Most underused MikroTik hardware and software features

OR

“*The path between fastpath and advanced features*”

MUM, Europe 2018

Objectives

- To help you understand and combine “FastPath” and “SlowPath” features.

- To allow MikroTik equipment to do more.

- Encourage not only to update RouterOS version, but also update existing configurations with the latest features.

- Reduce the amount of hardware performance issue emails to support@mikrotik.com!

Presentation plan

- This presentation will consist of most popular performance issues related to mistakes in
  - Hardware choice
  - Hardware usage
  - Layer-2 feature usage
  - Layer-3 feature usage
- We will cover the information needed to avoid such mistakes.

Know your hardware

- Improper use of hardware or using the wrong one for the job is by far the most popular mistake we see in the support requests.
- Each device made by MikroTik has it's specifics both in:
  - Structure (CPU cores, memory, port inter connections)
  - Performance (switching, bridging, routing, encrypting)

Meet Dave

- Dave is a smart and experienced network administrator, well certified in mainstream network equipment brands.
- There was a disaster, the main router died, and Dave needs to get at least something in the network working NOW!!
- The only spare equipment he can get his hands on is some strange “hEX” (RB750Gr3) board from someone called “MikroTik”, that a friend gave him to try out some time ago.
- Dave needs MPLS, L2TP+IPSec, firewall and routing.

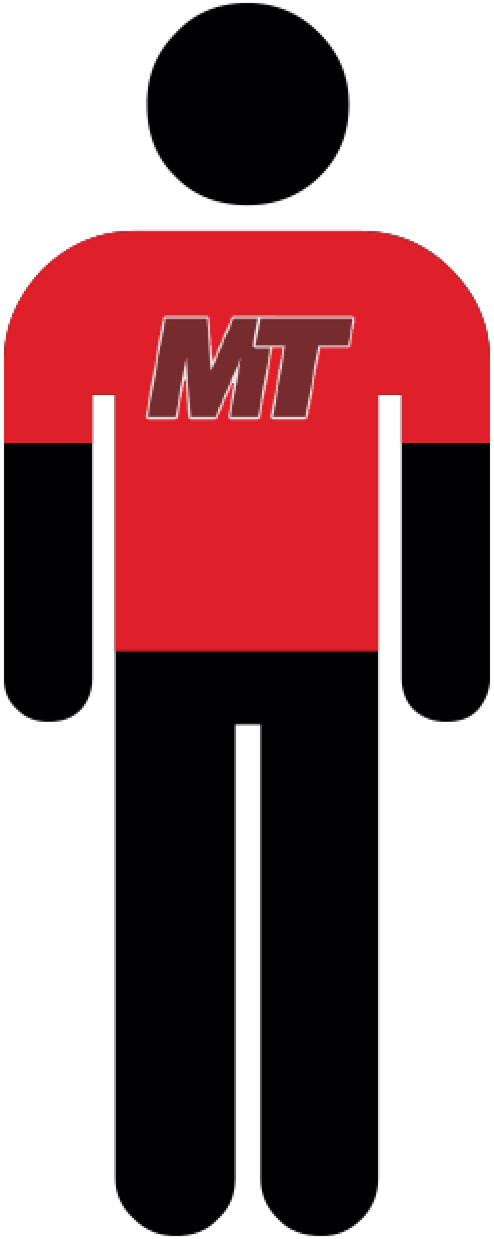

Few Days later

- Dave applied RB750Gr3 as a fix and got most of the services online.
- He is in shock how a $60 box was able to do all this.
- Dave has discovered RouterOS and MikroTik instantly becoming a MikroTik fanboy.
- He is sending lots of questions to support@mikrotik.com .

Analysis of the problem

- Dave’s problem #1:
  - The daily database exchange throughput is limited to 1 Gbps total, and CPU is not 100%, using routing with large packets.
- Diagnosis:
  - Block diagram for RB750Gr3.
- Reason:
  - Dave uses ether2 and ether4 ports for database exchange, both ports are on the same 1 Gbps line to CPU.

RB750Gr3 block diagram 1

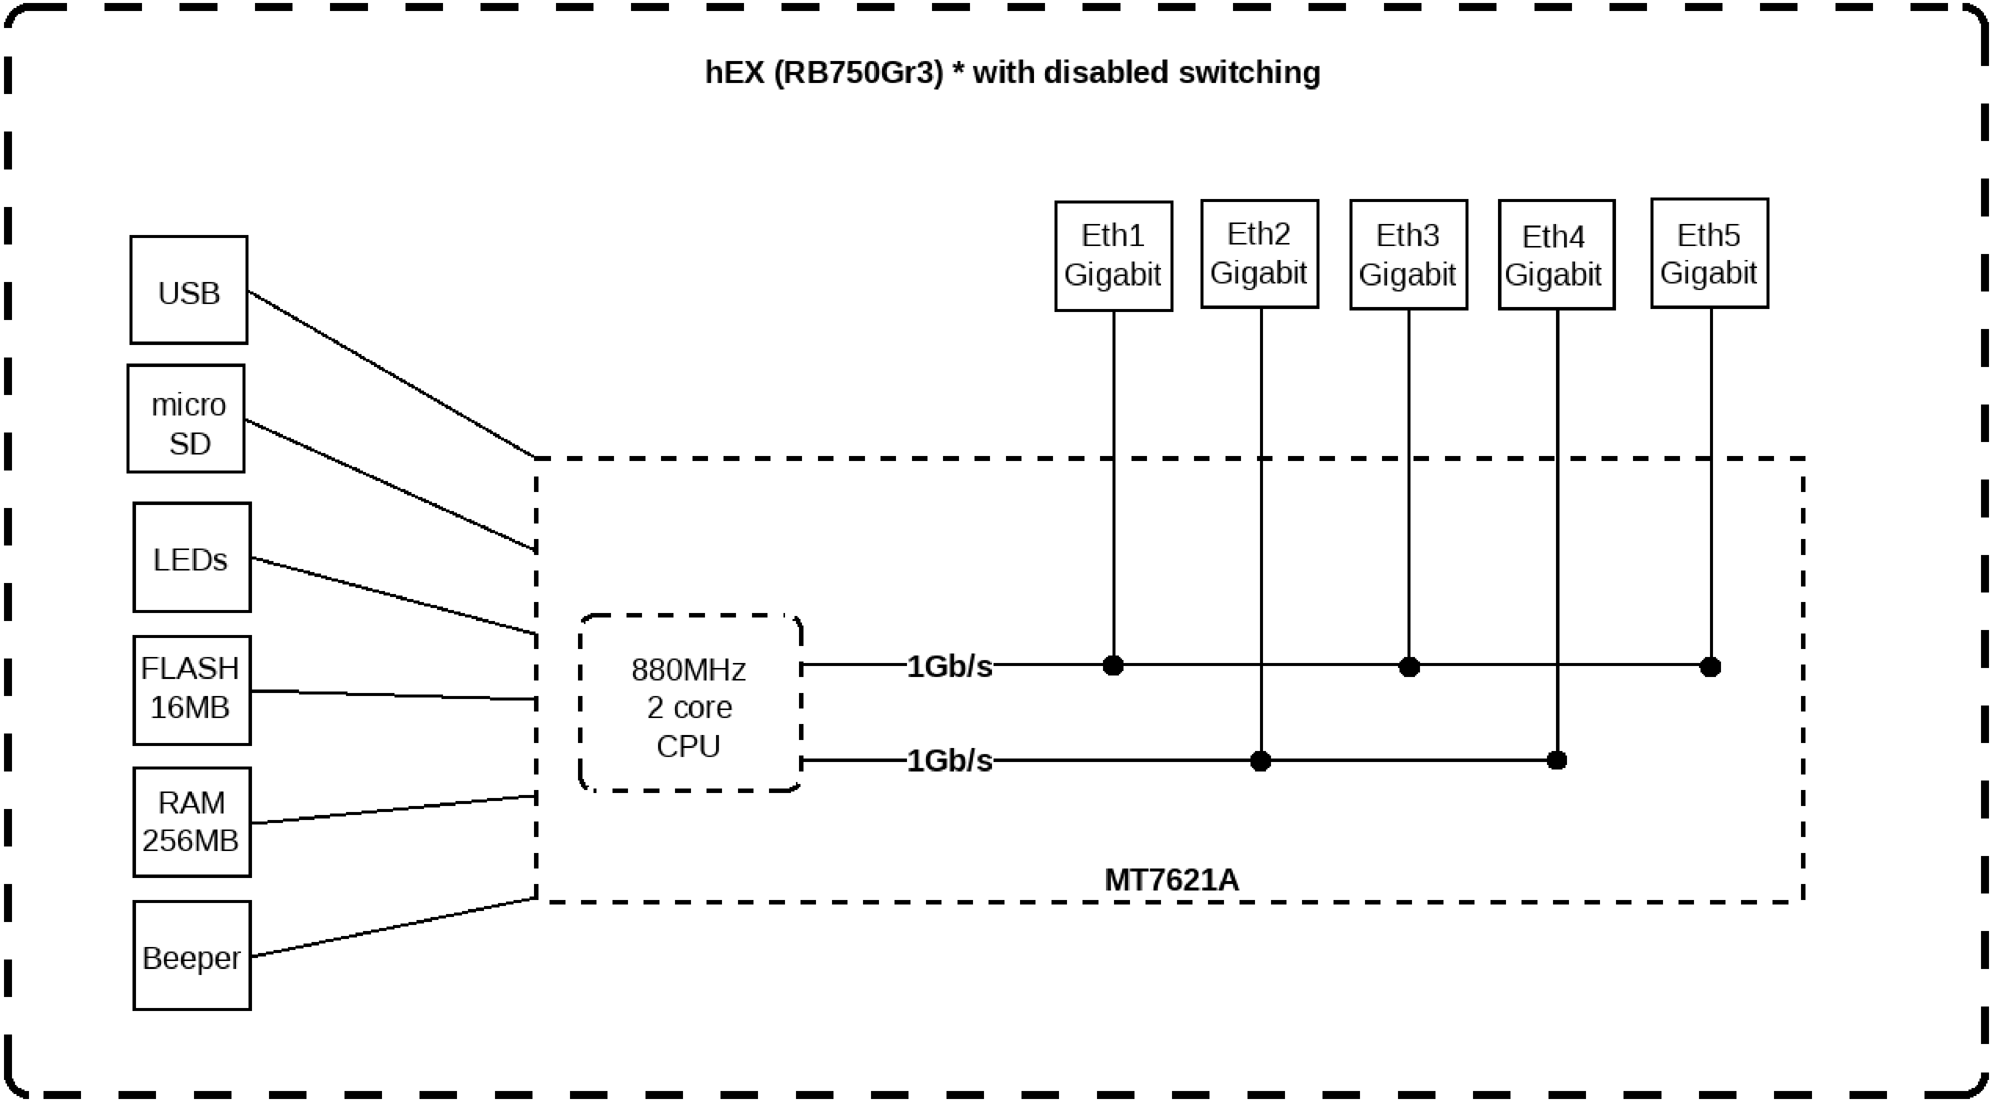

Analysis of the problem

- Dave’s problem #2:
  - Dave put two ports in a hardware bridge and suddenly is limited to 1 Gbps total again.
- Diagnosis:
  - Other block diagram for RB750Gr3.
- Reason:
  - Hardware bridge switches ports, CPU assigns one dedicated 1 Gbps line to switched ports.
  - Dave needs to use software bridging.

RB750Gr3 block diagram 2

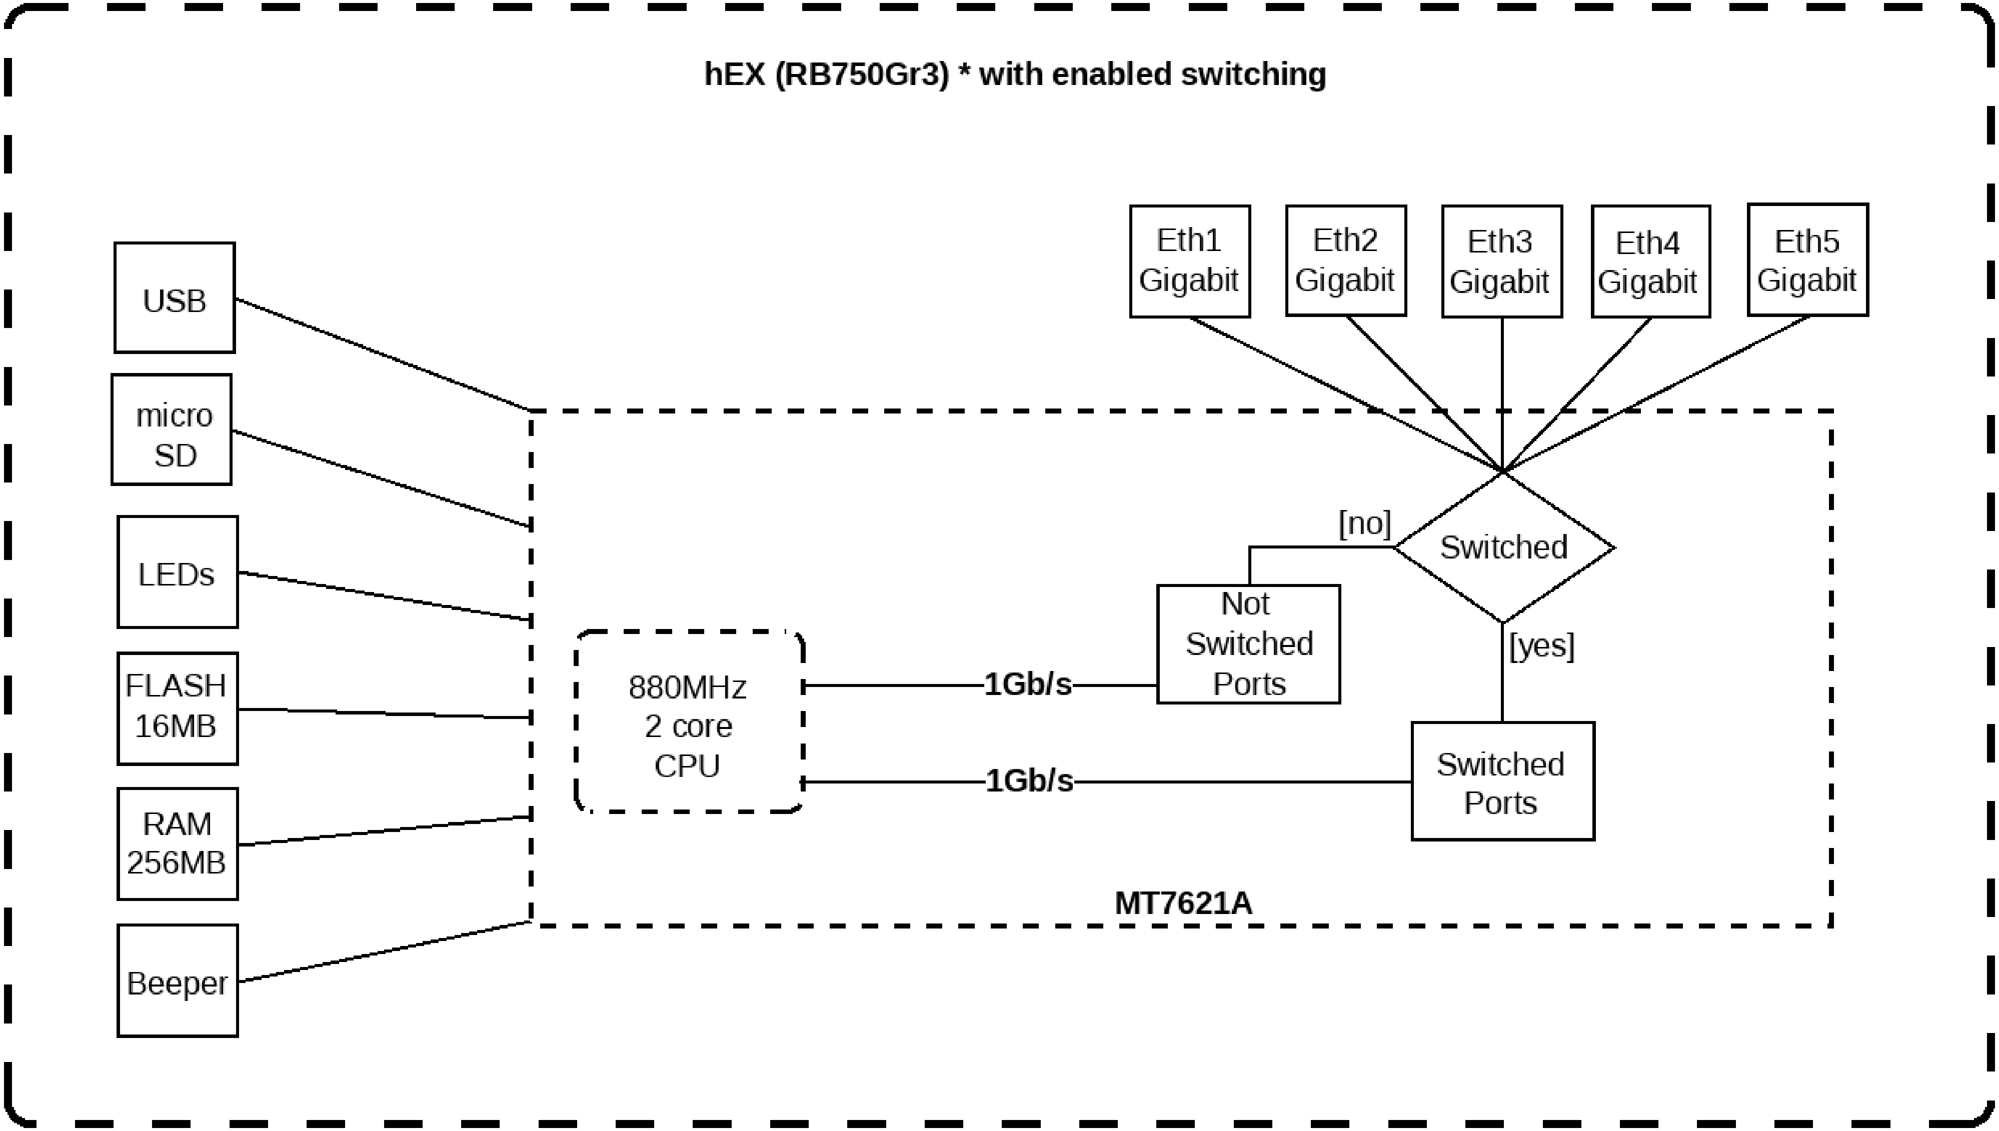

Improving the temporary fix

- Dave needed more ports, so he decided to replace the RB750Gr3 with an RB3011UiAS-RM.
- He needs more switching throughput so he examines RB3011UiAS-RM block diagram for bottlenecks.
- Dave is getting more enthusiastic about MikroTik (judging from mail frequency to support@mikrotik.com).

