

BirchenallHowden

information | communication | technology

Optimising your MikroTik Layer2 configuration

March 2019 © Jono Thompson
BirchenallHowden Ltd



Jono Thompson

- Networking background started as a Cisco Engineer
- Started using ROS June 2010
- MikroTik Consultant Since Dec 2014
- MikroTik Trainer since March 2017



BirchenallHowden Ltd

- Established in 2006
- 29 staff
- Based in Sheffield, UK and working throughout the UK and Europe
- Currently providing IT support for over 75 companies and 2800 users
- Currently have 2 MikroTik consultants



BirchenallHowden Ltd

- Services Provided
 - Wired and wireless network design and installation,
 - Desktop and server installation, support and maintenance
 - ISP Services, leased lines, connectivity
 - Telephony
 - Wireless installs
 - MikroTik Consultancy
 - MikroTik Training



• Visit www.birchenallhowden.co.uk

Presentation Objectives

- Since version 6.41 there have been some major changes to the Bridge configuration
- This presentation is designed to help and encourage you to use the new features.



Presentation Objectives

- This presentation will show some of the most common mistakes made with Layer 2 configurations.
- It will show some incorrect configurations, and then show the correct configurations along with an explanation of what was wrong.
 - This is not a step by step guide
- Most of these are taken from real setups we have had to fix


Presentation Objectives

- At UK MUM in Birmingham 2018, I did a presentation on step by step guide on the new Bridge VLAN filtering
 - “New bridge features in 6.43”



Configurations

- In the download PDF version the incorrect configurations will be marked like this



- Correct configurations will be marked like this:-




Meet Mike




- Last time we met Mike, he had just installed his new MikroTik Wireless
- He has finally got all his wireless working really well thanks to some really cool guy from the UK MUM
- Mike invites Dave around to see how good it is...

Mike and Dave

- They sit in Mike's office Kitchen while Mike shows Dave how fast the WiFi network now is...
- Dave looks up at Mike's Network patch cabinet



Mike's Network



Dave's Visit



- Dave suggests that Mike uses a PoE switch to power all his APs.
- Then he would be able to shut his cabinet door!

New Switch

- Mike does some searches the internet...
- Mike sees that MikroTik do PoE switches
- As they are cheap he buys one
- CRS328-24P-4S-RM



Neat Install



Mike looks at the features

- Mike is so excited about his new switch he Tweets about his new tidy install




Dave's Visit



- Dave replies and says "You should have a guest Wi-Fi to keep your visitors off your office network"

Mike's Network



CRS328

hAP ac^2

AP2 - 13



Guest Network - Router Configuration

- Dave sets about configuring a new bridge on his hAP ac² Router for his guest network

```
/interface bridge add name=bridge-guest
/interface bridge port add bridge=bridge-guest interface=ether4
/ip address add address=192.168.201.1/24 interface=bridge-guest network=192.168.201.0
/ip pool add name=dhcp_pool-guest ranges=192.168.201.2-192.168.201.254
/ip dhcp-server add address-pool=dhcp_pool-guest disabled=no interface=bridge-guest name=dhcp-guest
/ip dhcp-server network add address=192.168.201.0/24 dns-server=8.8.8.8 gateway=192.168.201.1

/interface bridge add name=bridge-lan
/interface bridge port add bridge=bridge-lan interface=ether2
/interface bridge port add bridge=bridge-lan interface=ether3
```

The Problem

- Mike notices that he now has really slow throughput from his wireless clients to his wired clients
- He also notices that his hAP ac² router has a high CPU load.
- He also notices that network traffic has huge latency
- He calls Dave.....



What's wrong?

- Dave tells Mike to look for the check hardware offloading is enabled and the H flag to show its being used.

The image displays three overlapping windows from Mikrotik WinBox:

- Bridge Table:** A table listing bridge ports. The first row is highlighted with a red box around the 'H' flag in the first column.

#	Interface	Bridge	Horizon	Trusted	Priority (h...	Path Cost	Role
H	ether4	bridge-guest		no	80	10	designated port
	ether2	bridge-lan		no	80	10	designated port
	ether3	bridge-lan		no	80	10	designated port

- Terminal:** Shows the command `interface bridge port print` and its output. The first line of output is highlighted with a red box.

```
[admin@MikroTik] > interface bridge port print
Flags: X - disabled, I - inactive, D - dynamic, H - hw-offload
#  INTERFACE  BRIDGE  HW  PVID  PRIORITY  PATH-COST  INTERNAL-PATH-COST  HORIZON
0  H ether4    bridge... yes  1    0x80    10           10           none
1  ether2    bridge... yes  1    0x80    10           10           none
2  ether3    bridge... yes  1    0x80    10           10           none
```

- Bridge Port <ether2>:** Shows the configuration for the ether2 interface. The 'Hardware Offload' checkbox is checked and highlighted with a red box. At the bottom, the 'Hw. Offload' status is also highlighted with a red box.

Multiple Bridges on a Single Switch Chip

Analysis:

- Only CRS1xx/2xx Support HW-offloading on more than 1 bridge
- Even though bridge port has HW-offloading enabled, the traffic for the 2nd bridge is going through the CPU. This has created a hardware limit for Mike



Multiple Bridges on a Single Switch Chip

Solution:

- You can control which bridge uses HW-offloading.

Preferred solution:

- Consider a network redesign to use hardware optimally



Bridge VLAN filtering

- Mike does some reading about VLANs and sees that since version 6.41 MikroTik supports VLANs with a single bridge
- Mike configures bridge VLAN filtering on his hAP ac² so he only needs to use one bridge and turns it on.....



Bridge VLAN filtering

```
/interface vlan
add interface=ether4 name=vlan11 vlan-id=11
add interface=ether4 name=vlan201 vlan-id=201
/ip dhcp-server
add address-pool=dhcp_pool-lan disabled=no interface=vlan11 name=dhcp-lan
add address-pool=dhcp_pool-guest disabled=no interface=vlan201 name=dhcp-guest

/interface bridge port
add bridge=bridge-lan interface=ether4
add bridge=bridge-lan interface=ether2 pvid=11
add bridge=bridge-lan interface=ether3 pvid=11
/interface bridge vlan
add bridge=bridge-lan untagged=ether2,ether3 tagged=ether4 vlan-ids=11
add bridge=bridge-lan tagged=ether4 vlan-ids=201
```


VLAN Interfaces

- Mike now notices that his clients don't always get a DHCP address, even though he has a DHCP server running on the VLAN....
- Mike emails MikroTik to point out that MikroTik don't support DHCP running in VLANs.



VLAN Interfaces

Analysis:

- Interfaces in a bridge are slave interfaces.
 - The bridge is the master.
- All traffic captured on the bridge is forwarded to the CPU using the bridge interface, not the physical interface



VLAN Interfaces

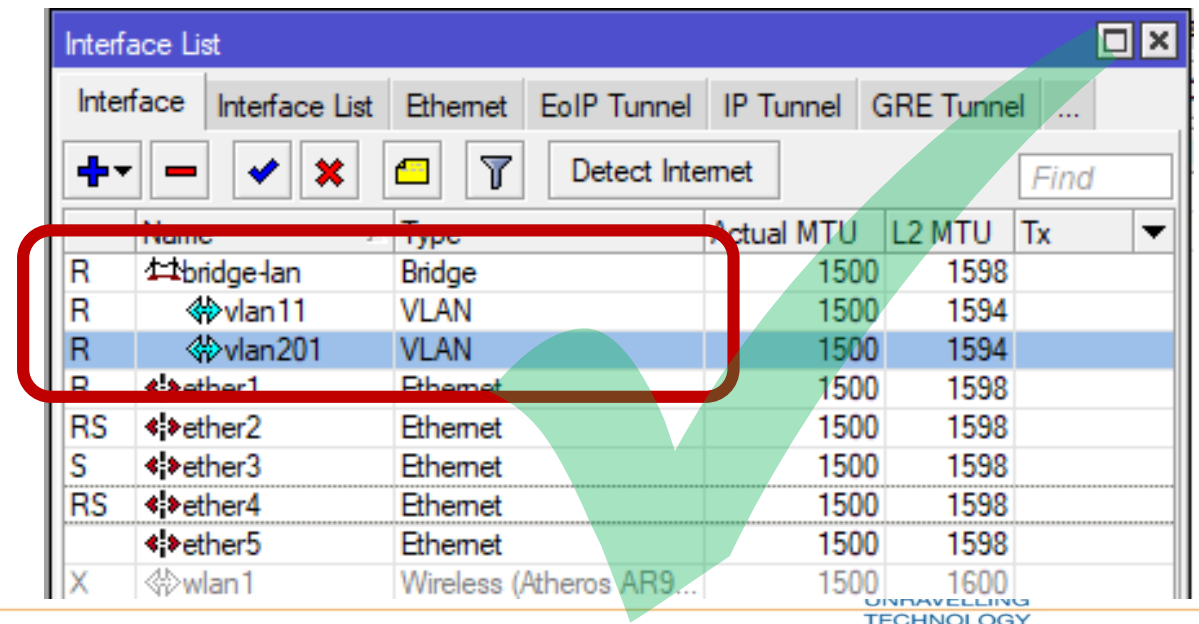
	Name	Type	Actual MTU	L2 MTU	Tx
R	bridge-lan	Bridge	1500	1598	
R	ether1	Ethernet	1500	1598	
RS	ether2	Ethernet	1500	1598	
S	ether3	Ethernet	1500	1598	
RS	ether4	Ethernet	1500	1598	
R	vlan11	VLAN	1500	1594	
R	vlan201	VLAN	1500	1594	
	ether5	Ethernet	1500	1598	
X	wlan1	Wireless (Atheros AR9...	1500	1600	
X	wlan2	Wireless (Atheros AR9...	1500	1600	



