



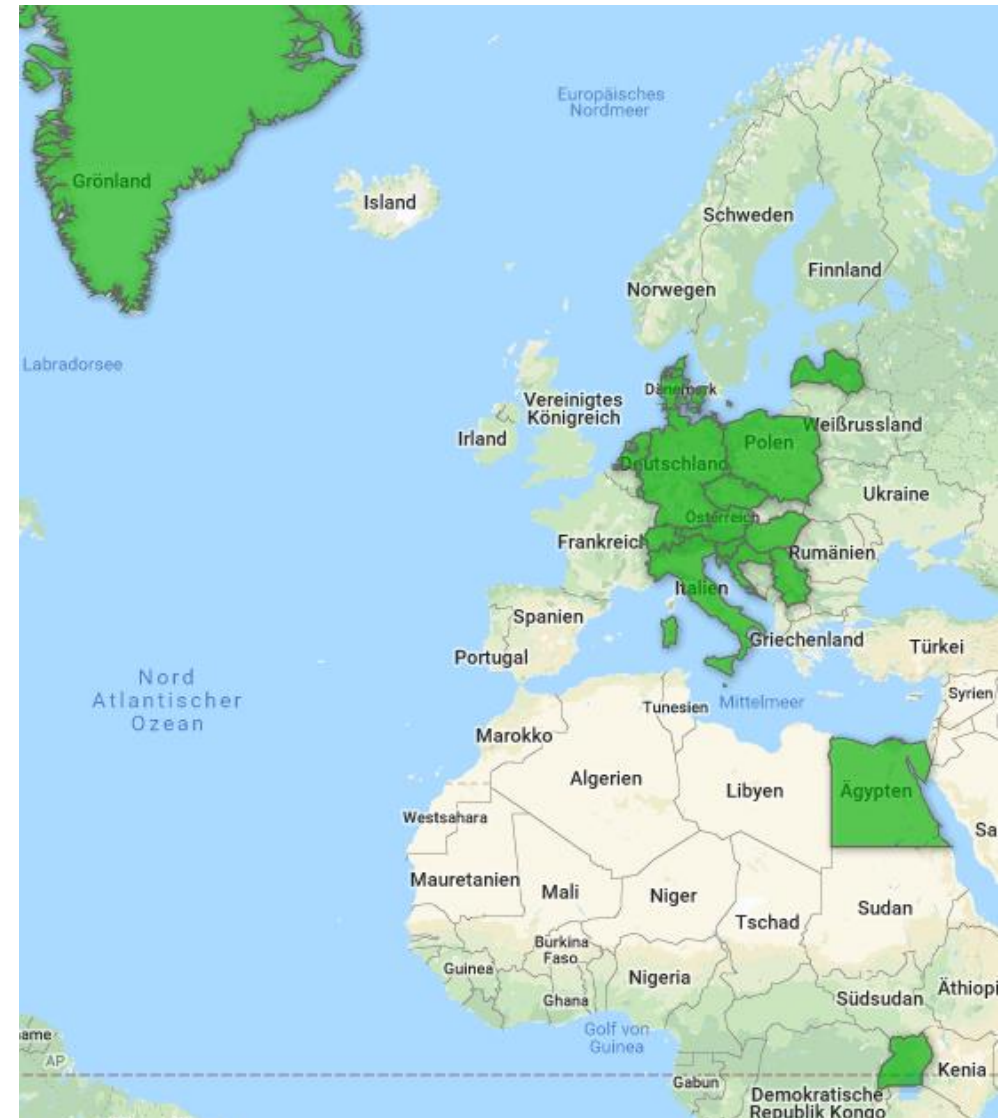
Let's take a look at the

Multiple Spanning Tree Protocol (MSTP)



About me

- Sebastian Inacker
<inacker@fmsweb.de>
- FMS Internetservice GmbH
Germany
- MikroTik Trainer
(TR0011, May 2007)
- MTCNA, MTCRE,
MTCTCE, MTCUME,
MTCWE, MTCIPv6E,
MTCINE, MTCSE





MikroTik trainings and workshops

- Own training center (south-west of Germany) and on site (Austria, Denmark, Germany, Greenland, Hungary, Luxembourg, Malta, Netherlands, Switzerland, Uganda)





Overview / big picture

“Implementing and running a RSTP (MSTP) network is easy.
As long as it is running well.”

Topics:

- Evolving network (improvements)
- Maybe more (R)STP than you expected (Background)
- MSTP for improved network



Overview / big picture

No topics:

- “New” style of bridge configuration
(See MUM presentations or the MikroTik wiki)
- Redundancy of routers
(see for example Patrik Schaub, 2018 Berlin)
- BPDU (Bridge Protocol Data Units)




The beginning



Existing setup

- Main site + one remote site
- One (reliable) 100 Mbit/s connection
- Local network 192.168.100.0/24
- Redundant internet connectivity

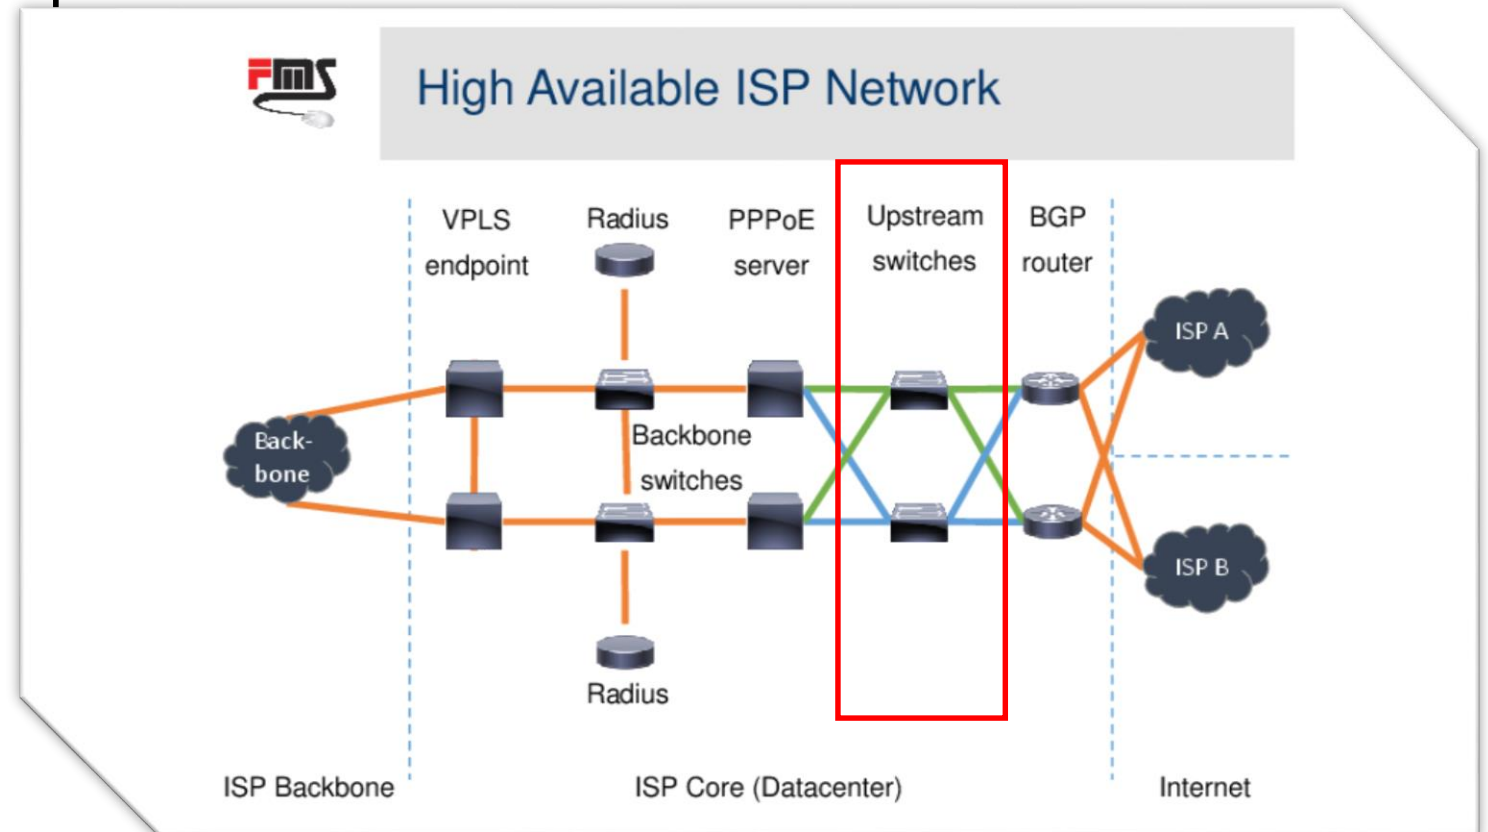
- 100 Mbit/s link will be replaced by new 1 GBps/s link





Redundant setup


Prior presentations about redundancy (by us, FMS Internetservice):
Skipped the switch setup.



Patrik Schaub,
MUM EU 2018 Berlin



Redundant setup





New setup (RSTP)

Used hardware:

- Router: CCR1009 (*example*, depending on network)
- Switch: CRS326 (CRS326-24G-2S+RM)

Software: RouterOS 6.43.12 / 6.44

The screenshot shows the RouterOS Bridge configuration window. The 'Ports' tab is selected, displaying a table of bridge ports. The table has columns for #, Interface, Bridge, Horizon, Trusted, Priority (hex), Path Cost, and Role. Two rows are visible, both with a red border around them:


#	Interface	Bridge	Horizon	Trusted	Priority (hex)	Path Cost	Role
0	H ether3	bridge		no	80	10	designated port
1	H ether4	bridge		no	80	10	designated port

At the bottom of the window, it says '2 items'.

CCR1009: HW offloading for RSTP on ether1-ether4



Traffic Flow





Redundant setup

Easy - just created bridges (with RSTP, default)

- Switches: All involved ports bridged
 - Routers: ether3, ether4 (faceing switches) bridged
- HW offloading on CCR1009 (e1-e4)

Issues with given setup:

- Only 1 IP-Subnet (no VLAN, no network seperation)
- Traffic uses 100 Mbit/s link.



Network separation (VLAN)



VLAN

- VLAN 100: Client network
- VLAN 42: Management network

Users separated from management network

Also possible: VLAN for

- VoIP / Video
- Guest network (wireless)
- Backup / (VMWare|SQL) replication
- ...



VLAN

Example config switch local 1

Bridge

Bridge Ports VLANs MSTIs Port MST Overrides Filters NAT Hosts MDB

+ - ✓ ✗ [icon] [icon] Find

#		Interface	Bridge	PVID	Comment
0	H	ether21	bridge		1 to router 1
1	H	ether22	bridge		1 to router 2
2	H	ether24	bridge		1 to remote site
3	H	sfp-sfplus 1	bridge		1 to neighbour switch
4	H	ether17	bridge	100	client laptop

5 items

Bridge

Bridge Ports VLANs MSTIs Port MST Overrides Filters NAT Hosts MDB

+ - ✓ ✗ [icon] [icon] Find

Bridge	VLAN IDs	Tagged	Untagged	Comment
bridge	42	ether21, ether22, ether24, sfp-sfplus 1, bridge		management
bridge	100	ether21, ether22, ether24, sfp-sfplus 1	ether17	client network

2 items



VLAN

Example config router 1

Bridge configuration window showing interface assignments:

#	Interface	Bridge	PVID	Comment
0	ether3	bridge		1 tagged to switch local 1
1	ether4	bridge		1 tagged to switch local 2

2 items

Bridge configuration window showing VLAN configuration:

Bridge	VLAN IDs	Tagged	Untagged	Comment
bridge	42	ether3, ether4, bridge		management
bridge	100	ether3, ether4		client network

2 items



VLAN

- HW offloading: missing on router 1 / router 2

The screenshot shows two windows from Mikrotik WinBox. The 'Bridge' window displays a table of bridge interfaces:

#	Interface	Bridge	PVID	Comment
0	ether3	bridge		1 tagged to swi
1	ether4	bridge		1 tagged to swi


The 'Interface <bridge>' window shows the configuration for a selected interface. The 'VLAN' tab is active, and the 'VLAN Filtering' checkbox is checked. A red box highlights this checkbox, and a red arrow points from it to the 'ether4' interface in the Bridge window table.

- VLAN-Filtering=yes → HW offloading CRS3xx only

But: No *switching traffic* through router, right? → *Wrong!*



Local traffic through router





Issues to solve

Switch to switch: Port
disabled by RSTP


- on local switch 1
- on remote switch 2

#	Interface	Bridge	PVID	Forwarding	Comment
0	H ether21	bridge		1 yes	to router 1
1	H ether22	bridge		1 yes	to router 2
2	H ether24	bridge		1 yes	to remote site
3	H sfppplus1	bridge		1 no	to neighbour switch
4	H ether17	bridge		100 yes	client laptop

- We have „switching traffic“ through router 1
- VLAN-Filtering required → No hw offloading on router → CPU
- Also: Slow link to remote site is used



Local switch to switch disabled





Path decision



Path decision on RSTP

How does RSTP selects path / non forwarding ports?

