ArchiTechs

MANAGED SERVICES

www.iparchitechs.com

ISP Design – Using MikroTik CHR as a BGP edge router

PRESENTED BY:

KEVIN MYERS,
NETWORK ARCHITECT

Profile: About Kevin Myers

Background:

- 19+ years in Networking
- Designed/Built Networks on 6 continents
- MikroTik Certified Trainer
- MikroTik, Cisco and Microsoft Certified

Community Involvement:

Packet Pushers (Podcast Guest / Blogger)

Group Contributor (RouterOS / WISP Talk and others)

Delegate/Roundtable contributor (NFD14)

MT Forum (Forum Veteran – Member since 2012)

Network Collective (Podcast Guest)

Profile: About IP ArchiTechs

IP ArchiTechs
MANAGED SERVICES

Expert Networking

Whitebox | ISP | Data Center | Enterprise

- ✓ Global Consulting
- ✓ Managed Networks
- ✓ Monitoring
- ✓ Load Testing
- ✓ Development

Locations in: US | Canada | South America

Call us at: +1 855-645-7684
E-mail: consulting@iparchitechs.com
Web: www.iparchitechs.com

Design: **Why use the CHR as a BGP edge router?**

Goal of this presentation: When the presentation is finished, hopefully you will have walked away with a few key concepts:

- Performance characteristics of the CHR on multiple hypervisors
- Best practices when deploying the CHR
- Benefits of using the CHR vs CCR as a BGP edge router

Design: **CHR vs. CCR for full BGP tables?**

- Which platform is better?
- Throughput capabilities?
- x86 CPU vs. Tilera?
- BGP Updates, Set BGP Attributes (Community, Localpref)

MikroTik CHR

VS.

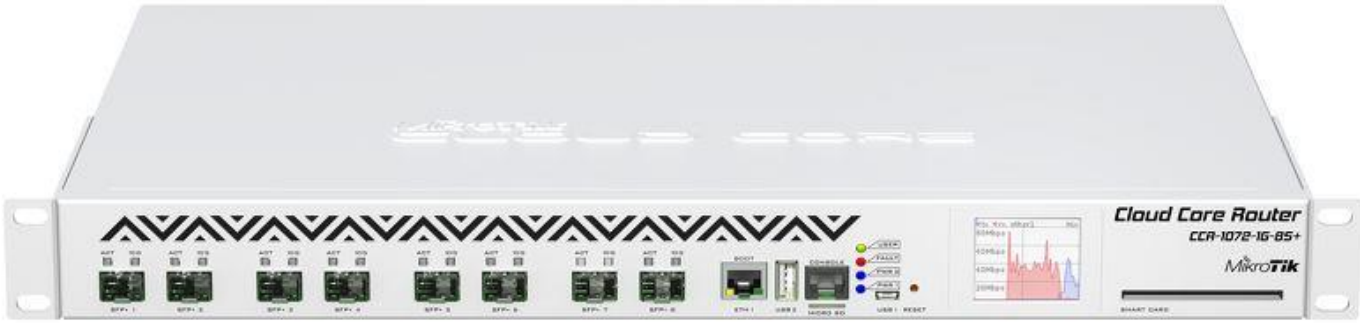

Design: CHR vs. CCR for full BGP tables?

Platform		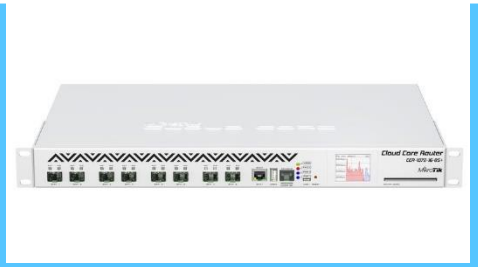
CPU BGP routers with full tables have a high computational requirement.	**x86** Better for heavy computational work. Higher power draw.	**Tilera** Optimized for packet transfer. Designed to be low power draw.
Throughput At 1500 bytes, 512 bytes and 64 bytes	**x86** More CPU and power is required to move data at the same speed as a CCR	**Tilera** Handles throughput at different frame sizes slightly better than x86
Routing table size Impact of multiple tables	**x86** Not limited by ASIC, x86 can process multiple tables more quickly than CCR if the HV host has a high clock speed CPU.	**Tilera** Not currently limited by ASIC. Limited to fixed speed single core performance