RB3011UiAS block diagram

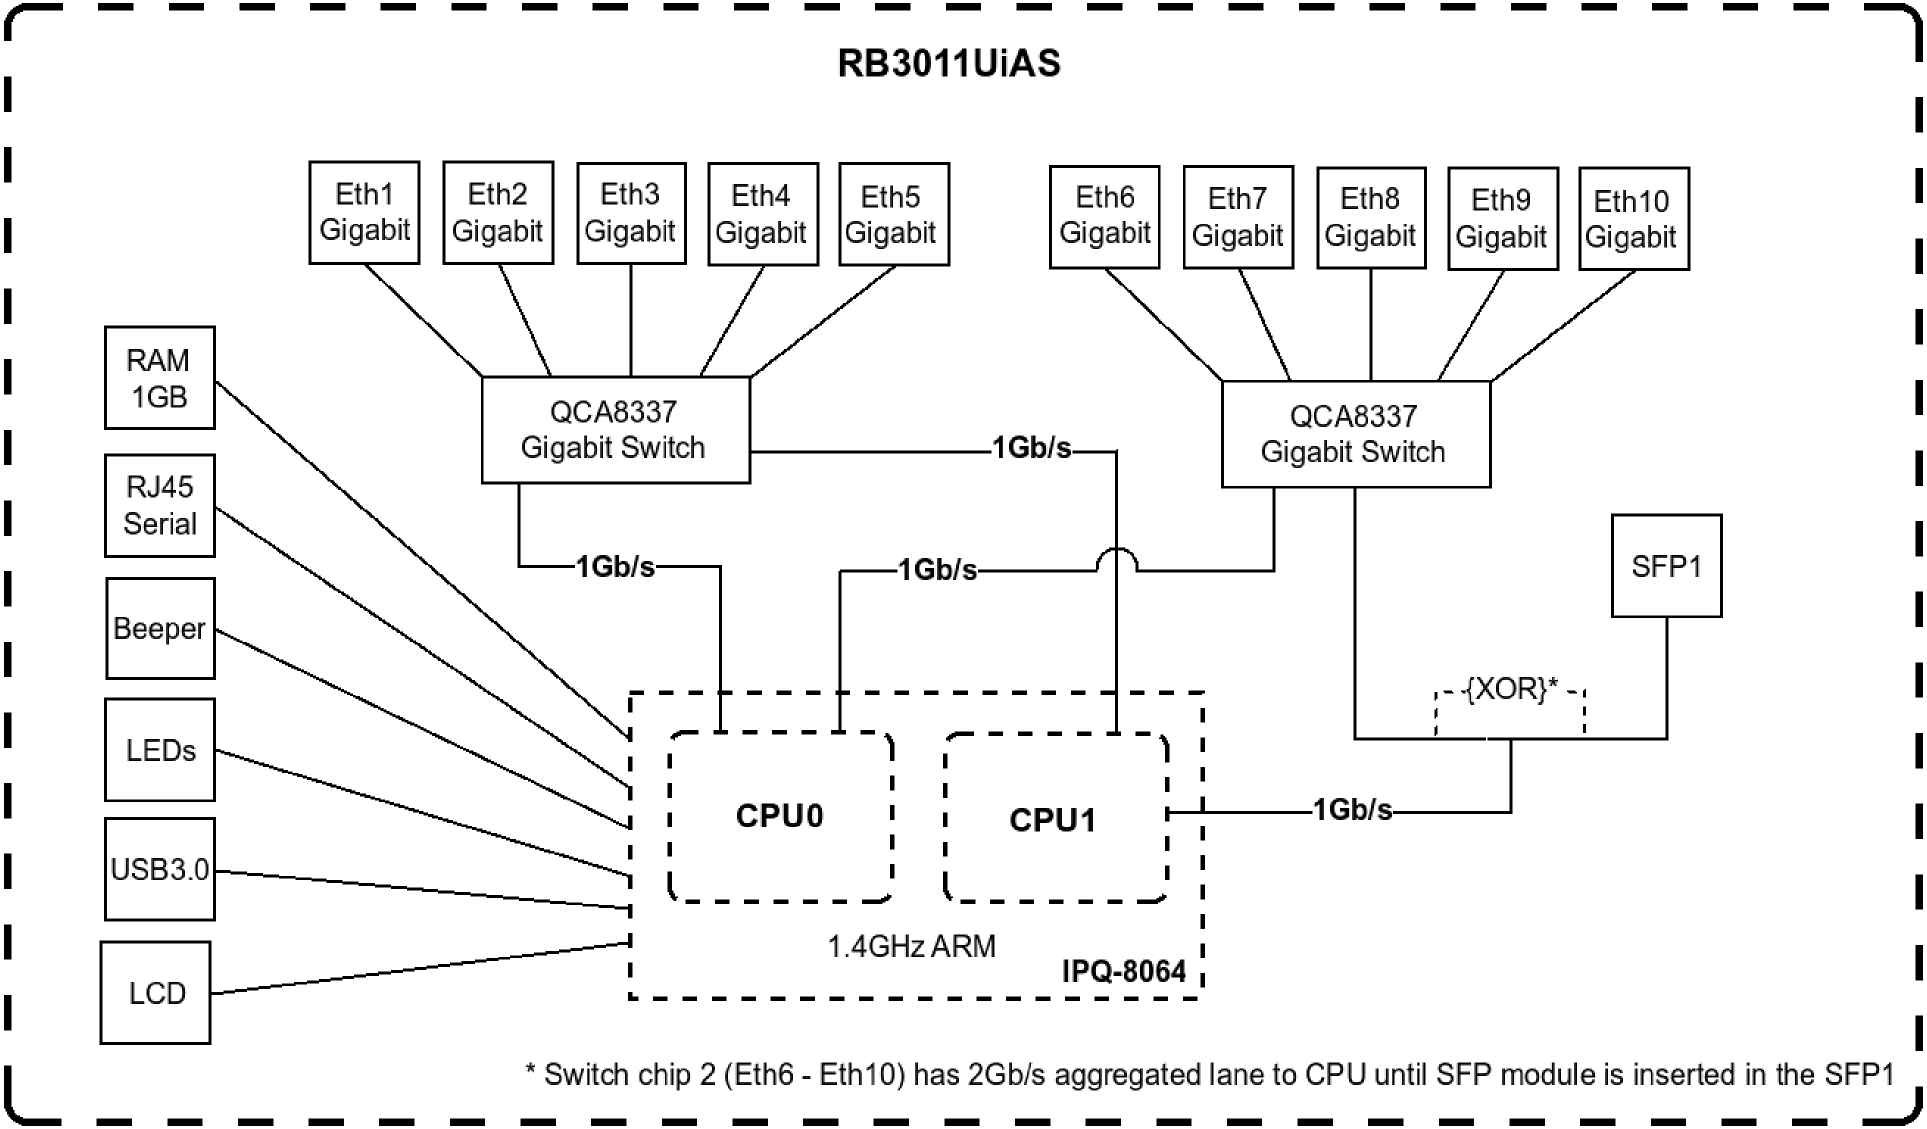

* Switch chip 2 (Eth6 - Eth10) has 2Gb/s aggregated lane to CPU until SFP module is inserted in the SFP1

Analysis of the problem

- Dave’s problem #3:
  - L2TP+IPSec connections are overloading the router, CPU is 100%, throughput is down, but RB3011 should be more powerful.
- Diagnosis:
  - Performance tables of both devices.
- Reason:
  - RB750Gr3 features hardware IPSec acceleration, but RB3011 doesn’t.

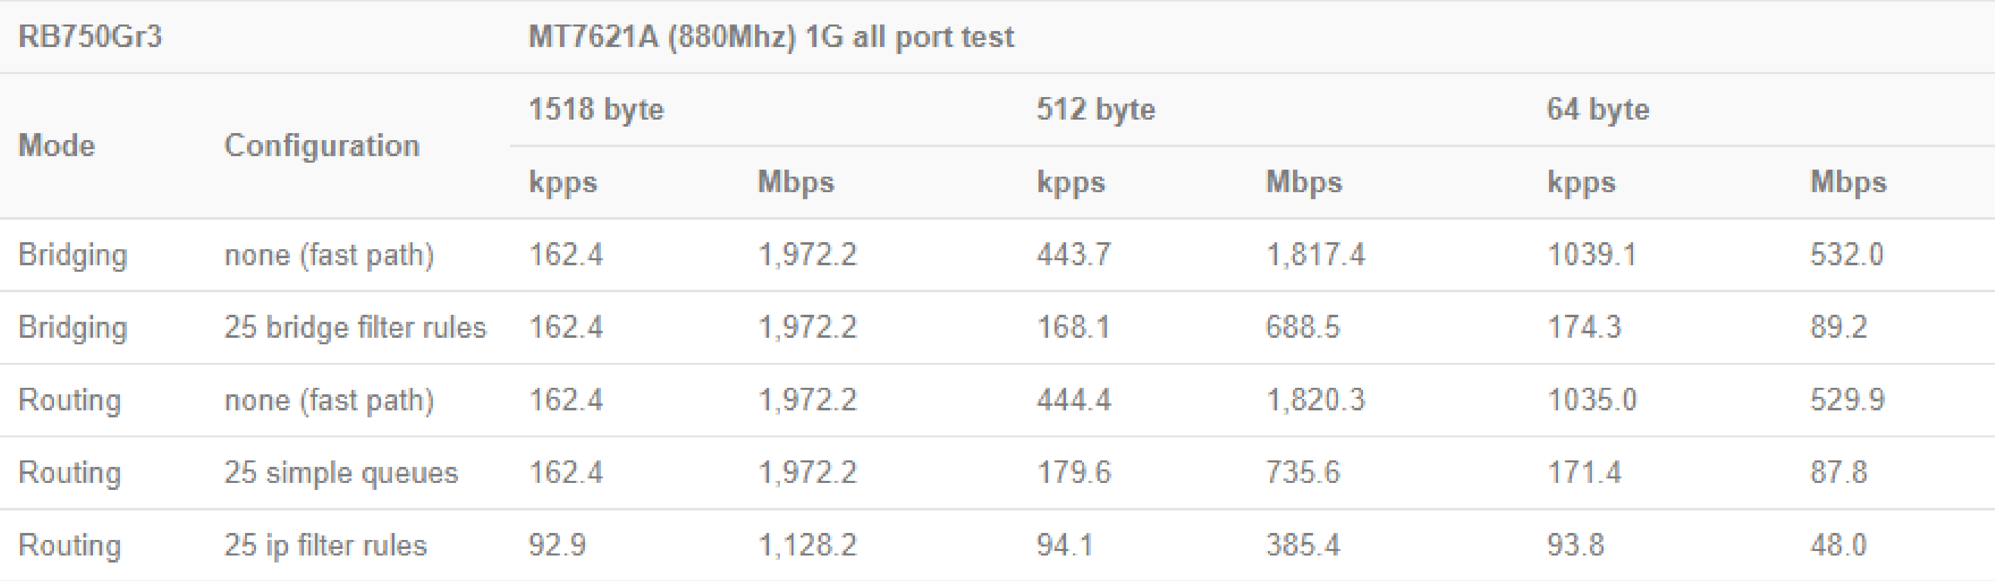

RB750Gr3		MT7621A (880Mhz) 1G all port test					
Mode	Configuration	1518 byte		512 byte		64 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Bridging	none (fast path)	162.4	1,972.2	443.7	1,817.4	1039.1	532.0
Bridging	25 bridge filter rules	162.4	1,972.2	168.1	688.5	174.3	89.2
Routing	none (fast path)	162.4	1,972.2	444.4	1,820.3	1035.0	529.9
Routing	25 simple queues	162.4	1,972.2	179.6	735.6	171.4	87.8
Routing	25 ip filter rules	92.9	1,128.2	94.1	385.4	93.8	48.0

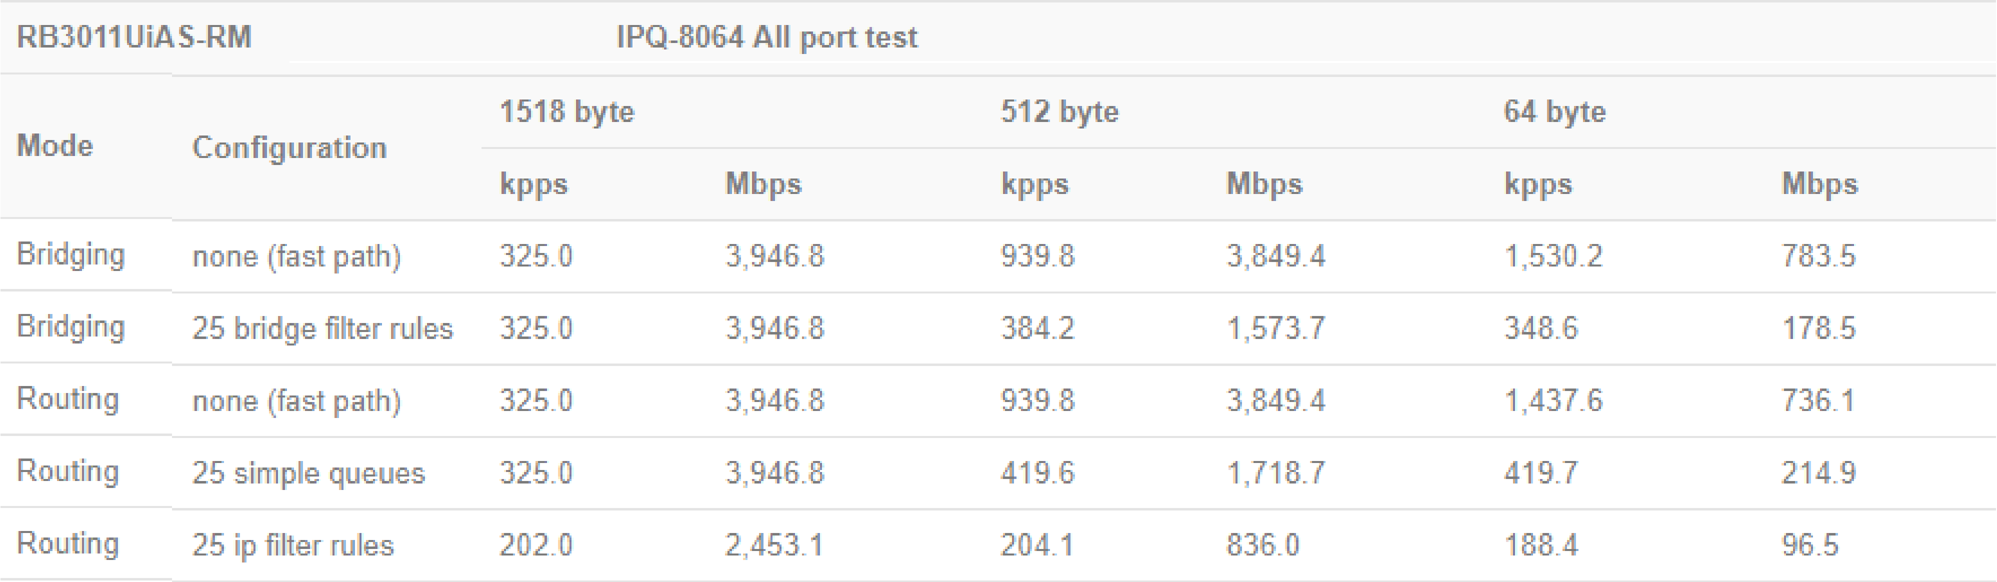

RB3011UiAS-RM		IPQ-8064 All port test					
Mode	Configuration	1518 byte		512 byte		64 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Bridging	none (fast path)	325.0	3,946.8	939.8	3,849.4	1,530.2	783.5
Bridging	25 bridge filter rules	325.0	3,946.8	384.2	1,573.7	348.6	178.5
Routing	none (fast path)	325.0	3,946.8	939.8	3,849.4	1,437.6	736.1
Routing	25 simple queues	325.0	3,946.8	419.6	1,718.7	419.7	214.9
Routing	25 ip filter rules	202.0	2,453.1	204.1	836.0	188.4	96.5

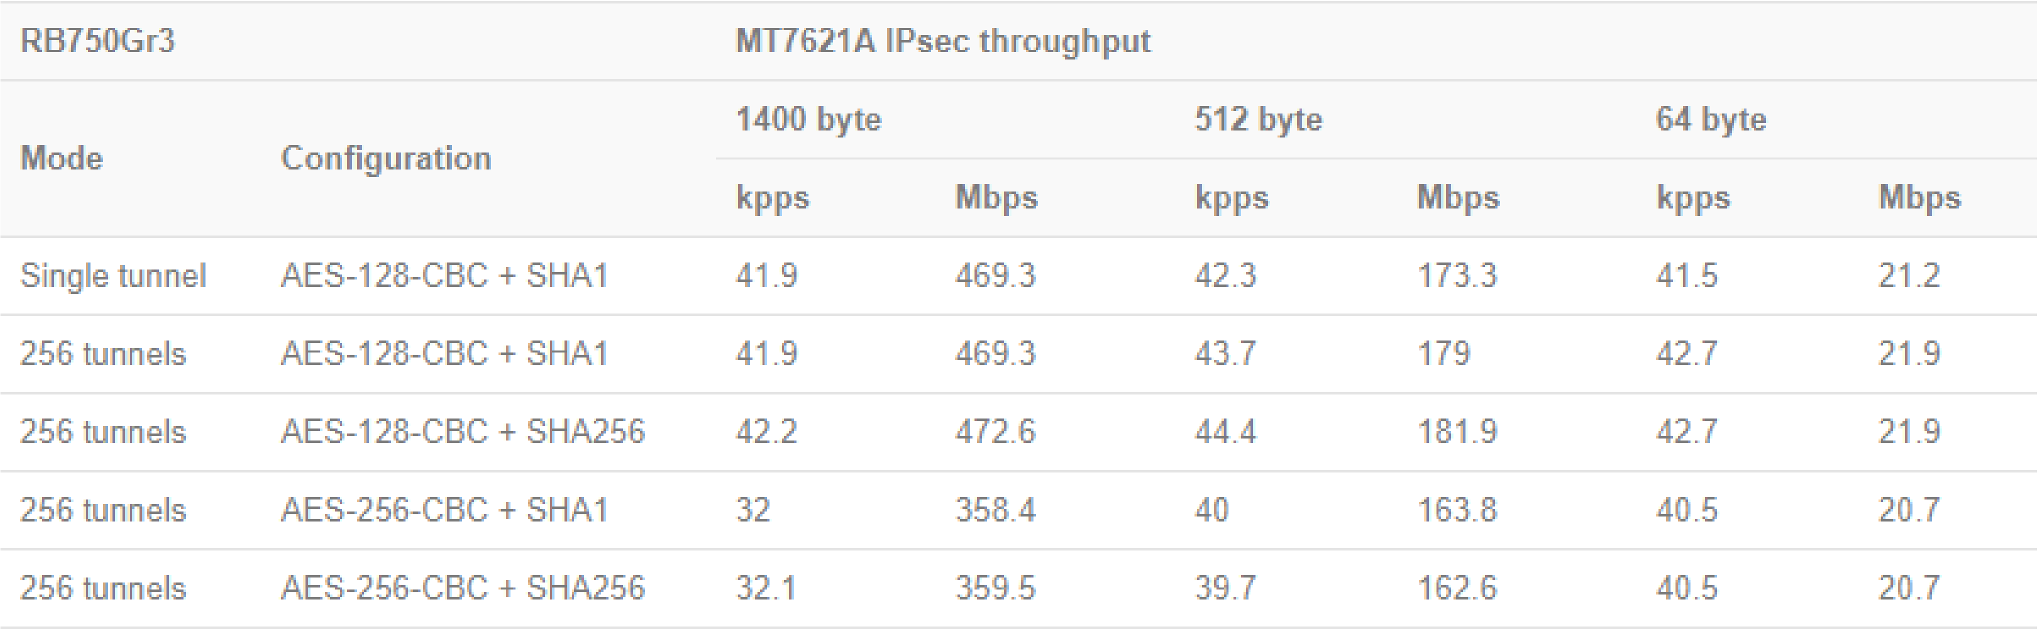

RB750Gr3		MT7621A IPsec throughput					
Mode	Configuration	1400 byte		512 byte		64 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Single tunnel	AES-128-CBC + SHA1	41.9	469.3	42.3	173.3	41.5	21.2
256 tunnels	AES-128-CBC + SHA1	41.9	469.3	43.7	179	42.7	21.9
256 tunnels	AES-128-CBC + SHA256	42.2	472.6	44.4	181.9	42.7	21.9
256 tunnels	AES-256-CBC + SHA1	32	358.4	40	163.8	40.5	20.7
256 tunnels	AES-256-CBC + SHA256	32.1	359.5	39.7	162.6	40.5	20.7

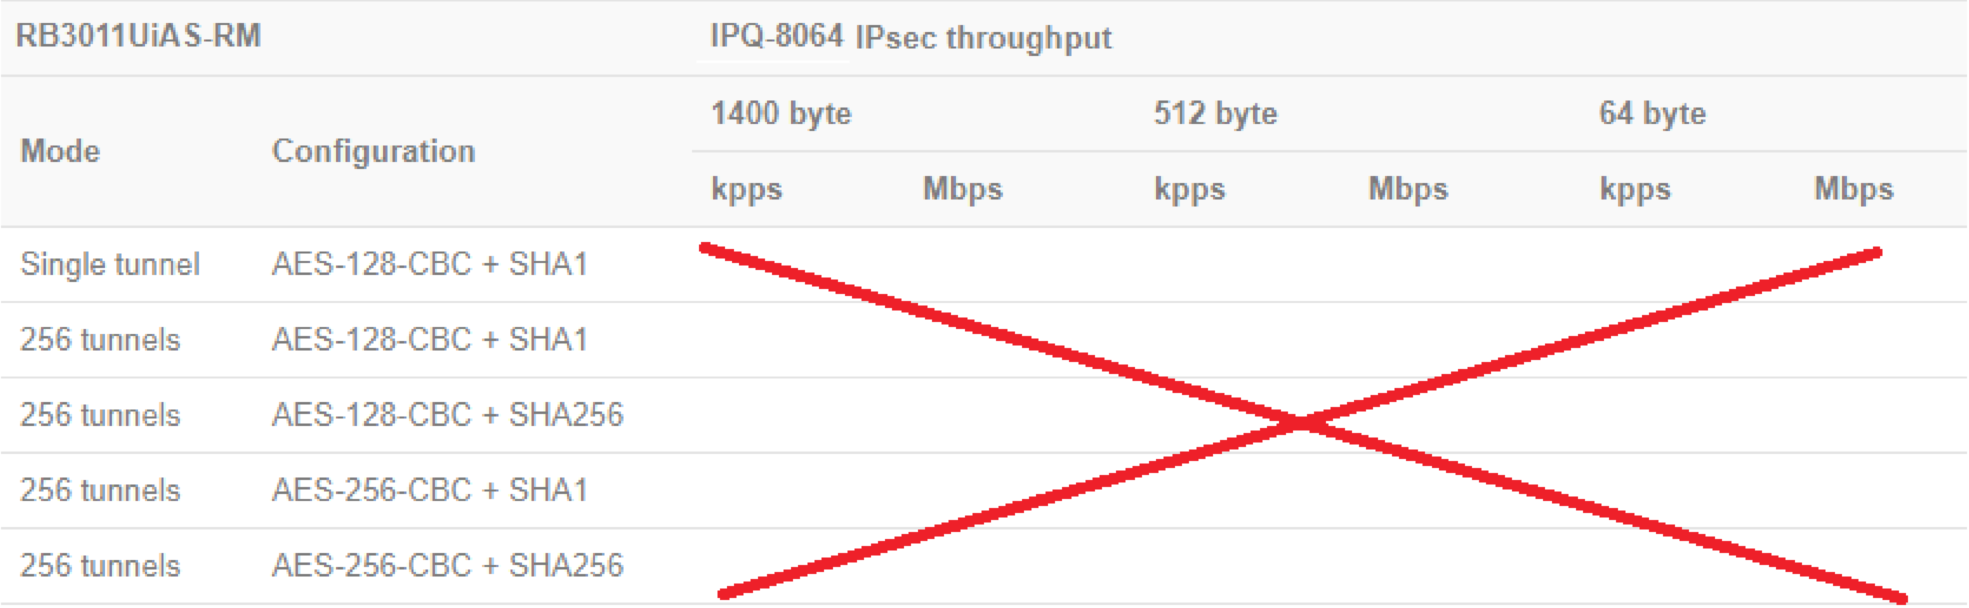

RB3011UiAS-RM		IPQ-8064 IPsec throughput					
Mode	Configuration	1400 byte		512 byte		64 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Single tunnel	AES-128-CBC + SHA1						
256 tunnels	AES-128-CBC + SHA1						
256 tunnels	AES-128-CBC + SHA256						
256 tunnels	AES-256-CBC + SHA1						
256 tunnels	AES-256-CBC + SHA256						

Improving the temporary fix

- Dave examines performance and IPSec hardware encryption performance numbers and decides to replace the RB750Gr3 with an RB1100AHx2.
- Dave examines block diagram for switching bottlenecks on RB1100AHx2 and decides to put most demanding throughput on ether11, ether12, ether13.