- Selection of „root bridge“
 1. Lowest bridge priority (default 0x8000 = 32768)
 2. Lowest bridge MAC address

Bridge ID: 0x<priority>.<mac_address>

Recommended priorities: Steps of 4096 (*required* for MSTP)

0x0000, 0x1000, 0x2000, ... 0x9000, 0xA000, ..., 0xF000



Path decision on RSTP

How does RSTP selects path / non forwarding ports?

- Selektion of port roles (forwarding: yes/no)
 1. Lowest port path costs (default 10) on ports *to root bridge*
 2. Lowest port priority (default 0x80) for ports *on root bridge*
 3. Lowest bridge port ID

Path costs: On ports facing *to root* bridge

Port priority: On ports facing *away from root* bridge



Port roles

Ports to root bridge:

- root-port: Forwarding traffic.
- alternate-port: Not forwarding traffic. (Backup for root-port)

Use *path costs* to select *path to root* bridge.

The screenshot shows a network configuration window titled "Bridge". It has several tabs: "Bridge", "Ports", "VLANs", "MSTIs", "Port MST Overrides", "Filters", "NAT", "Hosts", and "MDB". The "Ports" tab is selected. Below the tabs is a toolbar with icons for adding (+), removing (-), checking (✓), unchecking (✗), saving (floppy), and filtering (funnel). A "Find" search box is on the right. Below the toolbar is a table with the following data:

#	Interface	Bridge	Role	Forwarding
0	Ether1	rstp-bridge	root port	yes

The "Role" and "Forwarding" columns for the first row are highlighted with red boxes.



Port roles

Ports facing away from root bridge:

- designated-port: Forwarding traffic.
- backup-port: Not forwarding traffic.

Use *port priority* (if path costs are equal).

Additional port role:

- disabled-port: Disabled / inactive port (no link)



Root bridge

Please note:

- Root bridge (layer 2) is no default gateway (layer 3)
- Traffic is not forced to go through root bridge.
- Non root bridges will just know the path with lowest cost to root bridge

Selected root bridge will affect traffic distribution.



Path decision on STP

I think that I shall never see
A graph more lovely than a tree.
A tree whose crucial property
Is loop-free connectivity.
A tree that must be sure to span
So packets can reach every LAN.
First, the root must be selected.
By ID, it is elected.
Least-cost paths from root are traced.
In the tree, these paths are placed.
A mesh is made by folks like me,
Then bridges find a spanning tree.


Poem by Radia Perlman, who
invented STP.



Root bridge selection




Root bridge? I don't care...?!?



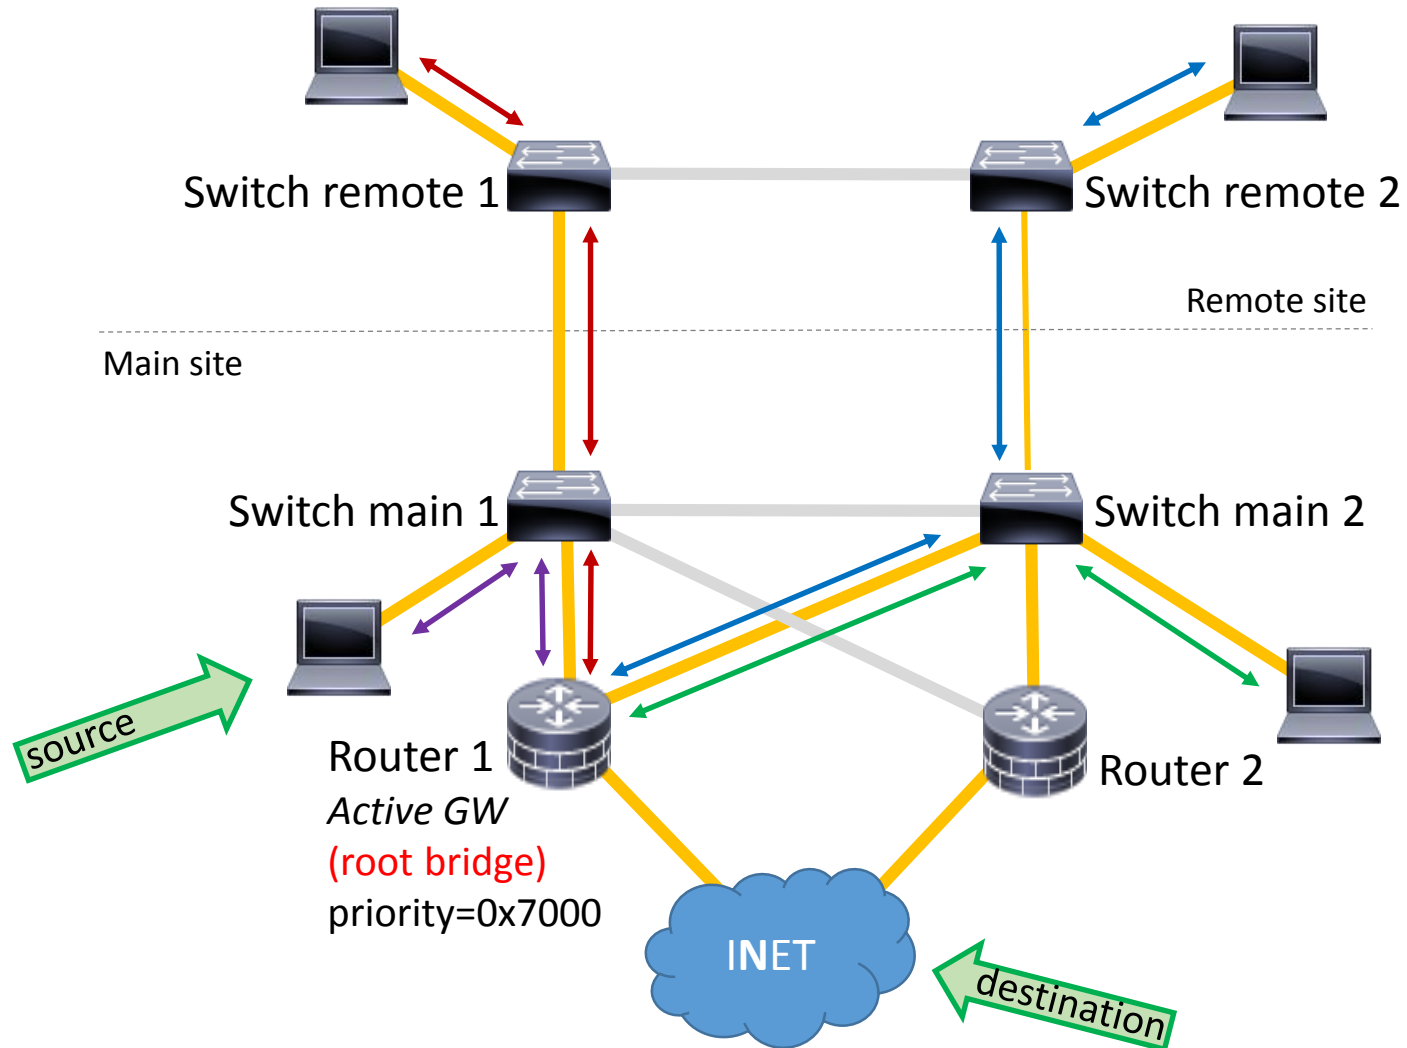


Root bridge? I don't care...?!?



