CloudCoreRouter and RouterOS v6.x

Tips and tricks!

Moscow

MUM Russia 2014

RouterOS v6 Tile architecture

- First 64bit RouterOS
- Multi Memory Channel support (faster RAM)
- Hardware Accelerated Multi-Threading (no RPS and IRQ needed)
- Hardware Accelerated Encryption

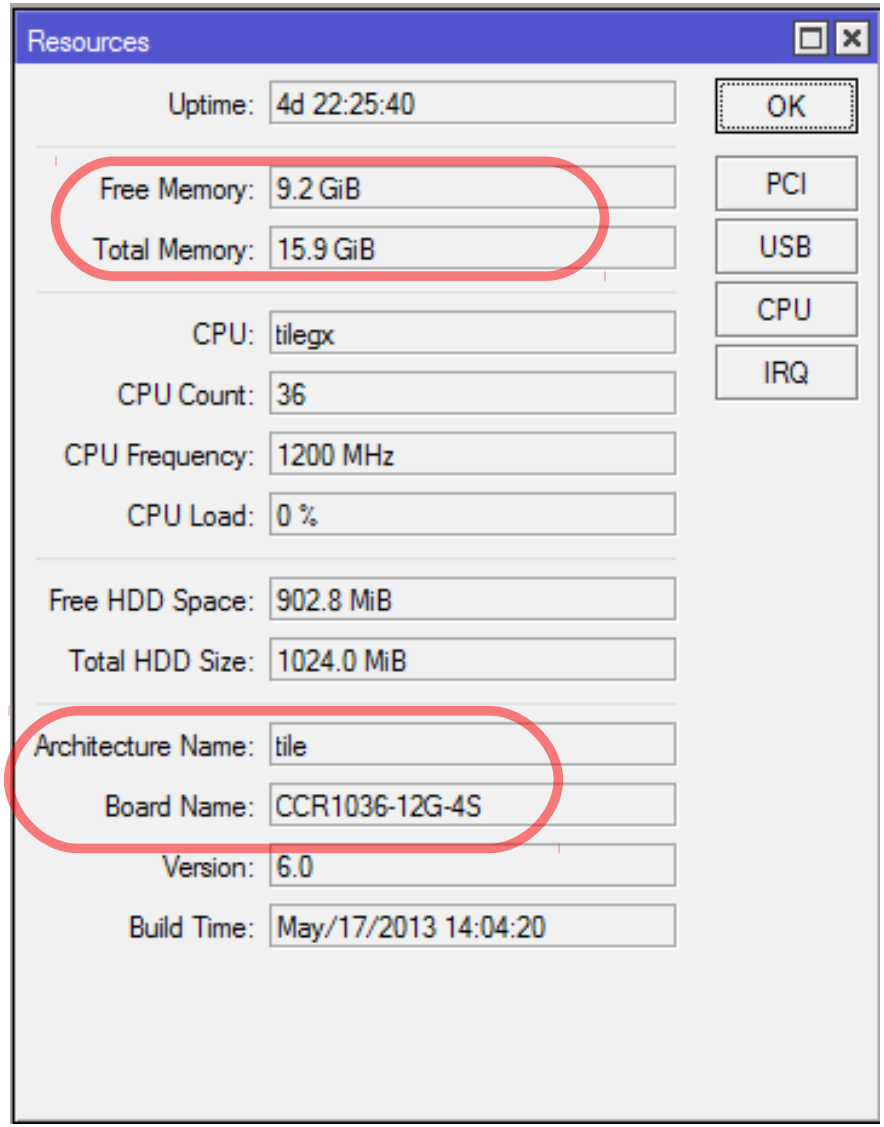

MikroTik RouterOS Packet Flow Diagram for version 6.x

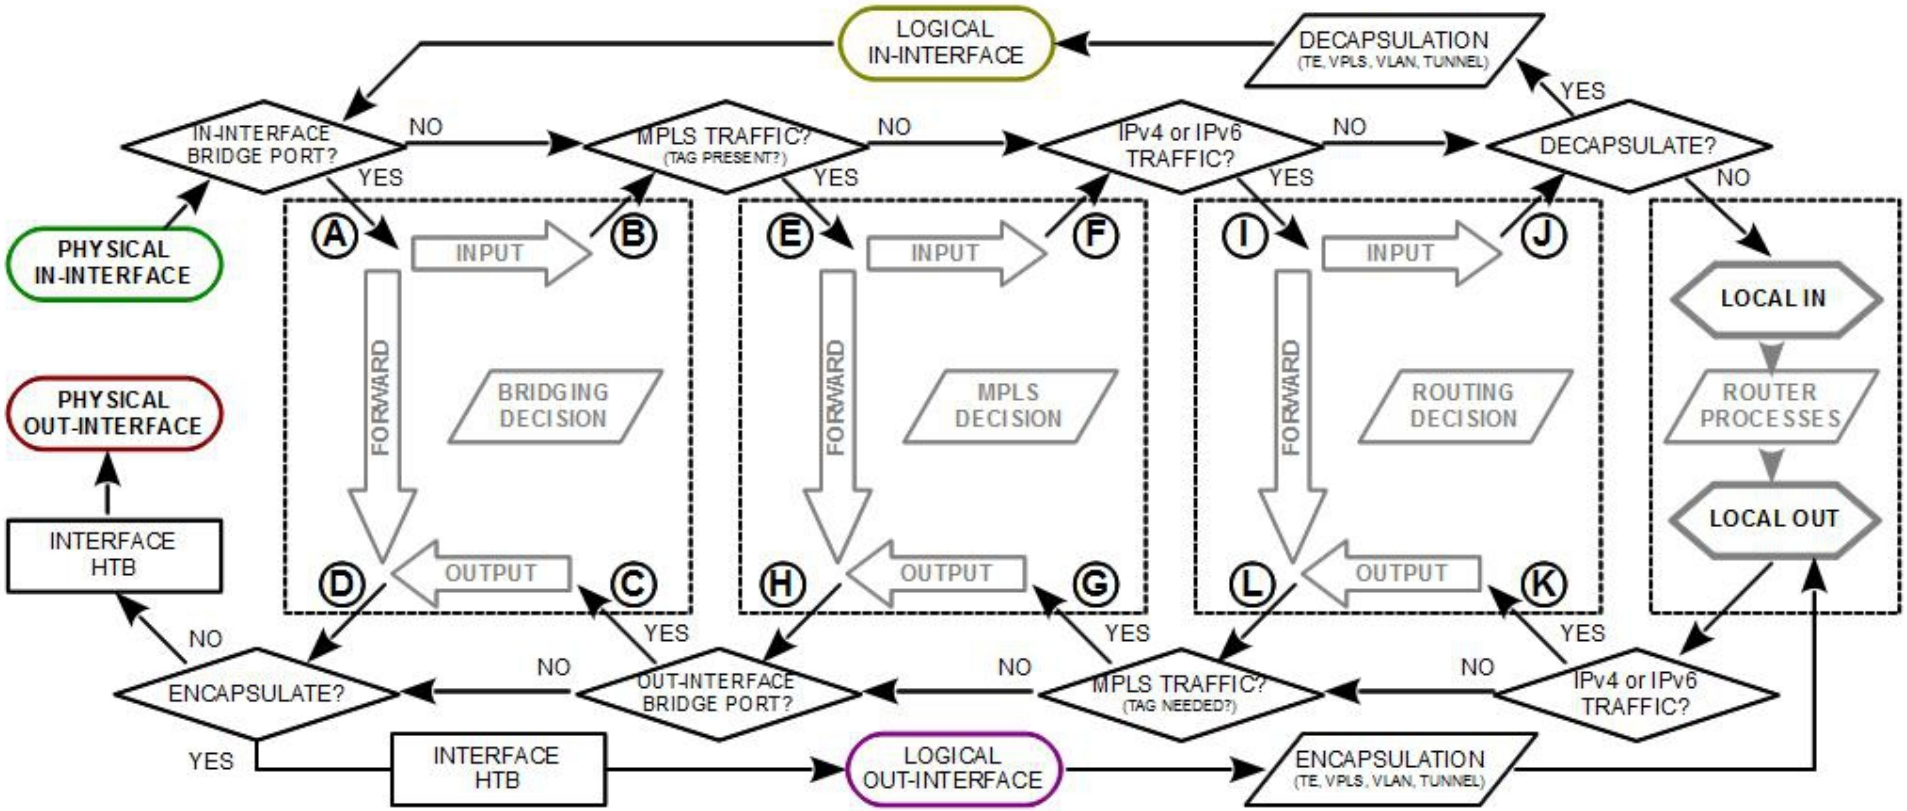

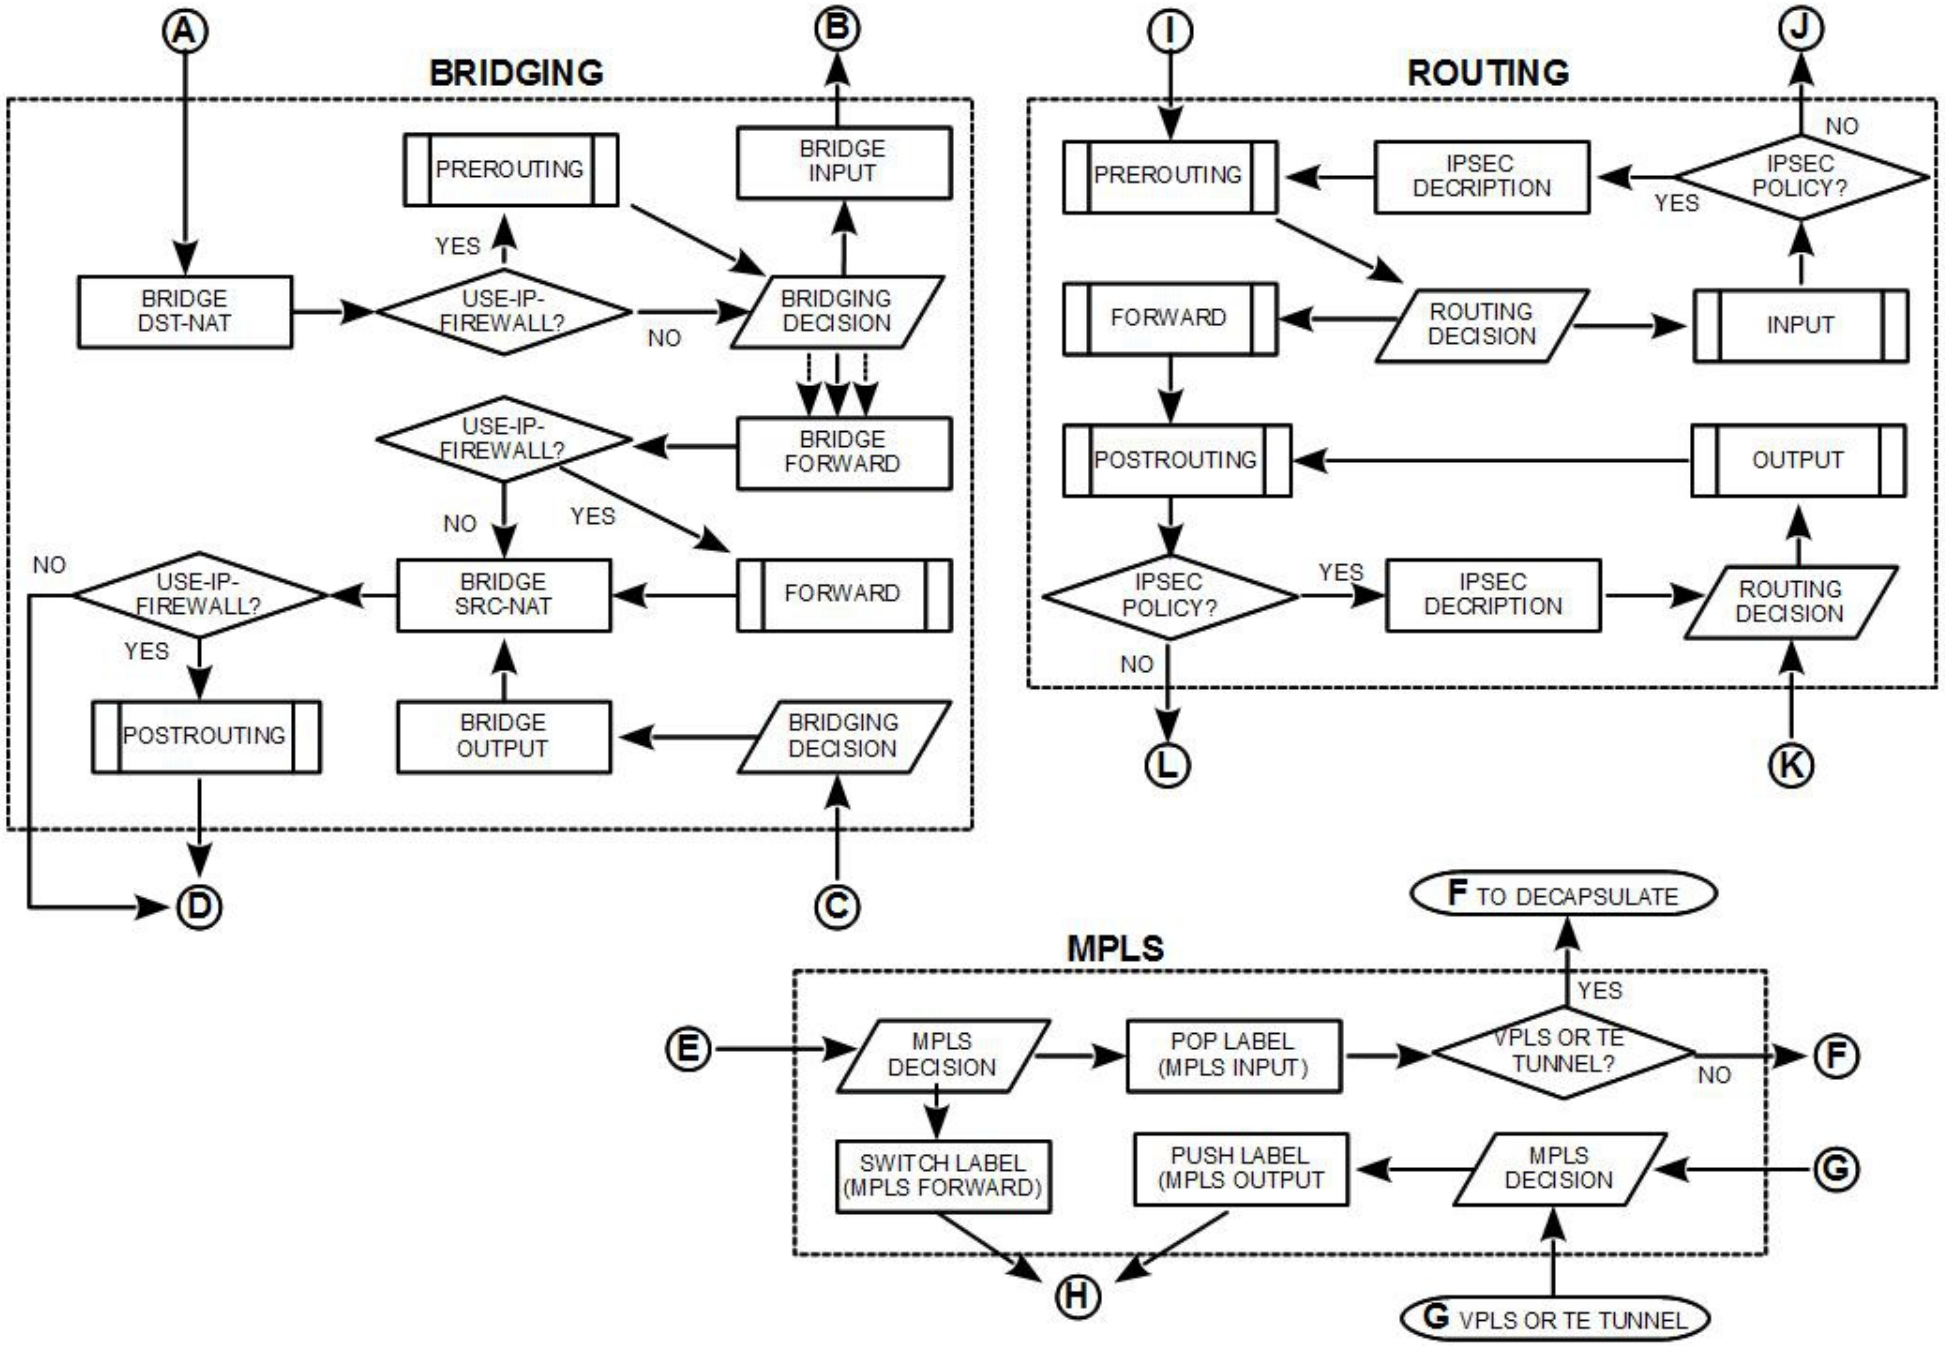
A
B
BRIDGING
PREROUTING
BRIDGE INPUT
BRIDGE DST-NAT
USE-IP-FIREWALL?
YES
NO
BRIDGING DECISION
USE-IP-FIREWALL?
BRIDGE FORWARD
NO
YES
NO
USE-IP-FIREWALL?
BRIDGE SRC-NAT
FORWARD
YES
POSTROUTING
BRIDGE OUTPUT
BRIDGING DECISION
D
C
I
J
ROUTING
PREROUTING
IPSEC DECRIPTION
NO
IPSEC POLICY?
YES
FORWARD
ROUTING DECISION
INPUT
POSTROUTING
OUTPUT
IPSEC POLICY?
YES
IPSEC DECRIPTION
ROUTING DECISION
NO
L
K
F TO DECAPSULATE
MPLS
YES
E
MPLS DECISION
POP LABEL (MPLS INPUT)
VPLS OR TE TUNNEL?
NO
F
SWITCH LABEL (MPLS FORWARD)
PUSH LABEL (MPLS OUTPUT
MPLS DECISION
G
H
G VPLS OR TE TUNNEL

Yes, still - Packet Flow Diagram (page 3)

