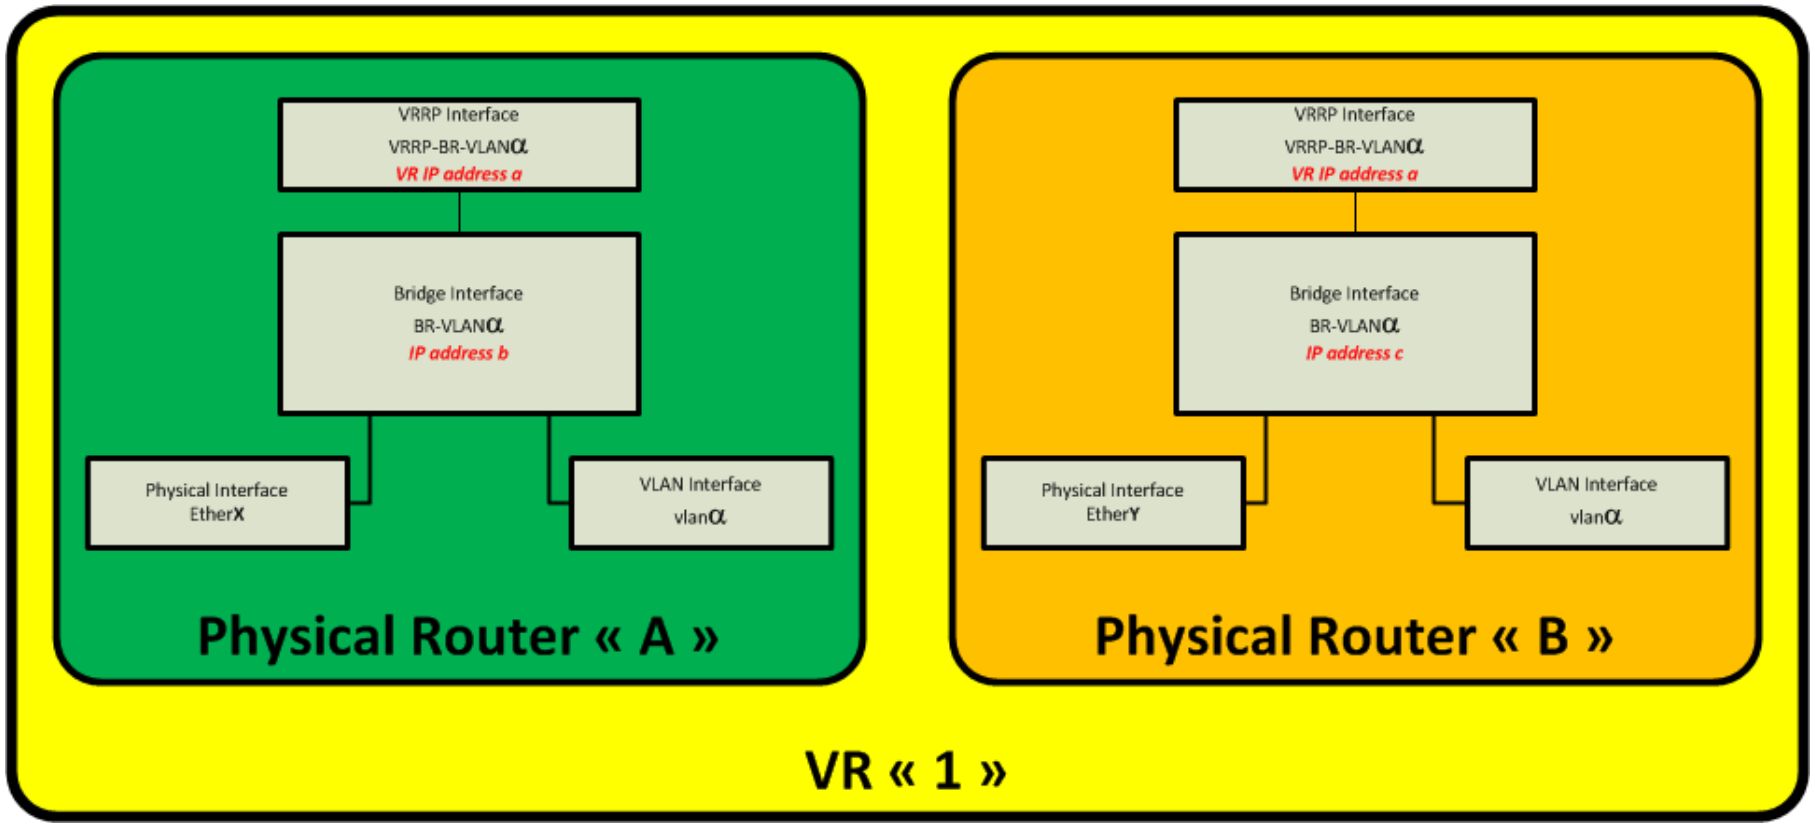

VRRP

- How to do it

Activating VRRP (step by step)