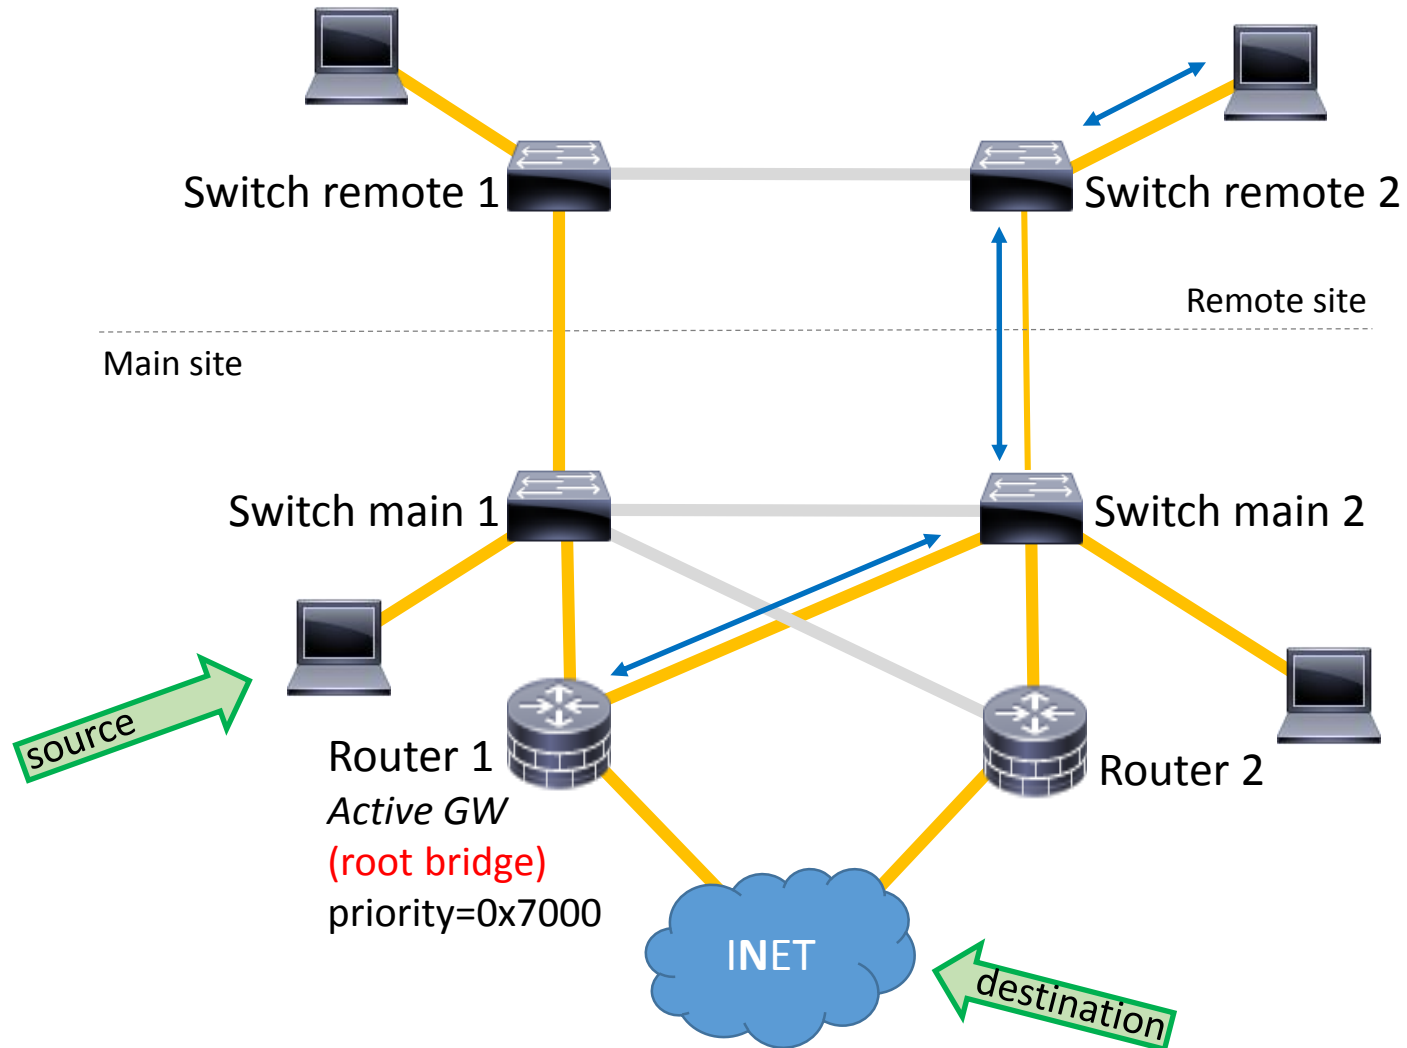


Traffic to internet is ok



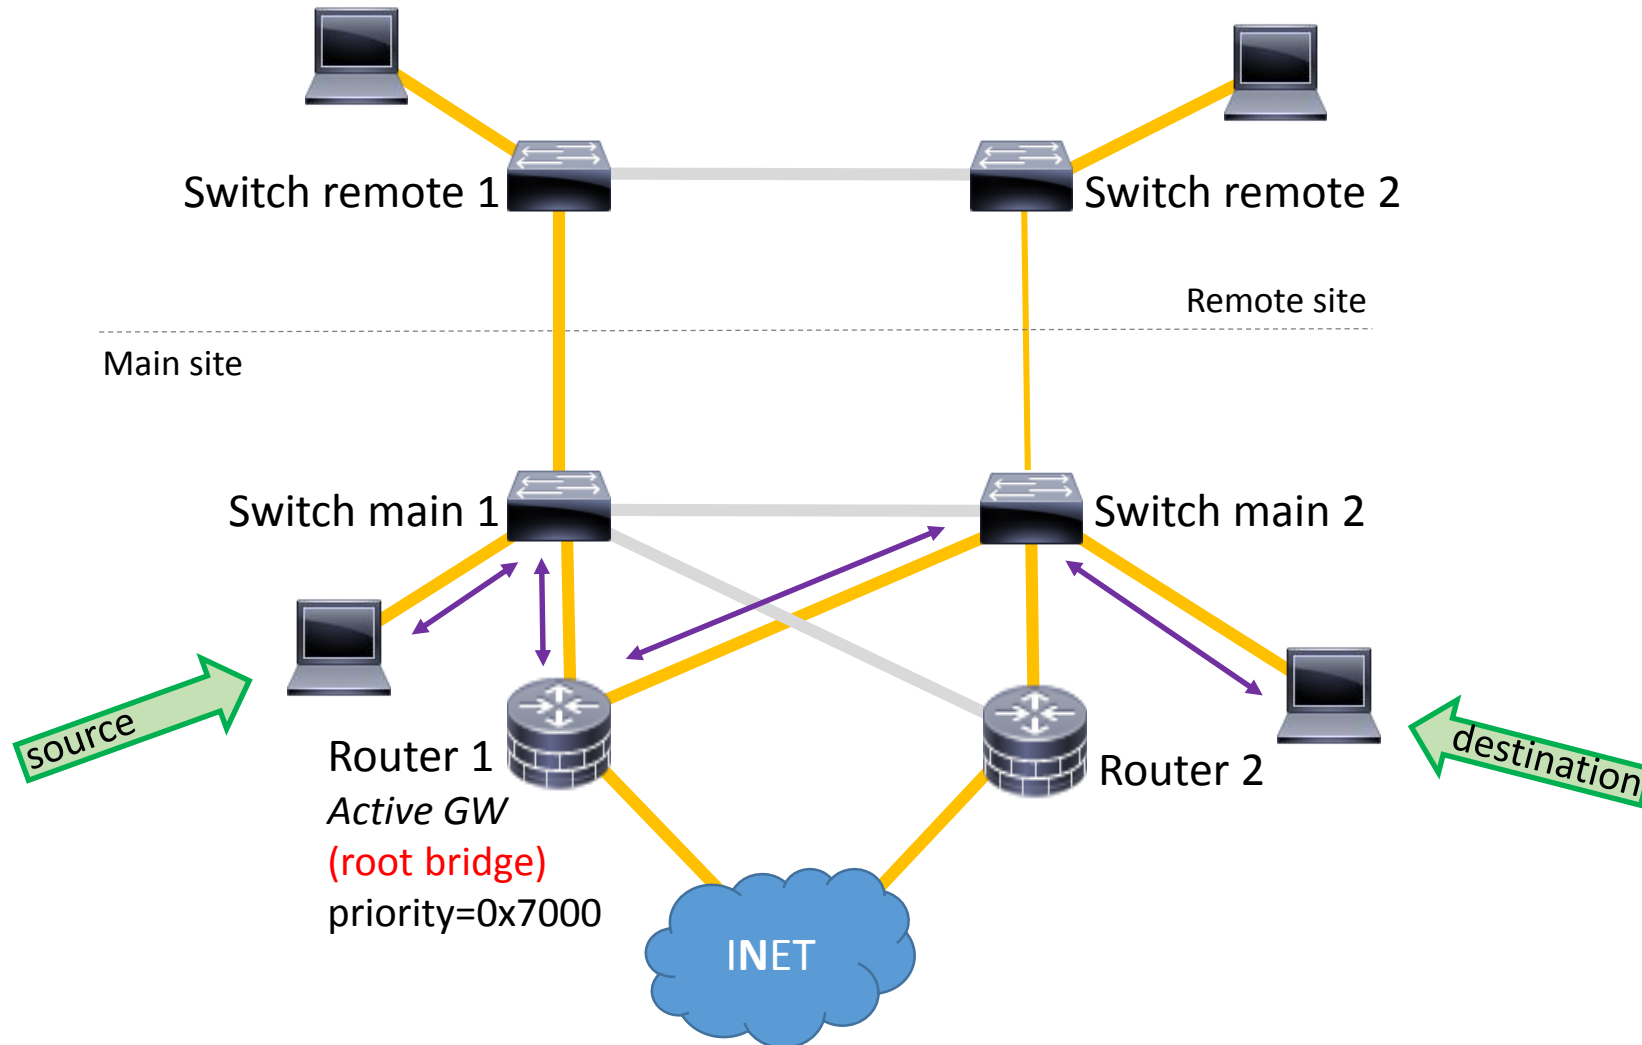


Slow link is used






Traffic through router





Traffic through router *and* other site





Path selection



Path selection

Intended traffic flow

- No traffic on slow link
- No traffic through router (no HW support)

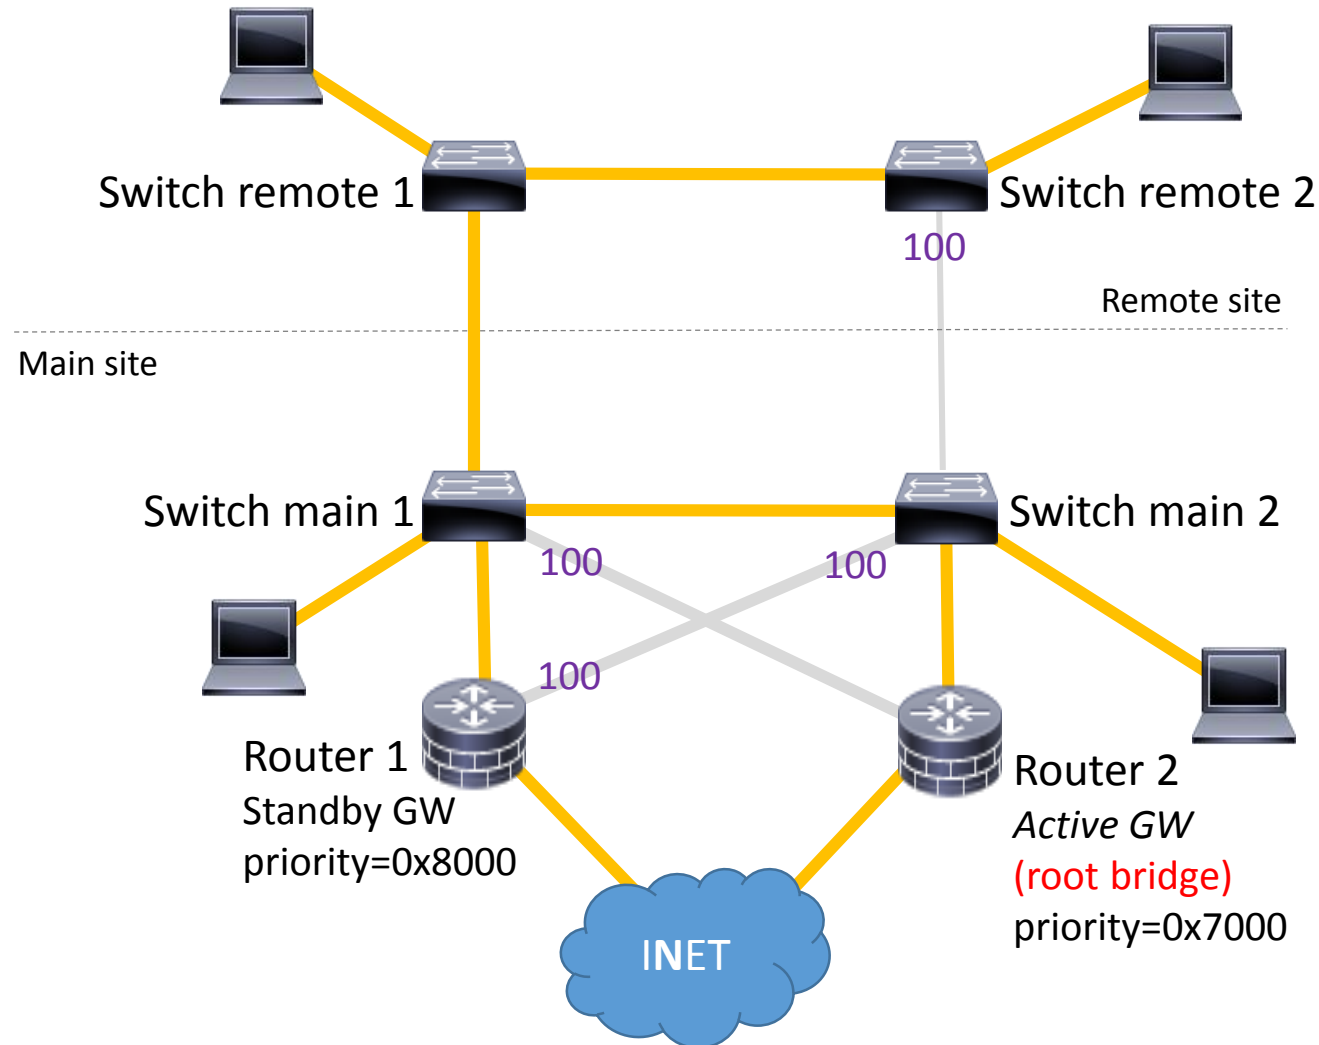
Root bridge will be on active router (VRRP)

- set priority=0x7000 on VRRP master and
set priority=0x8000 on VRRP backup

Bridge path costs will help disableing the correct ports



Traffic design (R2 active)



No traffic through router:

path-cost=100
on slow link
and crosslink

Rest: Default
cost (10)



Traffic design (R1 active)

Is this our final / perfect setup?

Depends on...