VLAN Interfaces

Solution:

- Change the VLAN interface to the bridge as the bridge is the new switch CPU and this is where all traffic will flow out of the bridge
- DHCP is now working for Mike

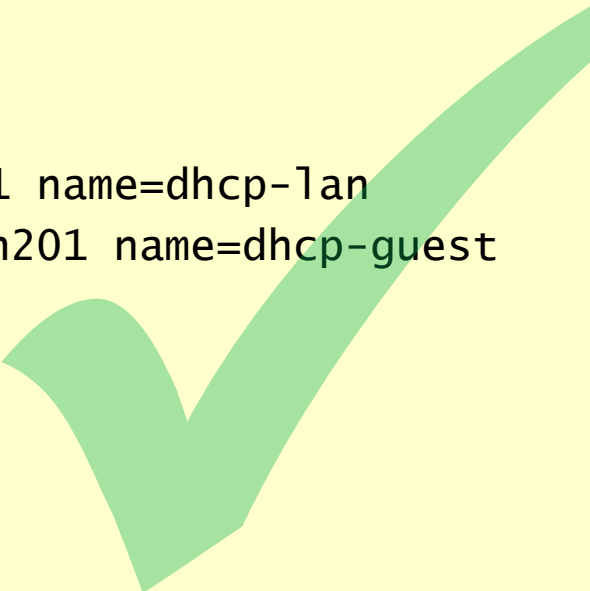


	Name	Type	Actual MTU	L2 MTU	Tx
R	bridge-lan	Bridge	1500	1598	
R	vlan11	VLAN	1500	1594	
R	vlan201	VLAN	1500	1594	
R	ether1	Ethernet	1500	1598	
RS	ether2	Ethernet	1500	1598	
S	ether3	Ethernet	1500	1598	
RS	ether4	Ethernet	1500	1598	
	ether5	Ethernet	1500	1598	
X	wlan1	Wireless (Atheros AR9...	1500	1600	

Bridge VLAN filtering – Correct Configuration

```
/interface vlan
add interface=bridge-lan name=vlan11 vlan-id=11
add interface=bridge-lan name=vlan201 vlan-id=201
/ip dhcp-server
add address-pool=dhcp_pool-lan disabled=no interface=vlan11 name=dhcp-lan
add address-pool=dhcp_pool-guest disabled=no interface=vlan201 name=dhcp-guest

/interface bridge port
add bridge=bridge-lan interface=ether4
add bridge=bridge-lan interface=ether2 pvid=11
add bridge=bridge-lan interface=ether3 pvid=11
/interface bridge vlan
add bridge=bridge-lan untagged=ether2,ether3 tagged=bridge-lan,ether4 vlan-ids=11
add bridge=bridge-lan tagged=bridge-lan,ether4 vlan-ids=201
```



Bridge VLAN filtering

- Mike tests his network and now see that he has slow throughput through both networks
- Mike also notices that when he runs a performance test he notices that his hAP ac² has a very high CPU load.



Bridge VLAN filtering

- Mike looks at his config on his hAP ac² and sees that even though hw-offload is enabled that there is still no H flag

The screenshot shows the Mikrotik WinBox interface for Bridge configuration. The 'Bridge' tab is active, displaying a table of bridge ports. The 'Hardware Offload' column is highlighted with a red box, showing 'yes' for all three interfaces. Below the WinBox window, a terminal window shows the command `interface bridge port print` and its output. The 'HW' column in the terminal output is also highlighted with a red box, showing 'yes' for all three interfaces. A large orange scribble is present in the bottom left corner of the image.

#	Interface	Bridge	Horizon	Trusted	Hardware Offload	Priority (h...	Path Cost	Role
0	ether4	bridge-lan		no	yes	80	10	designated port
1	ether2	bridge-lan		no	yes	80	10	designated port
2	ether3	bridge-lan		no	yes	80	10	designated port

```
[admin@MikroTik] > interface bridge port print
Flags: Y - disabled, I - inactive, D - dynamic, H - hw-offload
#  INTERFACE  BRIDGE  HW  PVID  PRIORITY  PATH-COST  INTERNAL-PATH-COST  HORIZON
0  ether4     bridge  .. yes   1      0x80      10                  10      none
1  ether2     bridge  .. yes   1      0x80      10                  10      none
2  ether3     bridge  .. yes   1      0x80      10                  10      none
```

Bridge VLAN filtering

- Mike can't understand why when he thinks he done everything right
 - He only has 1 bridge
 - He has hw-offloading enabled
- And yet the bridge does not show HW offloading
- Mike calls Dave



Bridge – HW offloading

- Dave tells Mike to look at the manual.....
- Depending on the model or the switch chip, only some features are supported with bridge HW offloading
- Use of unsupported features will disable HW-offloading



Bridge – HW offloading

Switch Chip	Model (example units)	STP/RSTP	MSTP	DHCP Snooping	VLAN Filtering	Bonding ³
	CRS3xx	✓	✓	✓	✓	✓
	CRS1xx/2xx	✓	✗	✓ ¹	✗	✗
QCA8337	hAP ac / hEX PoE / 3011 (1Gb)	✓	✗	✓ ²	✗	✗
AR8327	hAP ac ² /2011(1Gb)/1100AHx2	✓	✗	✓ ²	✗	✗
AR8227	hAP/hEX lite/2011 (100Mb)	✓	✗	✓ ²	✗	✗
AR8316		✓	✗	✓ ²	✗	✗
AR7240		✓	✗	✓ ²	✗	✗
MT7621	hEX (750Gr3)	✓	✗	✓	✗	✗
RTL8367	1100AHx4	✗	✗	✓	✗	✗
ICPlus175D		✗	✗	✓	✗	✗

1. Feature will not work properly in VLAN switching setups, you must make sure that required packets are sent out with the correct VLAN tag using ACL rules.
2. DHCP Snooping will not work properly with VLAN switching
3. Bridge hardware offloading only supported using 802.3ad bonding

Complete list https://wiki.mikrotik.com/wiki/Manual:Switch_Chip_Features#Bridge_Hardware_Offloading

Bridge – HW offloading

- Only a few devices are able to offload the traffic to the switch chip when using VLAN filtering
- When traffic is not offloaded to the switch chip the CPU is used to forward traffic. This can result in lower than expected throughput
- Dave suggest that he carries out a network redesign to use his hardware to its full potential


Bridge VLAN filtering

- Mike carries out a network redesign and no longer uses his hAP ac² as a switch in his network.
- Mike sees that his CRS328-24P-4S+ also supports VLAN filtering
- Mike now starts configuring his CRS328-24P-4S+ with VLANs so that his wireless access points can have both VLANs running on them



Mike's Network

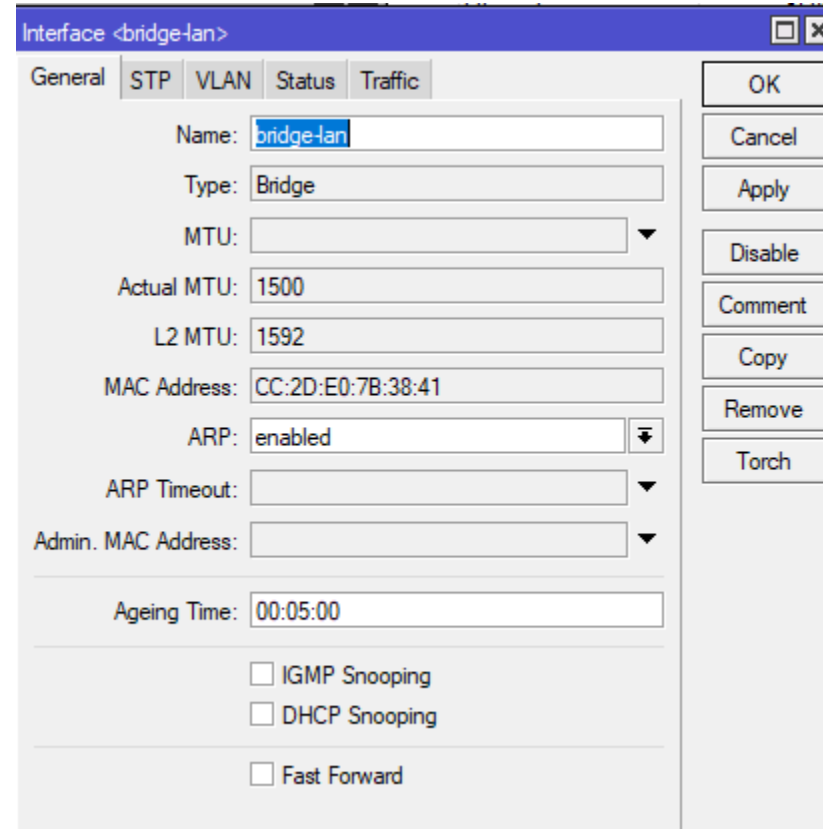
- New network topology



CRS326 VLAN Configuration

- Mike now sets about setting up his new CRS328-24P-4S+.