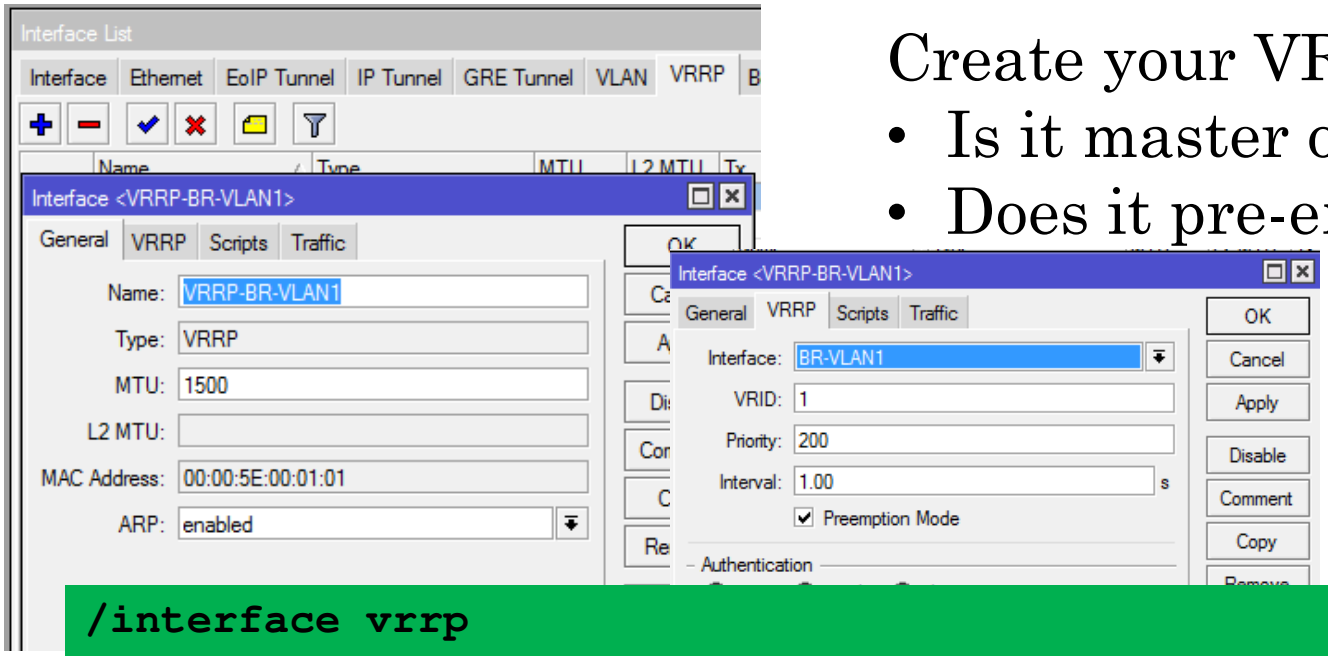

Create your VRRP interface.