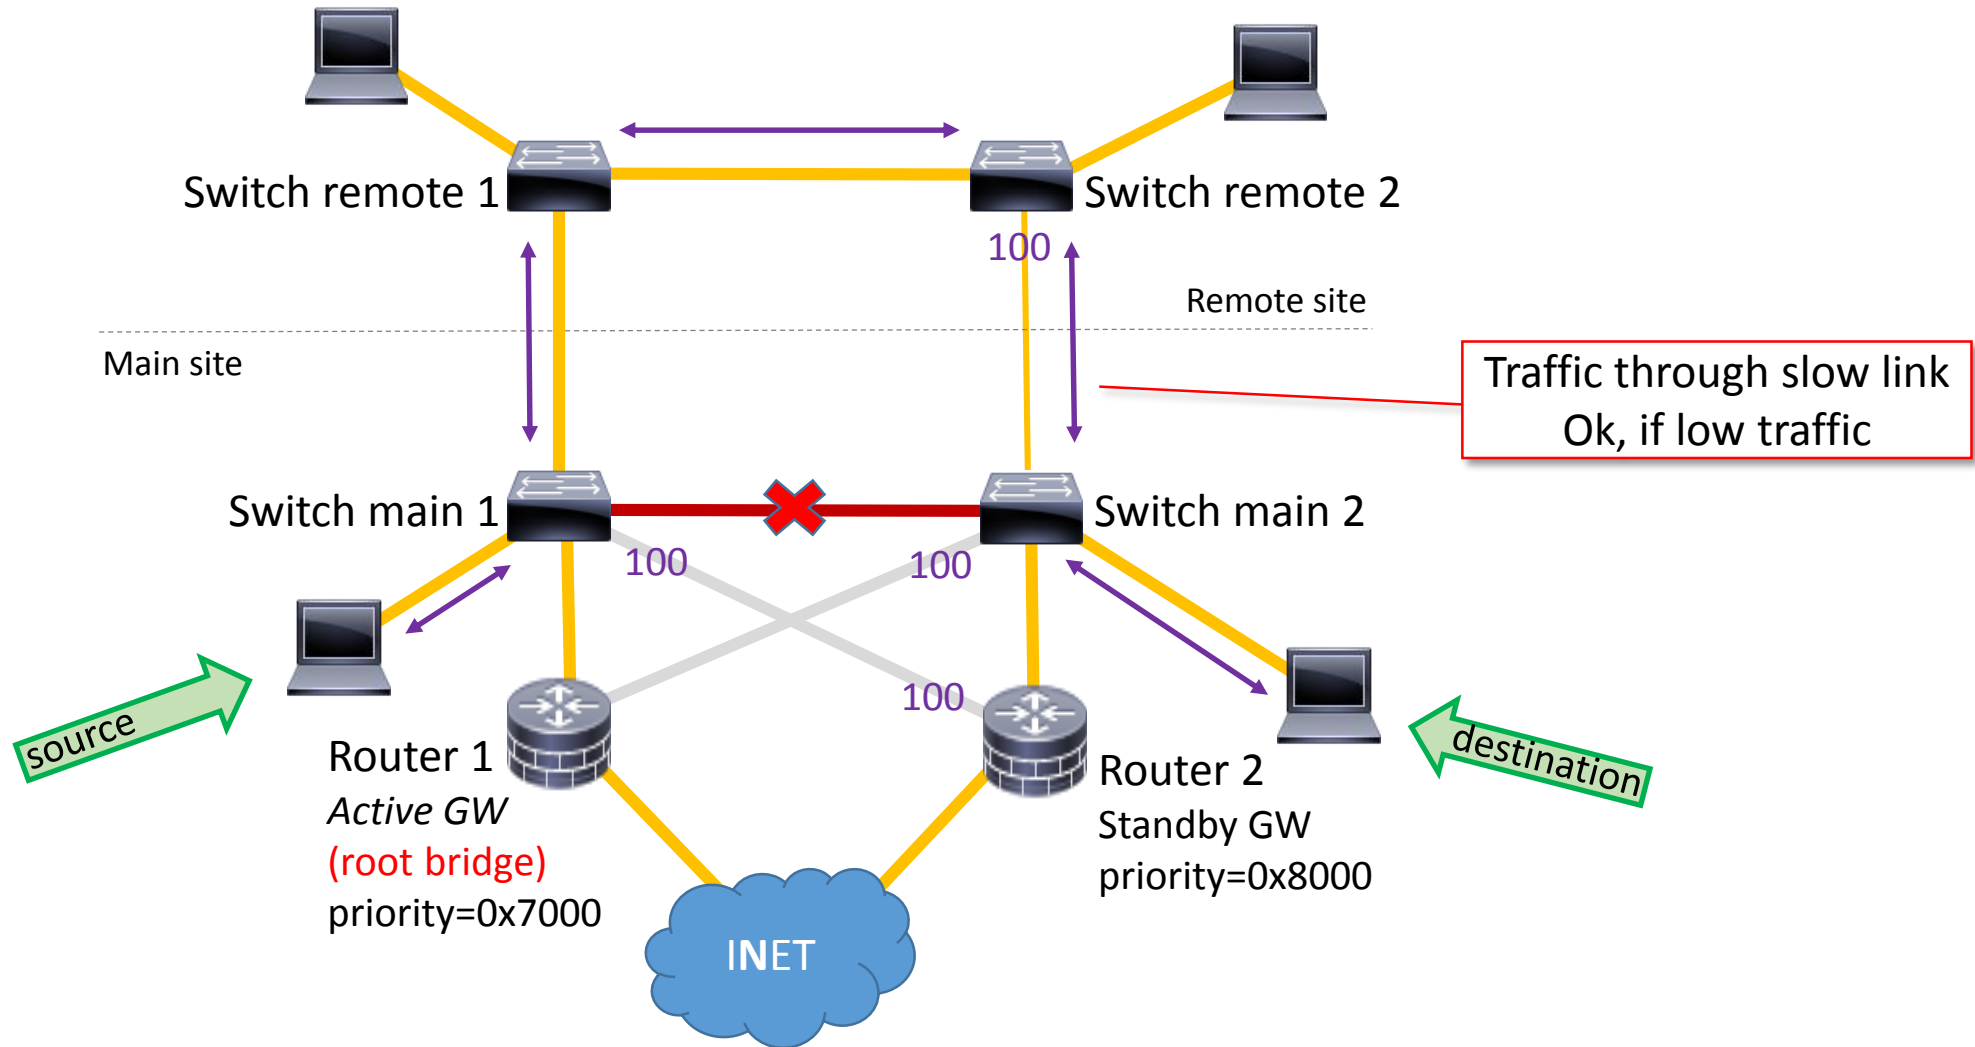
- devices talking to each other
(Device to device or device to internet only?)

Also: Let's check failure on main site.

(Failure of site to site link will just use slow link.)

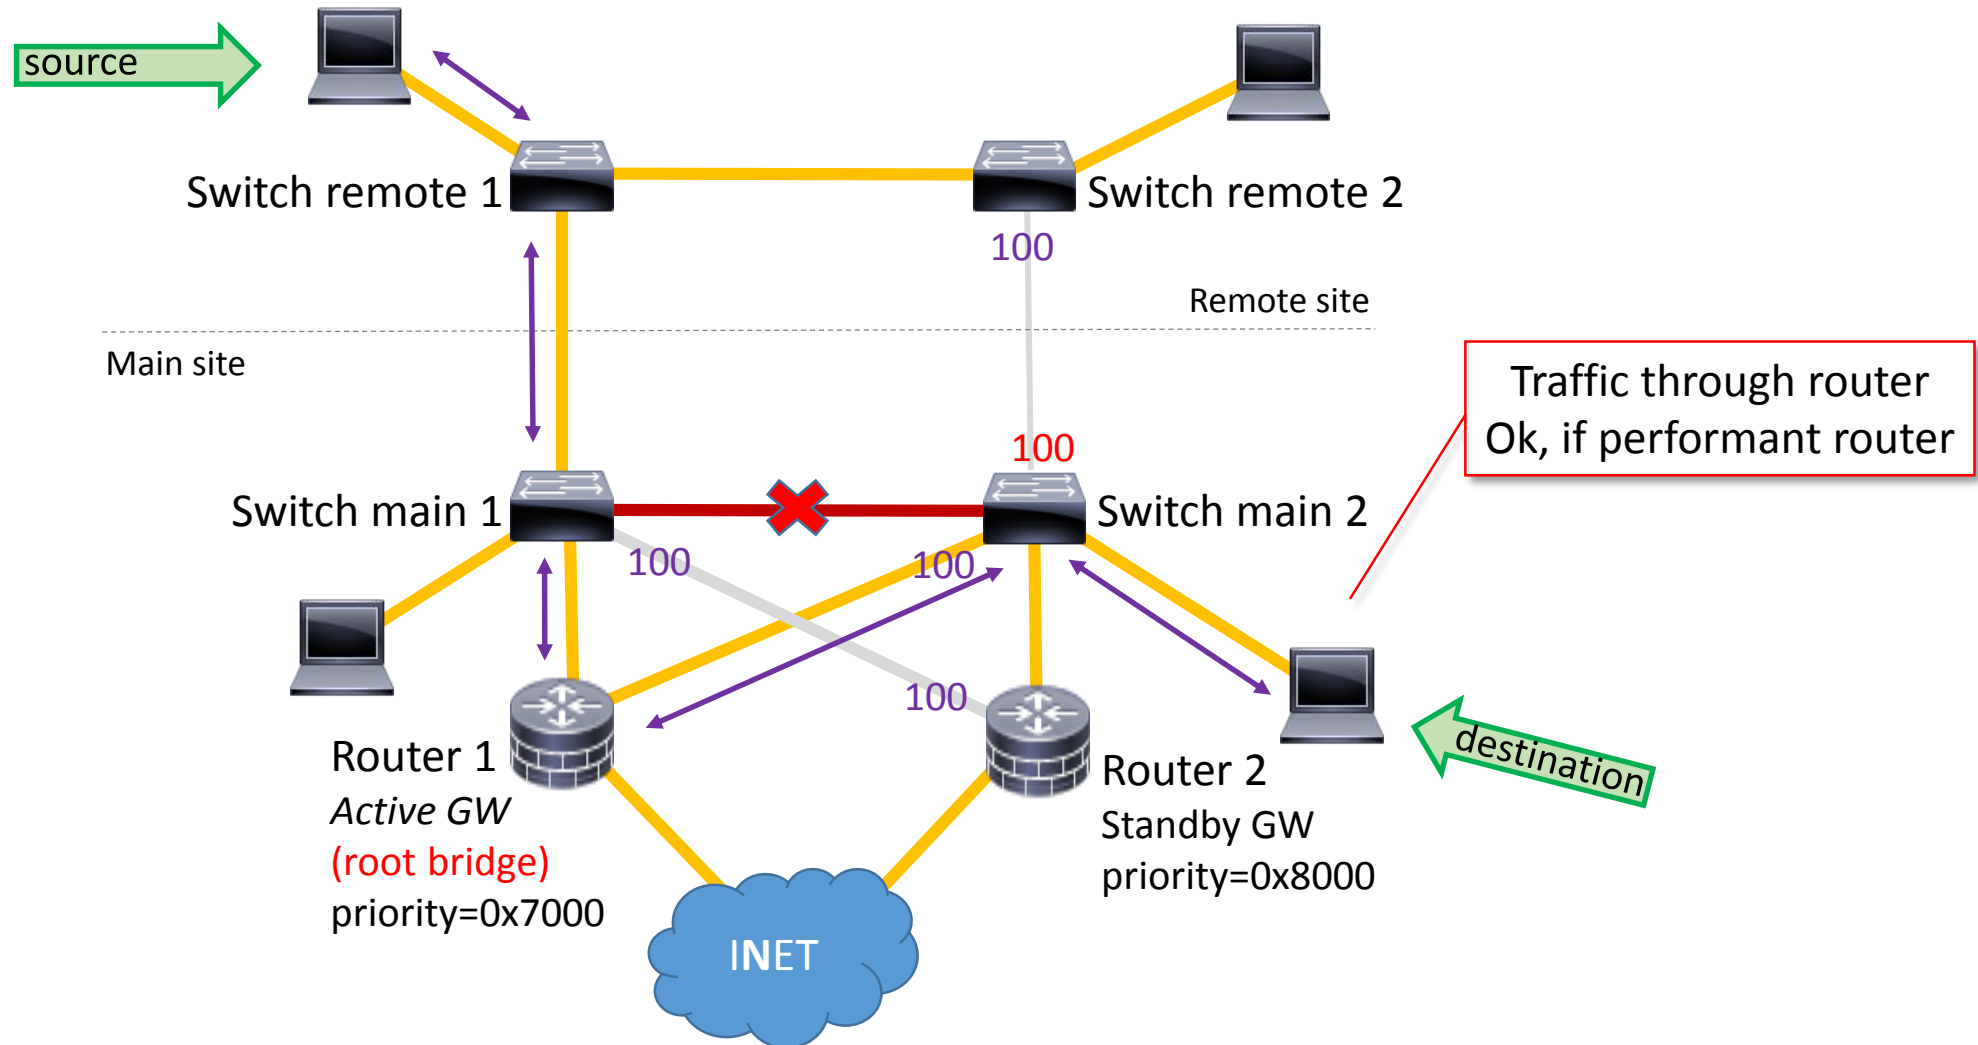


Failure (R1 active)





Failure (R1 active)