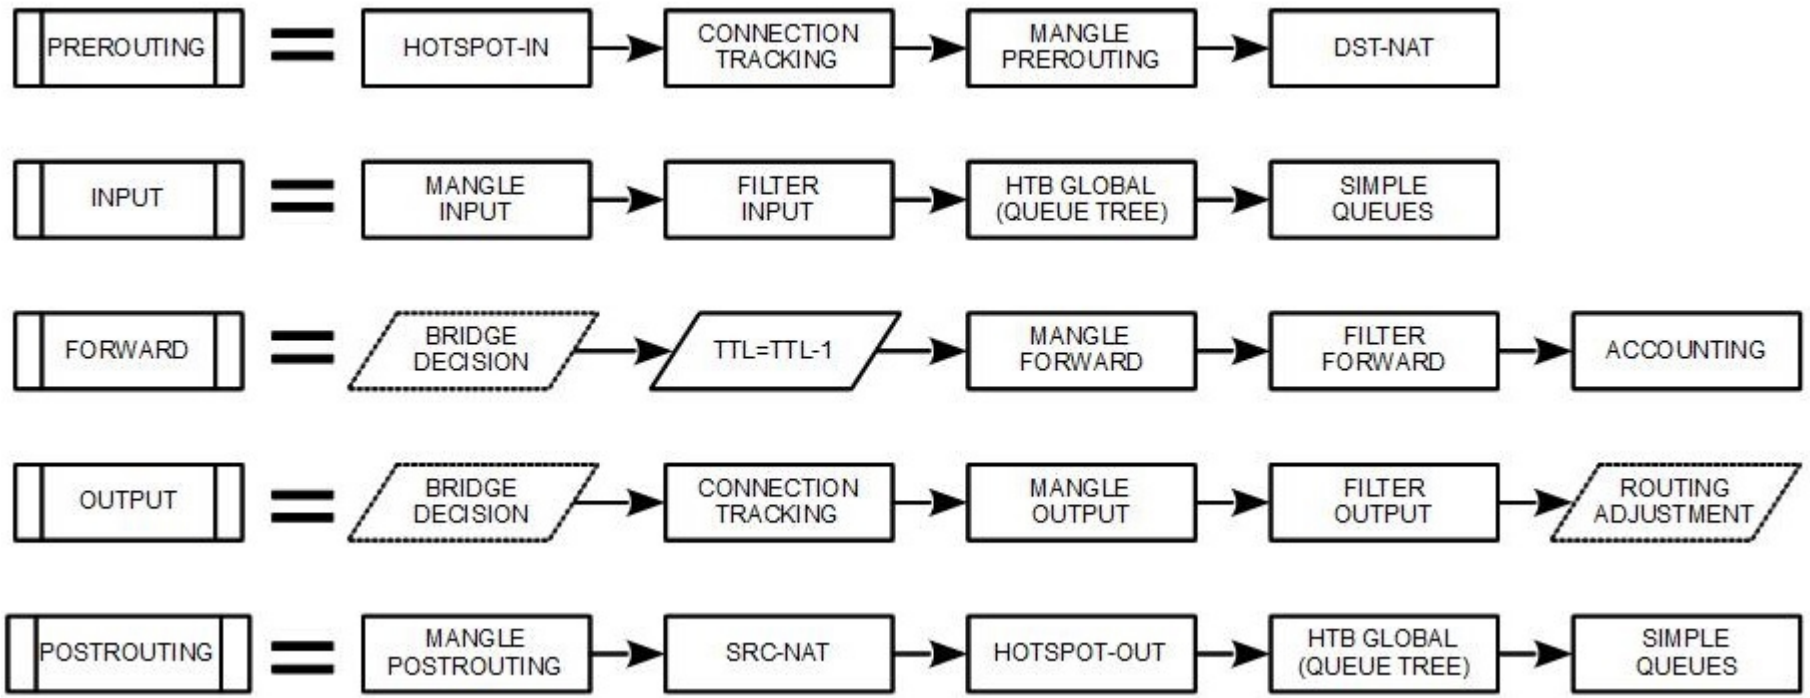

Packet Flow Scenario:

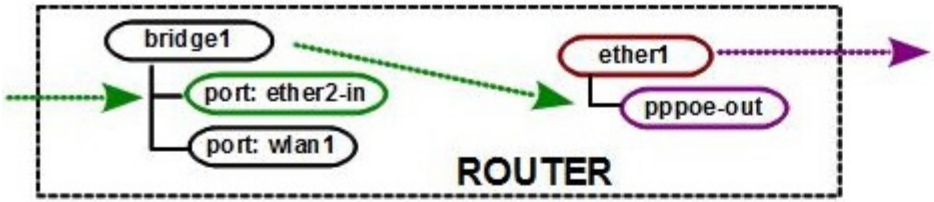

This Scenario in Packet Flow Diagram:

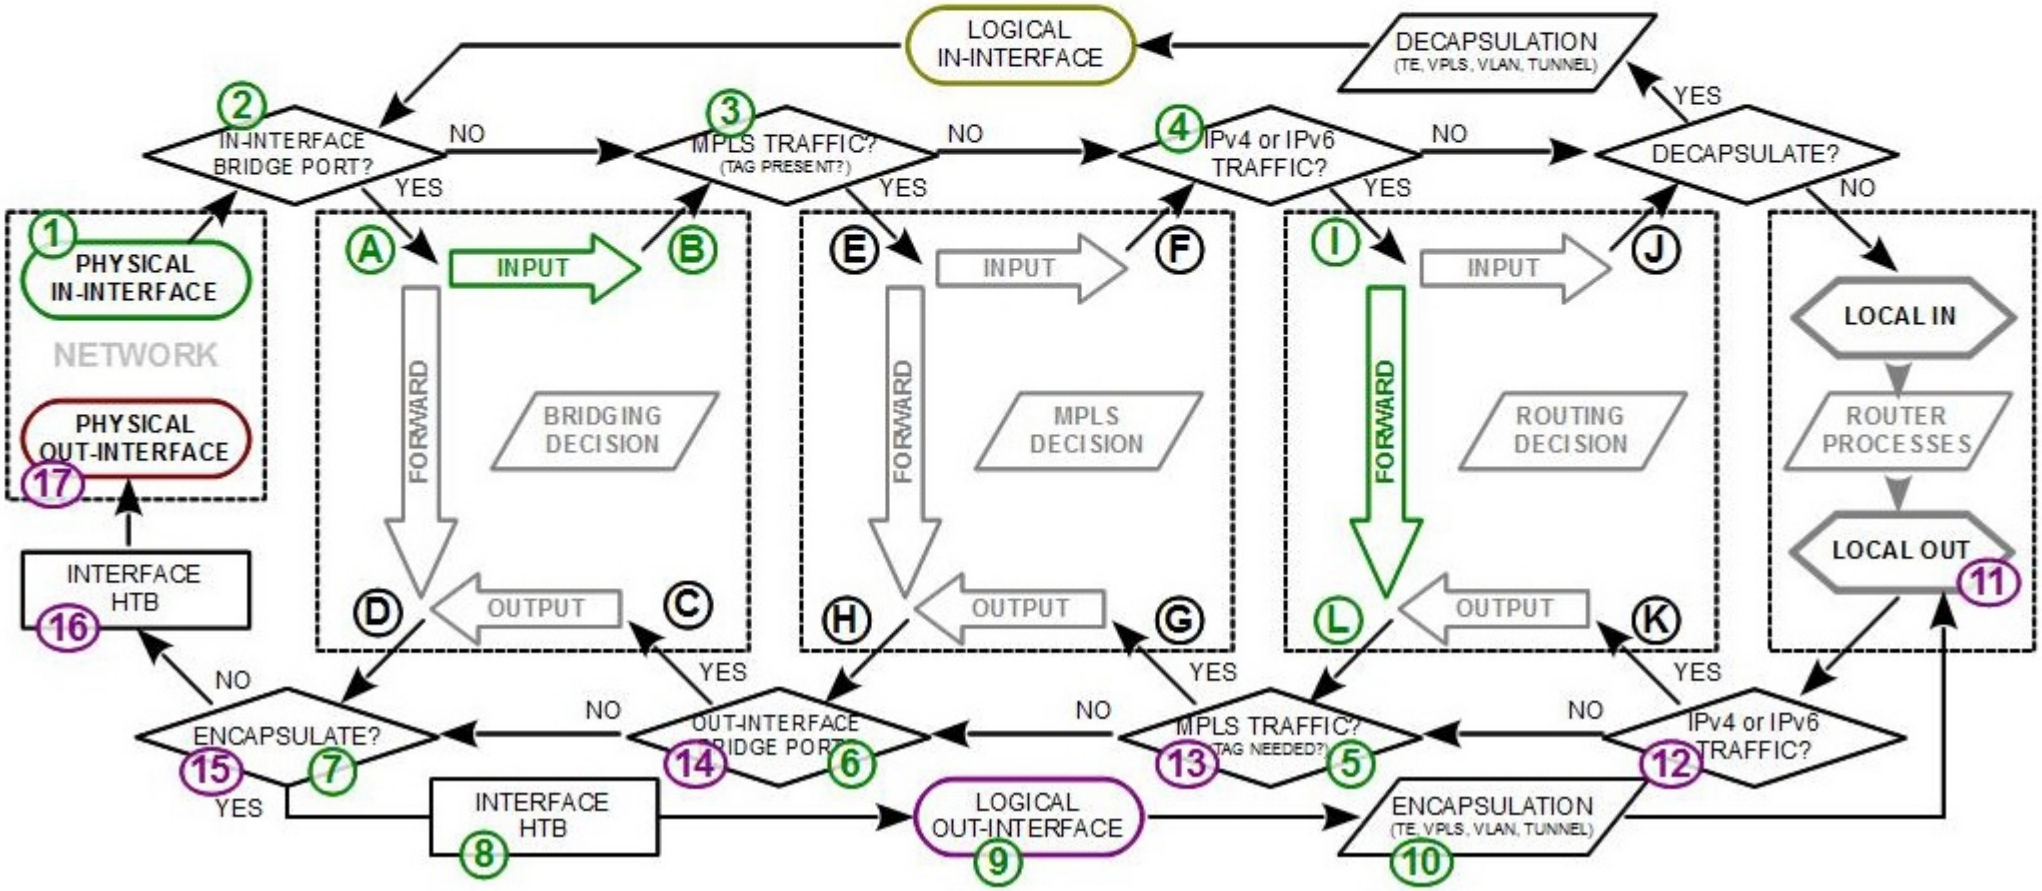

Packet Flow Scenario:

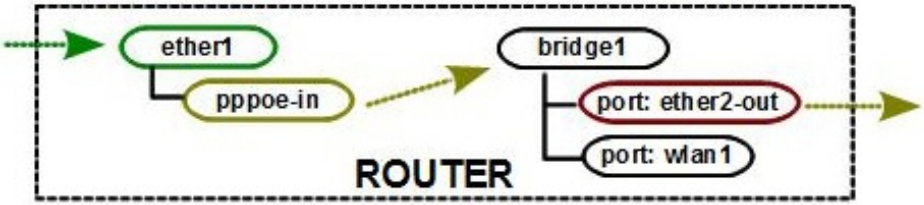

This Scenario in Packet Flow Diagram:

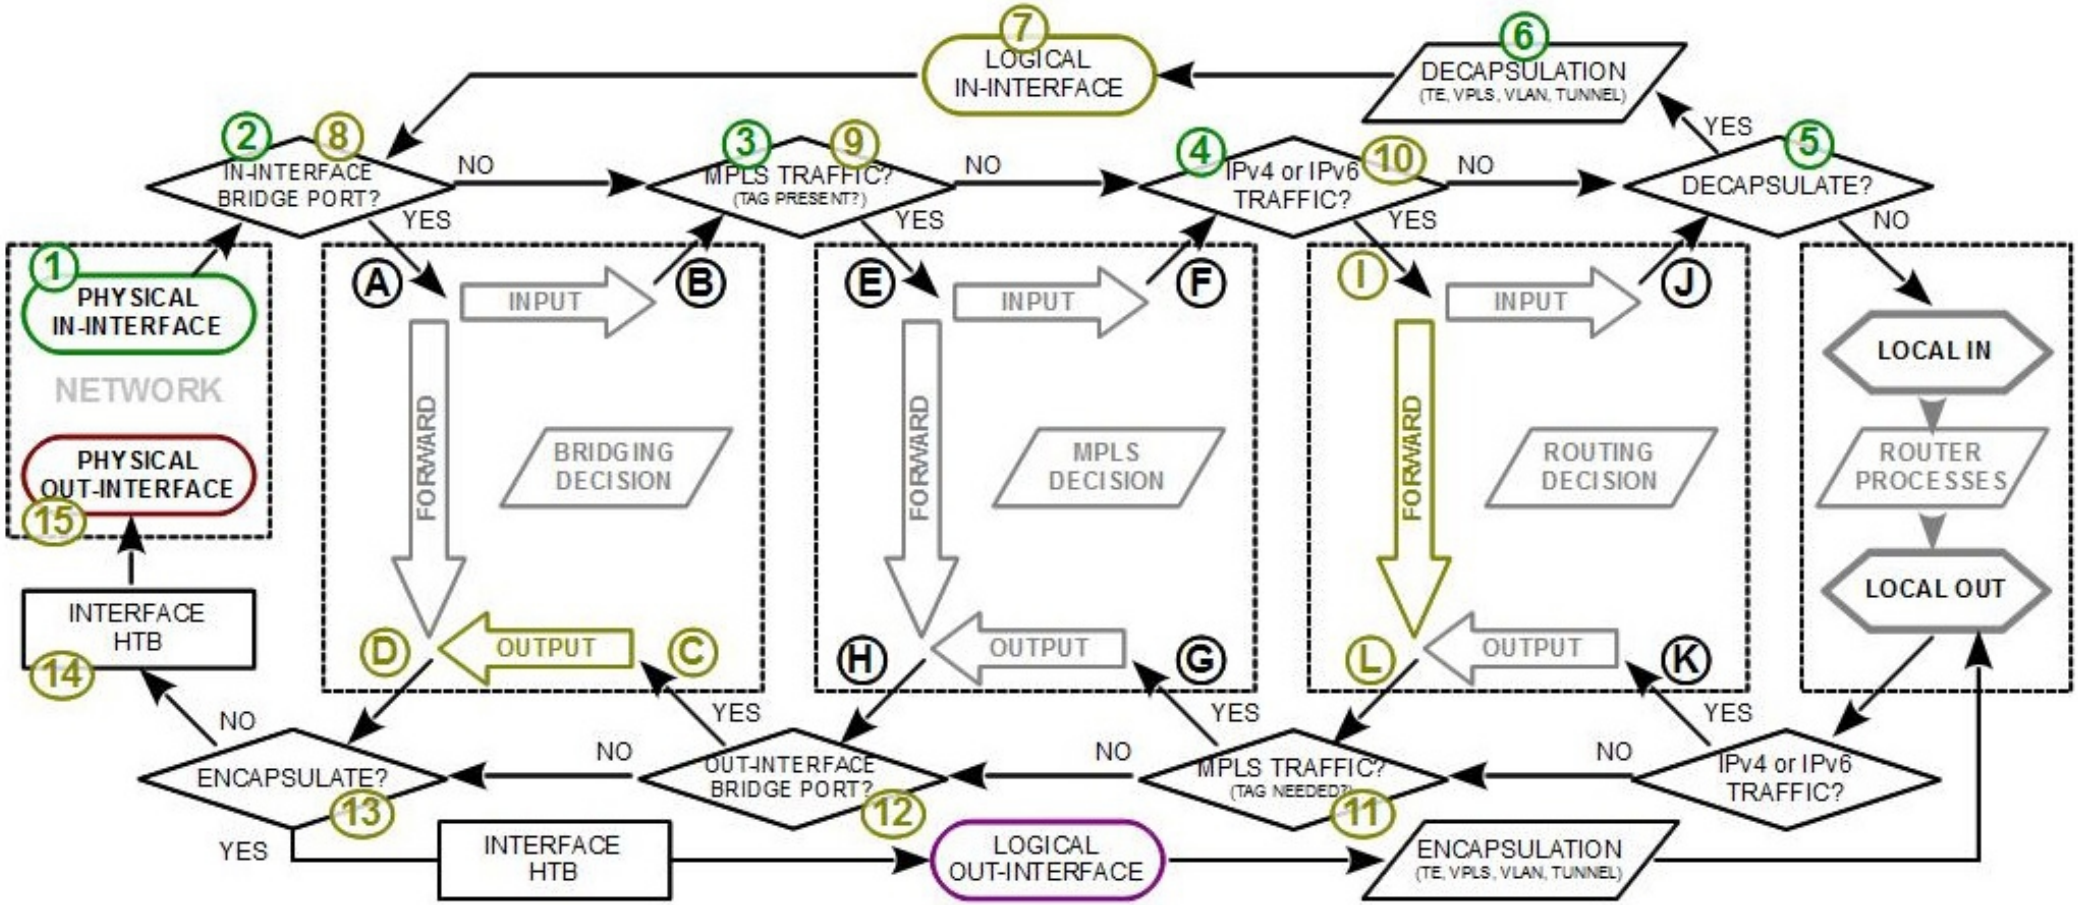

Packet Flow Scenario:

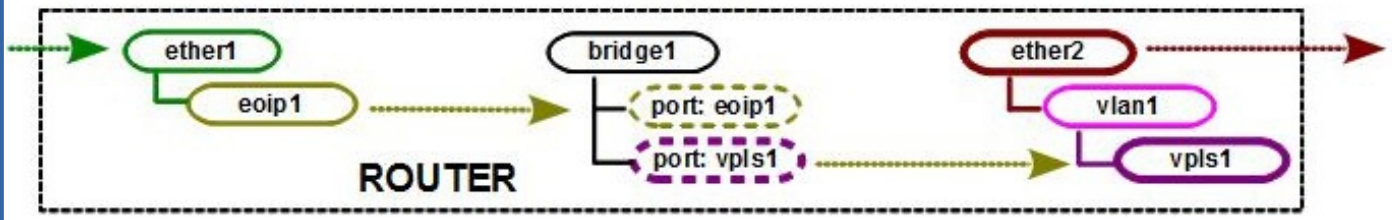

This Scenario in Packet Flow Diagram:

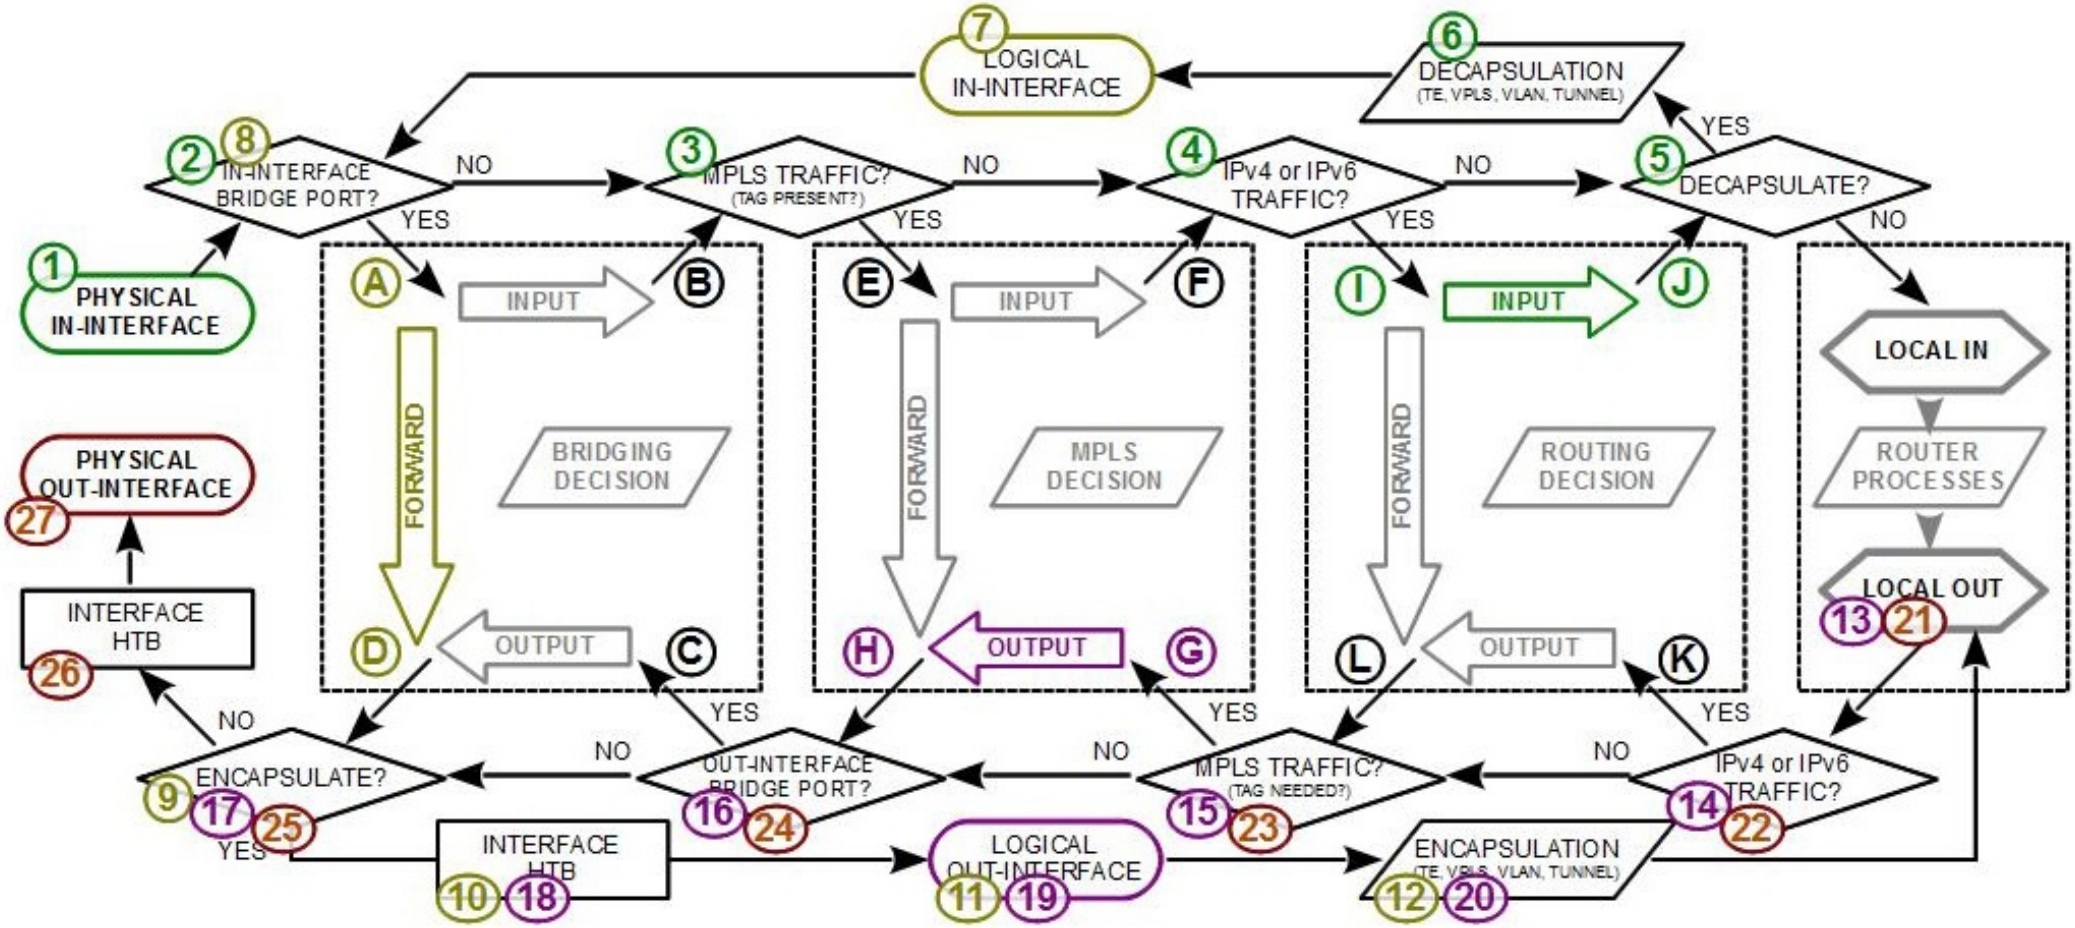

Multi-Core Packet Processing

- On receive packet gets assigned to a CPU core
- RouterOS is trying to keep packets from the same connection assigned to same CPU core
- If CPU core is overloaded single connection packets will be distributed between cores
- Re-assigning packet from one CPU core to another is very “expensive” process
- Processing packet on each separate CPU core might take different amount of time – packet order might change during the processing

Fast Path

- Fast Path allows to forward packets without additional processing in the Linux Kernel. It improves forwarding speeds significantly.
- Fast path requirements
  - Fast Path should be allowed in configuration
  - Interface driver must have support
  - Specific configuration conditions
- Currently RouterOS has fast path handlers for: ipv4 routing, traffic generator, mpls, bridge
- More handlers will be added in future

New Throughput test results

RB951G-2HnD		Gigabit Ethernet test (600Mhz)		RouterOS v6.0rc5			
Mode	Configuration	64 byte		512 byte		1518 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Bridging	none (fast path)	269.6	176.9	232	983.7	81	995.3
Bridging	25 Bridge filter rules	87.6	57.5	86	364.6	81	995.3
Routing	none (fast path)	226.9	148.8	210	890.4	81	995.3
Routing	25						
Routing							

RB1100Hx2				RouterOS v6.			
Mode	Configuration	64 byte		512 byte		1518 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Bridging	none (fast path)	1690	1108.6	704	2985.0	*406*	*4988.9*
Bridging	25 Bridge filter rules	412	270.3	396	1679.0	308	3784.7
Routing	none (fast path)	1496	980.7	704	2985.0	345	4239.4
Routing	25						
Routing							

CCR1036-12G-4S				RouterOS v6.			
Mode	Configuration	64 byte		512 byte		1518 byte	
		kpps	Mbps	kpps	Mbps	kpps	Mbps
Bridging	none (fast path)	*23808*	*15618.0*	*3759*	*15938.2*	*1300*	*15974.4*
Bridging	25 Bridge filter rules	7340	4815.0	*3759*	*15938.2*	*1300*	*15974.4*
Routing	none (fast path)	*23808*	*15618.0*	*3759*	*15938.2*	*1300*	*15974.4*
Routing	25 Simple Queues	7919	5194.9	*3759*	*15938.2*	*1300*	*15974.4*
Routing	25 IP filter rules	3127	2051.3	2998	12711.5	*1300*	*15974.4*