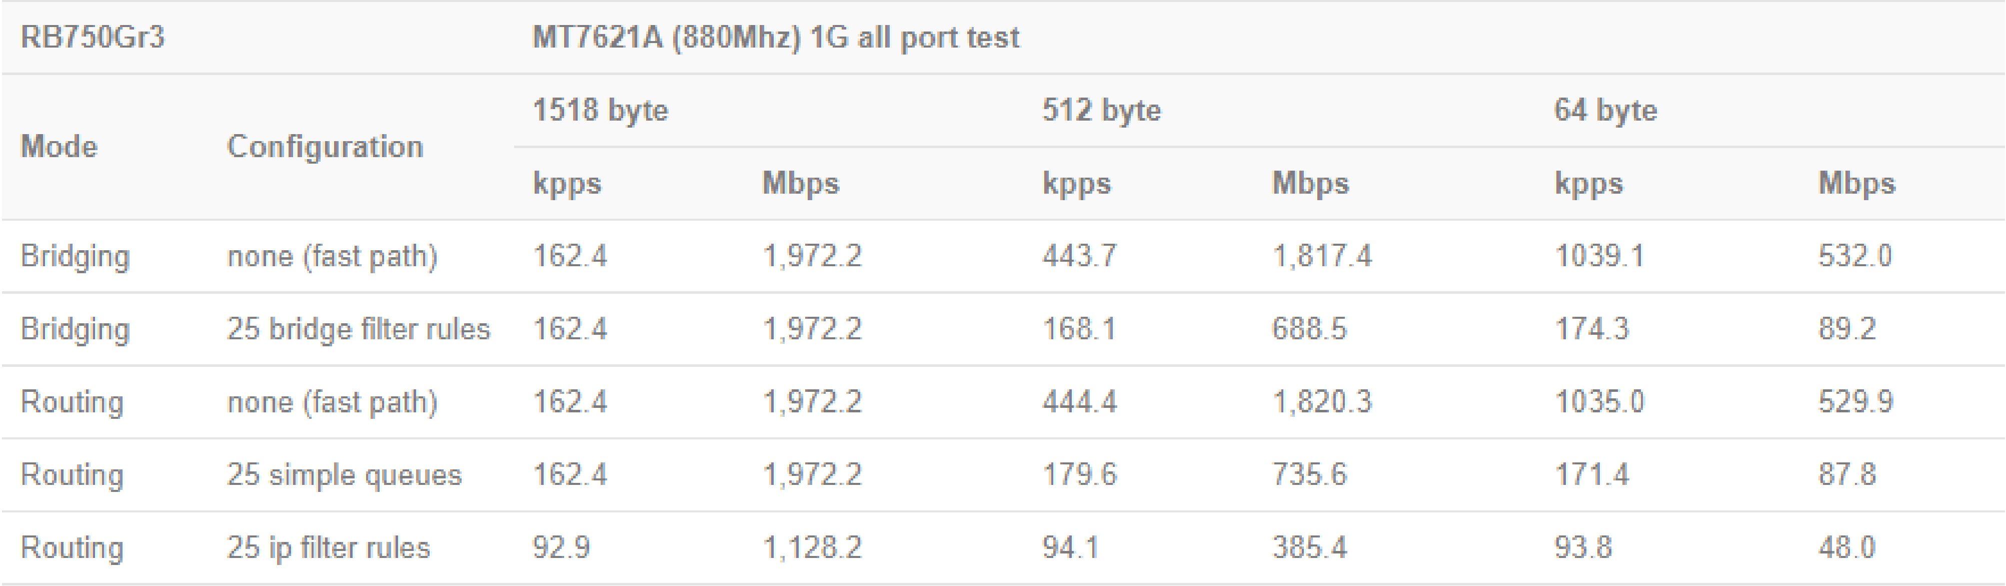

RB750Gr3		MT7621A (880Mhz) 1G all port test					
Mode	Configuration	1518 byte		512 byte		64 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Bridging	none (fast path)	162.4	1,972.2	443.7	1,817.4	1039.1	532.0
Bridging	25 bridge filter rules	162.4	1,972.2	168.1	688.5	174.3	89.2
Routing	none (fast path)	162.4	1,972.2	444.4	1,820.3	1035.0	529.9
Routing	25 simple queues	162.4	1,972.2	179.6	735.6	171.4	87.8
Routing	25 ip filter rules	92.9	1,128.2	94.1	385.4	93.8	48.0

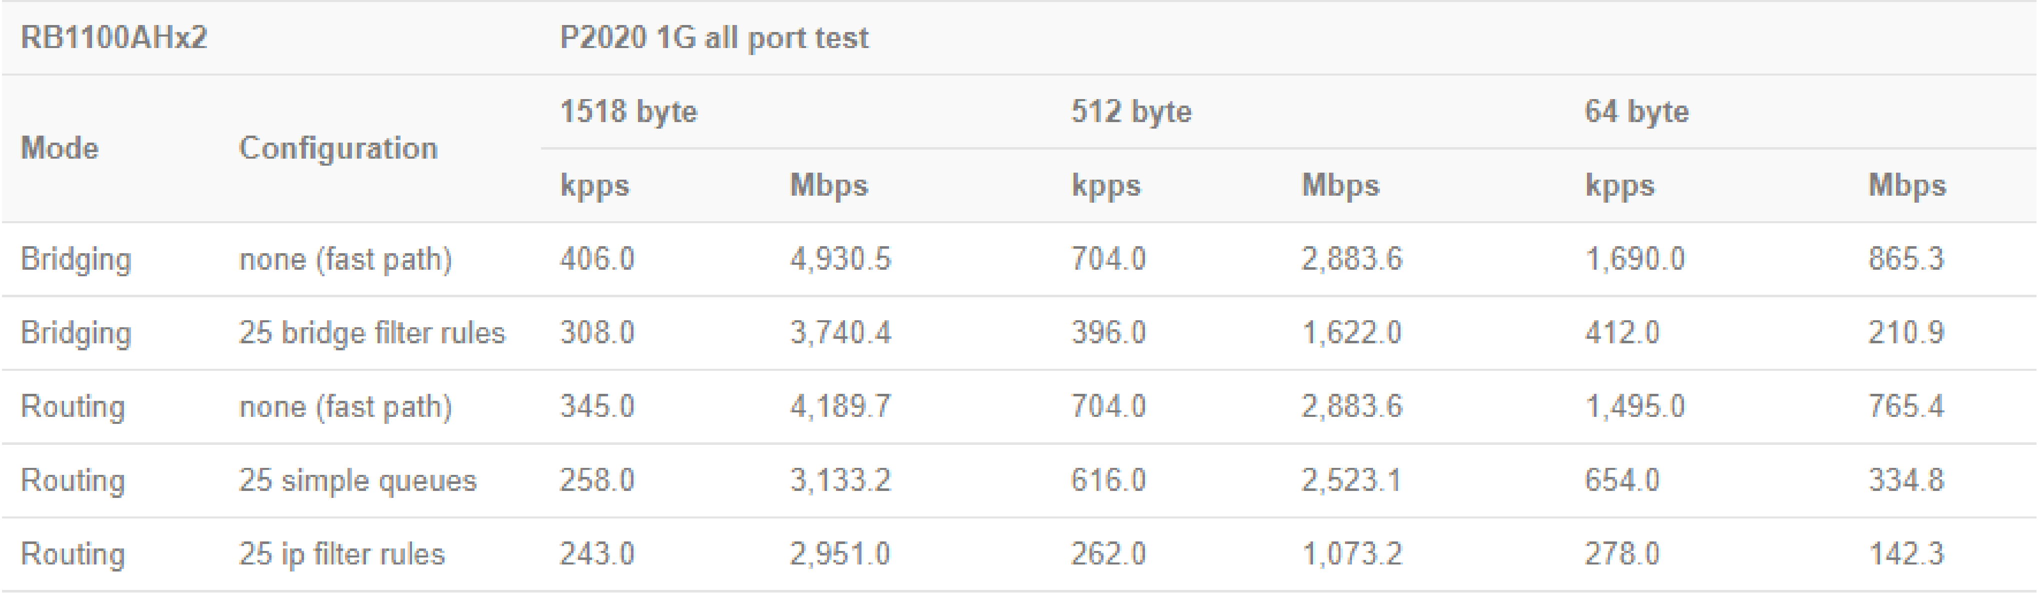

RB1100AHx2		P2020 1G all port test					
Mode	Configuration	1518 byte		512 byte		64 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Bridging	none (fast path)	406.0	4,930.5	704.0	2,883.6	1,690.0	865.3
Bridging	25 bridge filter rules	308.0	3,740.4	396.0	1,622.0	412.0	210.9
Routing	none (fast path)	345.0	4,189.7	704.0	2,883.6	1,495.0	765.4
Routing	25 simple queues	258.0	3,133.2	616.0	2,523.1	654.0	334.8
Routing	25 ip filter rules	243.0	2,951.0	262.0	1,073.2	278.0	142.3

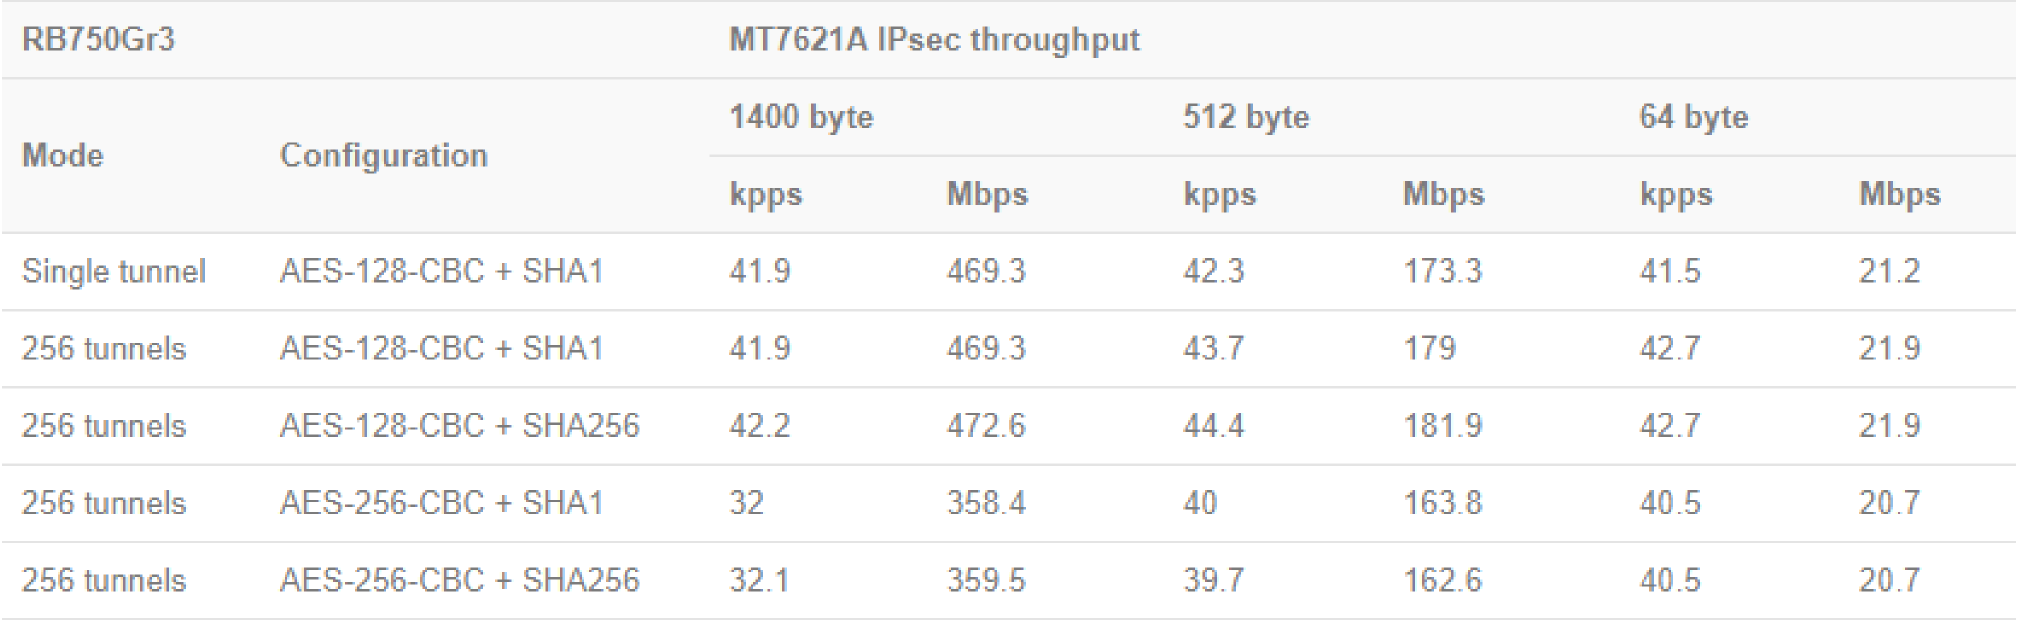

RB750Gr3		MT7621A IPsec throughput					
Mode	Configuration	1400 byte		512 byte		64 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Single tunnel	AES-128-CBC + SHA1	41.9	469.3	42.3	173.3	41.5	21.2
256 tunnels	AES-128-CBC + SHA1	41.9	469.3	43.7	179	42.7	21.9
256 tunnels	AES-128-CBC + SHA256	42.2	472.6	44.4	181.9	42.7	21.9
256 tunnels	AES-256-CBC + SHA1	32	358.4	40	163.8	40.5	20.7
256 tunnels	AES-256-CBC + SHA256	32.1	359.5	39.7	162.6	40.5	20.7

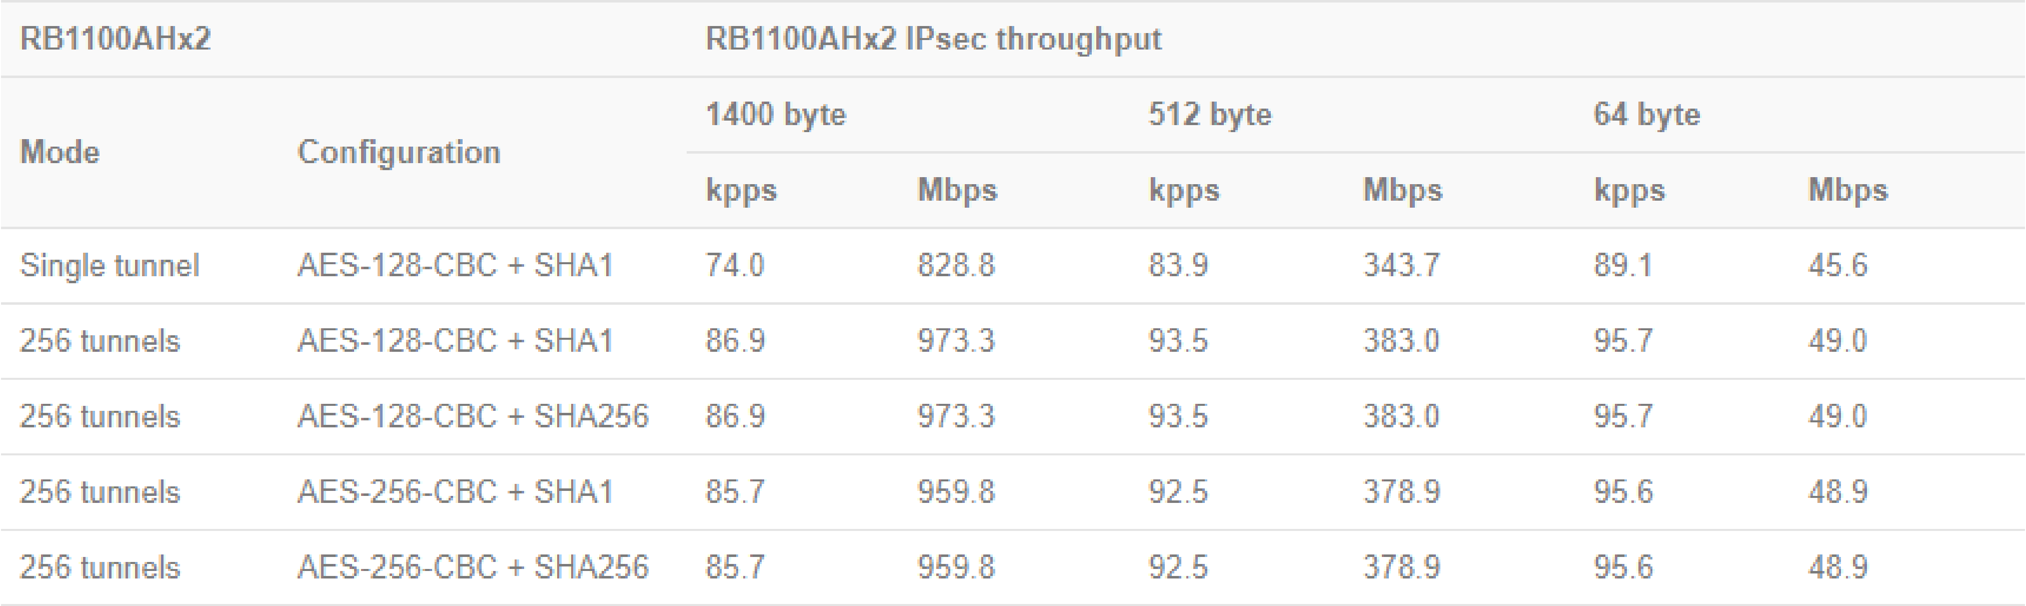

RB1100AHx2		RB1100AHx2 IPsec throughput					
Mode	Configuration	1400 byte		512 byte		64 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Single tunnel	AES-128-CBC + SHA1	74.0	828.8	83.9	343.7	89.1	45.6
256 tunnels	AES-128-CBC + SHA1	86.9	973.3	93.5	383.0	95.7	49.0
256 tunnels	AES-128-CBC + SHA256	86.9	973.3	93.5	383.0	95.7	49.0
256 tunnels	AES-256-CBC + SHA1	85.7	959.8	92.5	378.9	95.6	48.9
256 tunnels	AES-256-CBC + SHA256	85.7	959.8	92.5	378.9	95.6	48.9

Analysis of the problem

- Dave’s problem #4:
  - RB1100AHx2 doesn’t perform as expected, performance is not better, but even worse than RB750Gx3 on ether12 and ether13 ports.
- Diagnosis:
  - Block diagram of RB1100AHx2.
- Reason:
  - Management and emergency bypass port is used for main traffic management.