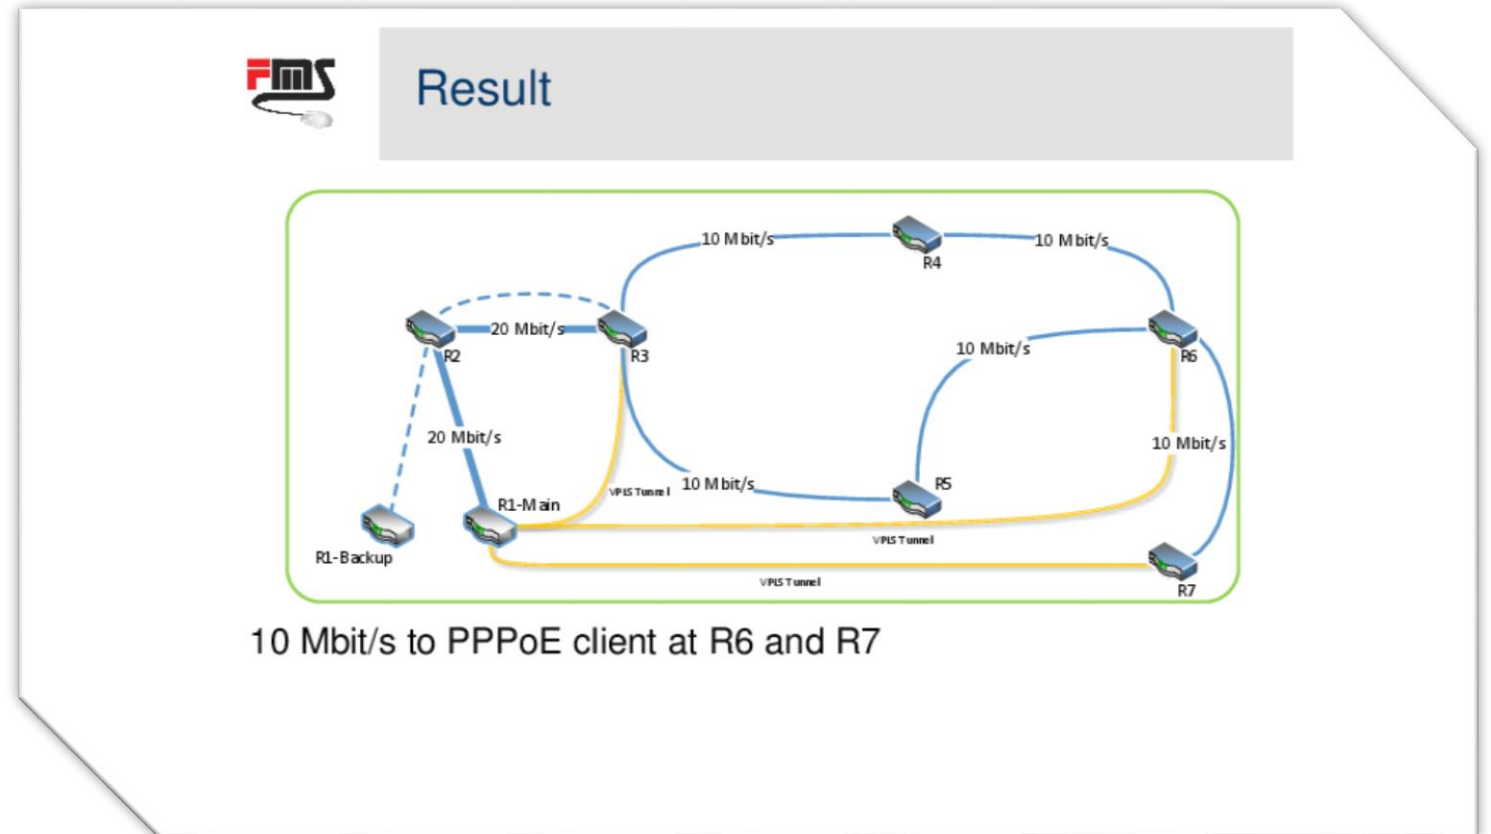


Improvements



Use backup link


- Wishlist: Use the backup link



Sebastian Inacker,
MUM EU 2018 Berlin



Client and management traffic



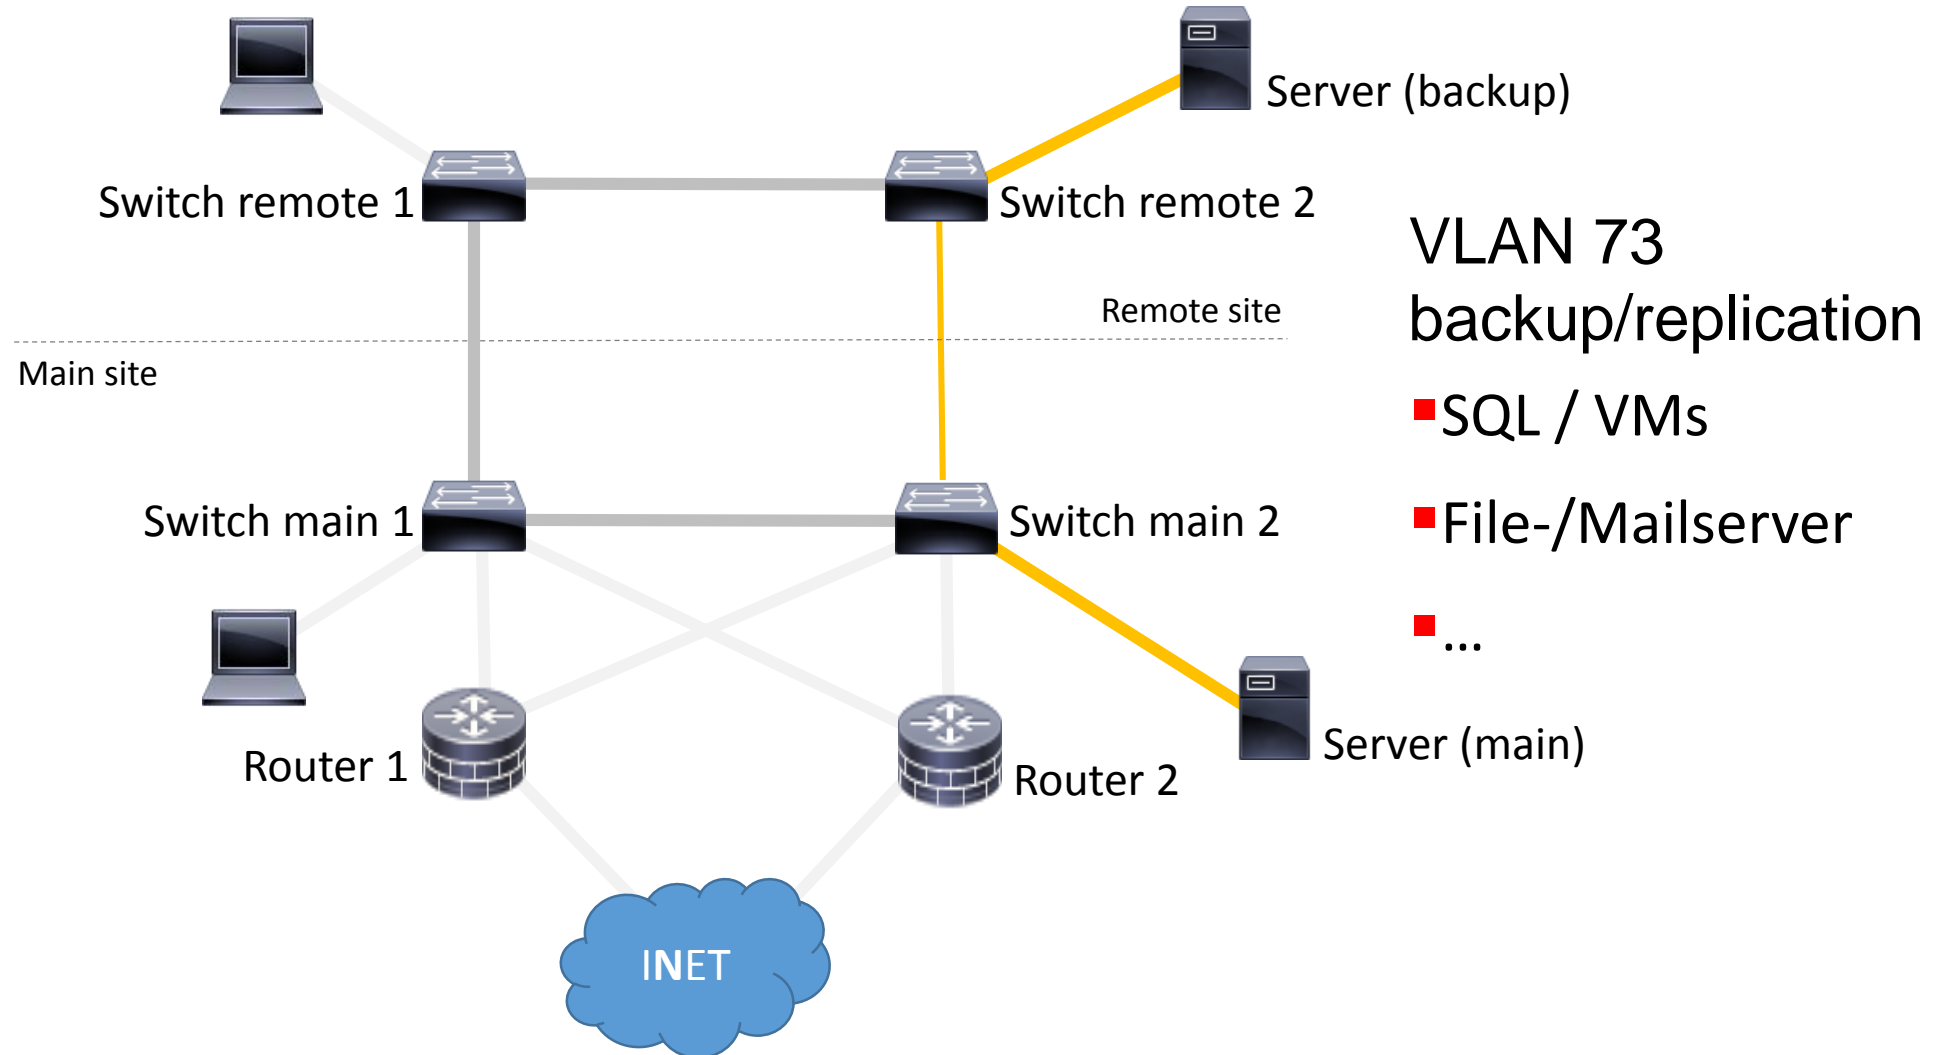
VLAN 100
client traffic

and

VLAN 42
management



Backup/replication traffic





Fact check

We have

- One bridge = One RSTP network
 - One root bridge
 - Same path costs for all VLANs

Would it be possible to have a root bridge and path costs for each VLAN (group of VLANs)?



Use backup link

Do not use multiple bridges, because

- No HW offloading, Possible Loops in network, ...
- See MUM presentations or MikroTik wiki

→ Manual:Layer2 misconfiguration

→ Bridged VLAN on physical interfaces

#	Interface	Bridge	Priority (hex)
0	ether4_vlan100	bridge_one	
1	ether5_vlan100	bridge_one	
2	ether4_vlan42	bridge_one	
3	ether5_vlan42	bridge_one	
4	ether4_vlan73	bridge_two	
5	ether5_vlan73	bridge_two	

6 items

Name	Type	L2 MTU	Tx	Rx	Tx Packet (p/s)
bridge_one	Bridge	1588	0 bps	0 bps	
bridge_two	Bridge	1588	0 bps	0 bps	



MSTP



MSTP

Multiple Spanning Tree Protocol (MSTP)

Available since RouterOS 6.41 (2017-12-22)

Multiple? Can we have...

- multiple STPs?
- multiple root bridges?
- Different paths on Layer 2?

Hardware offloading: Only on CRS3xx



Migration

Please note:

This is no guideline, how to migrate in a production network!

MSTP is compatible to RSTP (see MSTI0, later), but there can be loss of connectivity.



Enable MSTP

Use MSTP on routers and switches

The image displays two screenshots of the 'Interface <bridge>' configuration dialog box, illustrating the process of enabling MSTP. The top screenshot shows the 'STP' tab selected, with the 'Protocol Mode' set to 'RSTP' (highlighted with a red box). The bottom screenshot shows the 'Protocol Mode' set to 'MSTP' (also highlighted with a red box). A red arrow points from the 'RSTP' selection in the top screenshot to the 'MSTP' selection in the bottom screenshot, indicating the transition from RSTP to MSTP.

Interface <bridge>

General STP VLAN Status Traffic

Protocol Mode: none STP RSTP MSTP

Priority: 8000 hex

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

Interface <bridge>

General STP VLAN Status Traffic

Protocol Mode: none STP RSTP MSTP

Priority: 8000 hex

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



MSTI

MSTI – MST Instance

- One root bridge per MST Instance
- Different root bridges → different disabled ports
- Assign VLANs to Instances → different paths per VLAN (group)
- MST Instance Identifier: Number (1 – 31)

Sollution to not beeing able to use multiple bridges.



MSTI

MST Instances (on router 1, root bridge in MSTI1)

The screenshot displays the configuration interface for MST Instances on a bridge. The main window shows a table of MST instances, with the first instance (Identifier 1) highlighted. Two detailed configuration windows are open for this instance.