- He creates a bridge



The screenshot shows the Mikrotik WinBox configuration window for a new bridge interface. The window title is "Interface <bridge-lan>". The "General" tab is selected, and the "Name" field is set to "bridge-lan". The "Type" is set to "Bridge". The "MTU" is set to 1500, and the "Actual MTU" is 1500. The "L2 MTU" is set to 1592. The "MAC Address" is set to CC:2D:E0:7B:38:41. The "ARP" is set to "enabled". The "ARP Timeout" is set to 00:05:00. The "Admin. MAC Address" is set to 00:00:00:00:00:00. The "Ageing Time" is set to 00:05:00. There are three checkboxes at the bottom: "IGMP Snooping", "DHCP Snooping", and "Fast Forward", all of which are unchecked. On the right side of the window, there are buttons for "OK", "Cancel", "Apply", "Disable", "Comment", "Copy", "Remove", and "Torch".



CRS326 VLAN Configuration

Bridge Port <ether2>

General STP VLAN Status

Interface: ether2

Bridge: bridge-lan

Horizon:

Learn: auto

Unknown Unicast Flood

Unknown Multicast Flood

Broadcast Flood

Trusted

Hardware Offload

OK

Cancel

Apply

Disable

Comment

Copy

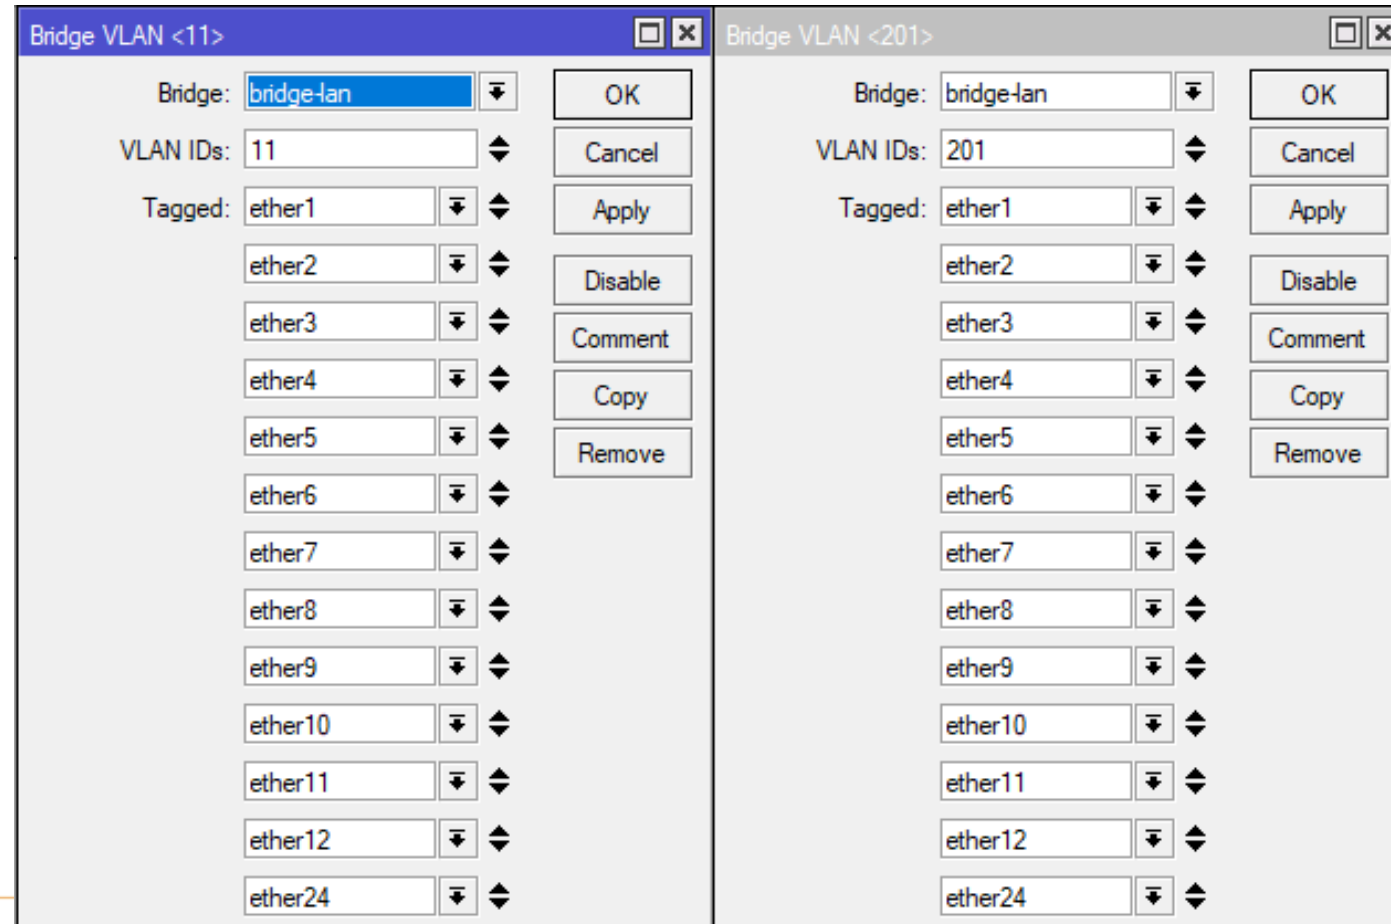
Remove

enabled inactive Hw. Offload

- Adds all Switch Ports to the Bridge
- Default is hardware offloaded enabled

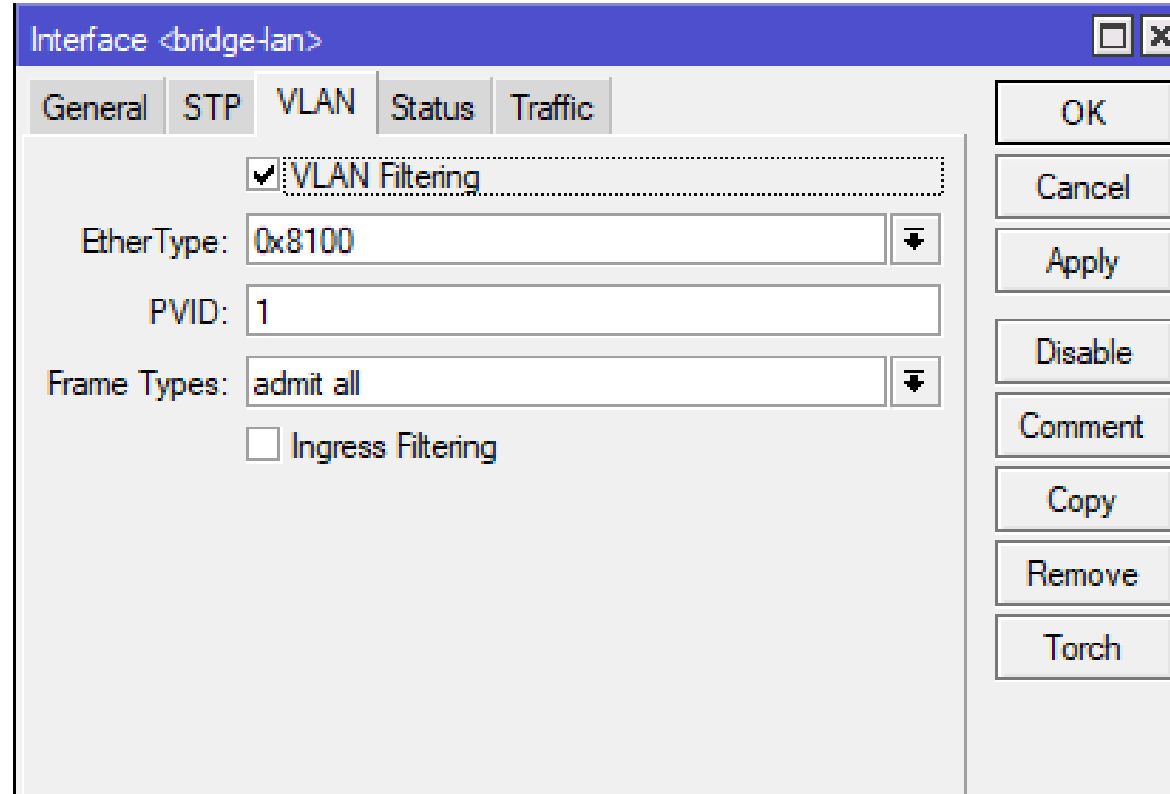
CRS326 VLAN Configuration

- Creates VLANs and adds the ports for the Wireless Access Points as Tagged Ports



CRS326 VLAN Configuration

- And finally Mike turns on VLAN filtering




The screenshot shows the 'Interface <bridge-lan>' configuration window with the 'VLAN' tab selected. The 'VLAN Filtering' checkbox is checked. Other settings include EtherType: 0x8100, PVID: 1, and Frame Types: admit all. The 'Ingress Filtering' checkbox is unchecked. On the right side, there are buttons for OK, Cancel, Apply, Disable, Comment, Copy, Remove, and Torch.