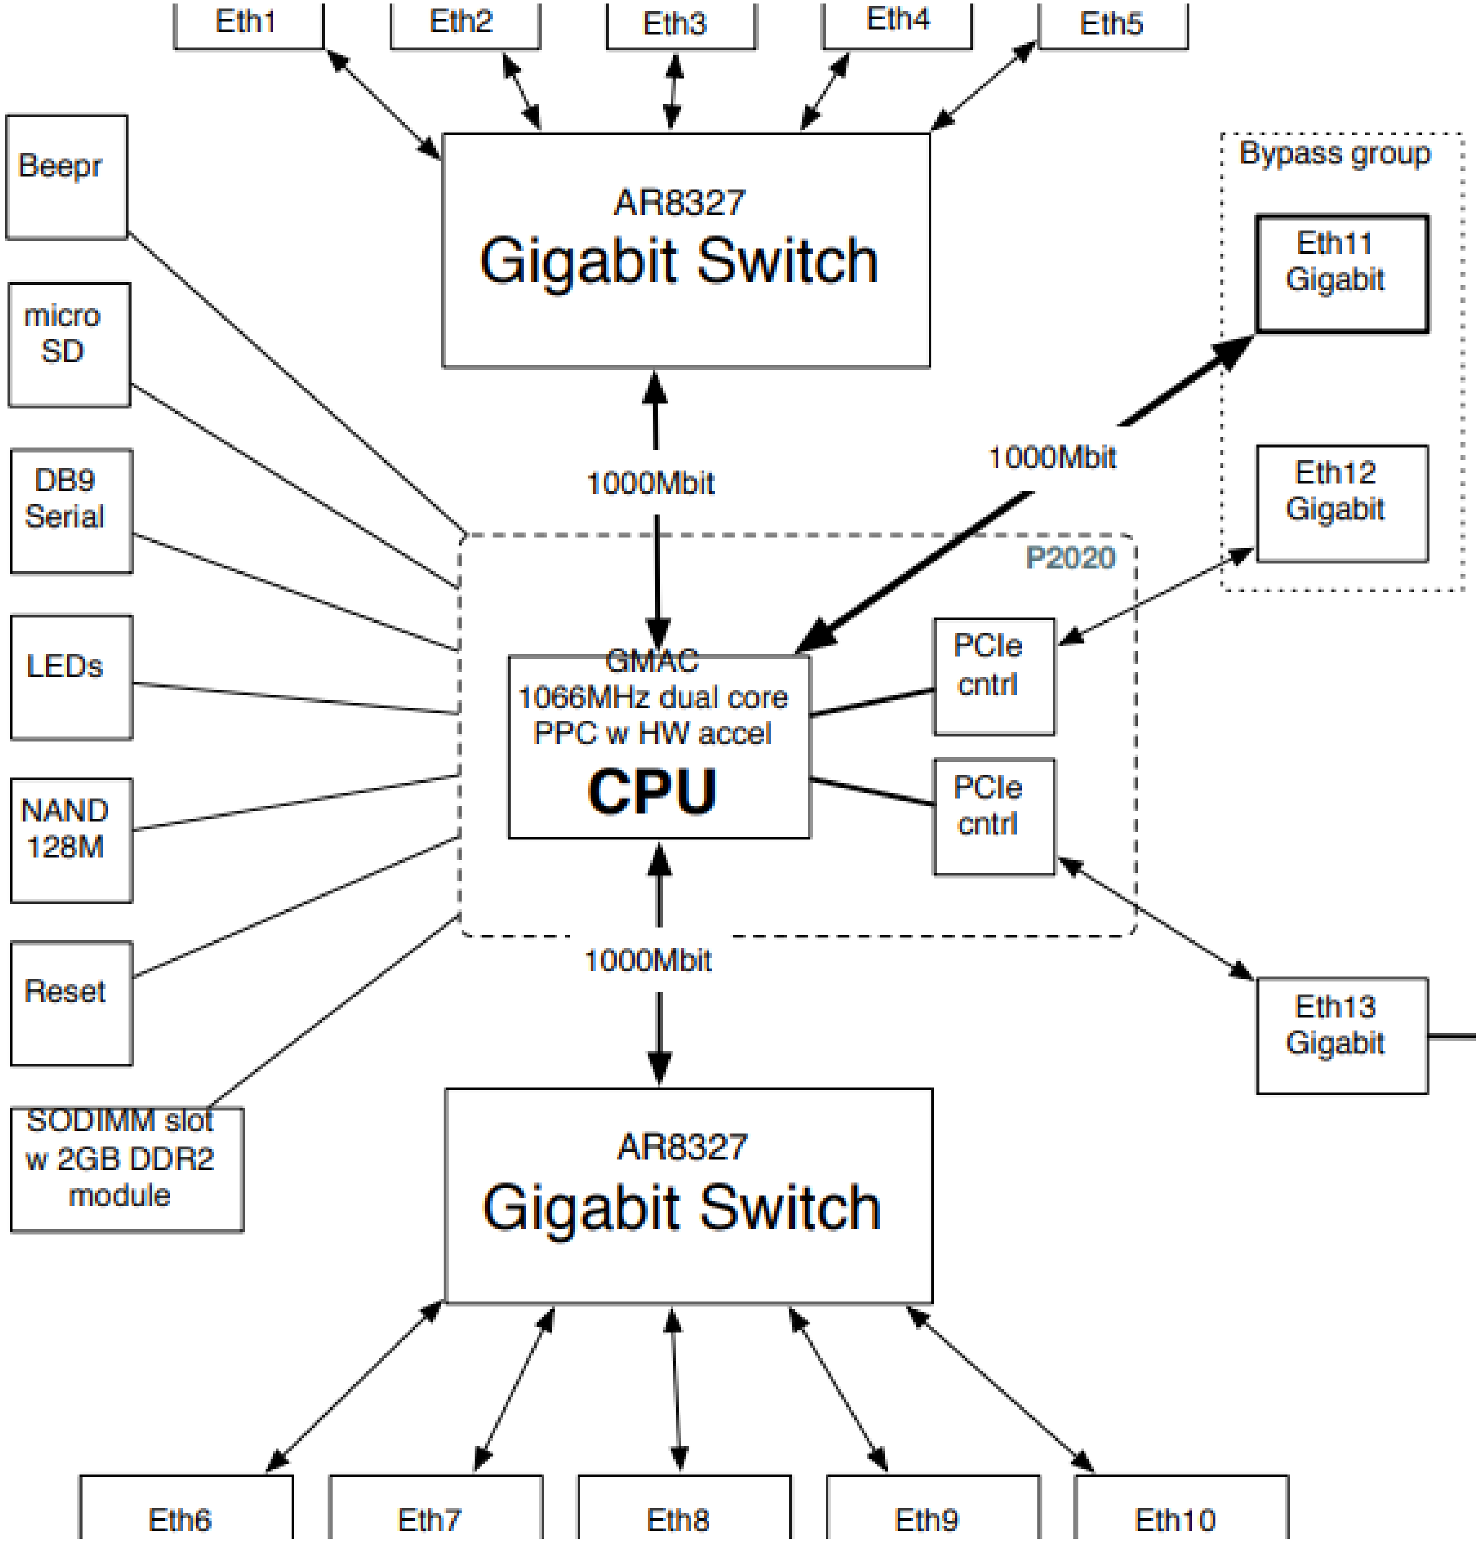
Eth1
Eth2
Eth3
Eth4
Eth5
Beepr
AR8327
Gigabit Switch
Bypass group
Eth11
Gigabit
micro
SD
1000Mbit
1000Mbit
DB9
Serial
Eth12
Gigabit
P2020
LEDs
GMAC
1066MHz dual core
PPC w HW accel
CPU
PCIe
cntrl
PCIe
cntrl
NAND
128M
Reset
1000Mbit
Eth13
Gigabit
SODIMM slot
w 2GB DDR2
module
AR8327
Gigabit Switch
Eth6
Eth7
Eth8
Eth9
Eth10

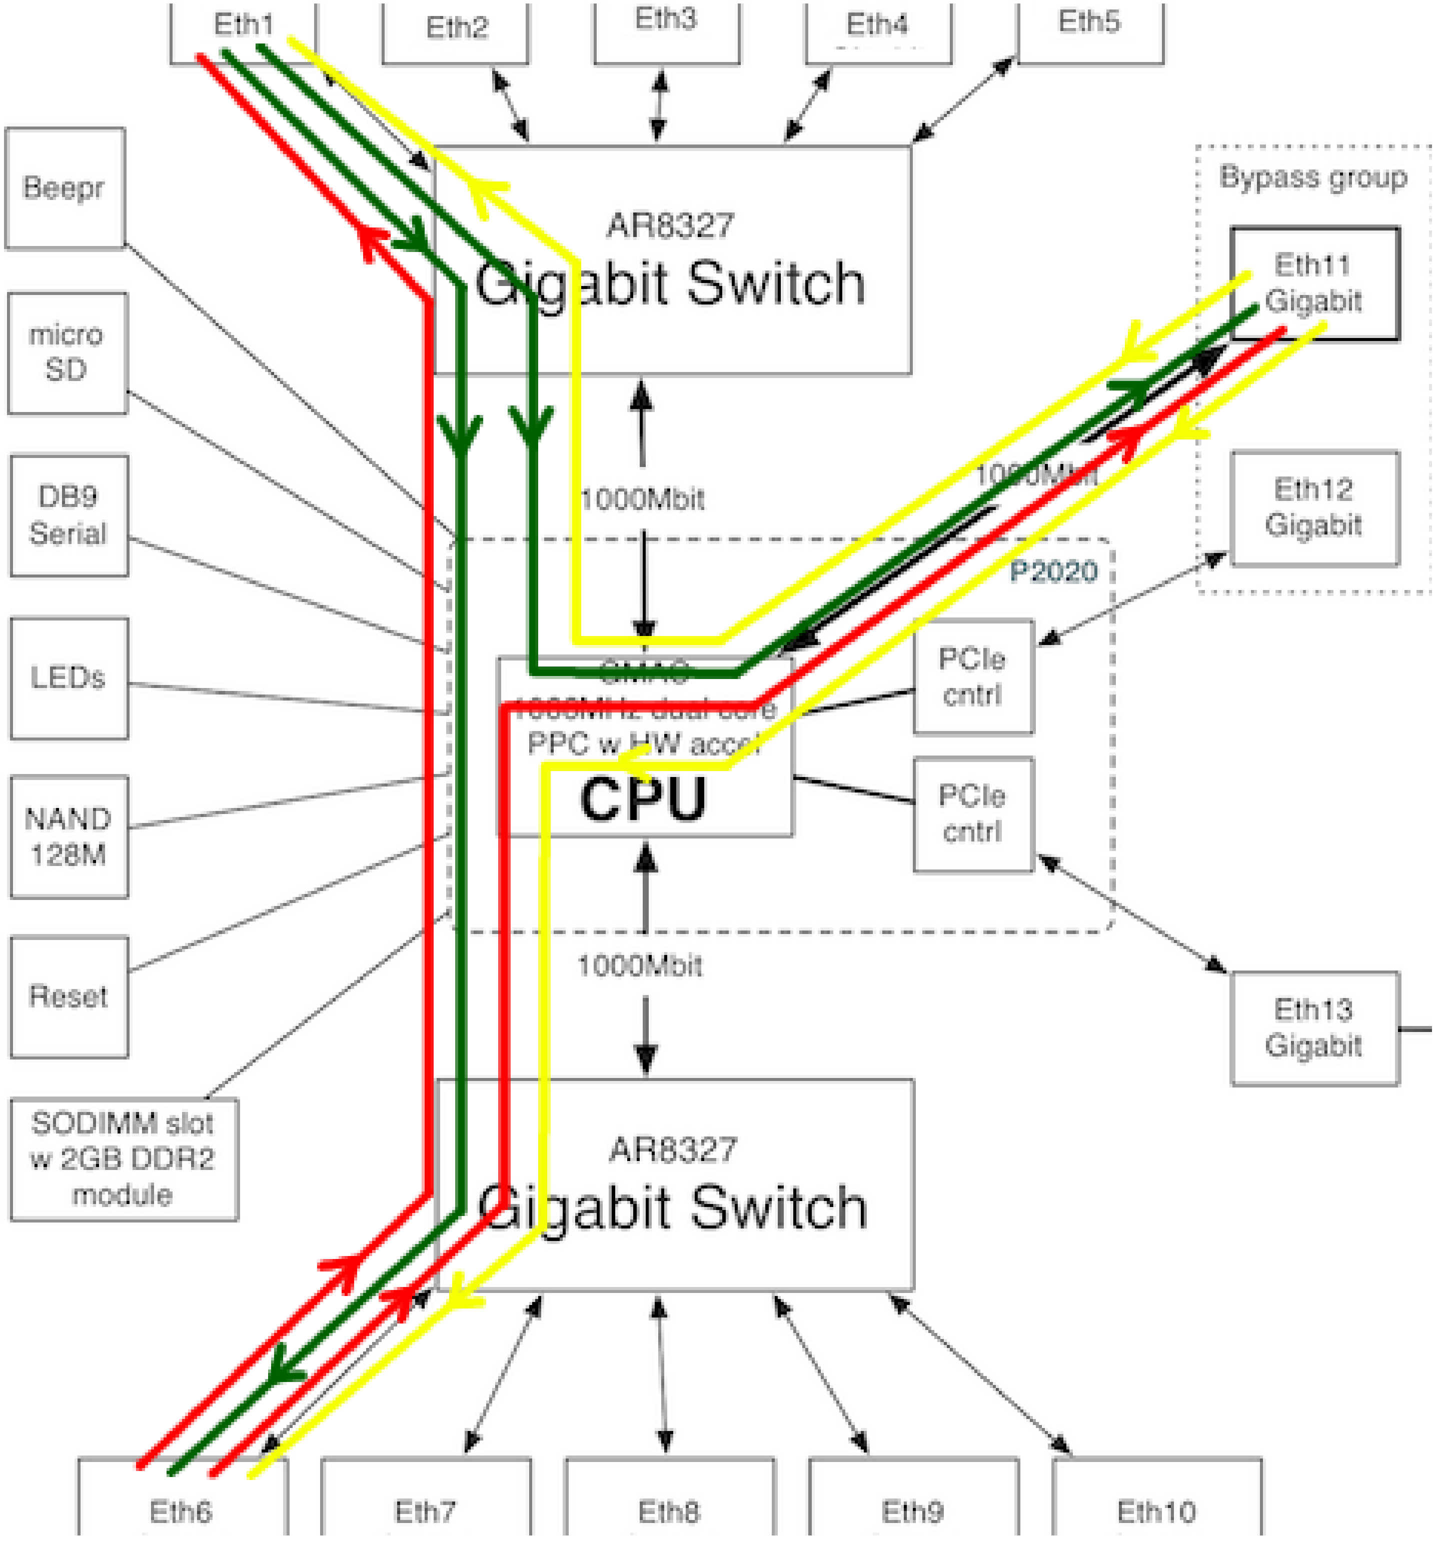
Eth1
Eth2
Eth3
Eth4
Eth5
Beepr
micro SD
DB9 Serial
LEDs
NAND 128M
Reset
SODIMM slot w 2GB DDR2 module
AR8327
Gigabit Switch
1000Mbit
Bypass group
Eth11 Gigabit
Eth12 Gigabit
1000Mbit
P2020
GMAC
1000MHz dual core PPC w HW accel
CPU
PCIe cntrl
PCIe cntrl
Eth13 Gigabit
1000Mbit
AR8327
Gigabit Switch
Eth6
Eth7
Eth8
Eth9
Eth10

Buying the right hardware

- Dave now uses all his experience and selects the perfect hardware for his permanent fix – RB1100AHx4.
- Dave starts to investigate other places where he can place MikroTik hardware in his network.
- Dave continues to write to support@mikrotik.com.