- Is it master or backup?
- Does it pre-empt?

```
/interface vrrp
add interface=BR-VLAN1 name=VRRP-BR-VLAN1 priority=200
add interface=BR-VLAN2 name=VRRP-BR-VLAN2 priority=200 vrid=2
```

```
/interface vrrp
add interface=BR-VLAN1 name=VRRP-BR-VLAN1 preemption-mode=no
add interface=BR-VLAN2 name=VRRP-BR-VLAN2 preemption-mode=no vrid=2
```

Activating VRRP (step by step)

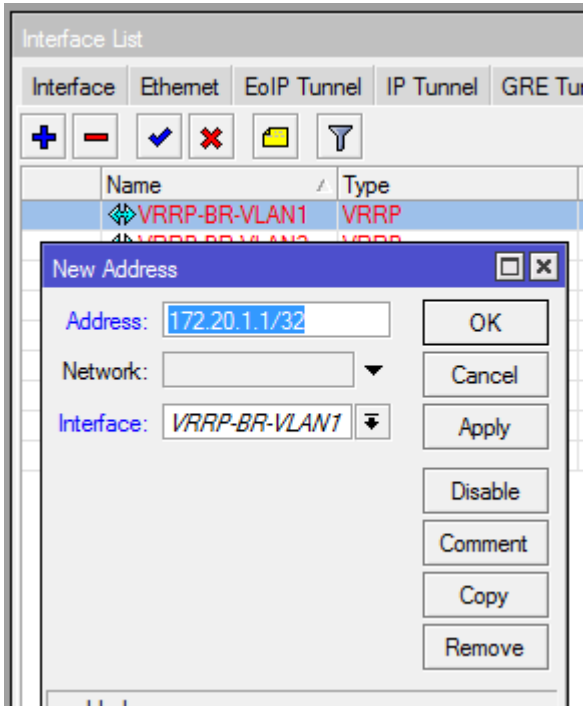

- Create an IP address for the VRRP interface.
  - Use free address
  - /32 mask
- If required, modify the IP address on the Bridge interface
- **Never have identical real & shared IP address**

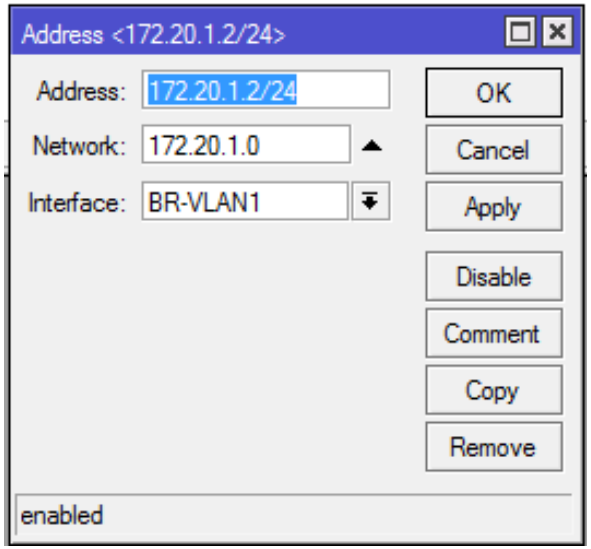

Optimization

- VRRP priorities
  - One core router is master for all VLANs
  - Other core is backup for all VLANs
- Why?
  - If the master core router fails, the backup quickly takes over
  - If the backup core router fails… not much happens

CSPE
Centre de services professionnels en éducation

VARIOUS

Optimization

- OSPF Interface priority
  - DR=high value (*ex. 100*)
  - BDR=low value (*ex. 50*)
  - Member=0
- Everything must be predictable!
  - DR, BDR, members