Tab	Value
VLAN Filtering	<input checked="" type="checkbox"/>
EtherType	0x8100
PVID	1
Frame Types	admit all
Ingress Filtering	<input type="checkbox"/>



CRS VLAN Configuration 1


```
/interface bridge add name=bridge-lan vlan-filtering=yes
/interface bridge port
add bridge=bridge-lan interface=ether1
add bridge=bridge-lan interface=ether2
add bridge=bridge-lan interface=ether3
add bridge=bridge-lan interface=ether4
add bridge=bridge-lan interface=ether5
add bridge=bridge-lan interface=ether6
add bridge=bridge-lan interface=ether7
add bridge=bridge-lan interface=ether8
add bridge=bridge-lan interface=ether9
add bridge=bridge-lan interface=ether10
add bridge=bridge-lan interface=ether11
add bridge=bridge-lan interface=ether12
add bridge=bridge-lan interface=ether24
/interface bridge vlan
add bridge=bridge-lan
tagged="ether1,ether2,ether3,ether4,ether5,ether6,ether7,ether8,ether9,ether10,ether11,ether12,ether24"
vlan-ids=201
add bridge=bridge-lan
tagged="ether1,ether2,ether3,ether4,ether5,ether6,ether7,ether8,ether9,ether10,ether11,ether12,ether24"
vlan-ids=11
```



VLAN in bridge with a physical interface

- Mike now adds in ether23 to link his old switch to the new switch.
- As his old switch has no VLANs and only needs traffic from the NEW VLAN11 he creates a VLAN on ether23 and adds this to the bridge

```
/interface vlan
add interface=ether23 name=vlan11-ether23 vlan-id=11
/interface bridge port
add bridge=bridge-lan interface=vlan11-ether23
/interface bridge vlan
add bridge=bridge-lan untagged=vlan11-ether23 vlan-ids=11
```



VLAN in bridge with a physical interface

Interface List

Interface	Interface List	Ethernet	EoIP Tunnel	IP Tunnel	GRE Tunnel	VLAN	VRRP	Bonding	LTE
	Name	Type	Actual MTU	L2 MTU	Tx	Rx			
R	ether13	Ethernet	1500	1592	0 bps				
R	ether14	Ethernet	1500	1592	0 bps				
	ether15	Ethernet	1500	1592	0 bps				
	ether16	Ethernet	1500	1592	0 bps				
	ether17	Ethernet	1500	1592	0 bps				
	ether18	Ethernet	1500	1592	0 bps				
	ether19	Ethernet	1500	1592	0 bps				
	ether20	Ethernet	1500	1592	0 bps				
	ether21	Ethernet	1500	1592	0 bps				
R	ether22	Ethernet	1500	1592	315.1 kbps				
R	ether23	Ethernet	1500	1592	0 bps				
R	vlan11-eth...	VLAN	1500	1588	0 bps				
RS	ether24	Ethernet	1500	1592	3.3 kbps				
	sfp-sfpplus1	Ethernet	1500	1592	0 bps				

32 items (1 selected)

Bridge VLAN <11>


Bridge: bridge-lan

VLAN IDs: 11

Tagged: ether1, ether2, ether3, ether4, ether5, ether6, ether7, ether8, ether9, ether10, ether11, ether12, ether24

Untagged: vlan11-ether23

Buttons: OK, Cancel, Apply, Disable, Comment, Copy, Remove



VLAN in bridge with a physical interface

- Mike is still having problems.
- He is seeing ports flapping on his old switch
- Spanning tree is doing strange things on his old switch
- Traffic just randomly stops

- He emails MikroTik to point out that their switches are incompatible with other vendors switches



VLAN in bridge with a physical interface

Analysis:

- Though this often works, this violates the 802.1w STP.
- The BPDU packets being sent on ether23 which is an untagged switch interface.
- The VLAN tagging is being applied in the CPU causing all packets being sent out as tagged traffic in VLAN11.
- Not all other vendors can understand tagged BPDU.
- In some more complex settings a switch may receive its own BPDU packet and trigger a loop detection when there is not one.

VLAN in bridge with a physical interface

Solution:

- The easiest solution is to simply disable STP/RSTP on the bridge

Preferred solution:

- Use bridge VLAN filtering correctly to tag the traffic



VLAN in bridge with a physical interface

- Mike decides he is going to use the preferred solution and makes changes to his bridge VLAN filtering
- Mike adds the physical port to the bridge and set the PVID and add the port as a untagged in the VLAN

```
/interface bridge port
add bridge=bridge-lan interface=ether23 pvid=11
/interface bridge vlan
add bridge=bridge-lan untagged=ether23 vlan-ids=11
```



VLAN in bridge with a physical interface

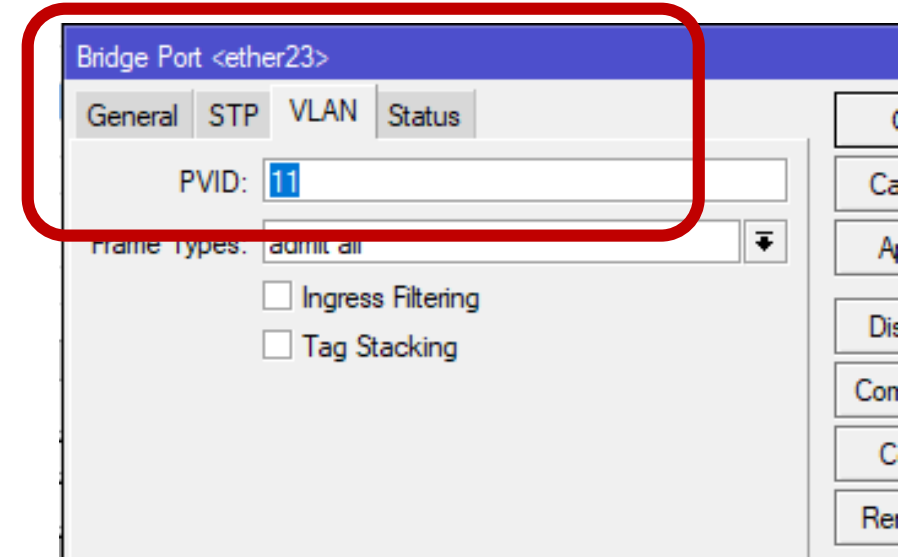
Bridge Port <ether23> (General tab):
Interface: ether23
Bridge: bridge-lan
Learn: auto
Unknown Unicast Flood:
Unknown Multicast Flood:
Broadcast Flood:
Trusted:
Hardware Offload:

Bridge Port <ether23> (VLAN tab):
PVID: 11
Frame Types: admit all
Ingress Filtering:
Tag Stacking:

Bridge VLAN <11>
Bridge: bridge-lan
VLAN IDs: 11
Tagged: ether1, ether2, ether3, ether4, ether5, ether6, ether7, ether8, ether9, ether10, ether11, ether12, ether24
Untagged: ether23
Current Tagged: ether1, ether2

Untagged Interfaces

- Setting a PVID on a bridge port once VLAN filtering is enabled will also set untagged interface in the VLAN in the VLAN table



Bridge

Bridge Ports VLANs MSTIs Port MST Overrides Filters NAT Hosts MDB

Bridge	VLAN IDs	Tagged	Untagged	Current Tagged	Current Untagged
bridge-lan	201	ether1, ether2, et...		ether1, ether2, ether3, et...	
bridge-lan	11	ether1, ether2, et...		ether1, ether2, ether3, et...	ether23
D bridge-lan	1				ether1, ether2, ether3, ether4, ether5, e...