RB1100AHx4 block diagram

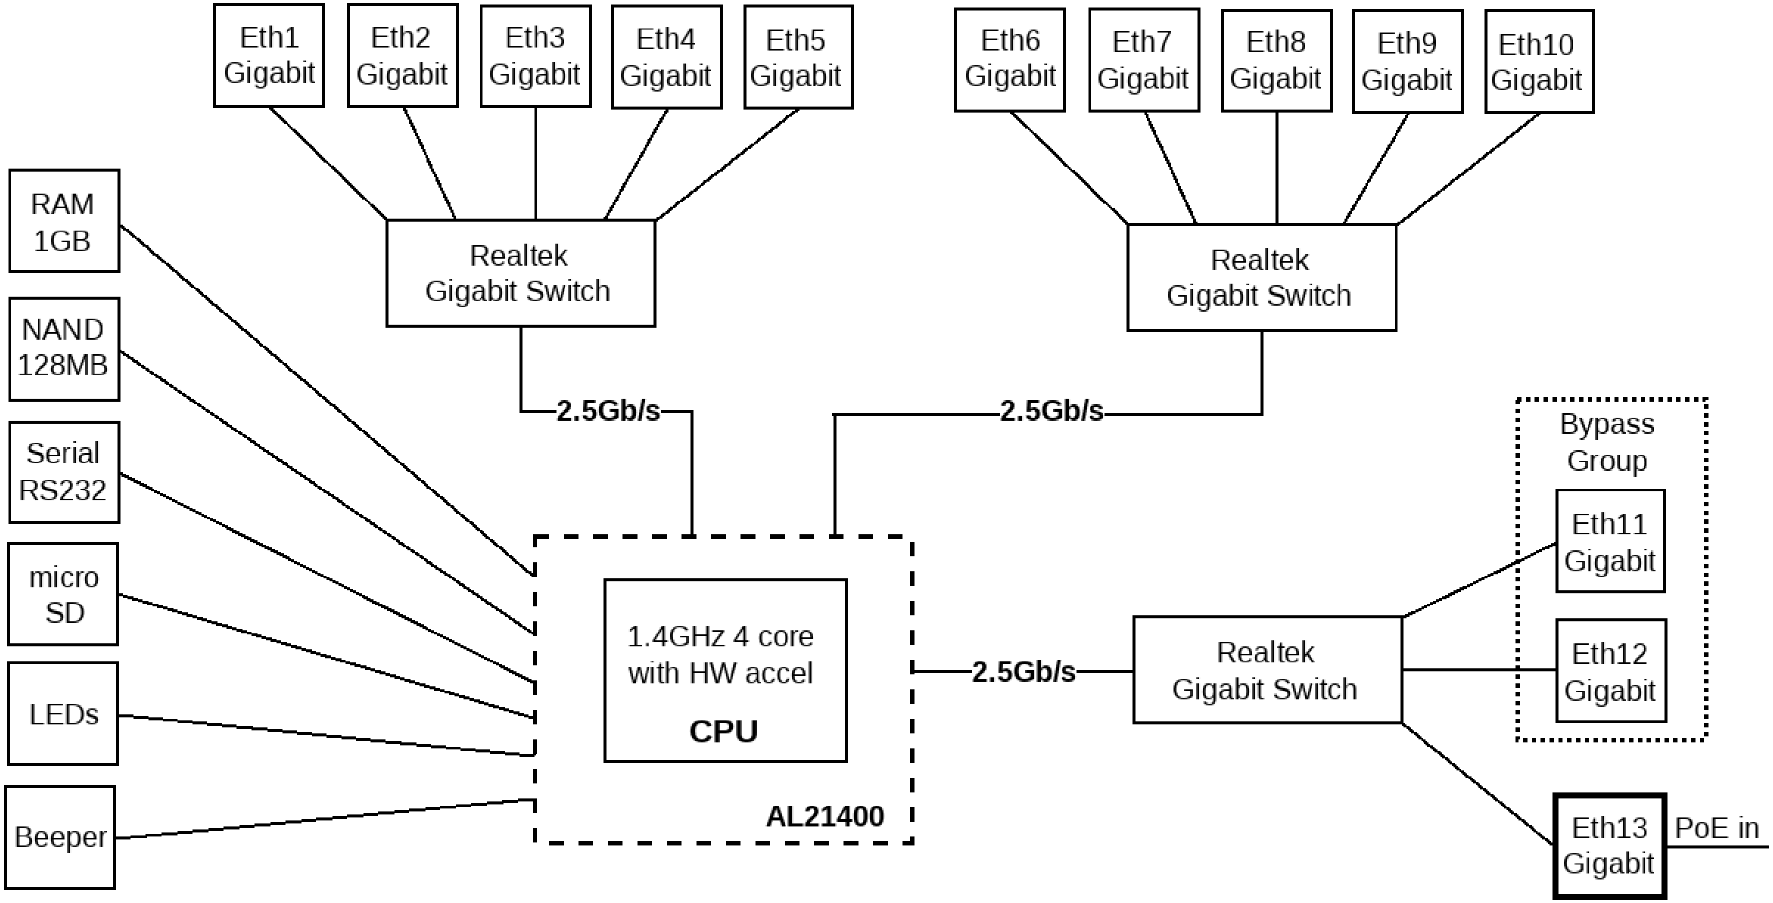

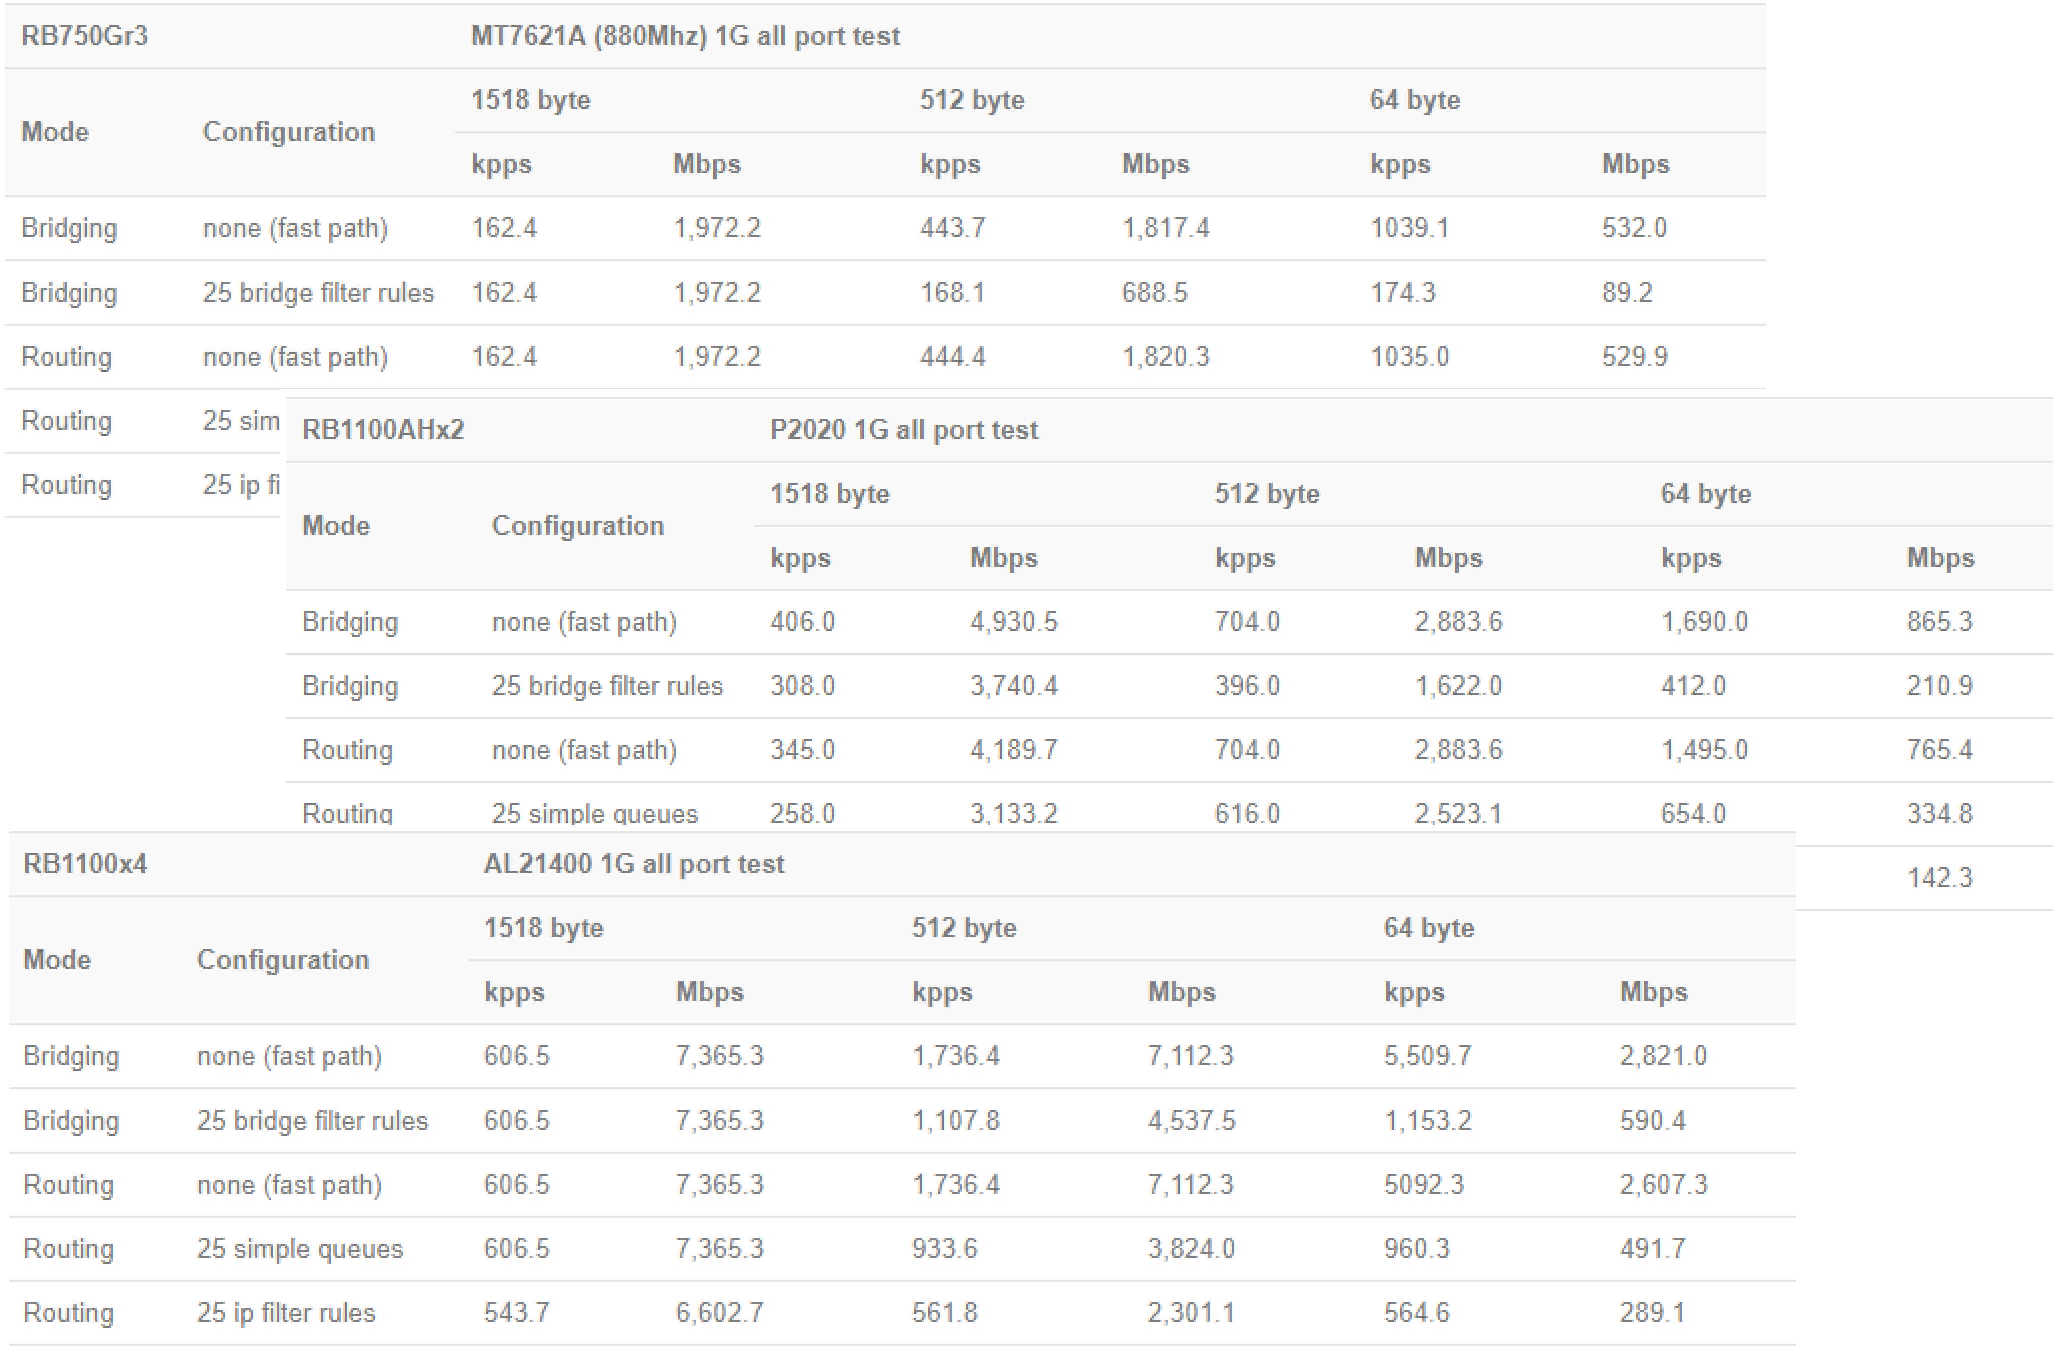

RB750Gr3		MT7621A (880Mhz) 1G all port test					
Mode	Configuration	1518 byte		512 byte		64 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Bridging	none (fast path)	162.4	1,972.2	443.7	1,817.4	1039.1	532.0
Bridging	25 bridge filter rules	162.4	1,972.2	168.1	688.5	174.3	89.2
Routing	none (fast path)	162.4	1,972.2	444.4	1,820.3	1035.0	529.9
Routing	25 sim						
Routing	25 ip fi						

RB1100AHx2		P2020 1G all port test					
Mode	Configuration	1518 byte		512 byte		64 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Bridging	none (fast path)	406.0	4,930.5	704.0	2,883.6	1,690.0	865.3
Bridging	25 bridge filter rules	308.0	3,740.4	396.0	1,622.0	412.0	210.9
Routing	none (fast path)	345.0	4,189.7	704.0	2,883.6	1,495.0	765.4
Routing	25 simple queues	258.0	3,133.2	616.0	2,523.1	654.0	334.8
							142.3

RB1100x4		AL21400 1G all port test					
Mode	Configuration	1518 byte		512 byte		64 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Bridging	none (fast path)	606.5	7,365.3	1,736.4	7,112.3	5,509.7	2,821.0
Bridging	25 bridge filter rules	606.5	7,365.3	1,107.8	4,537.5	1,153.2	590.4
Routing	none (fast path)	606.5	7,365.3	1,736.4	7,112.3	5092.3	2,607.3
Routing	25 simple queues	606.5	7,365.3	933.6	3,824.0	960.3	491.7
Routing	25 ip filter rules	543.7	6,602.7	561.8	2,301.1	564.6	289.1

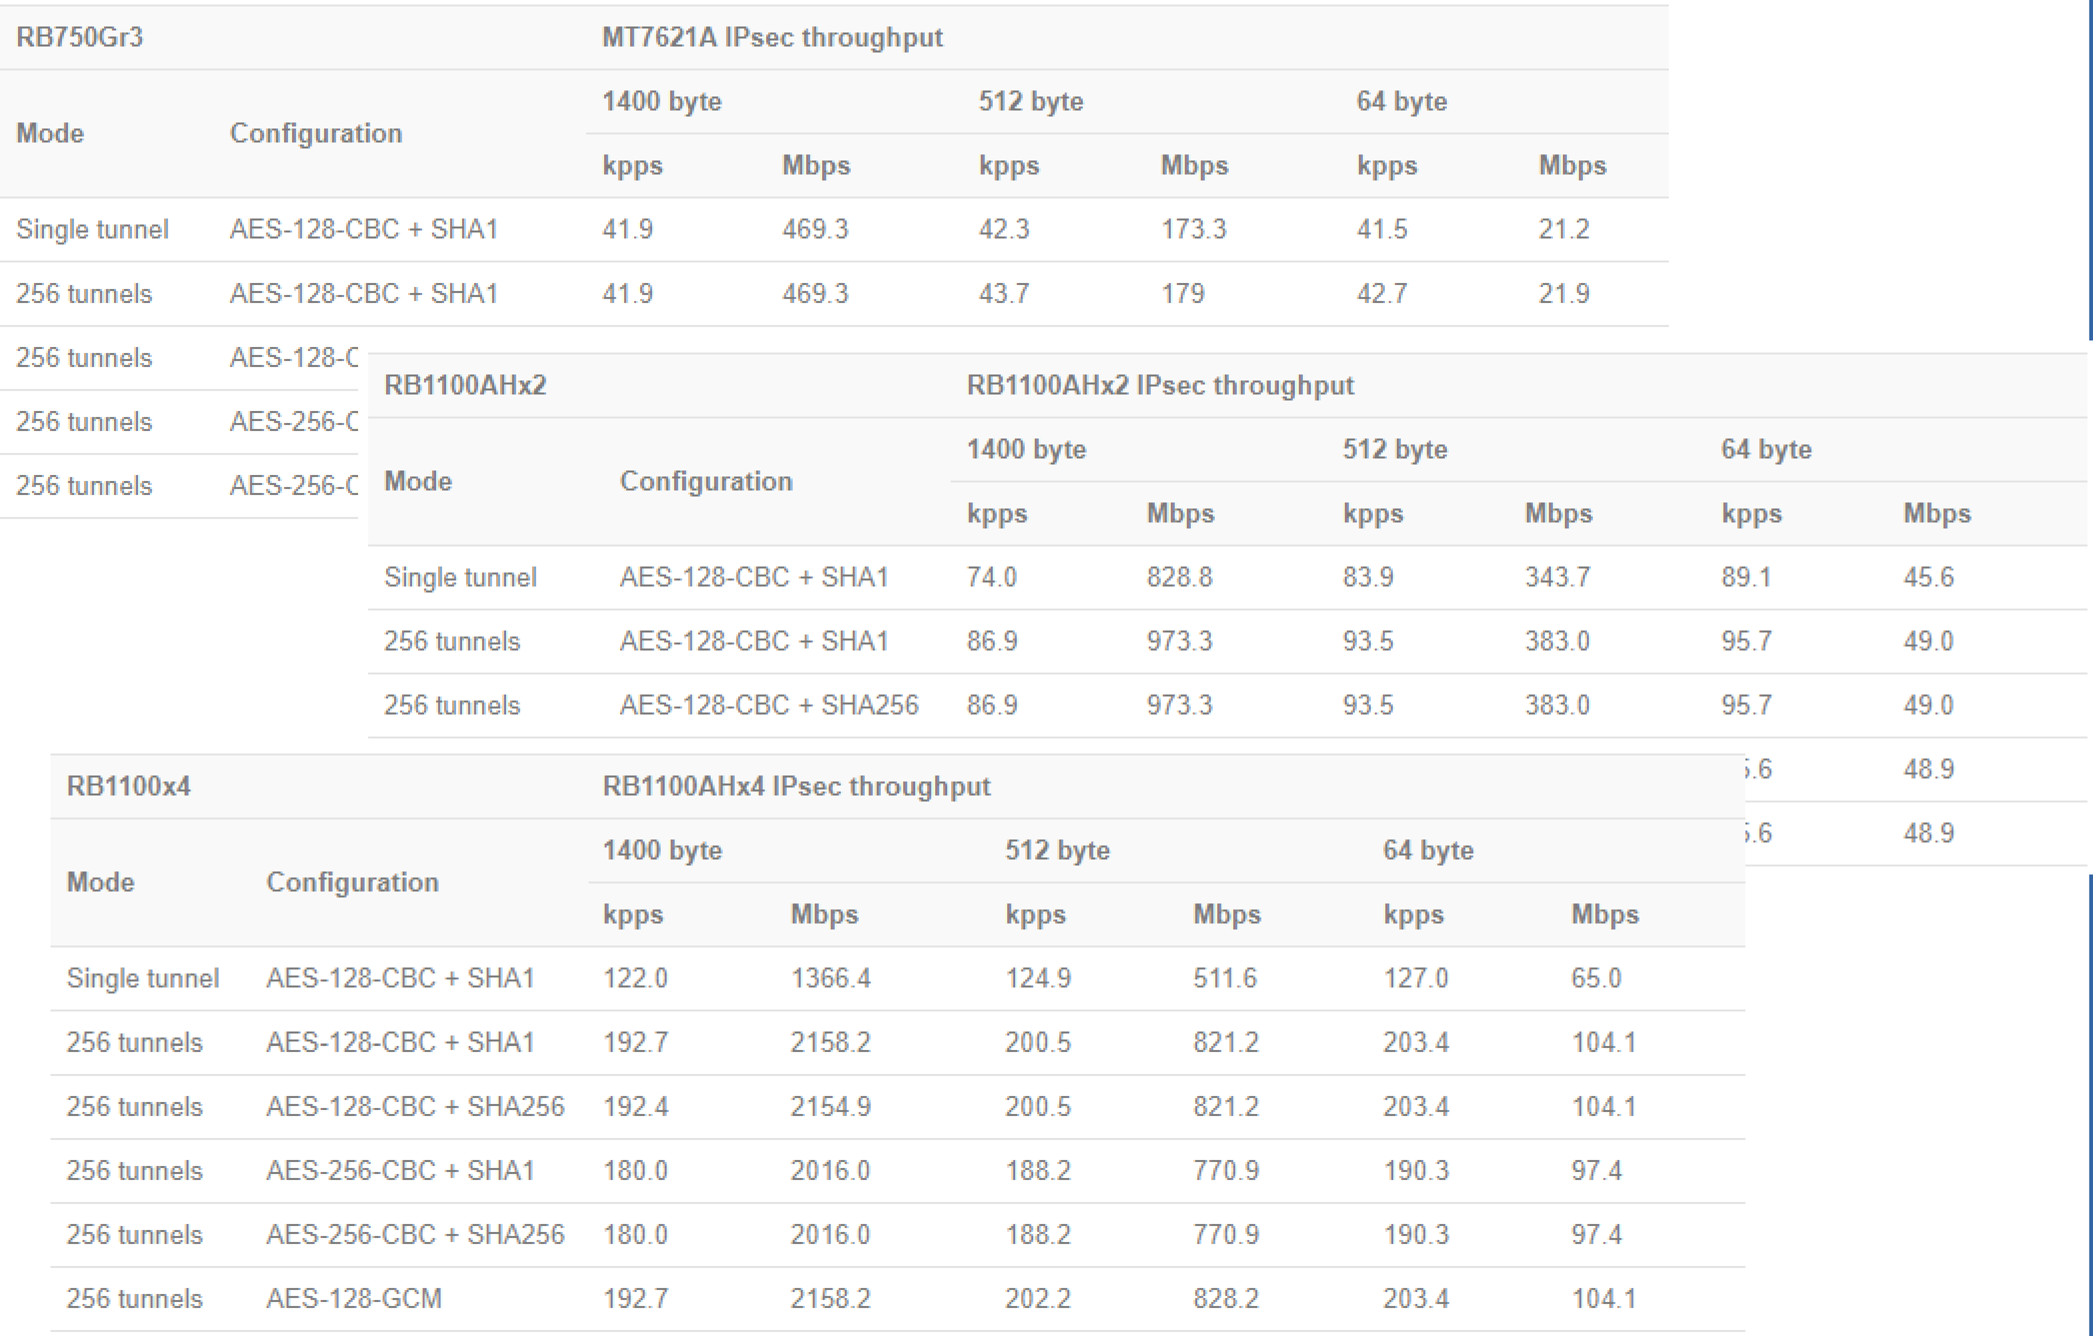

RB750Gr3		MT7621A IPsec throughput					
Mode	Configuration	1400 byte		512 byte		64 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Single tunnel	AES-128-CBC + SHA1	41.9	469.3	42.3	173.3	41.5	21.2
256 tunnels	AES-128-CBC + SHA1	41.9	469.3	43.7	179	42.7	21.9
256 tunnels	AES-128-C						
256 tunnels	AES-256-C						
256 tunnels	AES-256-C						