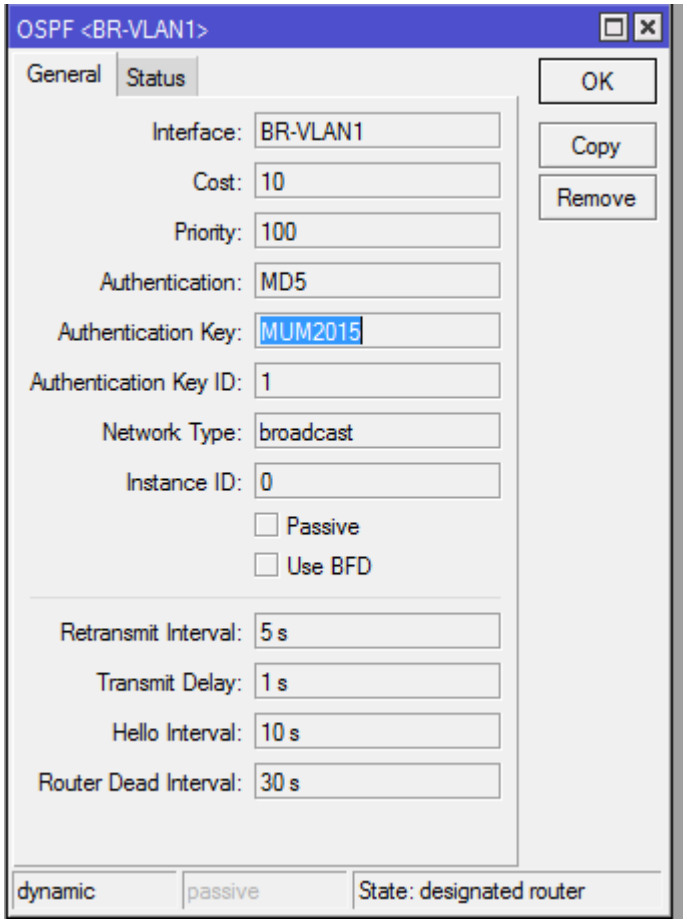

Optimization

- OSPF
  - Lower “hello-interval”
  - Lower “dead-interval”
- Beware of more OSPF traffic!

Pitfalls

- Configure one step at a time
  - Configure routers first
  - Plug after, otherwise spanning tree loops!
    - Traffic
    - CPU
    - You

Wish

- MSTP support
  - IEEE 802.1Q-2014, Chap. 13
- Why?
  - True Layer 2 and Layer 3 load-sharing
    - Load-shared VRRP config NOT recommended on two core routers without MSTP
  - In case of loss of a core router, only half affected (for a brief moment)
  - Easier migration from Cisco (PVSTP)

CSPE
Centre de services professionnels en éducation

DEMONSTRATIONS

Setup

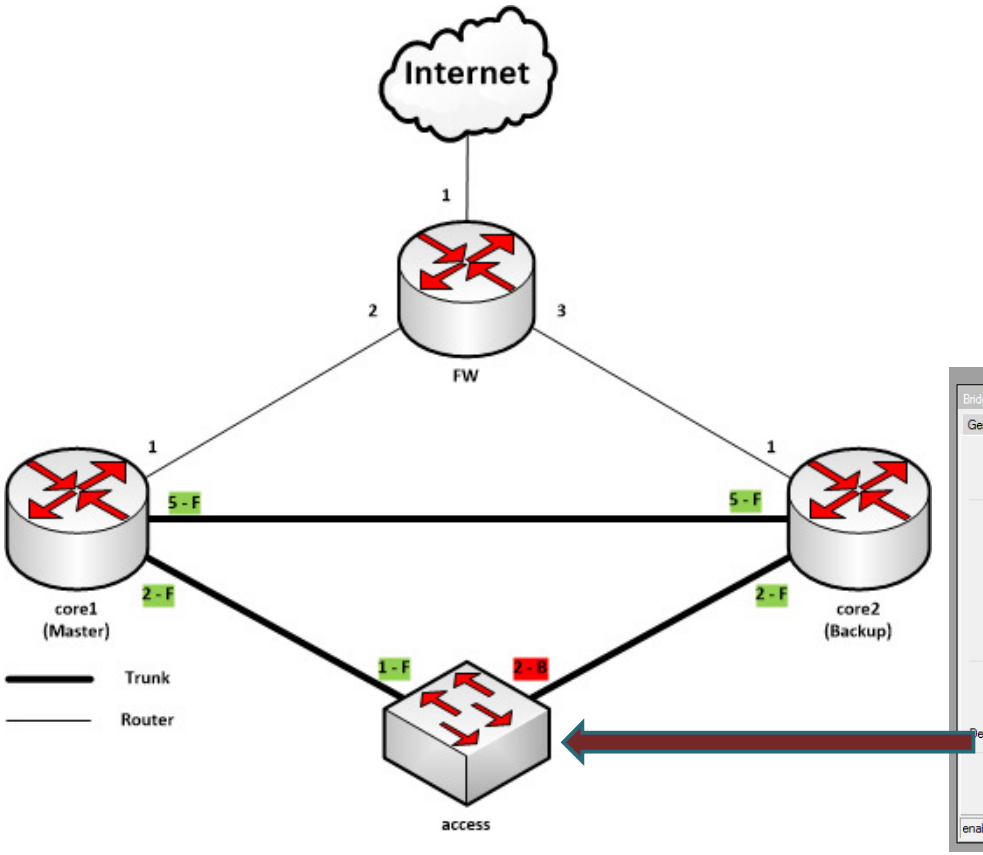
Internet
1
2
3
FW
1
1
5 - F
5 - F
2 - F
2 - F
core1
(Master)
core2
(Backup)
1 - F
2 - B
Trunk
Router
access