Management Interface


- Mike realises his new CRS328 switch needs to have an IP Address on it so that he can manage it, and also it can then get NTP time sync so logging is useful
- Mike plans to manage his switch from VLAN 11
- Mike doesn't want any of the other VLANs to access his switch



Management Interface 1

- Mike connects his laptop to ether22 and sets it untagged in VLAN11.
- Mike now sets an IP Address on ether22 so he can manage the switch

```
/interface bridge port
add bridge=bridge-lan interface=ether22 pvid=11
/interface bridge vlan
add bridge=bridge-lan untagged=ether22 vlan-ids=11
/ip address
add address=10.100.1.2/24 interface=ether22 network=10.100.1.0
```



Management Interface 1

The image displays three overlapping network configuration windows:

- Address <10.100.1.2/24>**: Shows configuration for an IP address. The **Address** field is set to `10.100.1.2/24`, **Network** to `10.100.1.0`, and **Interface** to `ether22`. The status is `enabled`.
- Bridge Port <ether22>**: Shows configuration for a bridge port. The **PVID** is set to `11`. **Frame Types** are set to `admit all`. **Ingress Filtering** and **Tag Stacking** are unchecked. The status is `enabled`, `inactive`, and `Hw. Offload`.
- Bridge VLAN <11>**: Shows configuration for a bridge VLAN. The **Bridge** is `bridge-lan` and **VLAN IDs** is `11`. The **Tagged** list includes `ether1` through `ether10`, `ether11`, `ether12`, and `ether24`. The **Untagged** field is set to `ether22`. A large red 'X' is drawn over the **Tagged** list.



Management Interface 1

- Mike is still unable to access his new CRS328 switch by IP Address

```
Pinging 10.100.1.2 with 32 bytes of data:  
Request timed out.  
Request timed out.  
Reply from 10.100.1.252: Destination host unreachable.  
Reply from 10.100.1.252: Destination host unreachable.  
  
Ping statistics for 10.100.1.2:  
    Packets: Sent = 4, Received = 2, Lost = 2 (50% loss),
```

- Mike calls Dave.....



Management Interface 1

Analysis:

- When you add an interface to the bridge it becomes a slave interface

R	ether21	Ethernet	1500	1592	329.6 kbps	19.8 kbps
RS	ether22	Ethernet	1500	1592	166.5 kbps	80.1 kbps
RS	ether23	Ethernet	1500	1592	16.9 kbps	0 bps

- All traffic captured on the interface is sent to the CPU on the bridge interface, not the physical ethernet port
- The slave interface never captures any traffic on the interface

Management Interface 1

Solution:


- Set the management interface VLAN to be on the master interface in this case the master is bridge-lan



Management Interface 2

- Mike add VLAN11 to the bridge with PVID11 and sets the VLAN interface as an untagged port in the bridge....

```
/interface vlan
add interface=bridge-lan name=vlan11 vlan-id=11
/interface bridge port
add bridge=bridge-lan interface=vlan11 pvid=11
/interface bridge vlan
add bridge=bridge-lan untagged=vlan11 vlan-ids=11
/ip address
add address=10.100.1.2/24 interface=vlan11 network=10.100.1.0
```



Management Interface 2

Interface List

Name	Type	Actual MTU	L2 MTU	Tx	Rx
bridge-lan	Bridge	1500	1592	0 bps	0
vlan11	VLAN	1500	1588	0 bps	0
ether1	Ethernet	1500	1592	128.2 kbps	
ether2	Ethernet				
ether3	Ethernet				
ether4	Ethernet				
ether5	Ethernet				
ether6	Ethernet				
ether7	Ethernet				
ether8	Ethernet				
ether9	Ethernet				
ether10	Ethernet				

Bridge Port <vlan11>

General STP VLAN Status

PVID: 11

Frame Types: admit all

Ingress Filtering

Tag Stacking

enabled | inactive | Hw. Offload

Bridge VLAN <11>

Bridge: bridge-lan

VLAN IDs: 11

Tagged: ether1, ether2, ether3, ether4, ether5, ether6, ether7, ether8, ether9, ether10, ether11, ether12, ether24

Untagged: vlan11

OK, Cancel, Apply, Disable, Comment, Copy, Remove



Management Interface 2

- Mike is still unable to access his new CRS328 switch by IP Address

```
Pinging 10.100.1.2 with 32 bytes of data:  
Request timed out.  
Request timed out.  
Reply from 10.100.1.252: Destination host unreachable.  
Reply from 10.100.1.252: Destination host unreachable.  
  
Ping statistics for 10.100.1.2:  
    Packets: Sent = 4, Received = 2, Lost = 2 (50% loss),
```

- Mike calls Dave



Management Interface 2

Analysis

- The only connection from the switch chip to the CPU is through the bridge interface.
- The bridge works like a Switch and this should not be mixed up with physical interfaces with VLAN interfaces on.



Management Interface 2

Solution:

1. Add the bridge as a tagged port into required VLANs
 2. Creating VLAN interfaces on the bridge
- This will allow tagged packets to enter the bridge from the switch chip and will be then passed through to the specific VLAN interface on the bridge.



Management Interface 2

Solution cont.:

- To secure management access further Dave tells Mike to use ingress filtering to allow only traffic on Specific VLAN ID



Management Interface 2

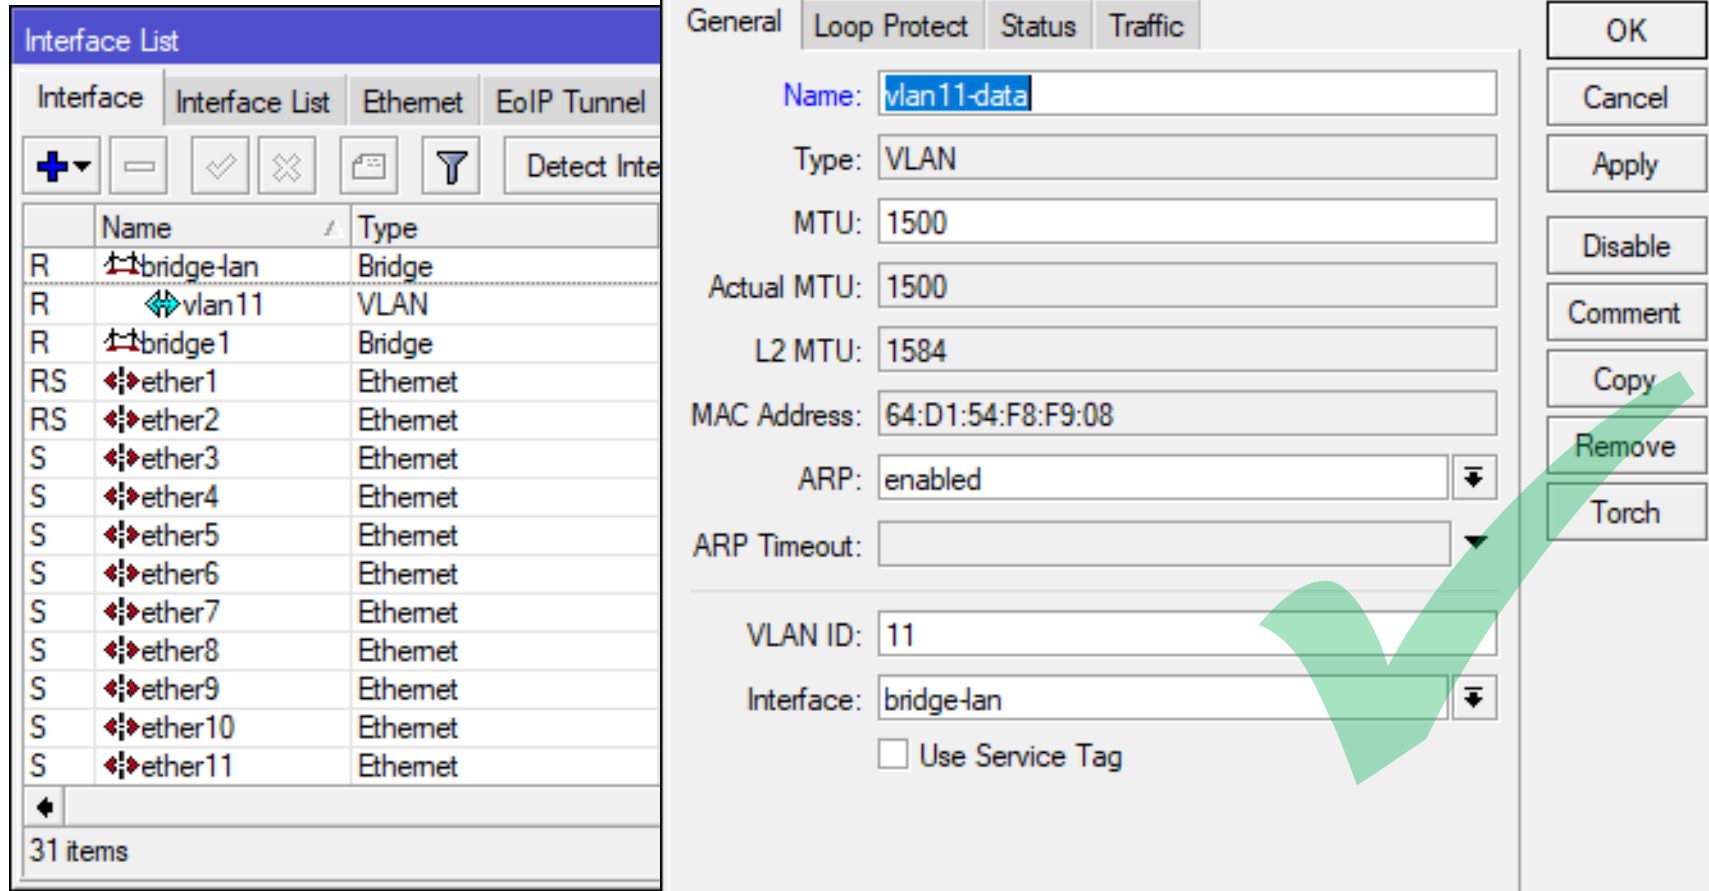
Solution cont.:

- Set the CPU/Bridge interface to pass tagged traffic.
- The CPU now receives tagged traffic and management access is allowed

```
/interface vlan add interface=bridge-lan name=vlan11 vlan-id=11  
/interface bridge vlan add bridge=bridge-lan tagged=bridge-lan vlan-ids=11  
/ip address add address=10.100.1.2/24 interface=vlan11 network=10.100.1.0  
/ip route add dst-address=0.0.0.0/0 gateway=10.100.1.1 distance=1
```

Management Interface - Correct

- Create a VLAN interface on the bridge interface



The image shows two windows from a network configuration application. The left window, titled 'Interface List', displays a table of network interfaces. The right window, titled 'Interface <vlan1>', shows the configuration for a new VLAN interface.