RB1100AHx2		RB1100AHx2 IPsec throughput					
Mode	Configuration	1400 byte		512 byte		64 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Single tunnel	AES-128-CBC + SHA1	74.0	828.8	83.9	343.7	89.1	45.6
256 tunnels	AES-128-CBC + SHA1	86.9	973.3	93.5	383.0	95.7	49.0
256 tunnels	AES-128-CBC + SHA256	86.9	973.3	93.5	383.0	95.7	49.0
						.6	48.9
						.6	48.9

RB1100x4		RB1100AHx4 IPsec throughput					
Mode	Configuration	1400 byte		512 byte		64 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Single tunnel	AES-128-CBC + SHA1	122.0	1366.4	124.9	511.6	127.0	65.0
256 tunnels	AES-128-CBC + SHA1	192.7	2158.2	200.5	821.2	203.4	104.1
256 tunnels	AES-128-CBC + SHA256	192.4	2154.9	200.5	821.2	203.4	104.1
256 tunnels	AES-256-CBC + SHA1	180.0	2016.0	188.2	770.9	190.3	97.4
256 tunnels	AES-256-CBC + SHA256	180.0	2016.0	188.2	770.9	190.3	97.4
256 tunnels	AES-128-GCM	192.7	2158.2	202.2	828.2	203.4	104.1

Meet Mike

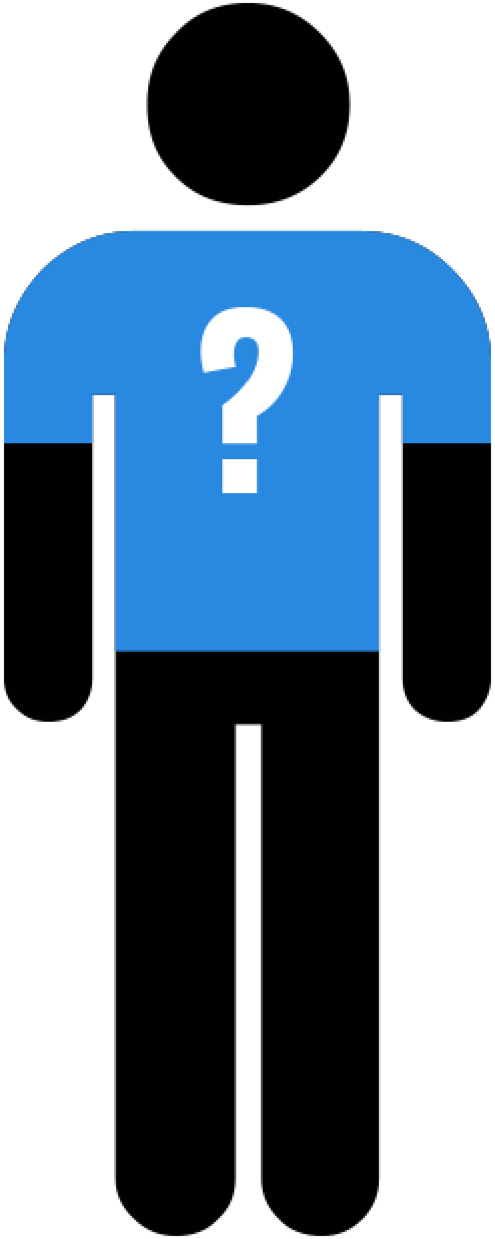

- Mike is a self made businessman with a small office that works with customers on site, several employees, few servers.
- Mike is strong believer in all-in-one solutions, he is looking for one network device that will satisfy all his needs.
- Mike needs an access point for office devices, guest network for customer access, 5 Ethernet ports to connect servers, Internet and a few PCs.
- Mike's friend Dave suggests to get MikroTik hAP ac² *(RBD52G-5HacD2HnD-TC)*.

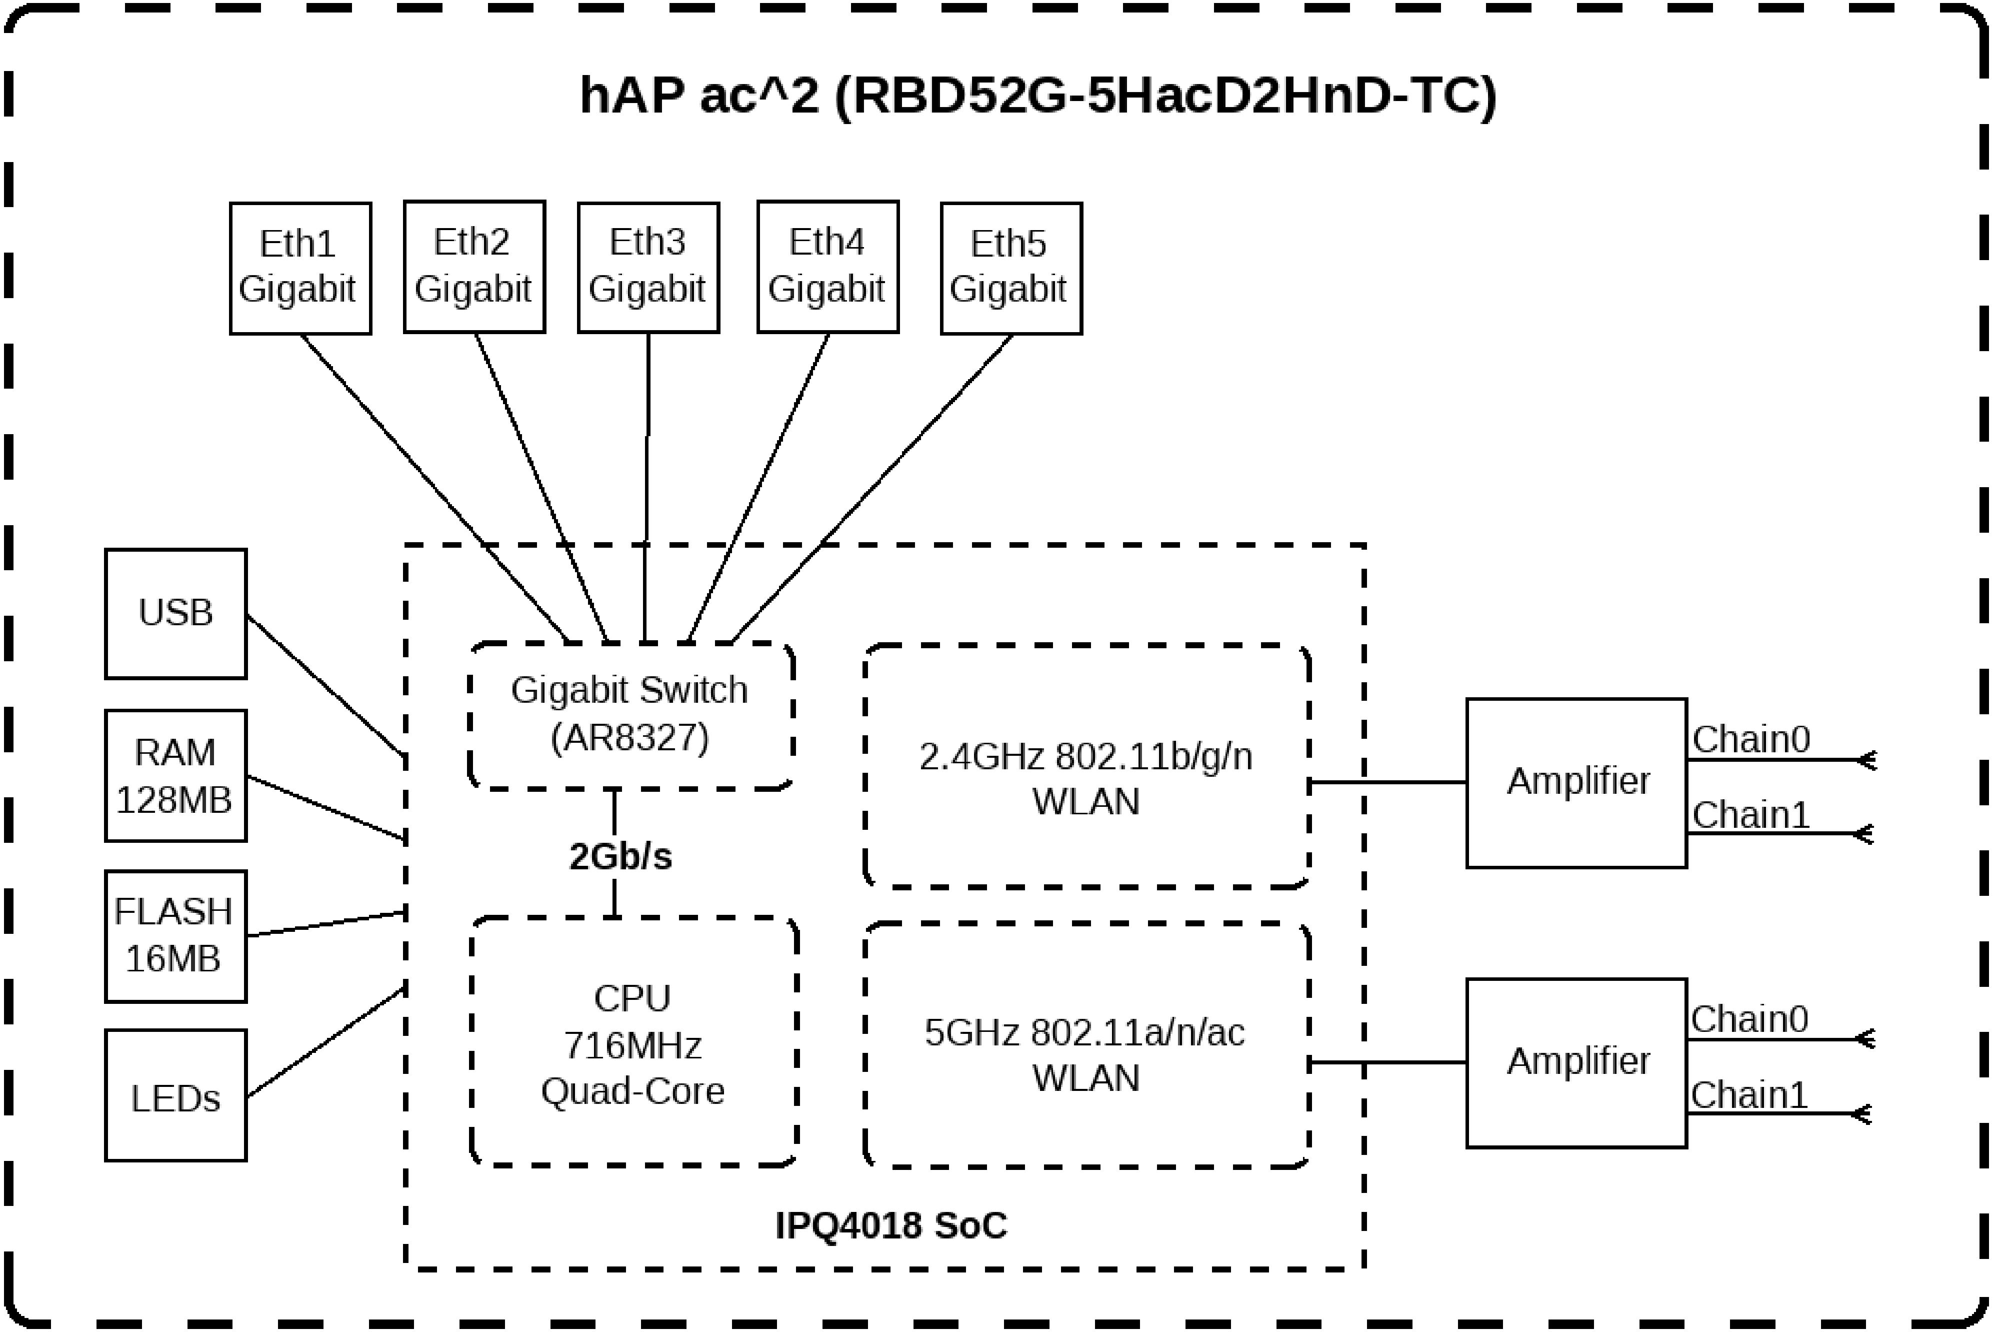

hAP ac^2 (RBD52G-5HacD2HnD-TC)

Eth1 Gigabit

Eth2 Gigabit

Eth3 Gigabit

Eth4 Gigabit

Eth5 Gigabit

USB

RAM 128MB

FLASH 16MB

LEDs

Gigabit Switch (AR8327)

2Gb/s

CPU 716MHz Quad-Core

2.4GHz 802.11b/g/n WLAN

5GHz 802.11a/n/ac WLAN

IPQ4018 SoC

Amplifier

Chain0

Chain1

Amplifier

Chain0

Chain1

“Slow bridging performance”

- RouterOS v6.40.5
- Internet port, All other ports bridged (wireless AP’s, virtual guest AP’s, other Ethernets)
- **/interface bridge filter**: to restrict guest access to servers

WRONG!!!

Analysis of the problem

- Mike’s problem #1:
  - Server to workstation speed on Ethernet not reaching 1 Gbps, CPU load high, Internet communication slowed down.
- Diagnosis:
  - “/tool profile” high bridging load.
- Reason:
  - All traffic is traveling through the bridge in “SlowPath”.
  - All bridge traffic is filtered in bridge filters.

New bridge implementation

- Starting from RouterOS v6.41, RouterOS switch functionality is included into the new bridge implementation that can hardware offload some of the bridge features.

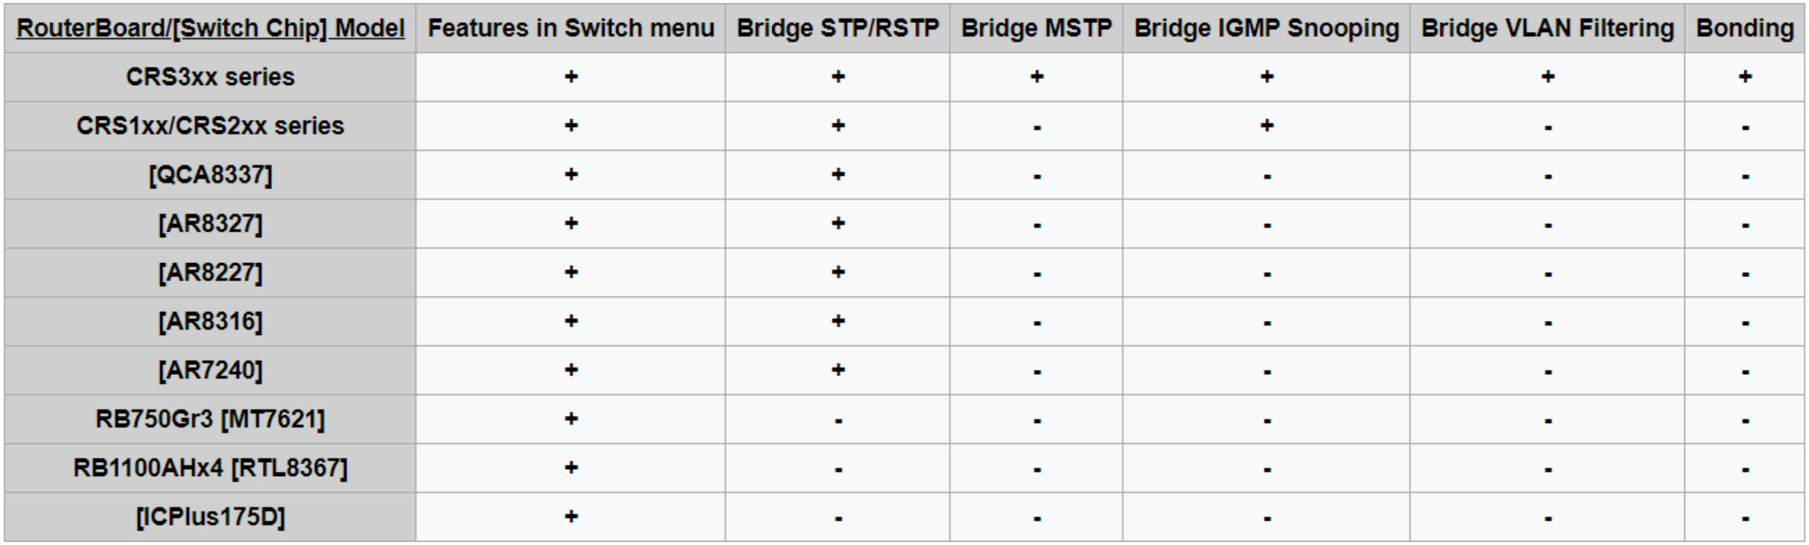

RouterBoard/[Switch Chip] Model	Features in Switch menu	Bridge STP/RSTP	Bridge MSTP	Bridge IGMP Snooping	Bridge VLAN Filtering	Bonding
CRS3xx series	+	+	+	+	+	+
CRS1xx/CRS2xx series	+	+	-	+	-	-
[QCA8337]	+	+	-	-	-	-
[AR8327]	+	+	-	-	-	-
[AR8227]	+	+	-	-	-	-
[AR8316]	+	+	-	-	-	-
[AR7240]	+	+	-	-	-	-
RB750Gr3 [MT7621]	+	-	-	-	-	-
RB1100AHx4 [RTL8367]	+	-	-	-	-	-
[ICPlus175D]	+	-	-	-	-	-

Hardware offload

- Each bridge port now has “hw” option that can enable hardware offload to switch for specific port or disable it, if port is attached to the switch chip.
- If in and out ports have “hw” enabled and are members of the same switch - traffic will skip all CPU processing and will be simply switched without causing any CPU load.
- Hardware offload can be used as a filter before the bridge filter, to reduce CPU load, without losing functionality.

Growing requirements

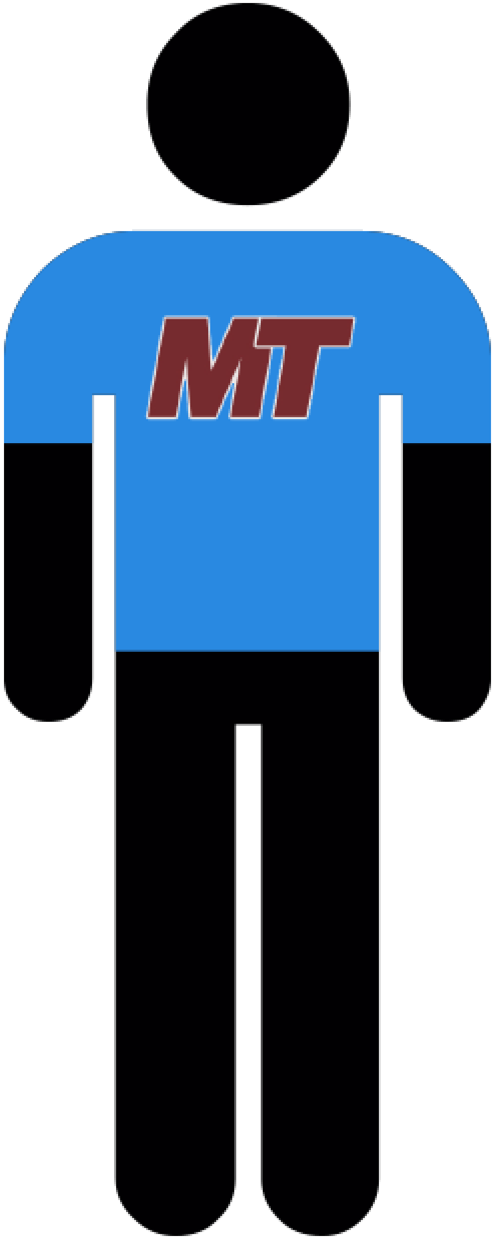

- Now internal devices work with server perfectly, load is normal.
- But Mike notices that some of the customers are abusing network privileges, by applying heavy download both on Mike’s server and Internet connection.
- Mike is about to implement some QoS.

“Transparent queuing”

- Same configuration as before
- Task is to apply limitation on guest network both for Internet and local server traffic.
- **/interface bridge settings**
  set use-ip-firewall=yes
- Simple queue for guest AP bridge port with PCQ queue type

WRONG!!!

Analysis of the problem

- Mike’s problem #2:
  - Queue doesn’t seem to work on all traffic, but causes additional load.
- Diagnosis:
  - “/tool profile”, packet flow diagram, firewall log rules.
- Reason:
  - Bridged traffic now travels through IP firewall including Connection Tracking.
  - From routing perspective guest traffic comes from bridge interface not bridge port interface.

Interface HTB

- There is one place where you can queue both bridged and routed traffic together – Interface HTB.