Throughput in millions pps

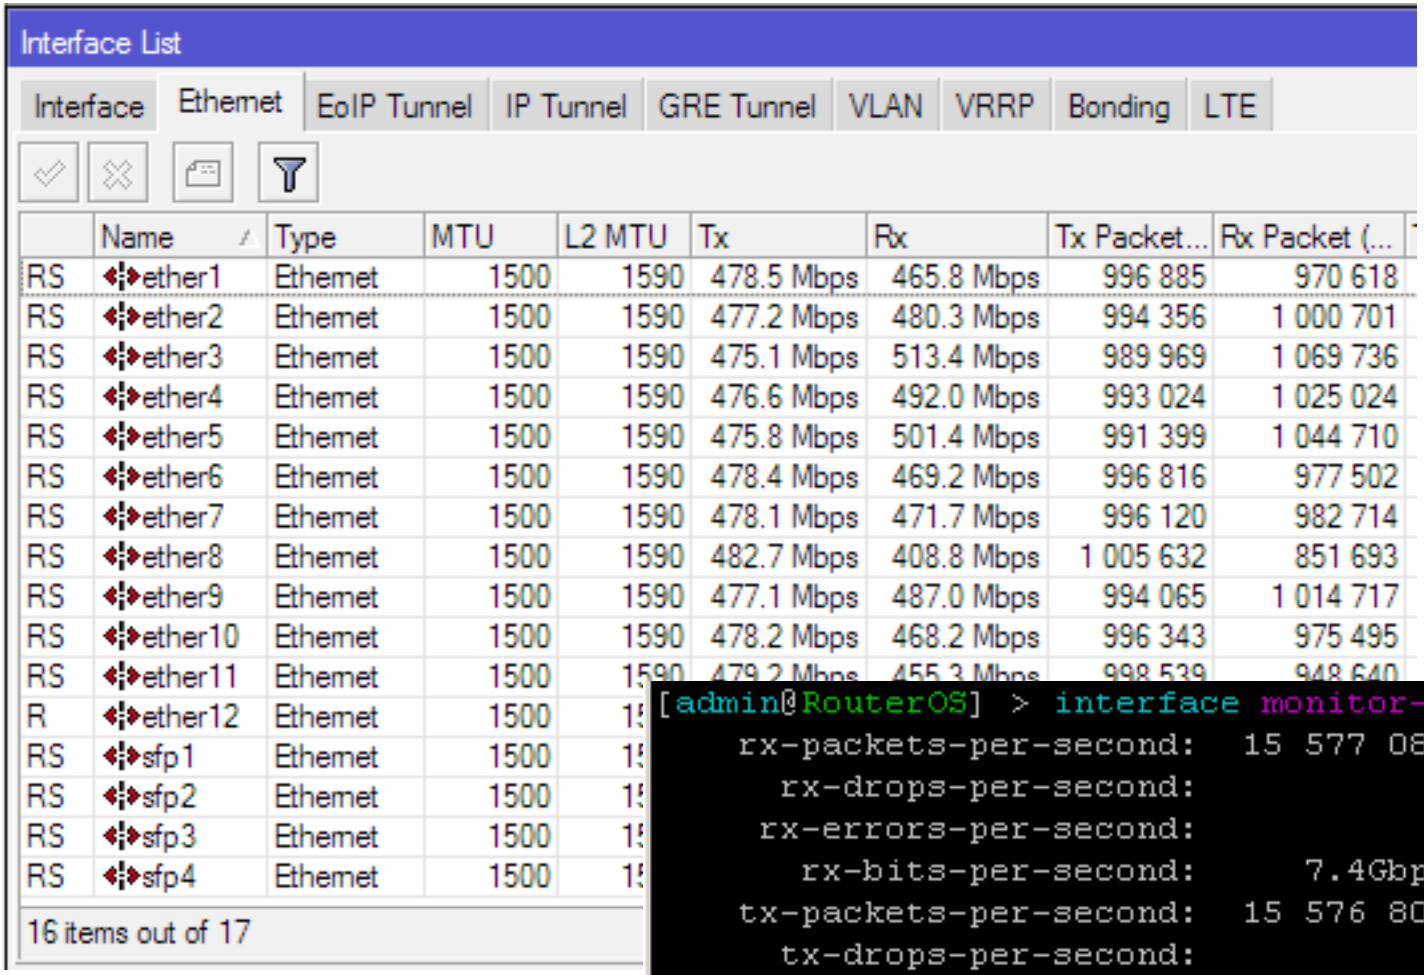

Traffic Generator Tool

- Traffic Generator is a bandwidth-tool evolution
- Traffic Generator can:
  - Determine transfer rates, packet loss
  - Detect out-of-order packets
  - Collect latency and jitter values
  - Inject and replay *.pcap file
  - Working on TCP protocol emulation
- “Quick” mode
- Full Winbox support (coming soon)

Queuing Changes

- Packets can be placed in queue by any number of CPU cores, but processed and taken out of queue only by a single CPU core
- In RouterOS v5.x there was several different places in packets “life-cycle” where it can be queued
- In RouterOS v6.x QoS system was redesigned so that queuing happens is the same place respectively to other processes in the router.
- Now all queuing happens at the very end of packet's “life-cycle” in the router