Interface	Name	Type
R	bridge-lan	Bridge
R	vlan11	VLAN
R	bridge1	Bridge
RS	ether1	Ethernet
RS	ether2	Ethernet
S	ether3	Ethernet
S	ether4	Ethernet
S	ether5	Ethernet
S	ether6	Ethernet
S	ether7	Ethernet
S	ether8	Ethernet
S	ether9	Ethernet
S	ether10	Ethernet
S	ether11	Ethernet

The 'Interface <vlan1>' configuration window shows the following settings:

- Name: vlan11-data
- Type: VLAN
- MTU: 1500
- Actual MTU: 1500
- L2 MTU: 1584
- MAC Address: 64:D1:54:F8:F9:08
- ARP: enabled
- ARP Timeout: (empty)
- VLAN ID: 11
- Interface: bridge-lan
- Use Service Tag

A large green checkmark is overlaid on the right side of the configuration window, indicating that the configuration is correct. The 'Interface List' window shows 31 items in total.

Management Interface - Correct

Bridge VLAN <11>

Bridge: bridge-lan

VLAN IDs: 11

Tagged: ether1, ether2, ether3, ether4, ether5, ether6, ether7, ether8, ether9, ether10, ether11, ether12, ether24, bridge-lan

Untagged: ether23

Buttons: OK, Cancel, Apply, Disable, Comment, Copy, Remove

- Add bridge as a Tagged Port on the VLAN11 – IMPORTANT
- Add an IP Address to the VLAN interface

New Address

Address: 10.100.1.2/24

Network: 10.100.1.0

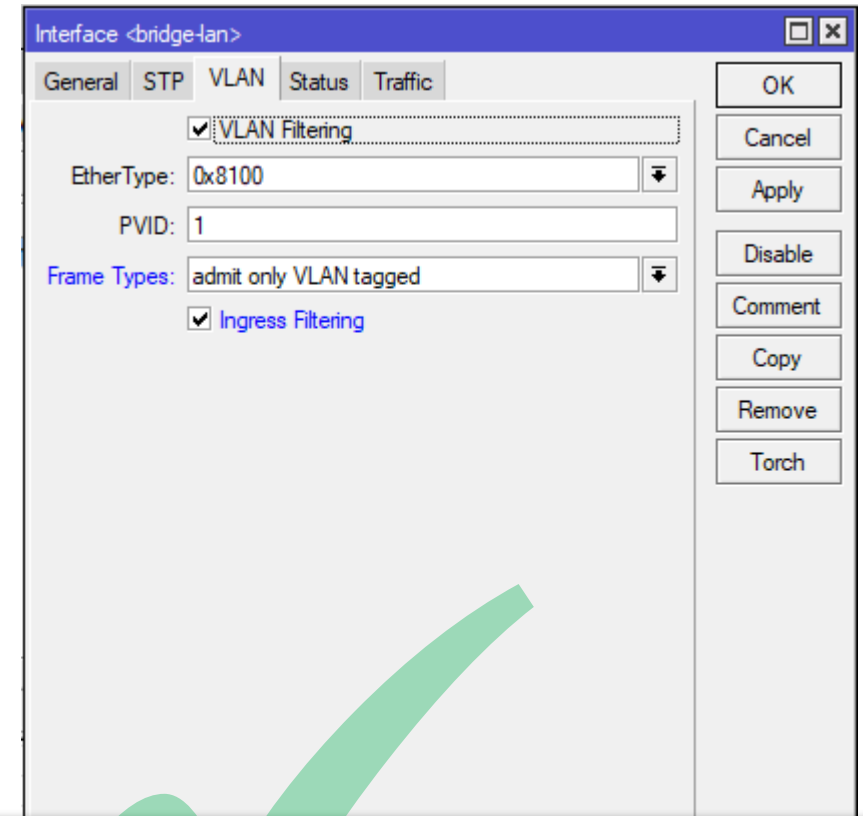
Interface: vlan11-data

Buttons: OK, Cancel, Apply, Disable, Comment, Copy, Remove

enabled

Ingress Filtering

- Mike also configures Ingress Filtering.
- Ingress filtering not only can be applied to a physical switch port, but also the bridge port / switch-cpu port



The screenshot shows a configuration window titled 'Interface <bridge-lan>' with tabs for General, STP, VLAN, Status, and Traffic. The VLAN tab is active, showing the following settings: 'VLAN Filtering' is checked, 'EtherType' is set to '0x8100', 'PVID' is set to '1', 'Frame Types' is set to 'admit only VLAN tagged', and 'Ingress Filtering' is checked. On the right side of the window, there are buttons for 'OK', 'Cancel', 'Apply', 'Disable', 'Comment', 'Copy', 'Remove', and 'Torch'. A large green checkmark is overlaid on the bottom right of the window.

```
/interface bridge  
add frame-types=admit-only-vlan-tagged ingress-filtering=yes name=bridge-lan vlan-filtering=yes
```

Egress Filtering

- When Bridge VLAN-filtering is enabled....
- By default the switch ports and the bridge port (switch CPU), filter on Egress based on the VLAN table Invalid VLANs are dropped on Egress.



Ingress Filtering

- Ingress Filtering can be used along with frame type to limit which packets are allowed to access the device, both from physical ports and also from the CPU.
- Ingress filtering will check the VLAN table and only allow VLANs. VLANs not specified in the VLAN table will be dropped



Ingress Filtering

- By applying `frame-type=VLAN-tagged-only` tagged packets will pass from the switch
- Applying `frame-type=VLAN-tagged-only` will also disable dynamic adding of untagged ports based on PVID



Bridge VLAN Filtering – Mike's Learnt



- Mike has learnt a few important things about the new Bridge-VLAN Filtering options



Bridge VLAN Filtering



1. Its important to check the features and use the choose the correct hardware



Bridge VLAN Filtering



2. The bridge works like a Switch and this should not be mixed up with physical interfaces with VLAN interfaces on.

Use only one method of VLAN configuration

Router or Switch

Bridge VLAN Filtering



3. The only connection from the switch to the CPU is the bridge interface. This needs adding as a tagged port in VLANs as needed and works like a switch port.

Bridge VLAN Filtering



- Mike is now a fan of VLANs on MikroTik switches!

Packet flow with hardware offloading


- Mike notices that a couple of his ports have high traffic
- Mike is curious as to what is happening. So he torches one of the interfaces
- He only sees a few packets. However he can see that there is a lot more traffic on that interface
- He calls Dave....

Packet flow with hardware offloading

- Dave immediately knows that Mike needs to look at the block diagram to understand how the hardware functions
- Dave spent lots of time last year looking at the block diagrams for different MikroTik hardware...
- Dave looks at the block diagram for the CRS328 with Mike and explains how it works....



Packet flow with hardware offloading



Packet flow with hardware offloading

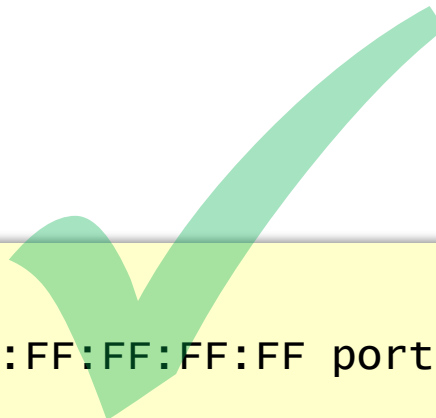
- Each ethernet port is connected to the switch chip
- The switch chip is connected the CPU (sometimes called switch-cpu port)
- Once hw-offloading has been enabled, the switch chip forwards packets between ports
- For packets to view in sniffer tools, they need to be sent to the CPU



Packet flow with hardware offloading

- If you know the traffic you are interested in then you can copy traffic to CPU.
- This is an example of how to send packets destined for EC:F4:BB:50:5E:CF

```
/interface ethernet switch rule  
add copy-to-cpu=yes dst-mac-address=EC:F4:BB:50:5E:CF/FF:FF:FF:FF:FF:FF ports=ether1  
switch=switch1
```




LAG / Bonding

- Mike wishes he could have more bandwidth between his servers.
- Mike also reads that MikroTik support bonded interfaces. His servers also support LACP so he bonds 2 ports together to increase throughput.