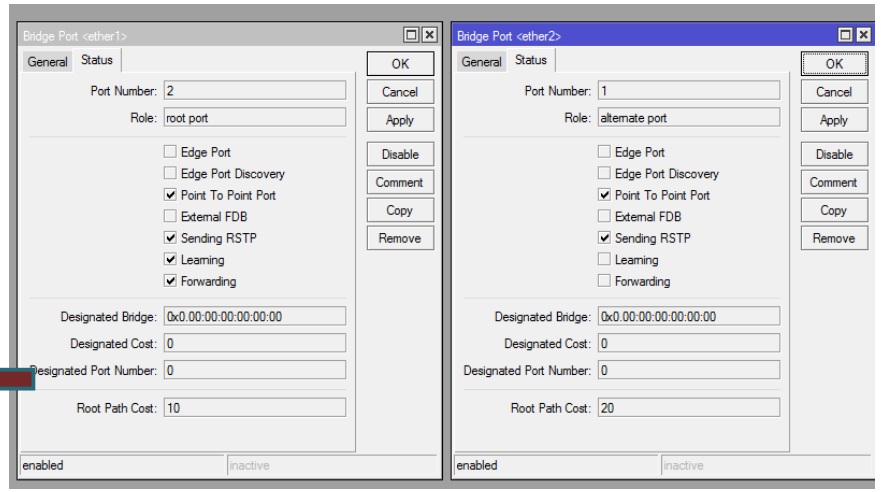
Bridge Port <ether1>
General
Status
Port Number: 2
Role: root port
Edge Port
Edge Port Discovery
Point To Point Port
External FDB
Sending RSTP
Learning
Forwarding
Designated Bridge: 0x0.00:00:00:00:00:00
Designated Cost: 0
Designated Port Number: 0
Root Path Cost: 10
OK
Cancel
Apply
Disable
Comment
Copy
Remove
enabled
inactive
Bridge Port <ether2>
General
Status
Port Number: 1
Role: alternate port
Edge Port
Edge Port Discovery
Point To Point Port
External FDB
Sending RSTP
Learning
Forwarding
Designated Bridge: 0x0.00:00:00:00:00:00
Designated Cost: 0
Designated Port Number: 0
Root Path Cost: 20
OK
Cancel
Apply
Disable
Comment
Copy
Remove
enabled
inactive

Outage of main core

Démo1-core1.mp4

Outage of backup core

Démo1-core2.mp4

CSPE
Centre de services professionnels en éducation

FINAL WORDS

What have we seen?

- Flexibility using VLANs
- Redundancy by doubling links and hardware and by using VRRP
- Certain things to be aware of

CSPE
Centre de services professionnels en éducation

QUESTION PERIOD

1st MUM in Canada!

Thank you MikroTik!!