Design: CHR testing lab setup

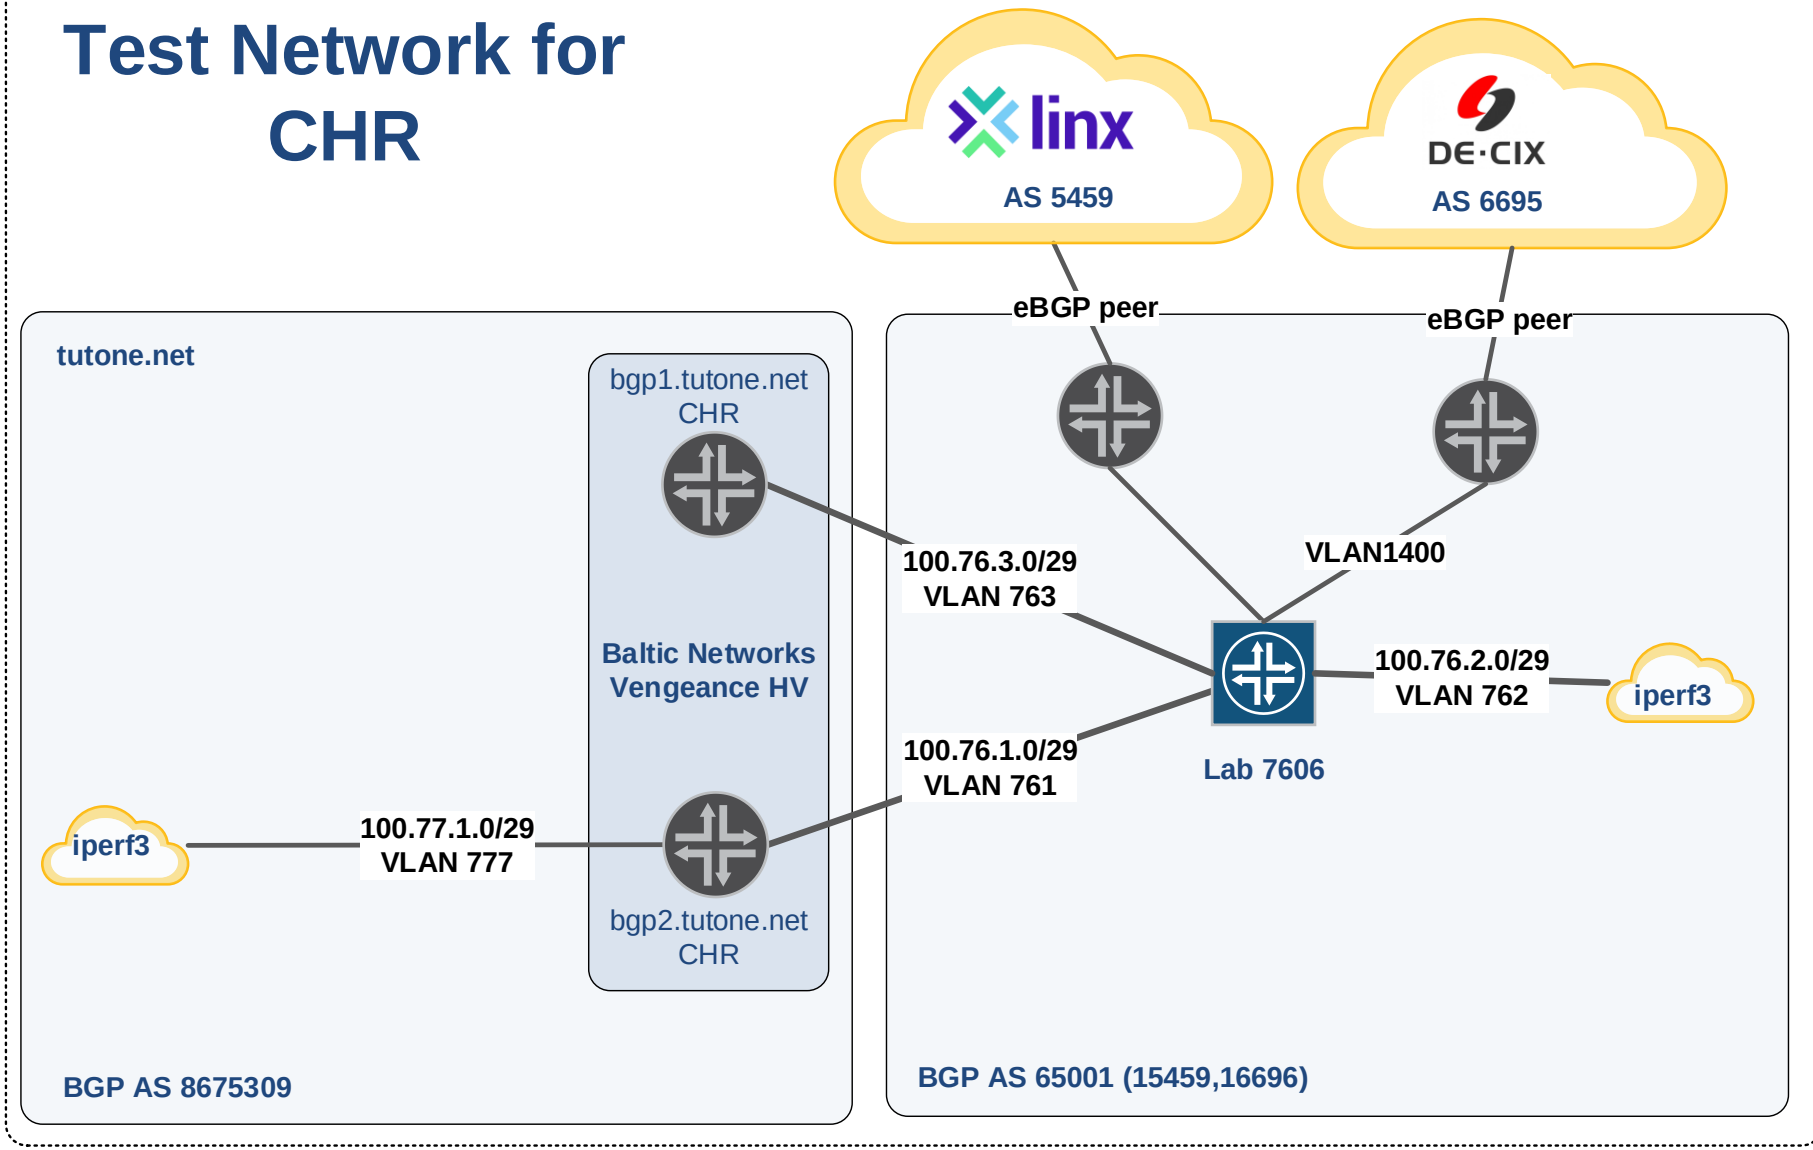

Design: CHR testing lab setup

Physical Test Network for CHR

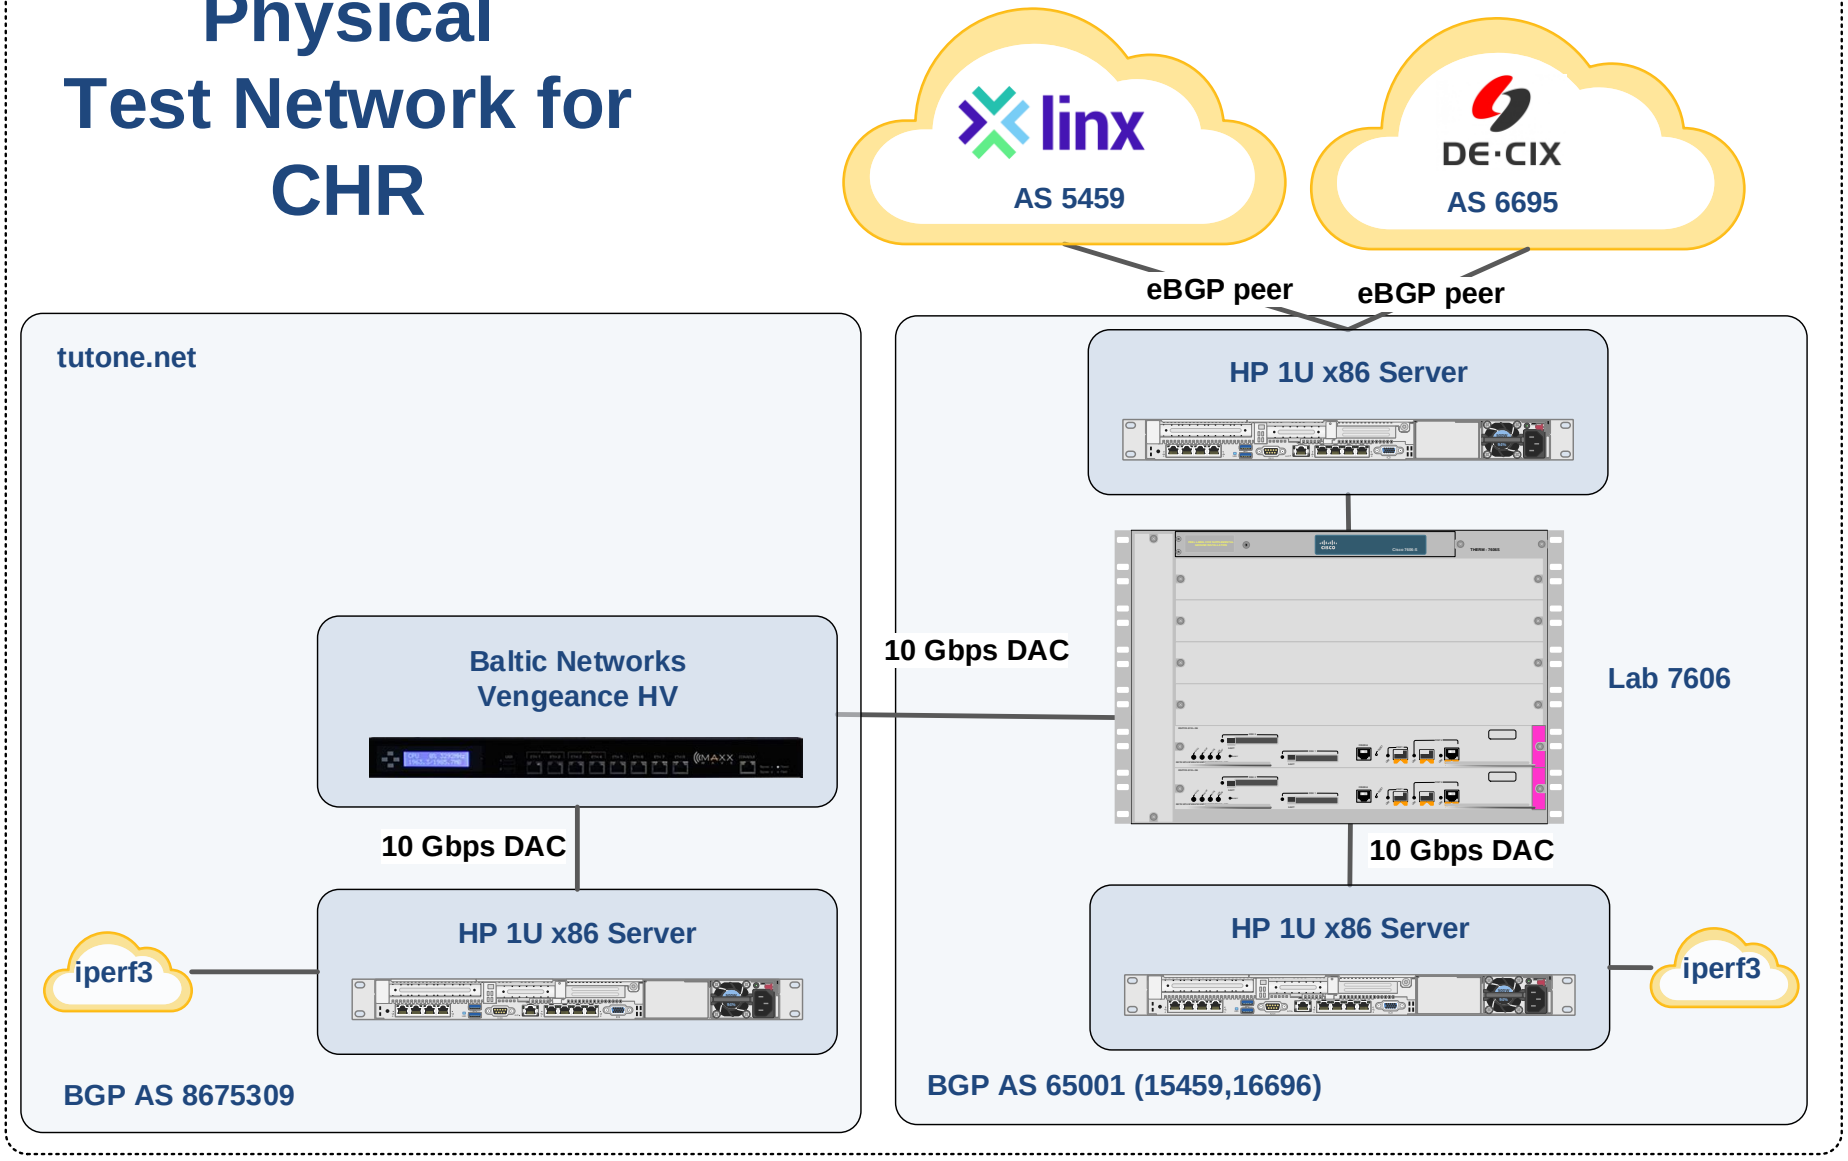

Design: **CHR testing lab setup**

- Hypervisor details – VM provisioning
- 8 vCPUs Total
  - Performance was better when using more sockets vs. more cores with a single socket
  - With BGP – single core performance is critical. You need to use a CPU with the most powerful single core performance that you can afford for the hypervisor host.
- 4096 MB RAM

- **Concept of testing**
  - Performance on 1 BGP full table
  - Performance on 2 BGP full tables
  - Performance at 1500 bytes, 512 bytes and 96 bytes
  - Throughput performance with convergence test
- **Performance settings**
  - Sockets vs CPUs – More sockets is better
  - TSO/LSO – Disable for best performance
  - Clock speed – Highest speed possible

Design: CHR performance on VMWARE ESXi (1 Full Table)