LAG / Bonding



LAG / Bonding

- Mike has also remembered to check if his CRS328 supports hw-offloading with bonded interfaces
- Mike sees that only 802.3ad (LACP) bonding is supported so he uses that bonding mode



LAG / Bonding

Switch Chip	Model (example units)	STP/RSTP	MSTP	DHCP Snooping	VLAN Filtering	Bonding ³
	CRS3xx	✓	✓	✓	✓	✓
	CRS1xx/2xx	✓	✗	✓ ¹	✗	✗
QCA8337	hAP ac / hEX PoE / 3011 (1gb)	✓	✗	✓ ²	✗	✗
AR8327	hAP ac ² /2011/1100AHx2 (1gb)	✓	✗	✓ ²	✗	✗
AR8227	hAP/hEX lite/2011 (100mb)	✓	✗	✓ ²	✗	✗
AR8316		✓	✗	✓ ²	✗	✗
AR7240		✓	✗	✓ ²	✗	✗
MT7621	hEX (750Gr3)	✓	✗	✓	✗	✗
RTL8367	1100AHx4	✗	✗	✓	✗	✗
ICPlus175D		✗	✗	✓	✗	✗

1. Feature will not work properly in VLAN switching setups, you must make sure that required packets are sent out with the correct VLAN tag using ACL rules.
2. DHCP snooping will not work properly with VLAN switching
3. Bridge hardware offloading only supported using 802.3ad bonding

Complete list https://wiki.mikrotik.com/wiki/Manual:Switch_Chip_Features#Bridge_Hardware_Offloading

LAG / Bonding

- Mike configures two bonding interfaces and sets them as untagged on VLAN11

```
/interface bonding
add mode=802.3ad name=bonding1 slaves=ether13,ether14
add mode=802.3ad name=bonding2 slaves=ether15,ether16
/interface bridge port
add bridge=bridge-lan interface=bonding1 pvid=11
add bridge=bridge-lan interface=bonding2 pvid=11
/interface bridge vlan
add bridge=bridge-lan untagged=bonding1,bonding2
```



LAG / Bonding

Interface <bonding1>

General Bonding Status Traffic

Slaves: ether13
ether14

Mode: 802.3ad

Primary: none

Link Monitoring: mii

Transmit Hash Policy: layer 2

Min. Links: 0

Down Delay: 0 ms

Up Delay: 0 ms

LACP Rate: 30 s

MII Interval: 100 ms

enabled running slave

Bridge Port <bonding1>

General STP VLAN Status

PVID: 11

Frame Types: admit all

Ingress Filtering

Tag Stacking

enabled inactive Hw. Offload



Bridge VLAN <11>

Bridge: bridge-lan

VLAN IDs: 11

Tagged: ether1
ether2
ether3
ether4
ether5
ether6
ether7
ether8
ether9
ether10
ether11
ether12
ether24
bridge-lan

Untagged: bonding1
bonding2




LAG / Bonding

- Mike notices that he still only gets 1Gb between his servers when testing with a well known network performance tool (iperf)
- He only sees the traffic on one of the bonded interface slave interfaces
- He checks the CPU and that it is not heavily loaded



LAG / Bonding



Interface List

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

+ - ✓ ✗ 📁 🔍 Detect Internet

Name	Type	Actual MTU	L2 MTU	Tx	Rx	Tx Packet (p/s)	Rx Packet (p/s)	FP Tx	FP R
RS ether5	Ethernet	1500	1592	25.6 kbps	2.8 kbps	19	3	0 bps	
RS ether6	Ethernet	1500	1592	27.6 kbps	0 bps	21	0	0 bps	
RS ether7	Ethernet	1500	1592	25.6 kbps	2.8 kbps	19	3	0 bps	
RS ether8	Ethernet	1500	1592	27.6 kbps	0 bps	21	0	0 bps	
RS ether9	Ethernet	1500	1592	27.6 kbps	0 bps	21	0	0 bps	
RS ether10	Ethernet	1500	1592	27.6 kbps	0 bps	21	0	0 bps	
RS ether11	Ethernet	1500	1592	27.6 kbps	0 bps	21	0	0 bps	
RS ether12	Ethernet	1500	1592	27.6 kbps	0 bps	21	0	0 bps	
RS ether13-bonding1	Ethernet	1500	1592	2.3 Mbps	81.6 kbps	4 500	12	432 bps	
RS ether14-bonding1	Ethernet	1500	1592	568 bps	948.1 Mbps	1	78 676	0 bps	
RS ether15-bonding2	Ethernet	1500	1592	960.9 Mbps	2.6 Mbps	79 762	5 070	0 bps	
RS ether16-bonding2	Ethernet	1500	1592	0 bps	0 bps	0	0	0 bps	
ether17	Ethernet	1500	1592	0 bps	0 bps	0	0	0 bps	
ether18	Ethernet	1500	1592	0 bps	0 bps	0	0	0 bps	

32 items (1 selected)

LAG / Bonding

- This time he remembers to check that the ports are hw-offloaded and the bonded interfaces are

The screenshot shows a network configuration interface with a 'Bridge' tab selected. The 'Ports' sub-tab is active, displaying a table of bridge ports. A red box highlights the bonded interfaces 'bonding1' and 'bonding2'. To the right, a 'Terminal' window displays the output of a command, showing details for various interfaces, including 'bonding1' and 'bonding2', which are also highlighted with a red box. The terminal output includes flags for hardware offload (HW) and path cost (PR).

#	Interface	Bridge	HW	PVID	PR	PATH-COST	INTERNA...	HORIZON
6	H ether7	bridge-lan	yes	1	0x	10	10	none
7	H ether8	bridge-lan	yes	1	0x	10	10	none
8	H ether9	bridge-lan	yes	1	0x	10	10	none
9	H ether10	bridge-lan	yes	1	0x	10	10	none
10	H ether11	bridge-lan	yes	1	0x	10	10	none
11	H ether12	bridge-lan	yes	1	0x	10	10	none
12	H ether24	bridge-lan	yes	1	0x	10	10	none
13	H ether23	bridge-lan	yes	1	0x	10	10	none
14	H bonding1	bridge-lan	yes	11	0x	10	10	none
15	H bonding2	bridge-lan	yes	11	0x	10	10	none

- He calls Dave

LAG / Bonding

Analysis:

- LACP (802.3ad) does not create 1 x 2Gbps interface but an interface that can transmit traffic over multiple slave interfaces
- LACP (802.3ad) using transmit hash policy (MAC, IP or Port).
- As the traffic is going to the same dst MAC and dst IP and dst Port, load balance between different members is not possible.
- This is correct for 802.3ad.

LAG / Bonding

Solution:

- Traffic going from multiple sources / destinations will load balance across LAG members
- Different transmit hash policies can increase single stream throughput. However these are not hw-offloaded in the bridge configuration so will be limited to CPU



DHCP Snooping

- Mike has a problem on his network. His users keep bring in their old routers from home and plugging them in to give them extra switch ports under their desks....
- This causes clients on his network to obtain incorrect IP information as these routers are also DHCP Servers.



DHCP Snooping

- Mike reads more about the features on his new switch.
- Mike sees that he can run DHCP Snooping to prevent his users plugging rogue routers into his network



DHCP Snooping

- Mike turns on DHCP Snooping on his bridge

```
/interface bridge
add dhcp-snooping=yes name=bridge-lan vlan-
filtering=yes
```

Interface <bridge-lan>

General STP VLAN Status Traffic

Name: bridge-lan

Type: Bridge

MTU: [dropdown]

Actual MTU: 1500

L2 MTU: 1592

MAC Address: CE:2D:E0:7B:38:41

ARP: enabled

ARP Timeout: [dropdown]

Admin. MAC Address: CE:2D:E0:7B:38:41

Ageing Time: 00:05:00

DHCP Snooping

Fast Forward

enabled running slave

DHCP Snooping

- Mike now has a problem...
- None of his client devices are getting a DHCP Address.
- He checks his switch and hw-offloading is still enabled
- He check his switch CPU and that is not maxed out...
- He Calls Dave....