Bridge	Identifier	Priority (hex)	VLAN Mapping	Regional Root Bridge ID	Root Port
bridge	1	7000	100, 42	0x7001.4C:5E:0C:C9:74:A5	none
bridge	2	8000	73	0x7002.B8:69:F4:90:0C:CA	ether4

Bridge MSTI <bridge:1> (General Tab):

- Bridge: bridge
- Identifier: 1
- Priority: 7000 (hex)
- VLAN Mapping: 100, 42

Bridge MSTI <bridge:1> (Status Tab):

- Current MAC Address: 4C:5E:0C:C9:74:A5
- Root Bridge
- Root Bridge ID: 0x7001.4C:5E:0C:C9:74:A5
- Regional Root Bridge ID: 0x7001.4C:5E:0C:C9:74:A5
- Root Path Cost: 0
- Root Port: none
- Port Count: 2
- Designated Port Count: 2



MST Override

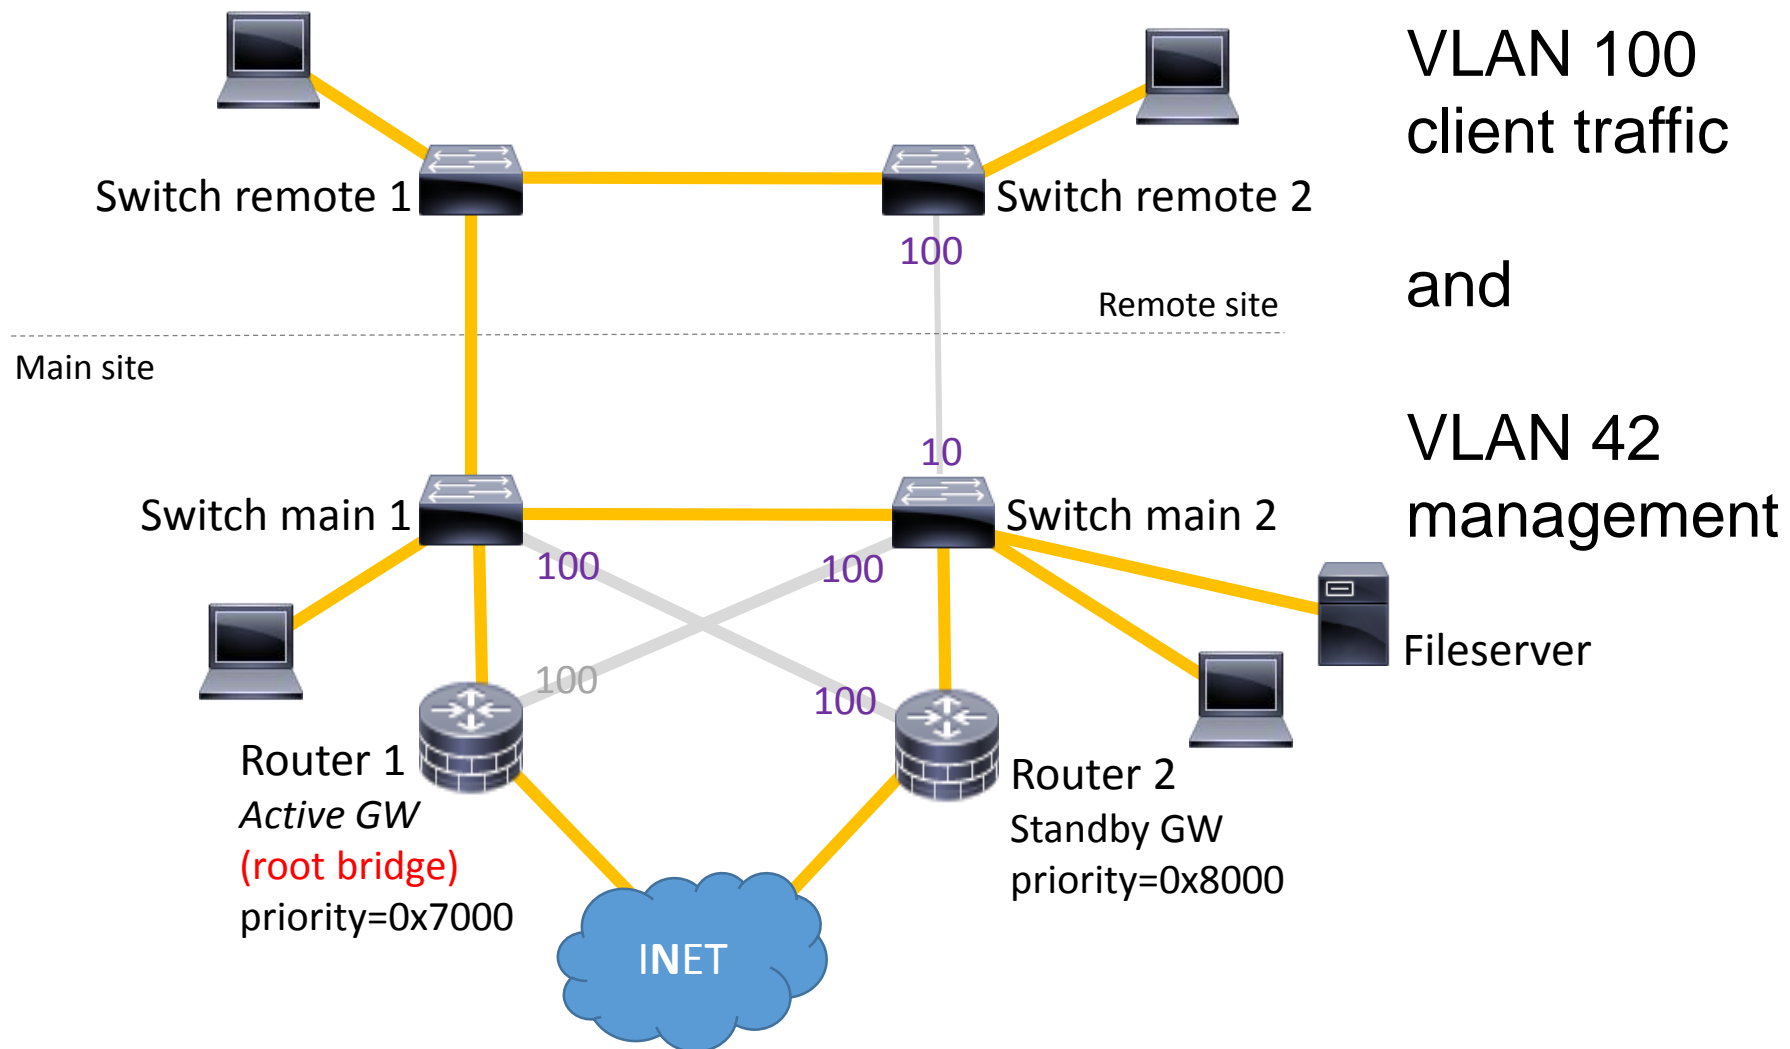
MST Override

- Different path costs, assigned to interfaces & Instances →
Selection of paths to be disabled
- Path costs (MSTI1, MSTI2, ...) will *not* be taken from bridge ports!

For MSTI1 (VLANs for user traffic and management) we will
configure path costs of next slide



MSTI 1 (VLAN 42, 100)





MST Override, MSTI1 (VLAN 100, 42)

Switch main 2, interface facing router 1

Bridge

Bridge Ports VLANs MSTIs Port MST Overrides Filters NAT Hosts MDB

Dynamic is no

Interface	Identifier	Priority (hex)	Internal Path Cost	Role	Forwarding
ether21	1	80	100	alternate port	no

1 item out of 8

Bridge MST Override <ether21>

General Status

Interface: ether21

Identifier: 1

Priority: 80 hex

Internal Path Cost: 100

enabled

Switch remote 2, interface facing slow link

Bridge

Bridge Ports VLANs MSTIs Port MST Overrides Filters NAT Hosts MDB

Dynamic is no

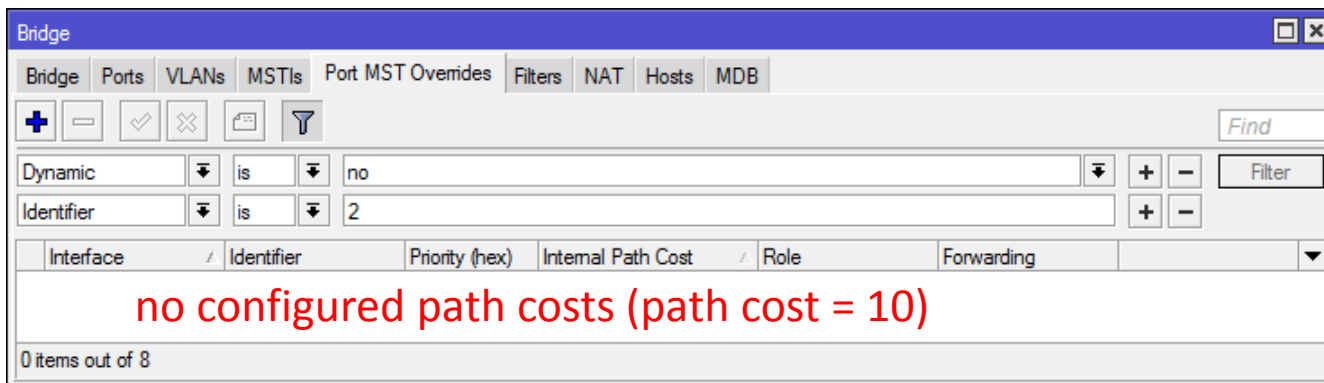
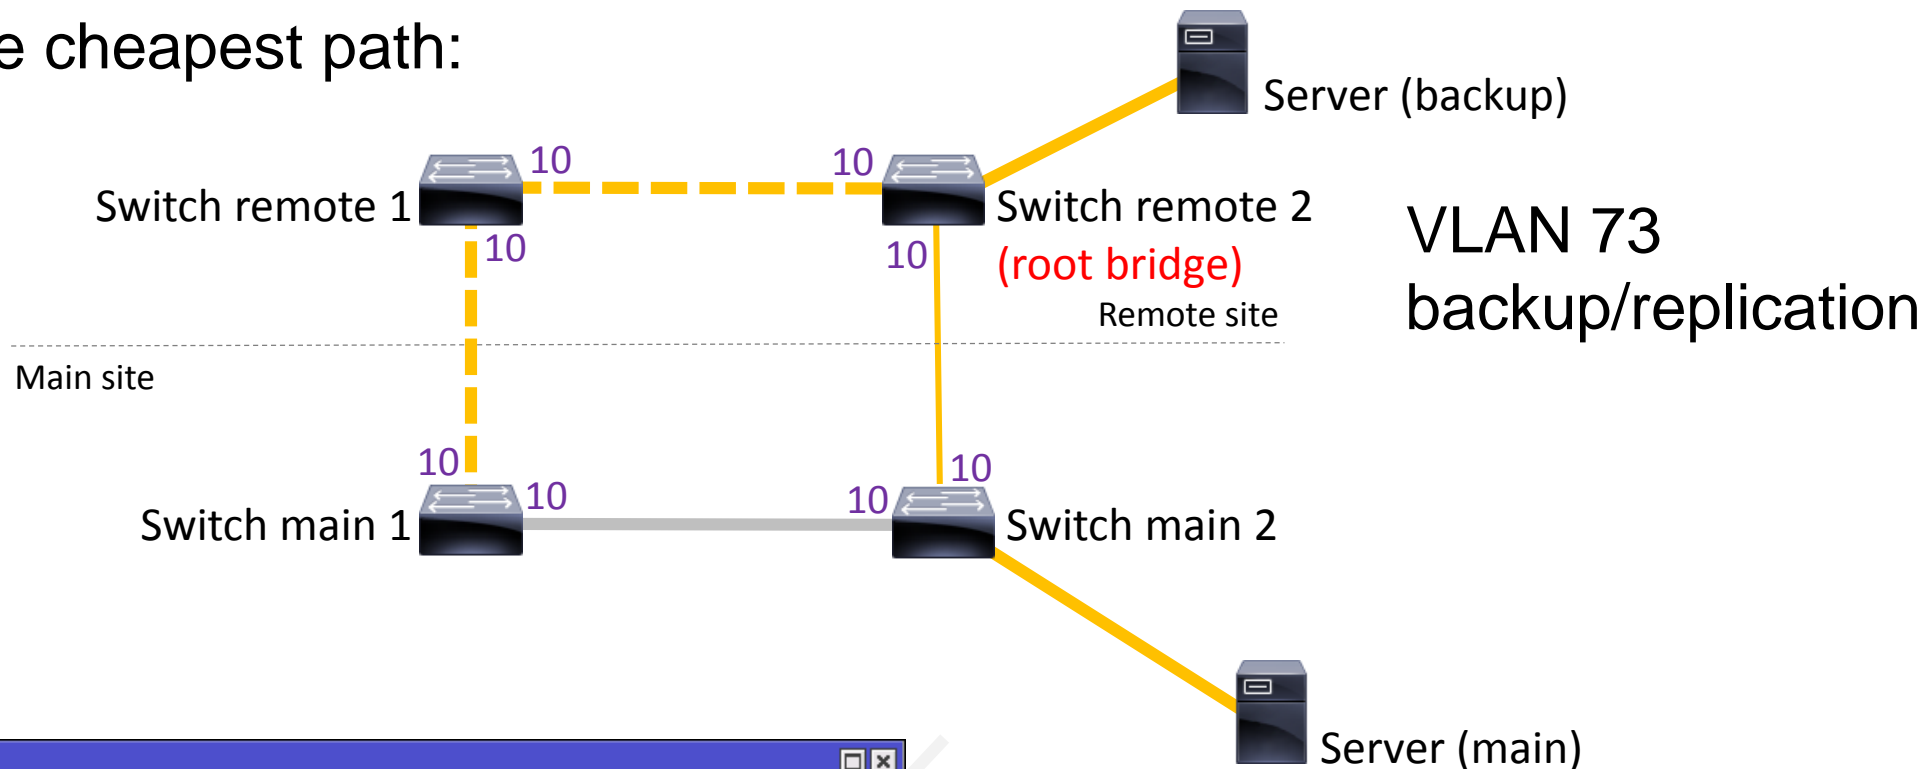
Interface	Identifier	Priority (hex)	Internal Path Cost	Role	Forwarding
ether24	1	80	100	alternate port	no

1 item out of 6



MSTI 2 (VLAN 73)

Use cheapest path:





MSTP Regions



MSTP Regions

MikroTik wiki: First subchapter for MSTP is „Regions“.