HTB in RouterOS v5

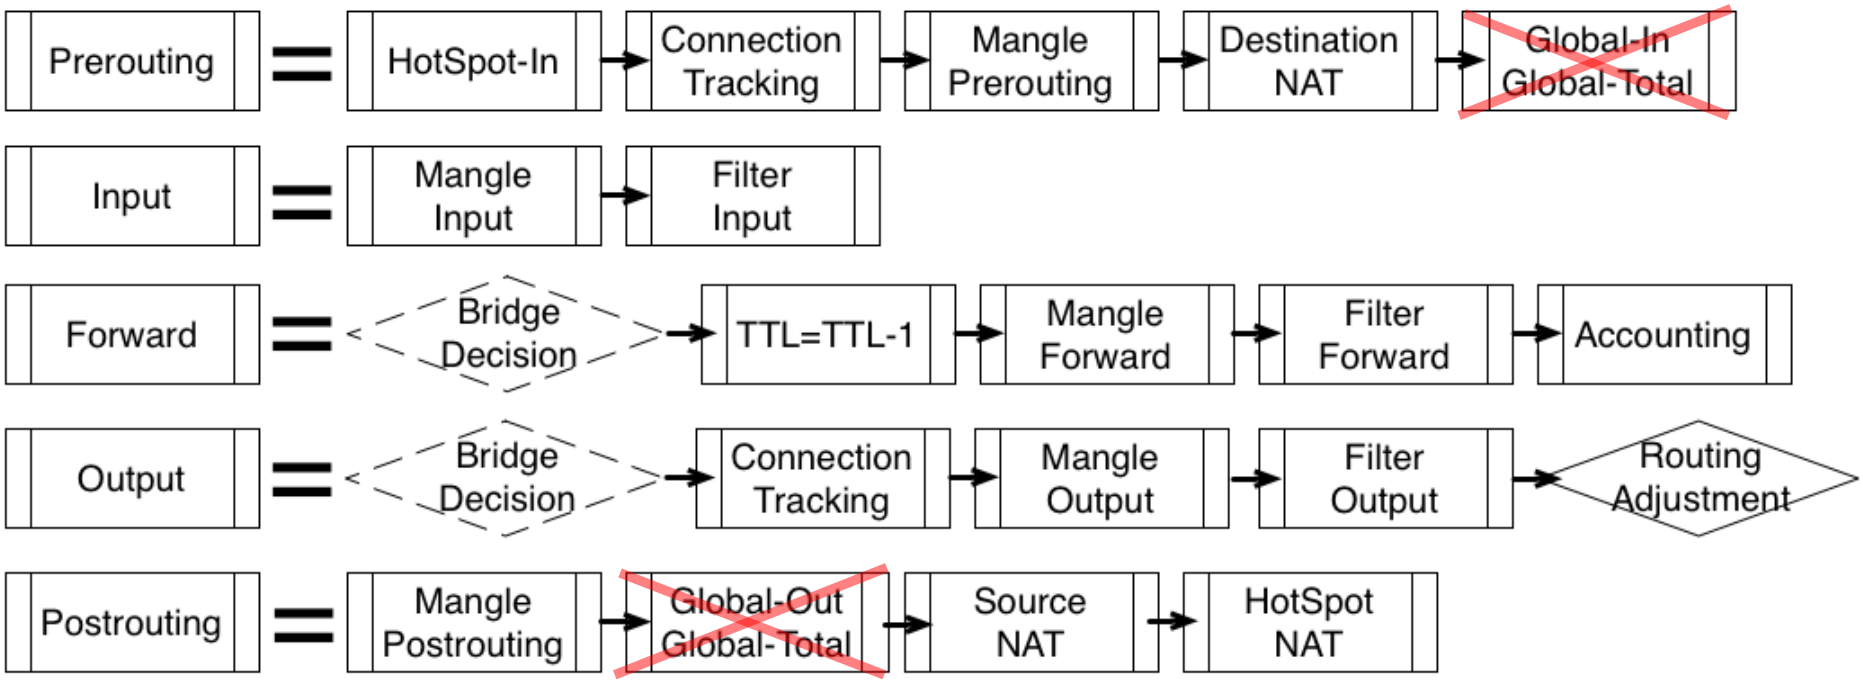

HTB in RouterOS v6

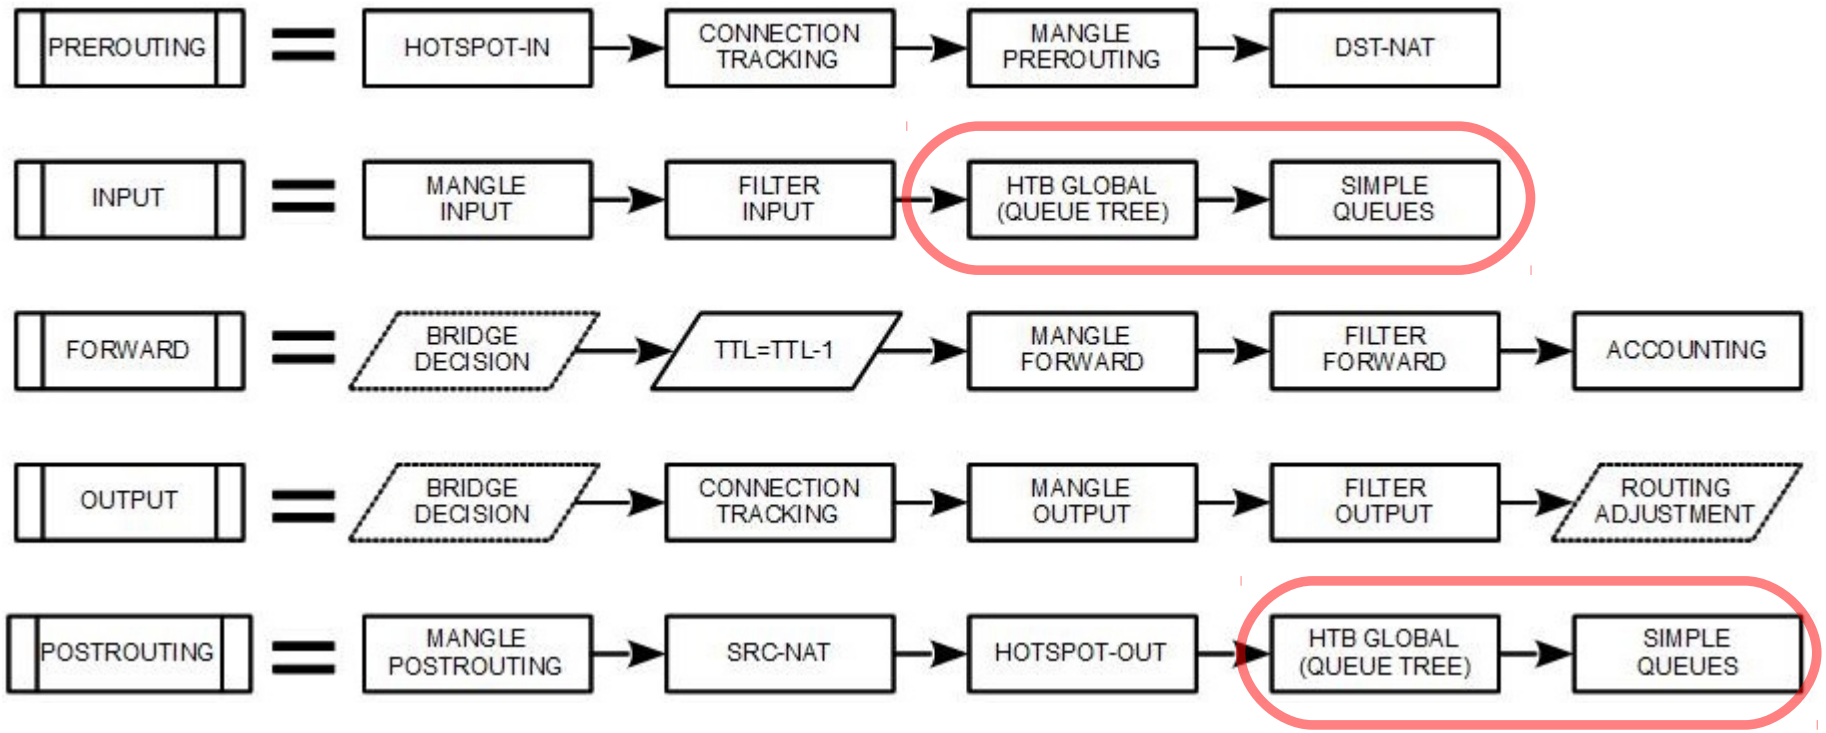

Simple Queues

- Matching algorithm has been updated
  - based on hash
  - faster miss-matches
- At least 32 top level queues are necessary to fully utilize CCR1036 potential (~9x faster than single queue)

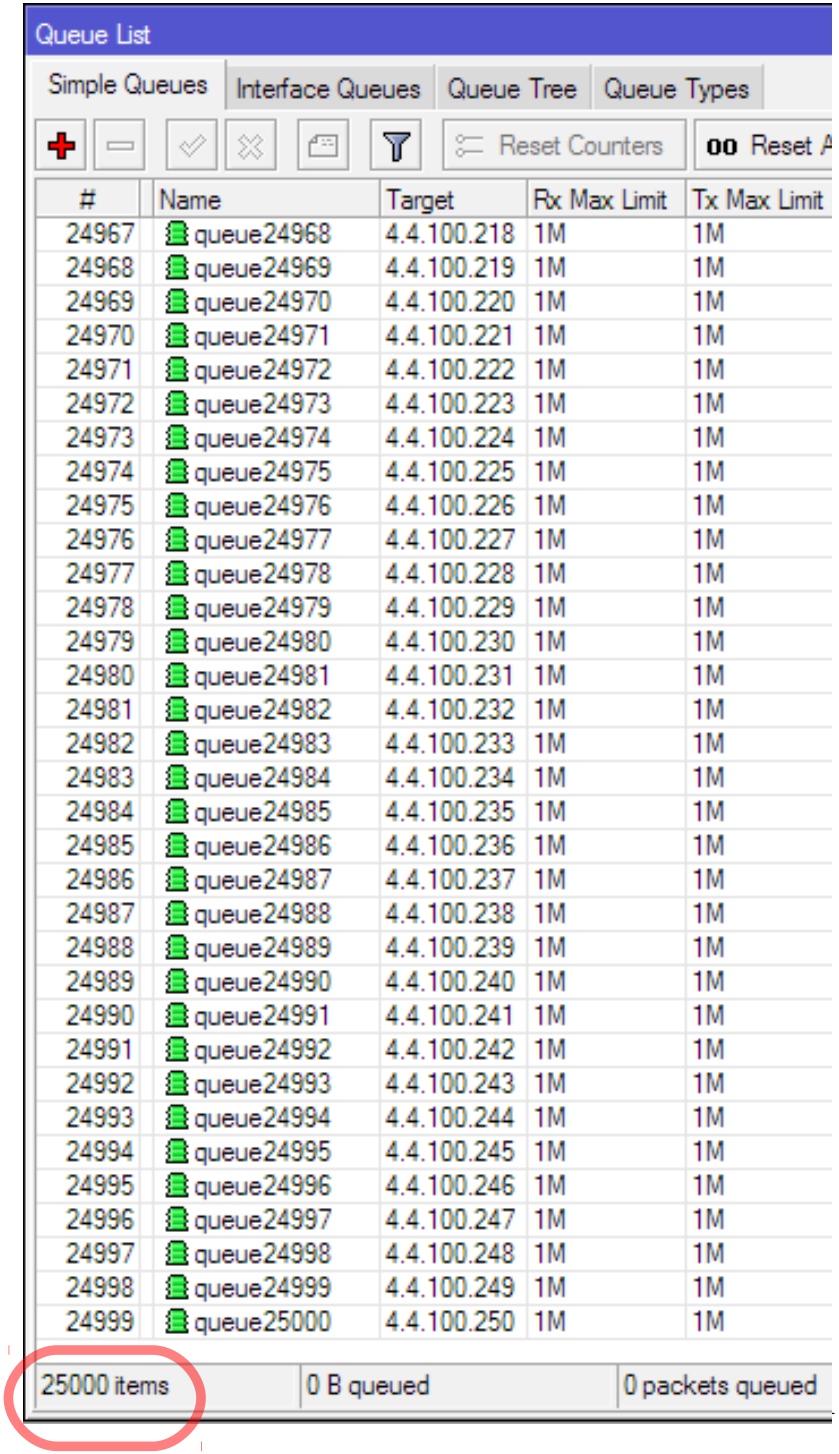

Queue Tree and CCR

- Currently (RouterOS v6.11) only one CPU core can take packets out from one HTB tree
- We are working on possible update of HTB algorithm, or introducing completely new method instead of HTB
- Suggestions:
  - Use Interface HTB as much as possible to offload traffic from HTB “global”
  - Use simple queues

PPTP,L2TP and PPPoE on CCR

- Changes introduced in v6.8:
  - kernel drivers for ppp, pppoe, pptp, l2tp now are lock-less on transmit & receive
  - all ppp packets (except discovery packets) now can be handled by multiple cores
  - MPPE driver now can handle up to 256 out-of-order packets (Previously even single out-of-order packet was dropped)
  - roughly doubled MPPE driver encryption performance

Single PPTP Tunnel Performance on CCR1036

in packets per second with 0,01% loss tolerance

Encryption	Conntrack	Version	64 byte	512 byte	1024 byte	1280 byte	1518 byte
No	No	6.7	135,788	143,454	148,728	147,046	71,400
No	No	6.8rc1	2,123,150	2,336,314	2,305,772	1,865,984	958,942
No	Yes	6.7	98,818	88,294	89,442	100,090	36,826
No	Yes	6.8rc1	1,191,870	1,215,884	1,201,644	1,179,104	32,738
Yes	No	6.7	47,432	38,180	30,830	27,042	13,206
Yes	No	6.8rc1	190,894	233,676	239,462	190,820	80,828
Yes	Yes	6.7	39,062	29,884	22,560	21,062	12,826
Yes	Yes	6.8rc1	189,266	233,676	239,462	190,820	14,192

Single L2TP Tunnel Performance on CCR1036

in packets per second with 0,01% loss tolerance

Encryption	Conntrack	Version	64 byte	512 byte	1024 byte	1280 byte	1518 byte
No	No	6.7	120,906	123,428	197,486	197,846	84,290
No	No	6.8rc1	3,708,612	3,522,140	2,312,554	1,868,990	1,214,036
No	Yes	6.7	98,120	105,438	103,230	101,506	50,160
No	Yes	6.8rc1	1,687,126	1,580,948	1,382,294	1,302,770	53,170
Yes	No	6.7	62,894	47,598	35,628	32,038	19,238
Yes	No	6.8rc1	212,052	234,962	226,366	215,688	91,830
Yes	Yes	6.7	47,432	37,134	29,408	26,614	16,464
Yes	Yes	6.8rc1	212,052	231,364	206,372	208,552	22,782

Single PPPoE Tunnel Performance on CCR1036

in packets per second with 0,01% loss tolerance

Encryption	Conntrack	Version	64 byte	512 byte	1024 byte	1280 byte	1518 byte
No	**No**	**6.7**	294,828	305,358	302,276	305,326	N/A
No	**No**	**6.8rc1**	5,519,320	4,633,852	2,376,862	1,912,372	N/A
No	**Yes**	**6.7**	277,156	260,386	192,272	183,856	83,844
No	**Yes**	**6.8rc1**	2,730,596	2,462,744	2,103,608	1,910,588	80,028
Yes	**No**	**6.7**	71,730	51,084	37,940	33,894	N/A
Yes	**No**	**6.8rc1**	212,052	239,322	228,588	208,552	N/A
Yes	**Yes**	**6.7**	56,286	43,412	33,318	29,754	19,358
Yes	**Yes**	**6.8rc1**	234,840	234,962	226,366	208,552	29,004

CCR and Packet Fragment

- Currently (in RouterOS v6.11) Connection Tracking required packet to be re-assembled before further processing
- It is impossible to ensure that all fragments of the packet is received by the same CPU core
- Process that stores and waits for fragments to re-assemble nullifies all multi-core benefits
- We plan to
  - add full support to Path MTU Discovery to all tunnels and interfaces
  - Update Connection Tracking to handle fragments.