DHCP Snooping

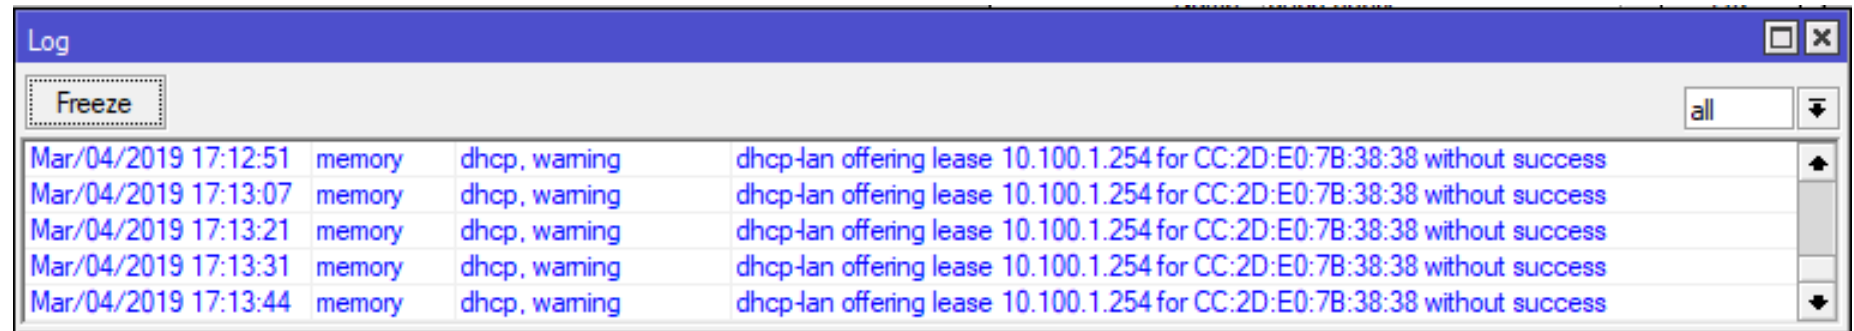
- Dave tells Mike to check the logs on both his CRS328 and his hAP ac²

- CRS328



Time	Category	Source	Message
Mar/04/2019 17:12:46	memory	interface, warning	ether24: received DHCP server message on untrusted port
Mar/04/2019 17:12:58	memory	interface, warning	ether24: received DHCP server message on untrusted port
Mar/04/2019 17:13:13	memory	interface, warning	ether24: received DHCP server message on untrusted port
Mar/04/2019 17:13:22	memory	interface, warning	ether24: received DHCP server message on untrusted port
Mar/04/2019 17:13:29	memory	interface, warning	ether24: received DHCP server message on untrusted port
Mar/04/2019 17:13:41	memory	interface, warning	ether24: received DHCP server message on untrusted port
Mar/04/2019 17:13:51	memory	interface, warning	ether24: received DHCP server message on untrusted port

- hAP ac²



Time	Category	Source	Message
Mar/04/2019 17:12:51	memory	dhcp, warning	dhcp-lan offering lease 10.100.1.254 for CC:2D:E0:7B:38:38 without success
Mar/04/2019 17:13:07	memory	dhcp, warning	dhcp-lan offering lease 10.100.1.254 for CC:2D:E0:7B:38:38 without success
Mar/04/2019 17:13:21	memory	dhcp, warning	dhcp-lan offering lease 10.100.1.254 for CC:2D:E0:7B:38:38 without success
Mar/04/2019 17:13:31	memory	dhcp, warning	dhcp-lan offering lease 10.100.1.254 for CC:2D:E0:7B:38:38 without success
Mar/04/2019 17:13:44	memory	dhcp, warning	dhcp-lan offering lease 10.100.1.254 for CC:2D:E0:7B:38:38 without success



DHCP Snooping

Analysis:

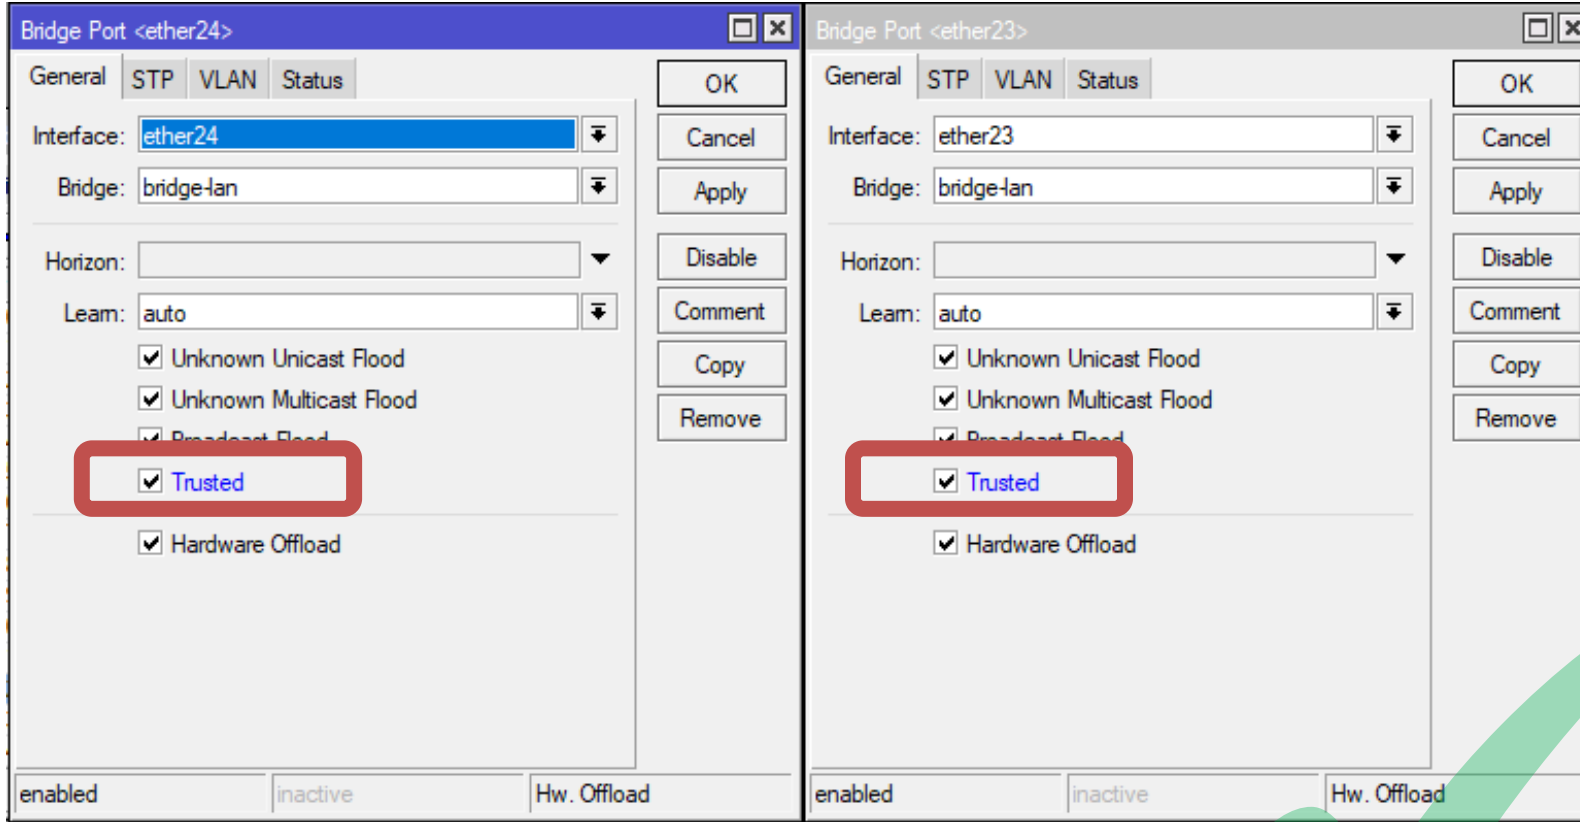
- DHCP Snooping is a Layer 2 Security feature
- This limits the ports on which DHCP messages are received
- Mike has successfully drop DHCP messages from the rogue routers on his network
- Mike has also blocked DHCP messages from his legitimate DHCP Server too!

DHCP Snooping


Solution:

- Dave tells Mike that he needs to allow DHCP messages to be forward on the port with his server and the port facing other switches.
- Mike sets the ports as trusted
- Mike's Clients now get an IP Address

DHCP Snooping



```
/interface bridge port  
add bridge=bridge-lan interface=ether24 trusted=yes  
add bridge=bridge-lan interface=ether23 trusted=yes
```



Thank you for Listening



References

- Visio Templates – Mikrotik Forum user FernandoSuperGG
<https://forum.mikrotik.com/viewtopic.php?f=2&t=120957>
- MikroTik Manual
https://wiki.mikrotik.com/wiki/Manual:CRS_Router#CRS3xx_series_switches
https://wiki.mikrotik.com/wiki/Manual:CRS3xx_series_switches
https://wiki.mikrotik.com/wiki/Manual:Layer2_misconfiguration
<https://wiki.mikrotik.com/wiki/Manual:Interface/Bridge>
https://wiki.mikrotik.com/wiki/Manual:Switch_Chip_Features#Bridge_Hardware_Offloading
https://wiki.mikrotik.com/wiki/Manual:Bridge_VLAN_Table