We ignored regions so far...!

Our Region Name: Empty

Region Revision: 0 (default)

Interface <bridge>

General STP VLAN Status Traffic

Protocol Mode: none STP RSTP MSTP

Priority: 8000 hex

Region Name:

Region Revision: 0

Max Message Age: 00:00:20

Forward Delay: 00:00:15

Transmit Hold Count: 6

Max Hops: 20

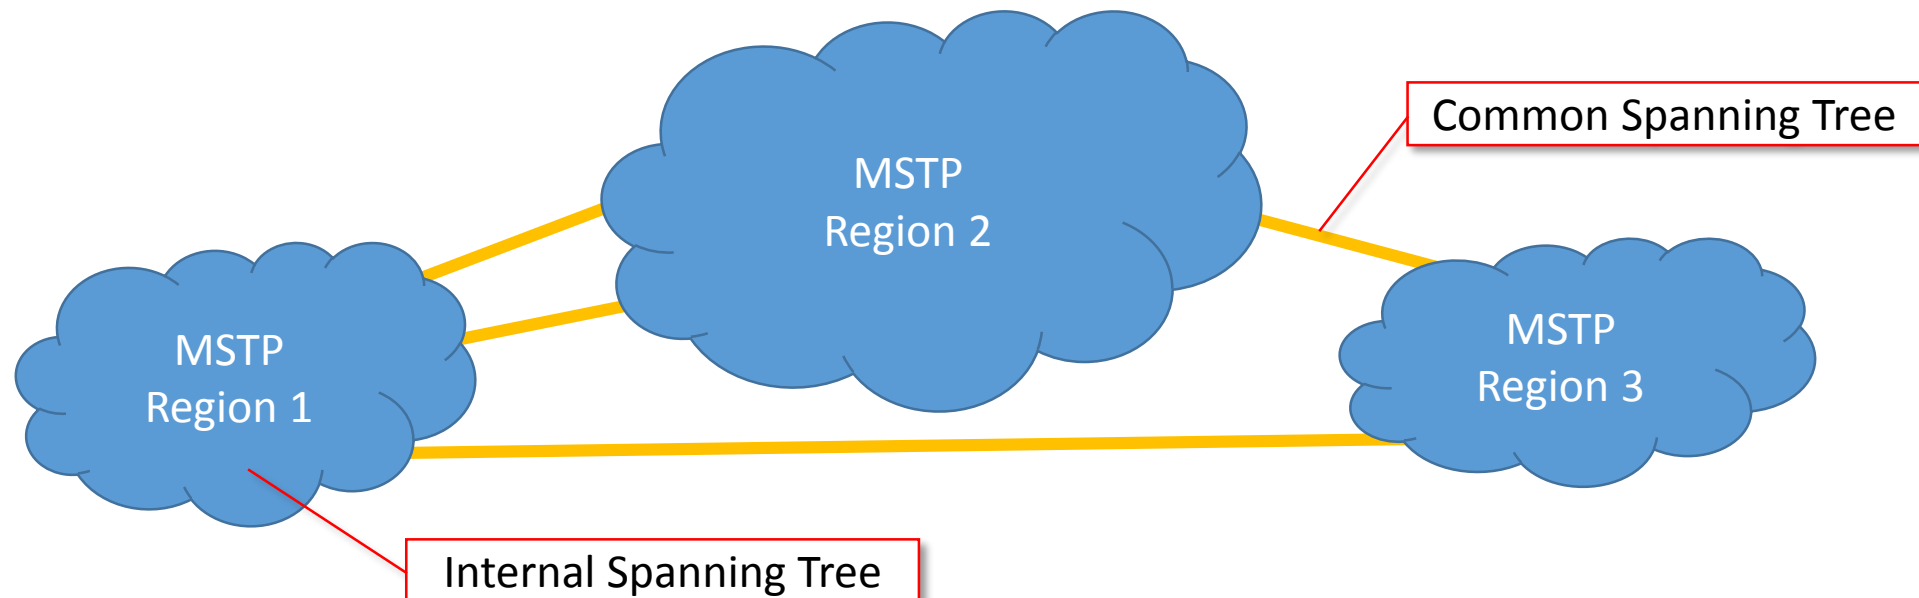
OK
Cancel
Apply
Disable
Comment
Copy
Remove
Torch



MSTP Regions

With regions it's possible to have

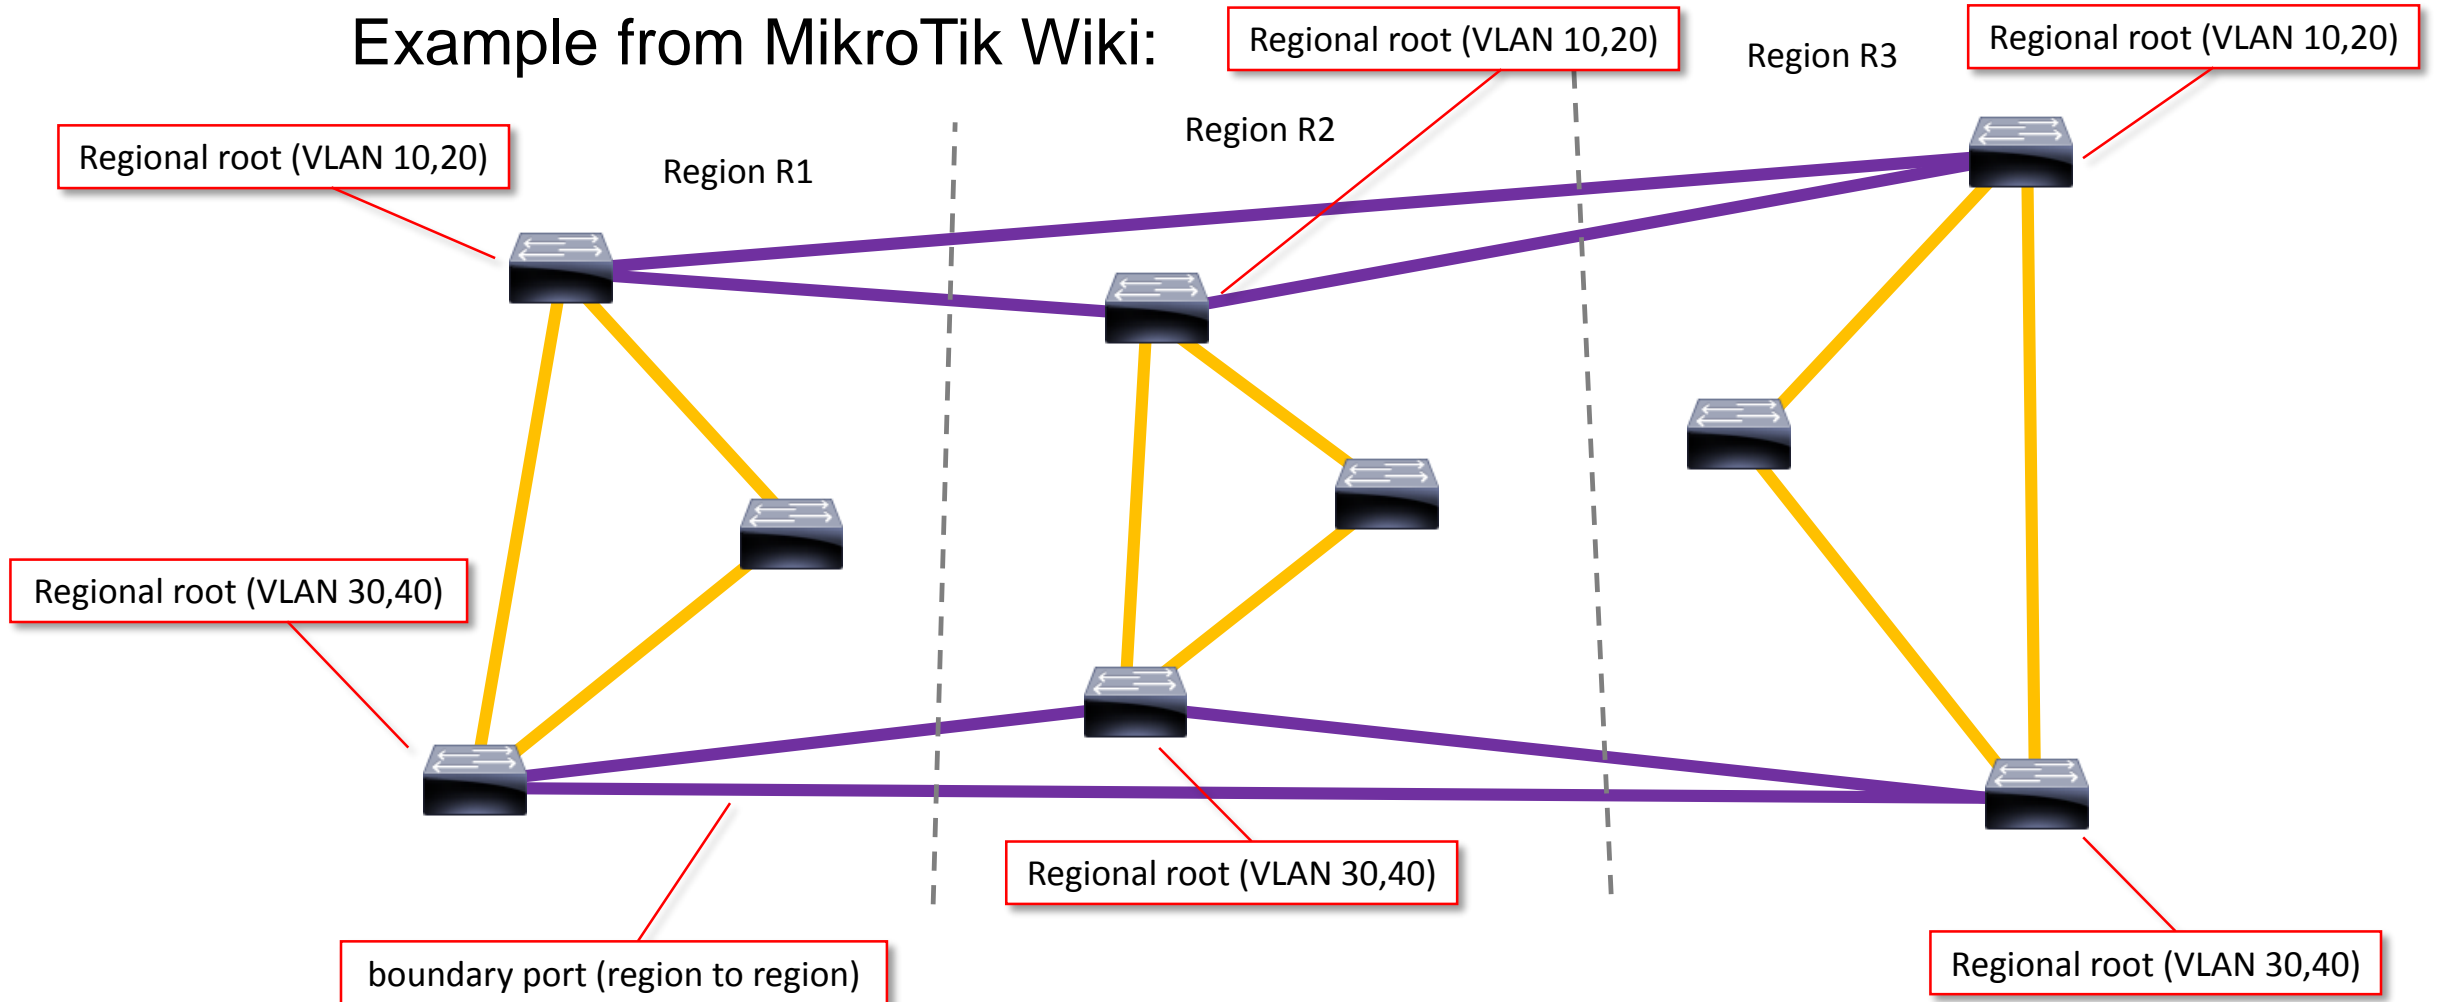
- One *regional* root bridge per *region* and VLAN group
- STP running in each region (Internal Spanning Tree, IST)
- STP between regions (Common Spanning Tree, CST)





MSTP Regions

Example from MikroTik Wiki:





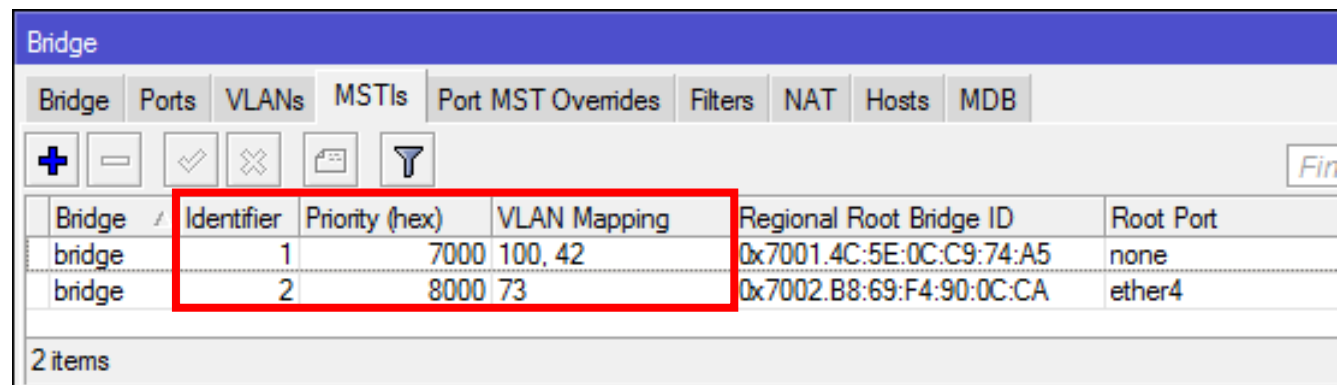
MSTP Regions

One region is fine *for us*.