Platform	Hypervisor	CHR
Baltic Vengeance	**VM Ware ESXi 6.5**	**6.41.3**

DE-CIX March 2018 Routing Table: **Time - 26 seconds**

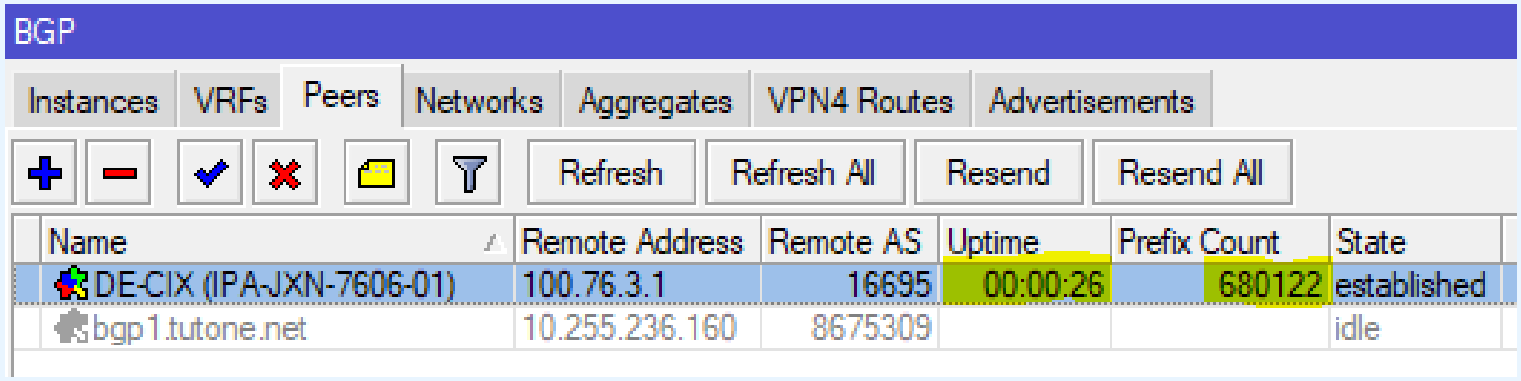

Design: CHR performance on VMWARE ESXi (2 Full Tables)

Platform	Hypervisor	CHR
Baltic Vengeance	VM Ware ESXi 6.5	6.41.3

DE-CIX and LINX March 2018 Routing Table: **Time – 4:46 seconds**

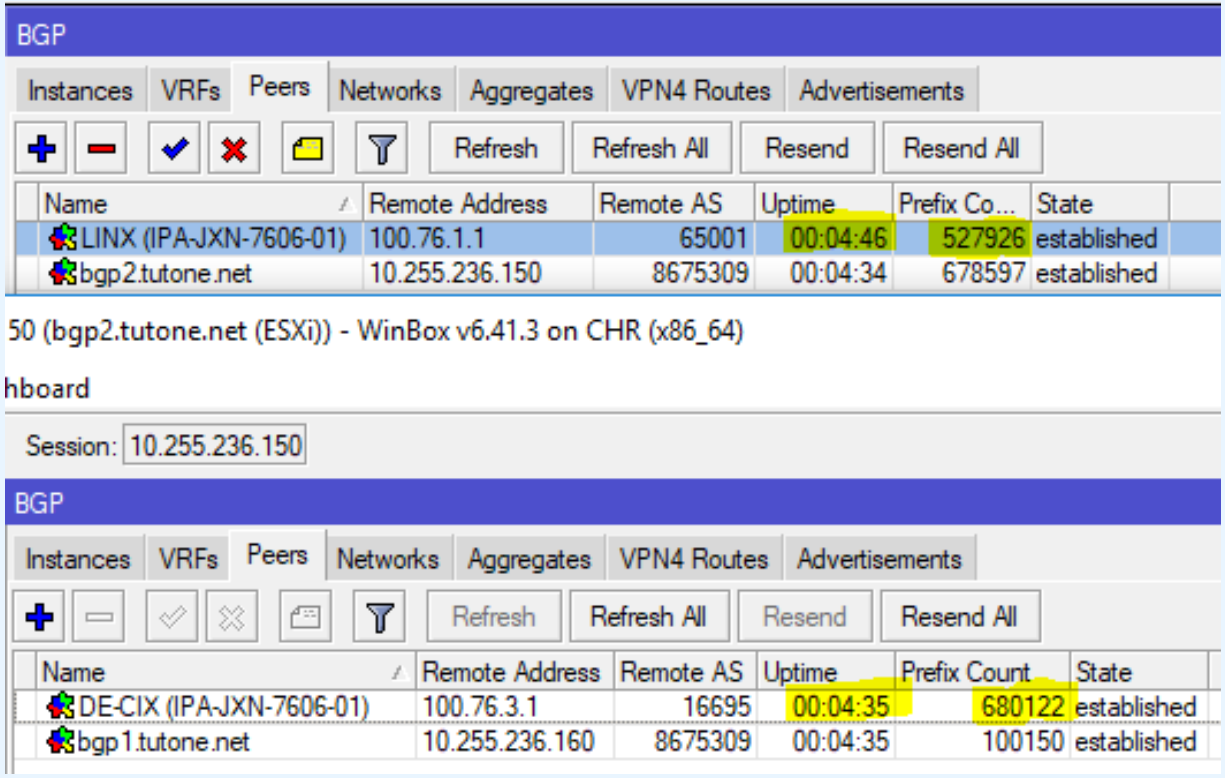

Design: CHR performance on VMWARE ESXi (1500 byte MTU)