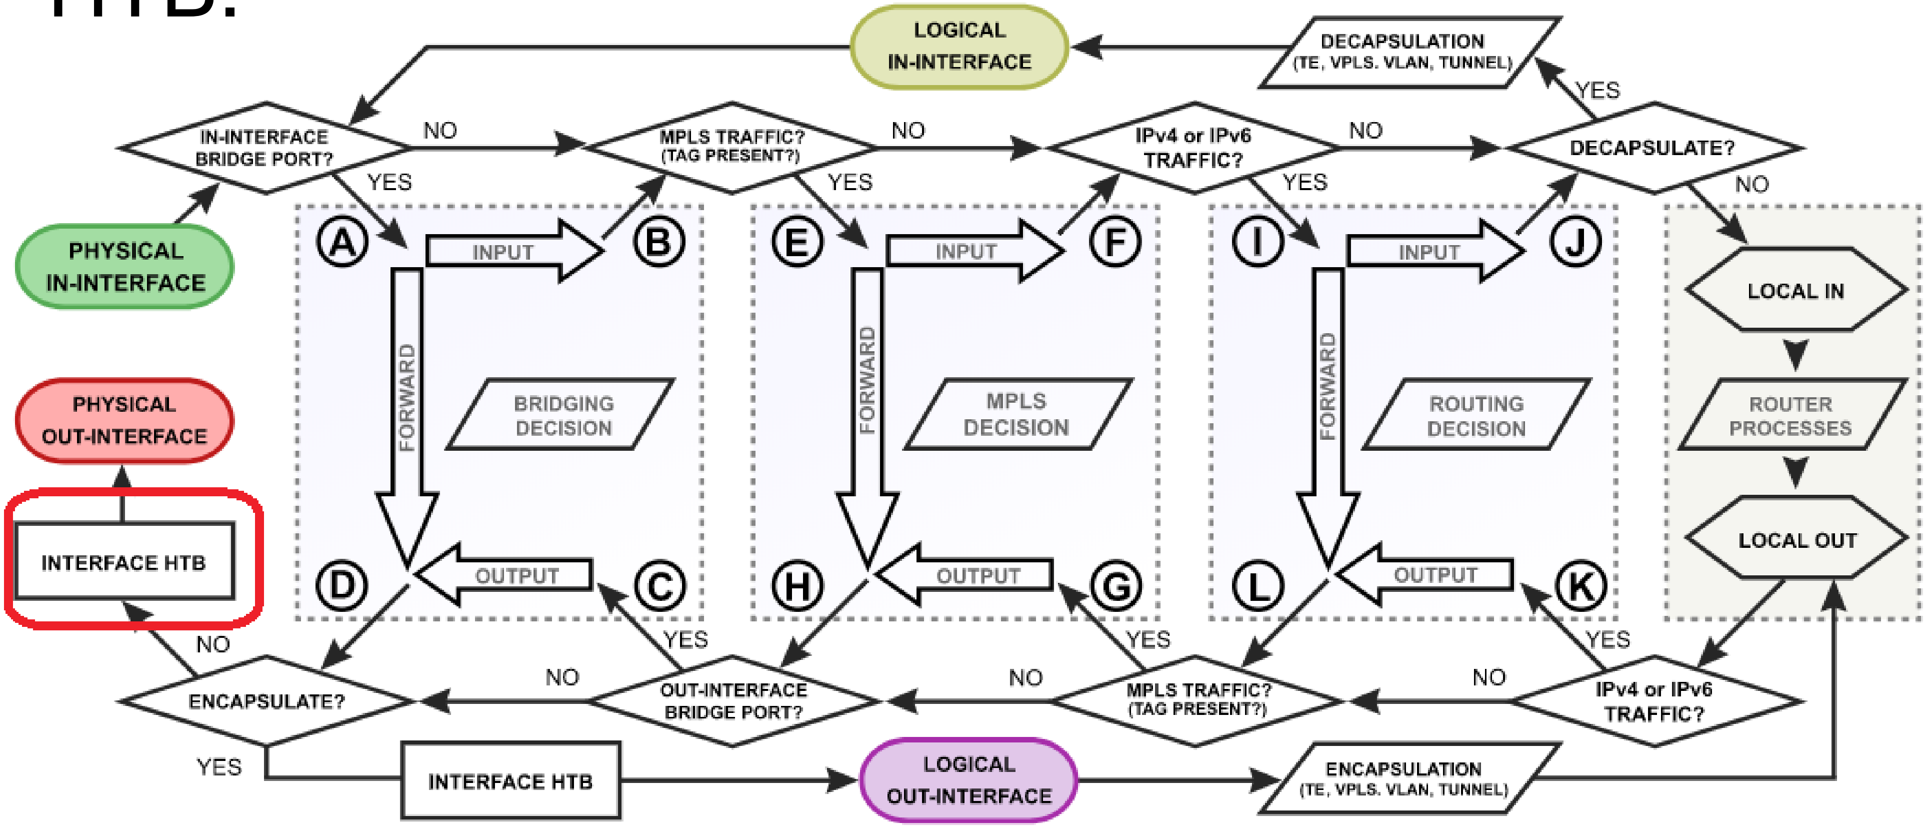

Solution

- Both “/interface bridge filter” and “/ip firewall mangle” have “packet-mark” options, to mark the traffic.
- Use packet-mark in Queue Tree placed on specific bridge port.
- This queue tree will override default interface queue from:

  ```
  /queue interface
  ```
- No need for “use-ip-firewall” anymore.

Business specific issues

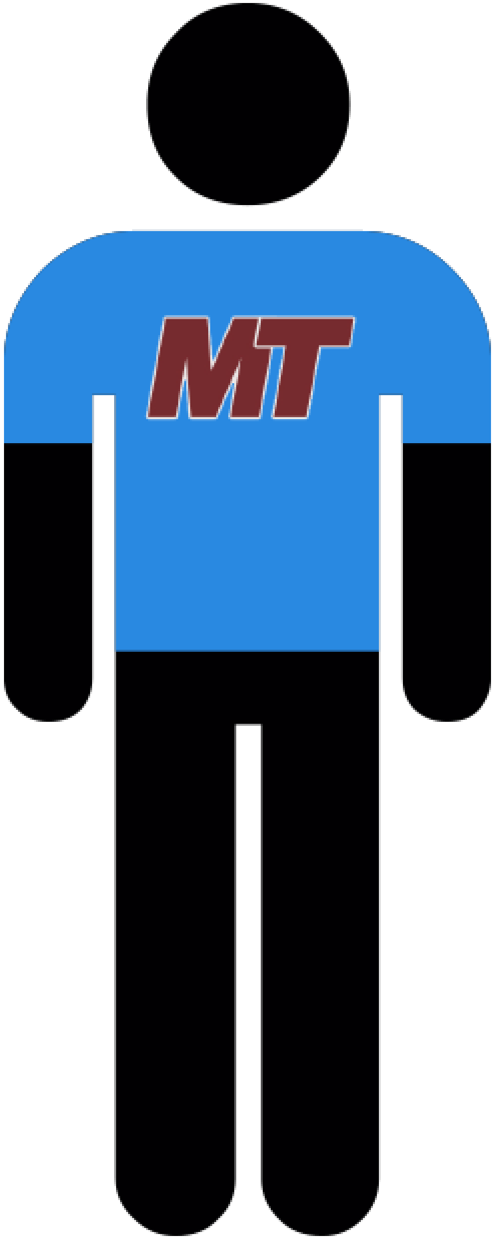

- With the guests limited to certain speed, Mike looks into what exactly customers are using his network for?
- Mikes notices in DNS cache that customers are browsing his competitor webpages most likely to compare prices.
- Mike investigates the way to restrict access to those pages and while at it, how to restrict YouTube and Facebook for the employees

“High Layer7 load”

-
  ```
  /ip firewall layer7-protocol
     add name=youtube regexp="^.+(youtube).*\$"
     add name=facebook regexp="^.+(facebook).*\$"
  ```

-
  ```
  /ip firewall filter
     add action=drop chain=forward layer7-protocol=facebook
     add action=drop chain=forward layer7-protocol=youtube
  ```

WRONG!!!

Analysis of the problem

- Mike’s problem #3:
  - High CPU load, increased latency, packet loss, jitter, YouTube and Facebook is not blocked.
- Diagnosis:
  - “/tool profile” high Layer7 load.
- Reason:
  - Each connection is rechecked over and over again.
  - Layer7 is checked in the wrong place and against all traffic.

Layer7

- Layer7-protocol is a method of searching for patterns in **ICMP/TCP/UDP** streams.

- On trigger Layer7 collects the next 10 packets or 2 KB of a connection and searches for the pattern in the collected data.

- All Layer7 patterns available on the Internet are designed to work only for the first 10 packets or 2 KB of a connection.

Correct implementation (old)

-
  ```
  /ip firewall mangle
  add action=mark-connection chain=prerouting protocol=udp dst-port=53 connection-mark=no-mark layer7-protocol=youtube new-connection-mark=youtube_conn passthrough=yes

  add action=mark-packet chain=prerouting connection-mark=youtube_conn new-packet-mark=youtube_packet
  ```

-
  ```
  /ip firewall filter
  add action=drop chain=forward packet-mark=youtube_packet
  add action=drop chain=input packet-mark=youtube_packet
  ```

(and same set for Facebook)

TLS-Host

- Since most of the Internet now uses HTTPS, it has become much harder to filter specific WWW content.
- For this reason, RouterOS 6.41 introduces a new firewall option that allows you to block HTTPS websites (TLS traffic).
- Based on the TLS SNI extension, called “TLS-Host”. The new parameter supports GLOB-style patterns.

Correct implementation (new)

-
  ```
  /ip firewall filter
       add chain=forward dst-port=443
  protocol=tcp tls-host=*.facebook.com
  action=reject
       add chain=forward dst-port=443
  protocol=tcp tls-host=*.youtube.com
  action=reject
  ```

Using the latest features

- Mike is happy with the device, but it is running little bit high on the load with all the Layer7 and TLS-Host filters, he reads up on the ways to improve performance.
- Mike discovers FastTrack.

FastTracked

- Connection tracking entries now have “FastTracked” flag.
- Packets from “FastTracked” connections are allowed to travel in “FastPath”.
- Works only with IPv4/TCP and IPv4/UDP.
- Traffic traveling in “FastPath” will be invisible to other router facilities (firewall, queues, etc.)
- Some packets will still follow the regular path to maintain Connection Tracking entries.

“Layer7 and TLS-Host stopped”

- Implemented as “fasttrack-connection” action for firewall filter/mangle like this:

  -
    ```
    /ip firewall filter
    add chain=forward action=fasttrack-connection
    connection-state=established,related
    add chain=forward action=accept connection-
    state=established,related
    ```

WRONG!!!

Analysis of the problem

- Mike’s problem #4:
  - Rules with Layer7 and TLS-Host options stopped capturing traffic.
- Diagnosis:
  - Counters on firewall rules, and fasttrack-connection rule.
- Reason:
  - Layer7 and TLS-Host options require several packets from connection to work, Fasttrack configuration only lets one packet to get to them.

Correct implementation

-
  ```
  /ip firewall filter
        add chain=forward action=fasttrack-connection connection-bytes=10000-0
        add chain=forward action=accept connection-bytes=10000-0
  ```

Growing

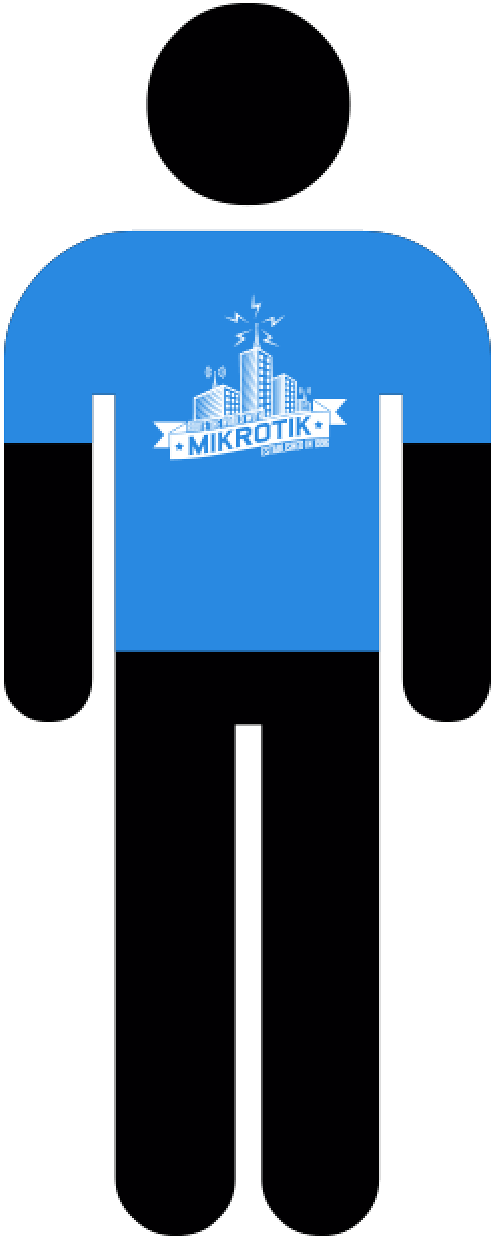

- Mike’s business is booming, he opens up a few more stores, deploys MikroTik devices in them.
- He now needs to interconnect offices with VPNs:
  - Securely.
  - So that devices would be in the same subnet.
  - With high throughput.

“IPSec tunnel doesn’t work”

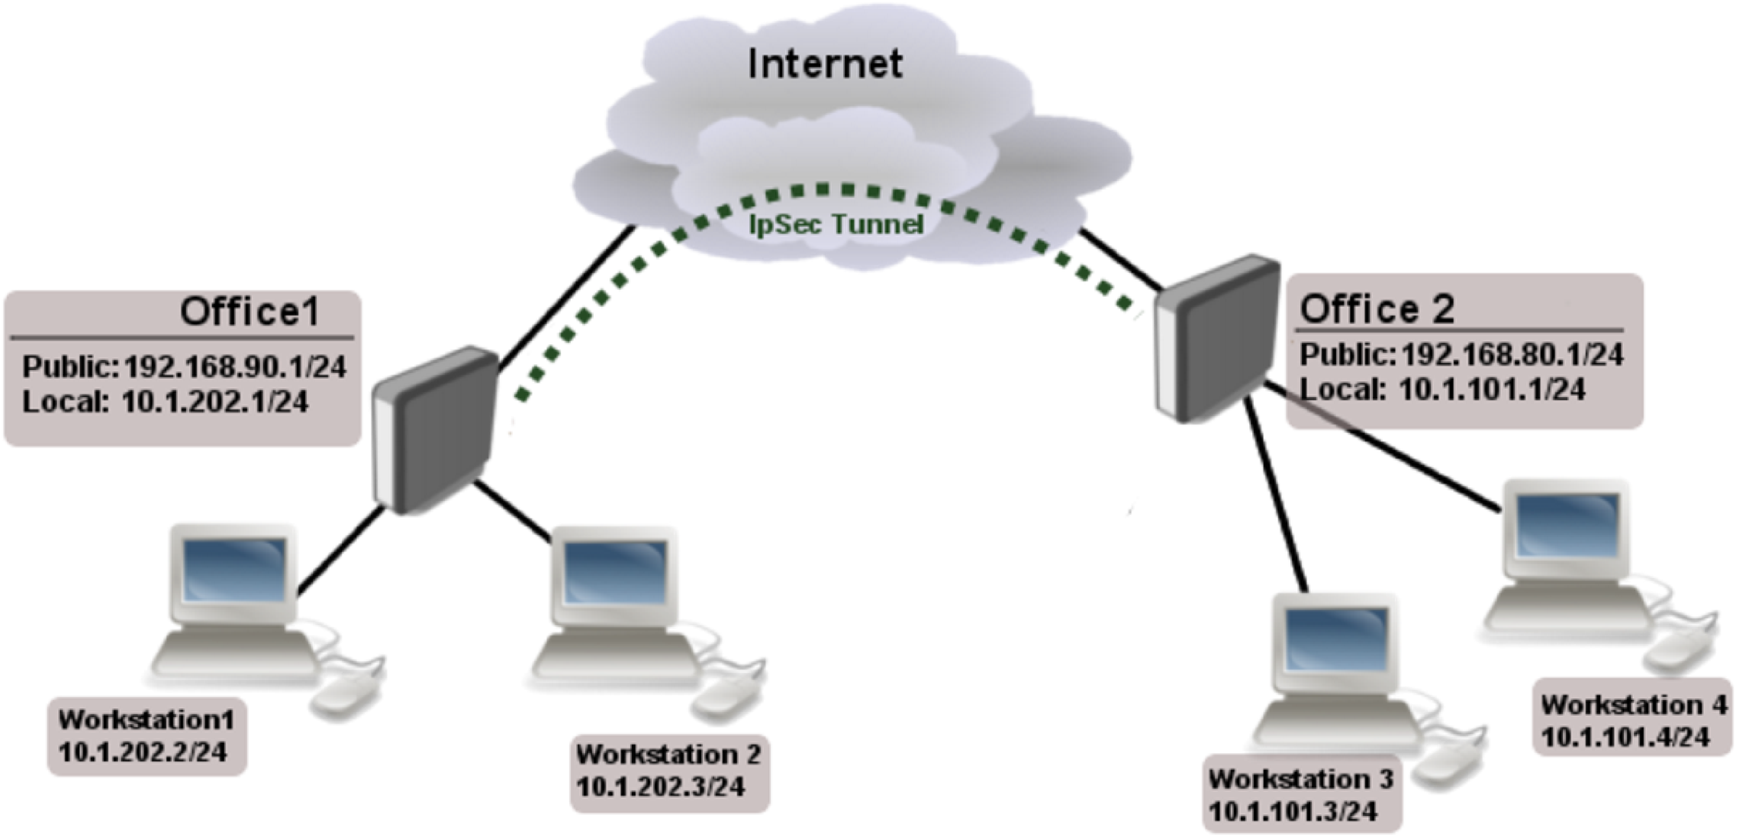

- Simple masquerade on both routers

WRONG!!!

Analysis of the problem

- Mike’s problem #5:
  - IPSec packets are rejected, tunnel cannot be established.
- Diagnosis:
  - /tool sniffer
- Reason:
  - NAT rules are changing “src-address” of encrypted packets, “src-address” doesn’t correspond to IPSec policy on opposite end.

Raw table

- Firewall RAW table allows to selectively bypass or drop packets before connection tracking thus significantly reducing the load on the CPU.
- If packet is marked to bypass connection tracking:
  - Packet de-fragmentation will not occur.
  - NAT will be skipped.
  - Options that depend on connection tracking will not trigger (fasttrack-connection, mark-connection, layer7 etc.)
  - Will have connection-state=untracked.

Correct implementation