Firewall Efficiency

- Each Firewall rule in RouterOS takes a dedicated place in system memory (RAM)
- CPU need to process a packet through all rules that packet passes before it is captured by a rule
- Reducing Average number of rules that packet need to pass before it is captured can significantly improve your firewall performance
- Make use of action=jump
- Simplify rules.

Changes in the Firewall

- Firewall now has “all-ether”,”all-wireless”,”all-vlan”,”all-ppp” as possibilities in interface matching
- Only 2 dynamic “change-mss” mangle rules are created for “all-ppp” interfaces
- New Mangle Actions “snif-tzsp”,”snif-pc” to send packet stream to remote sniffer.

Layer-7

- Layer-7 is the most “expensive” firewall option, it takes a lot of memory and processing power to match each connection to regexp string.
- Layer-7 should be used only on traffic that can't be identified any different way
- Layer-7 should be used only as trigger - use connection-mark or address-list to keep track of related packets or connections
- Do not use direct action (like accept, drop) in Layer-7 rule

Routing and CCR

- Packet routing can utilize all cores
- All dynamic routing protocols (more precisely - routing table updates and protocol calculations) in RouterOS v6.x are limited to a single core.
  - One BGP full feed will take 1-3min to load on CCR
  - Two BGP full feeds will take 6min to load on CCR
- Try to avoid configurations that continuously updates routing table
- All routing protocols will be updated to multi-core for RouterOS v7

IPSec and CCR

- Hardware acceleration support for aes-cbc + md5|sha1|sha256 Authenticated Encryption with Associated Data (AEAD) was added on CCR in RouterOS v6.8
- Now CCR1036 can handle 3,2Gbps encrypted IPSec traffic
  - Maintaining ~80% CPU load
  - No fragmentation (1470byte packets)
  - Many peers (100 separate tunnels)
  - AES128 was used

Tools

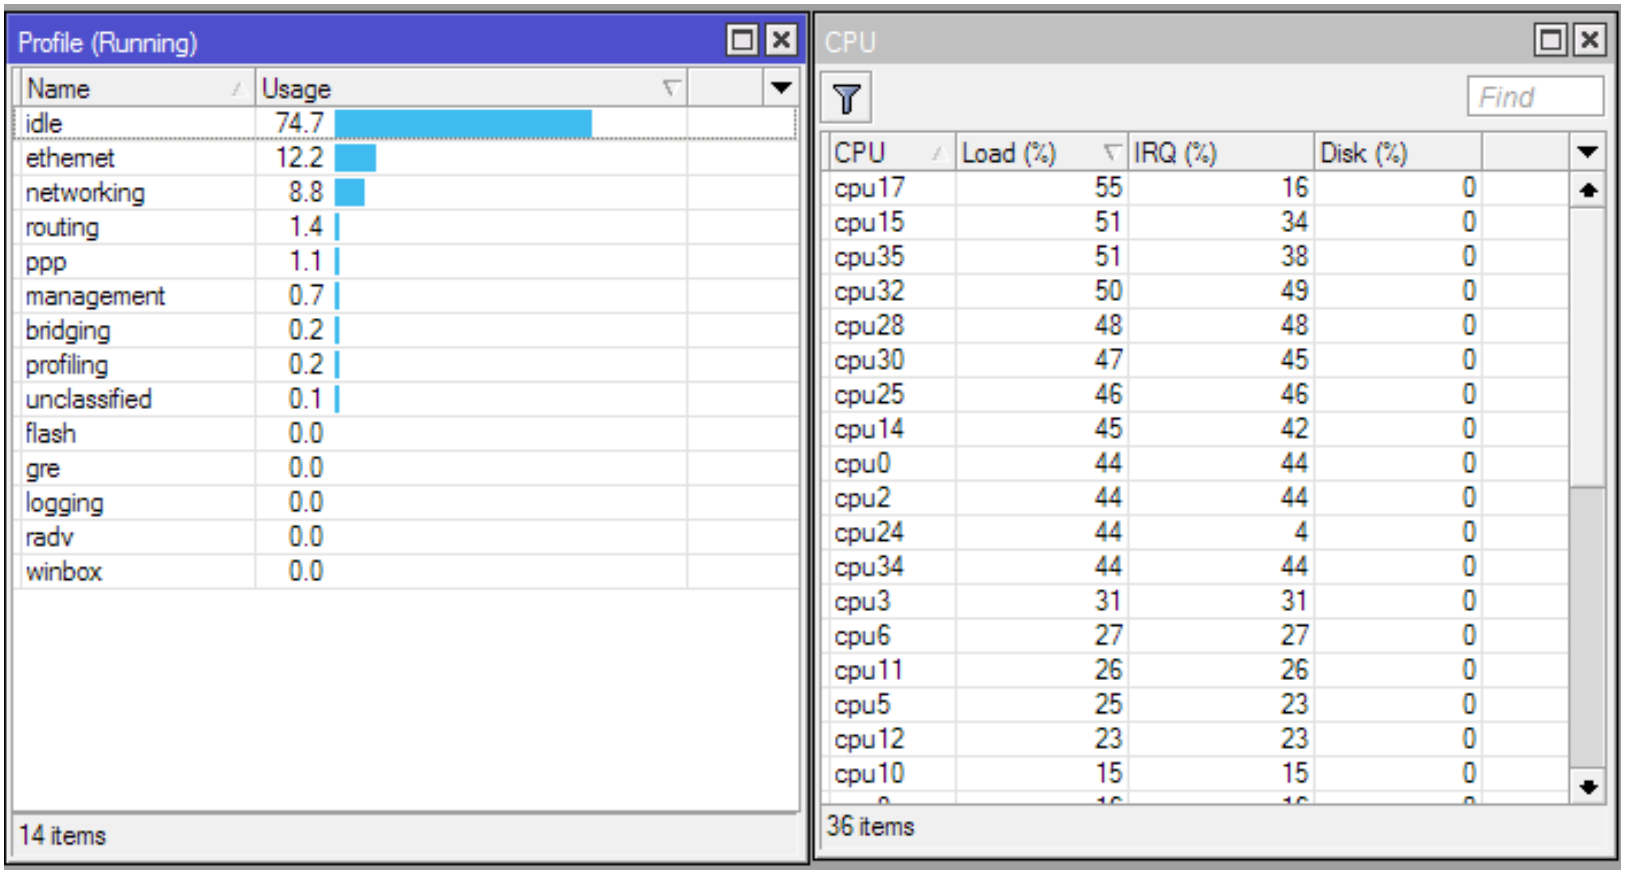

- /system resources cpu
- /tool profile

Partitions

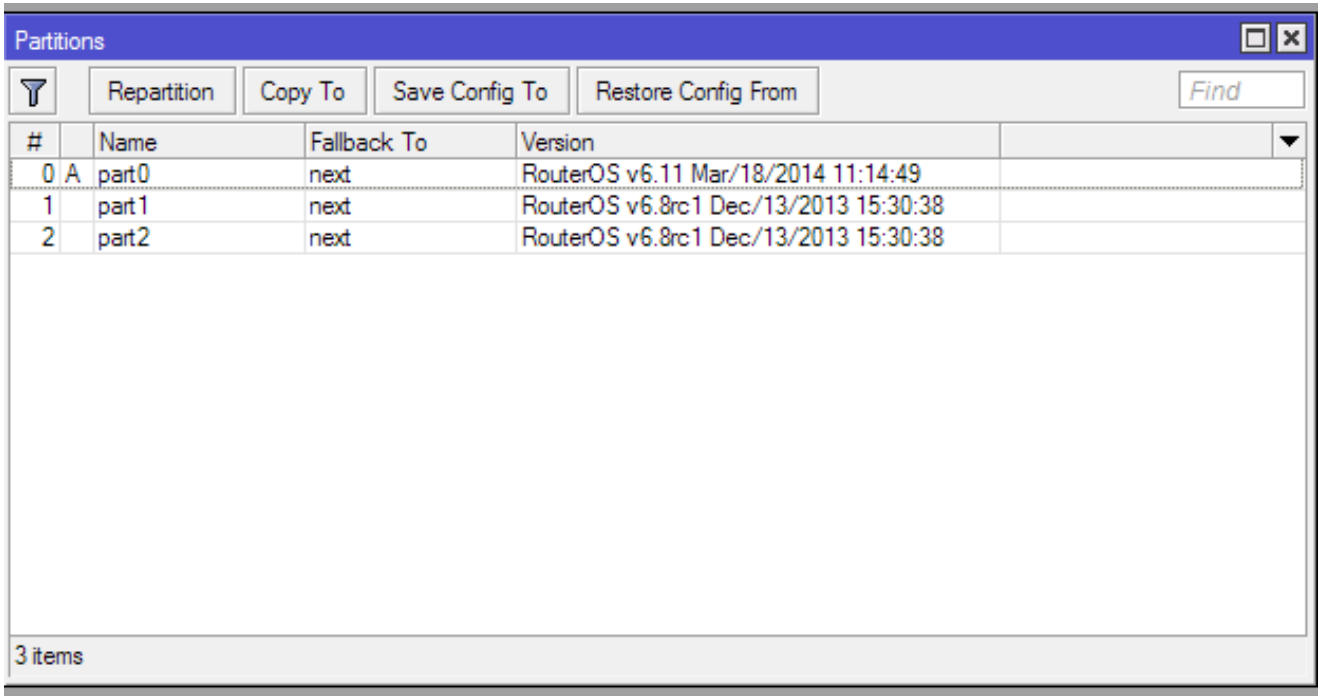

- Partition will always allow you to keep one working copy of RouterOS just one reboot away and backup configuration before mayor changes

Questions!!!