Use *same values for*

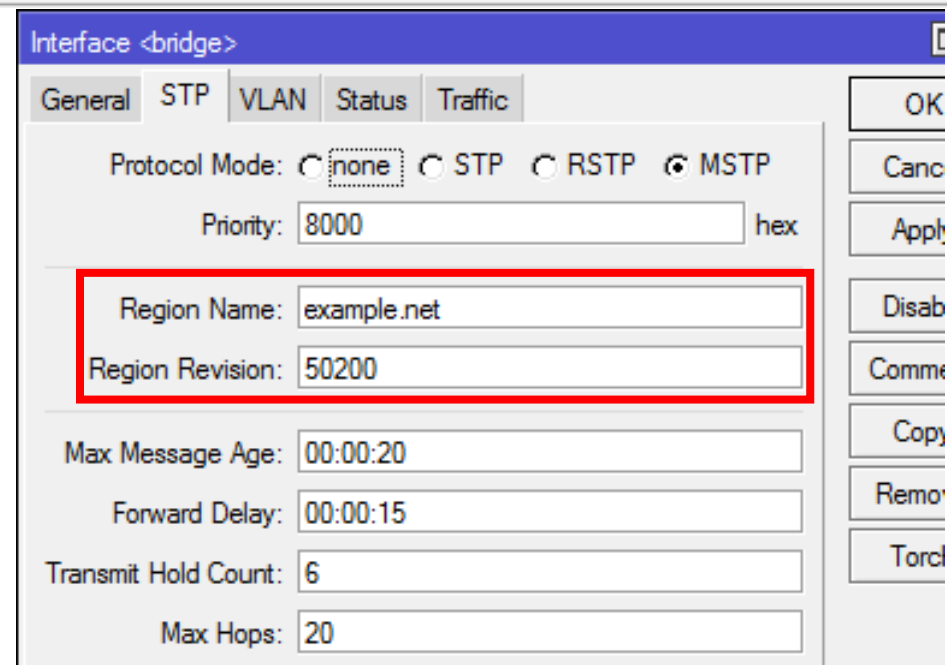
- Region Name
 - Region Revision
 - VLAN mappings to MSTI IDs
- on all MSTP devices (for BPDU packets).

If not: MSTP will assume multiple regions!



The screenshot shows the 'Bridge' configuration window with the 'MSTIs' tab selected. A table lists two MSTI entries. The first entry has Identifier 1, Priority (hex) 7000, and VLAN Mapping 100, 42. The second entry has Identifier 2, Priority (hex) 8000, and VLAN Mapping 73. The 'Regional Root Bridge ID' and 'Root Port' are also visible for each entry.

Bridge	Identifier	Priority (hex)	VLAN Mapping	Regional Root Bridge ID	Root Port
bridge	1	7000	100, 42	0x7001.4C:5E:0C:C9:74:A5	none
bridge	2	8000	73	0x7002.B8:69:F4:90:0C:CA	ether4



The screenshot shows the 'Interface <bridge>' configuration window with the 'STP' tab selected. The 'Protocol Mode' is set to 'MSTP'. The 'Priority' is set to '8000' in hex. The 'Region Name' is 'example.net' and the 'Region Revision' is '50200'. Other settings like 'Max Message Age', 'Forward Delay', 'Transmit Hold Count', and 'Max Hops' are also visible.

Protocol Mode: none STP RSTP MSTP

Priority: 8000 hex

Region Name: example.net

Region Revision: 50200

Max Message Age: 00:00:20

Forward Delay: 00:00:15

Transmit Hold Count: 6

Max Hops: 20



MST Instance 0



MSTI 0

We have internal path costs for our (single) region for each MST instance (VLAN group).

Do we need path costs on *bridge ports* (not at MST Override)?

Depends.

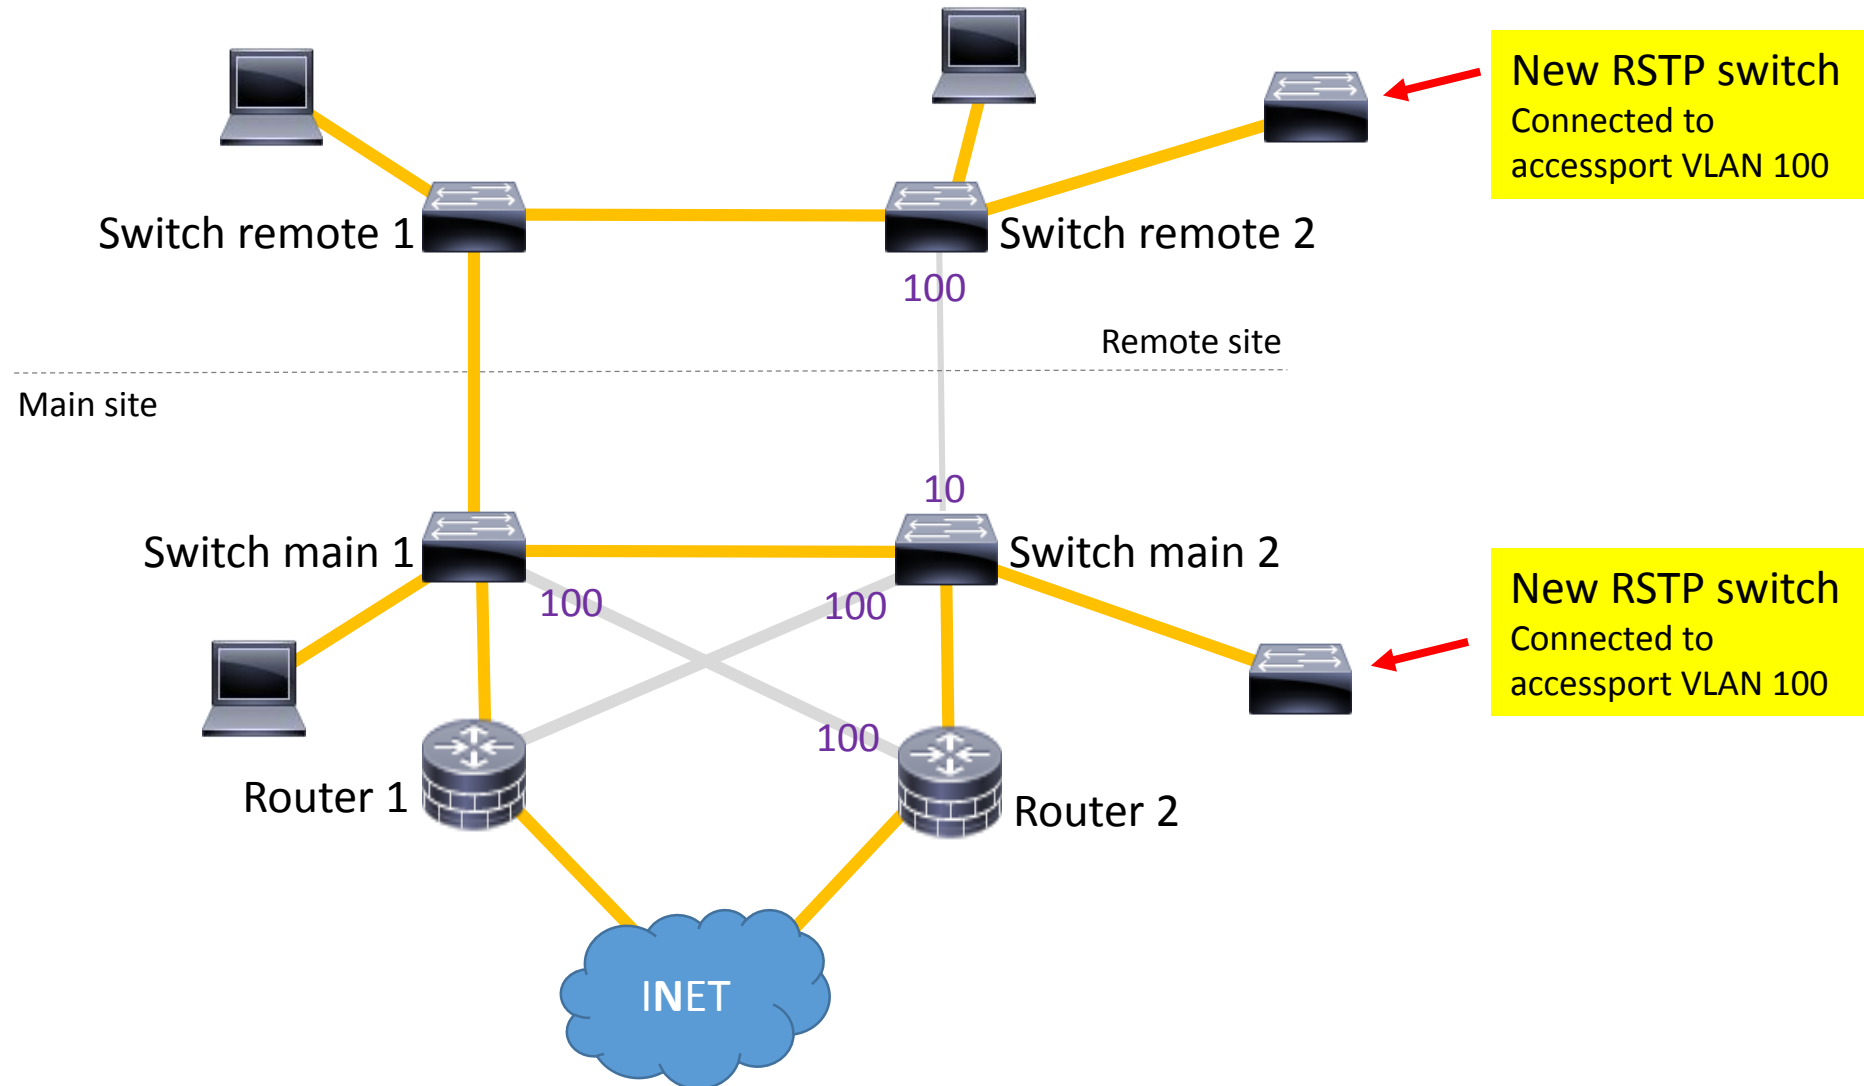
There is a MST Instance Zero for all VLANs not assigned to a MST Instance. *MSTI0 will use bridge port costs.*



MSTP between RSTP segments



MSTP between RSTP segments





MSTP between RSTP segments

What do you expect?

- Who is root bridge *on new RSTP switches*?
 - Router 1 (it's root for MSTI1, VLAN 100)
 - One „random“ switch/router



MSTP between RSTP segments

What do you expect?


- Who is root bridge on new RSTP switches?
 - Router 1 (it's root for MSTI1, VLAN 100)
 - One „random“ switch/router

Expect the unexpected

- It's one „random“ switch (here: switch main 2) – depending on priority/MAC of */interface bridge* config
- Traffic is going the (correct) MSTI1 (VLAN 100) path.
(also the *bridge / port* is *not forwarding*.)



Disabled *bridge* ports (at *MST10*)





Thank you!