-
  ```
  /ip firewall raw
  add action=notrack chain=prerouting src-address=10.1.101.0/24 dst-address=10.1.202.0/24

  add action=notrack chain=prerouting src-address=10.1.202.0/24 dst-address=10.1.101.0/24
  ```

“Securely bridge two local networks”

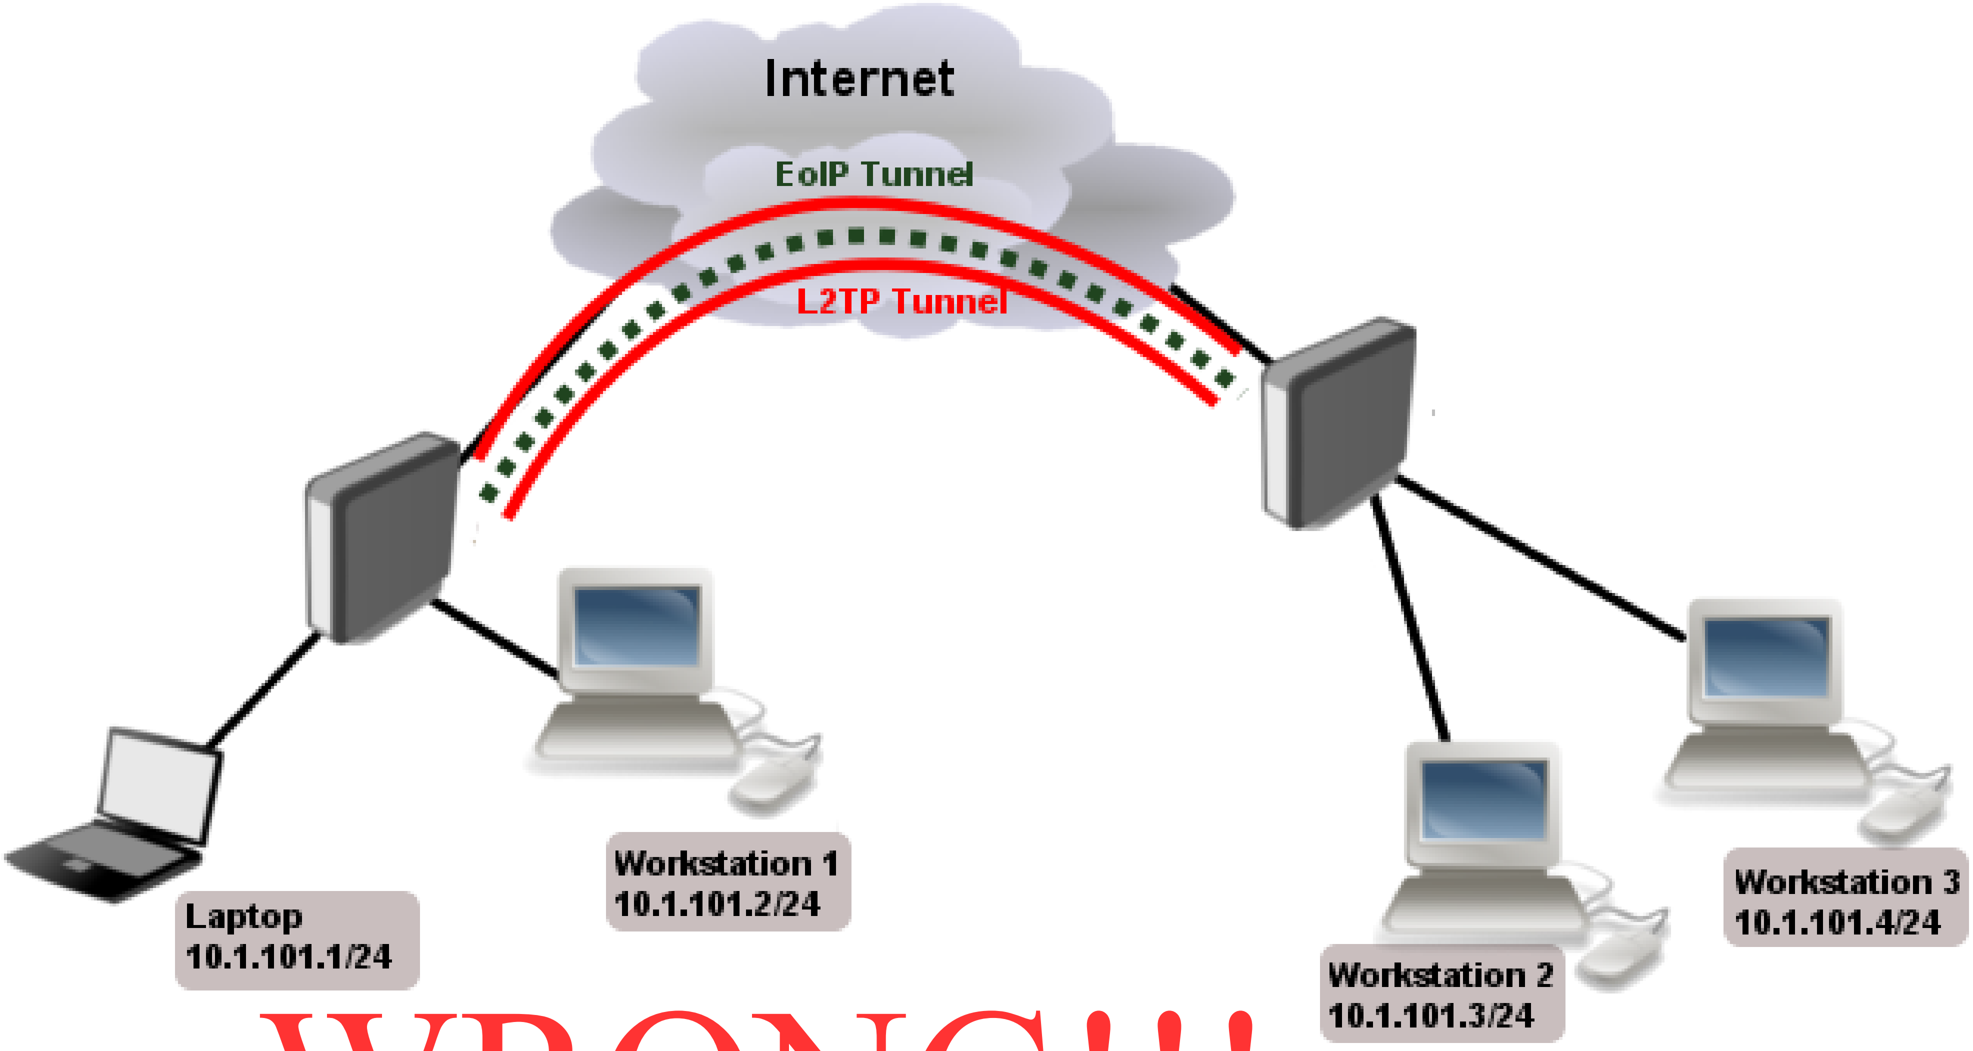

WRONG!!!

Analysis of the problem

- Mike’s problem #6:
  - Webpages very slow to open, slow download speeds, strange suspicion that competition knows your secret information :)
- Diagnosis:
  - /tool bandwidth-test, /tool ping with different packet sizes.
- Reason:
  - PPTP/L2TP is not secure anymore, severe packet overhead from two tunnel overheads, fragmentation, because of reduced MTU.

Correct implementation

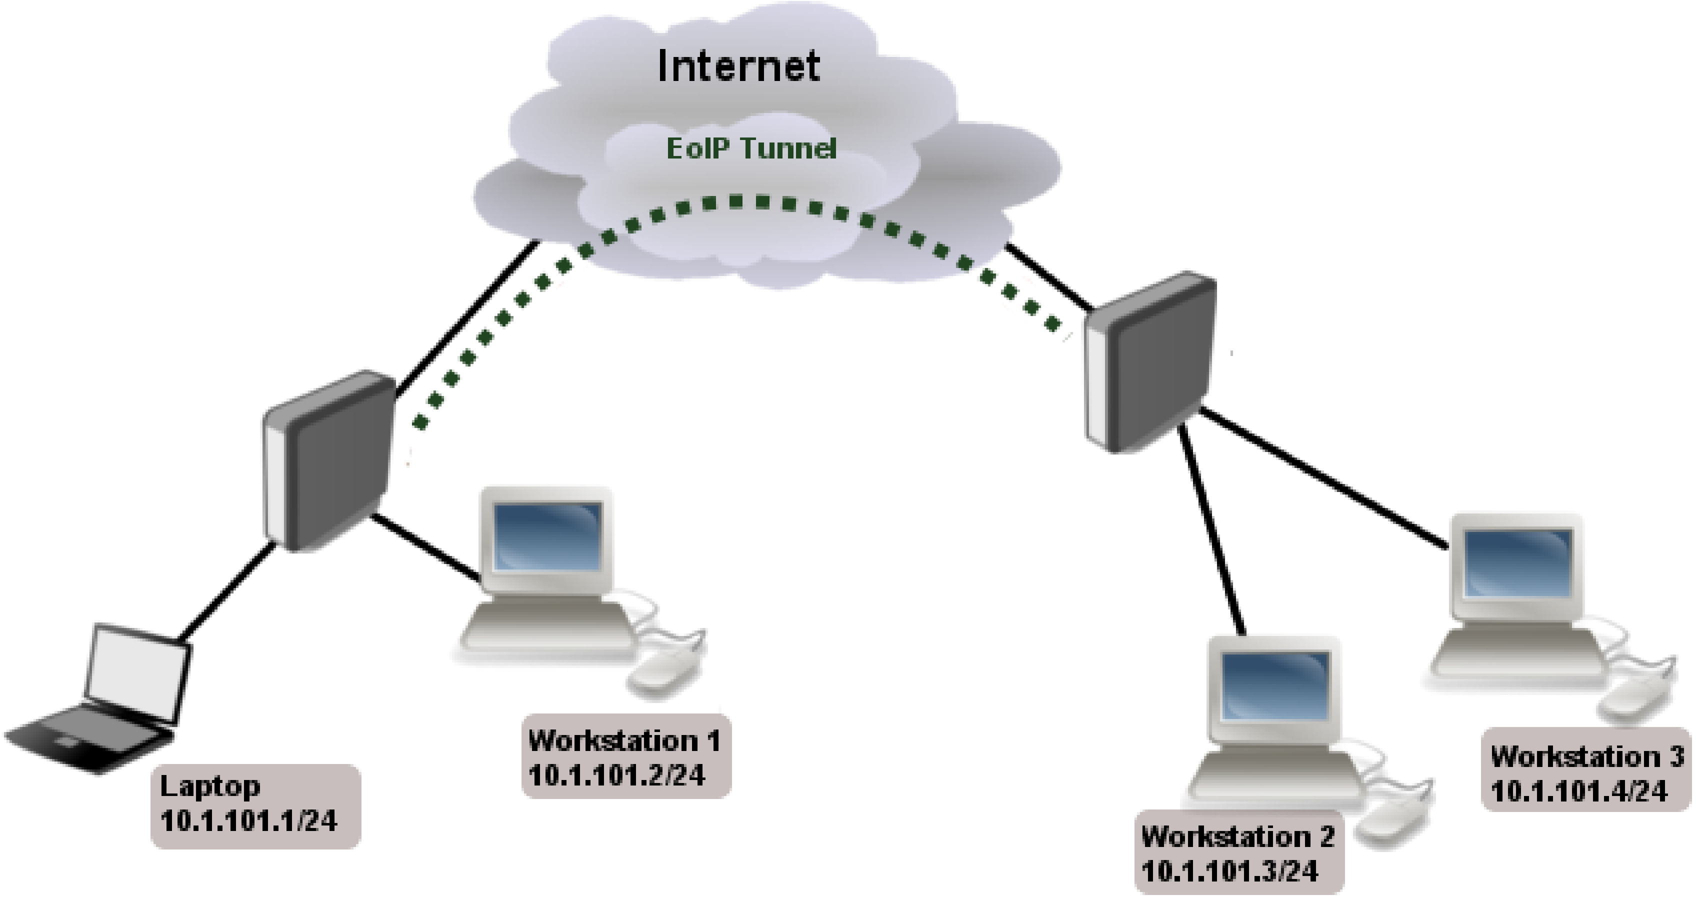

- `/interface eoip set ipsec-secret=`

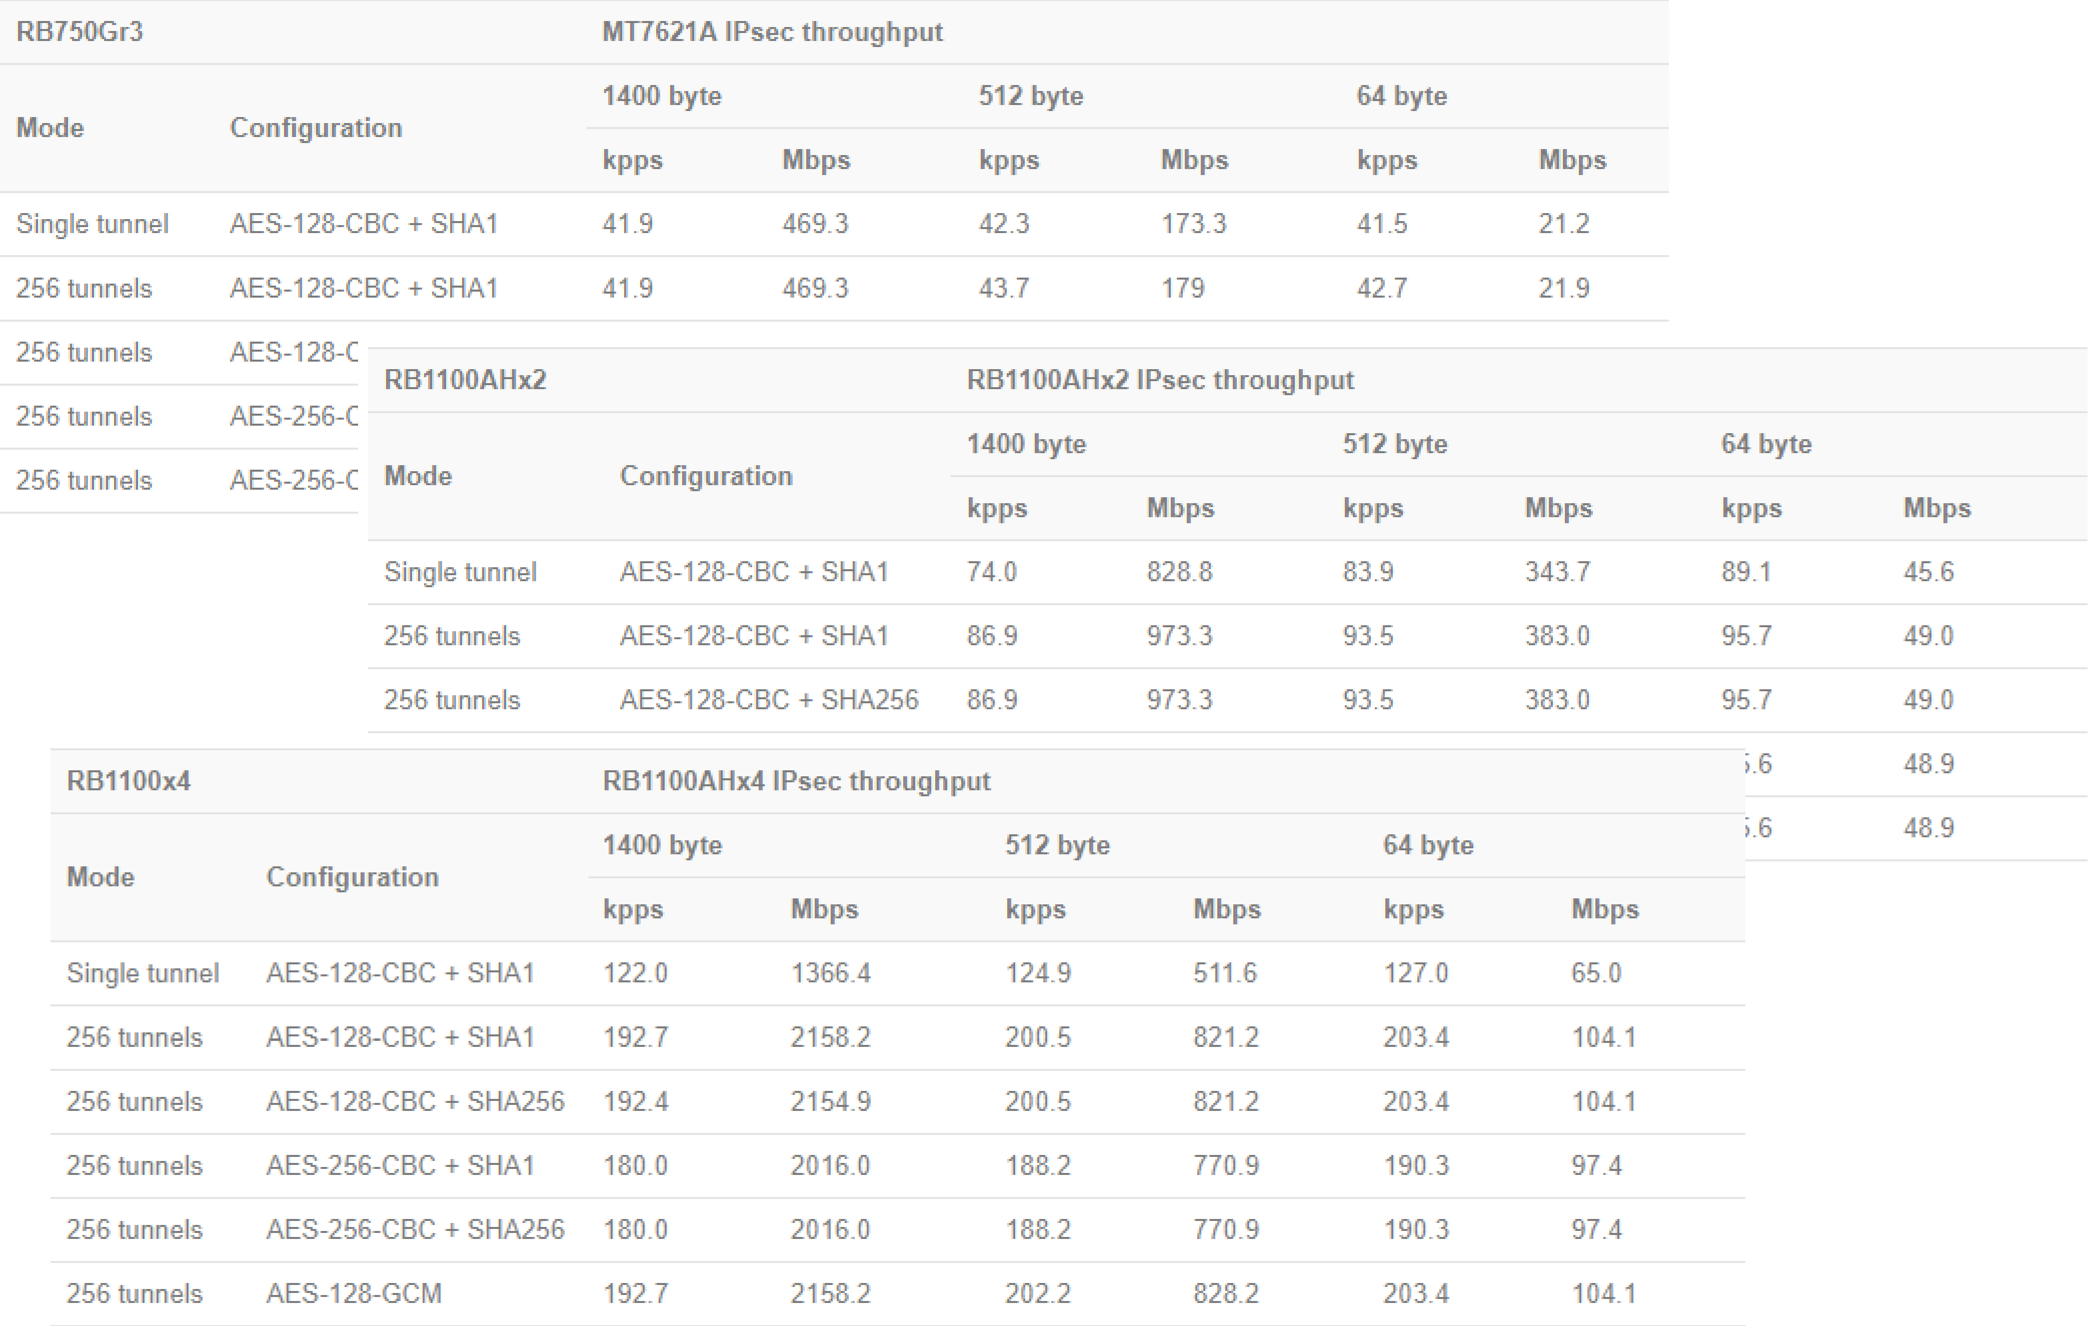

RB750Gr3		MT7621A IPsec throughput					
Mode	Configuration	1400 byte		512 byte		64 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Single tunnel	AES-128-CBC + SHA1	41.9	469.3	42.3	173.3	41.5	21.2
256 tunnels	AES-128-CBC + SHA1	41.9	469.3	43.7	179	42.7	21.9
256 tunnels	AES-128-C						
256 tunnels	AES-256-C						
256 tunnels	AES-256-C						

RB1100AHx2		RB1100AHx2 IPsec throughput					
Mode	Configuration	1400 byte		512 byte		64 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Single tunnel	AES-128-CBC + SHA1	74.0	828.8	83.9	343.7	89.1	45.6
256 tunnels	AES-128-CBC + SHA1	86.9	973.3	93.5	383.0	95.7	49.0
256 tunnels	AES-128-CBC + SHA256	86.9	973.3	93.5	383.0	95.7	49.0
						.6	48.9
						.6	48.9

RB1100x4		RB1100AHx4 IPsec throughput					
Mode	Configuration	1400 byte		512 byte		64 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Single tunnel	AES-128-CBC + SHA1	122.0	1366.4	124.9	511.6	127.0	65.0
256 tunnels	AES-128-CBC + SHA1	192.7	2158.2	200.5	821.2	203.4	104.1
256 tunnels	AES-128-CBC + SHA256	192.4	2154.9	200.5	821.2	203.4	104.1
256 tunnels	AES-256-CBC + SHA1	180.0	2016.0	188.2	770.9	190.3	97.4
256 tunnels	AES-256-CBC + SHA256	180.0	2016.0	188.2	770.9	190.3	97.4
256 tunnels	AES-128-GCM	192.7	2158.2	202.2	828.2	203.4	104.1

Questions?