Platform	Hypervisor	CHR
Baltic Vengeance	VM Ware ESXi 6.5	6.41.3

Throughput: **5.3 Gbps** Peak VM CPU: **28%**

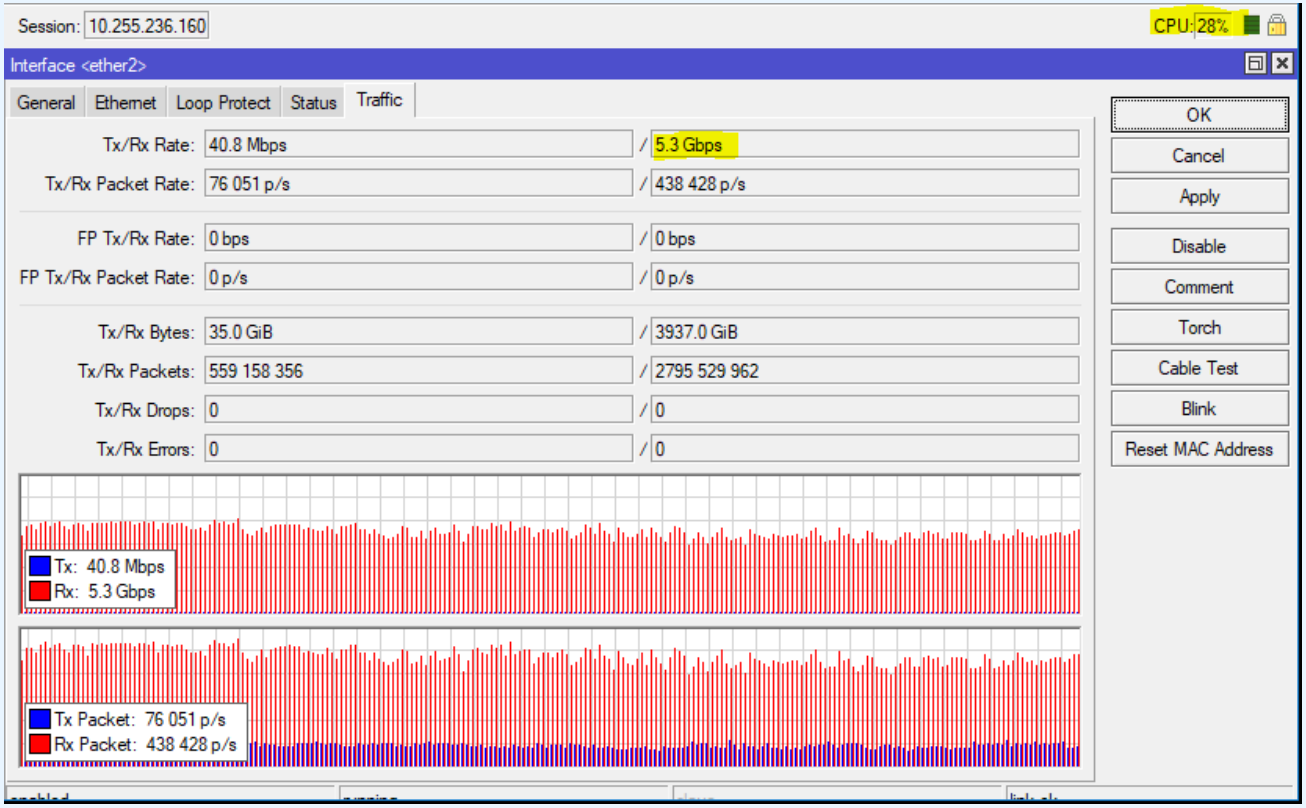

Design: CHR performance on VMWARE ESXi (512 byte MTU)

Platform	Hypervisor	CHR
Baltic Vengeance	VM Ware ESXi 6.5	6.41.3

Throughput: **3.0 Gbps** Peak VM CPU: **38%**

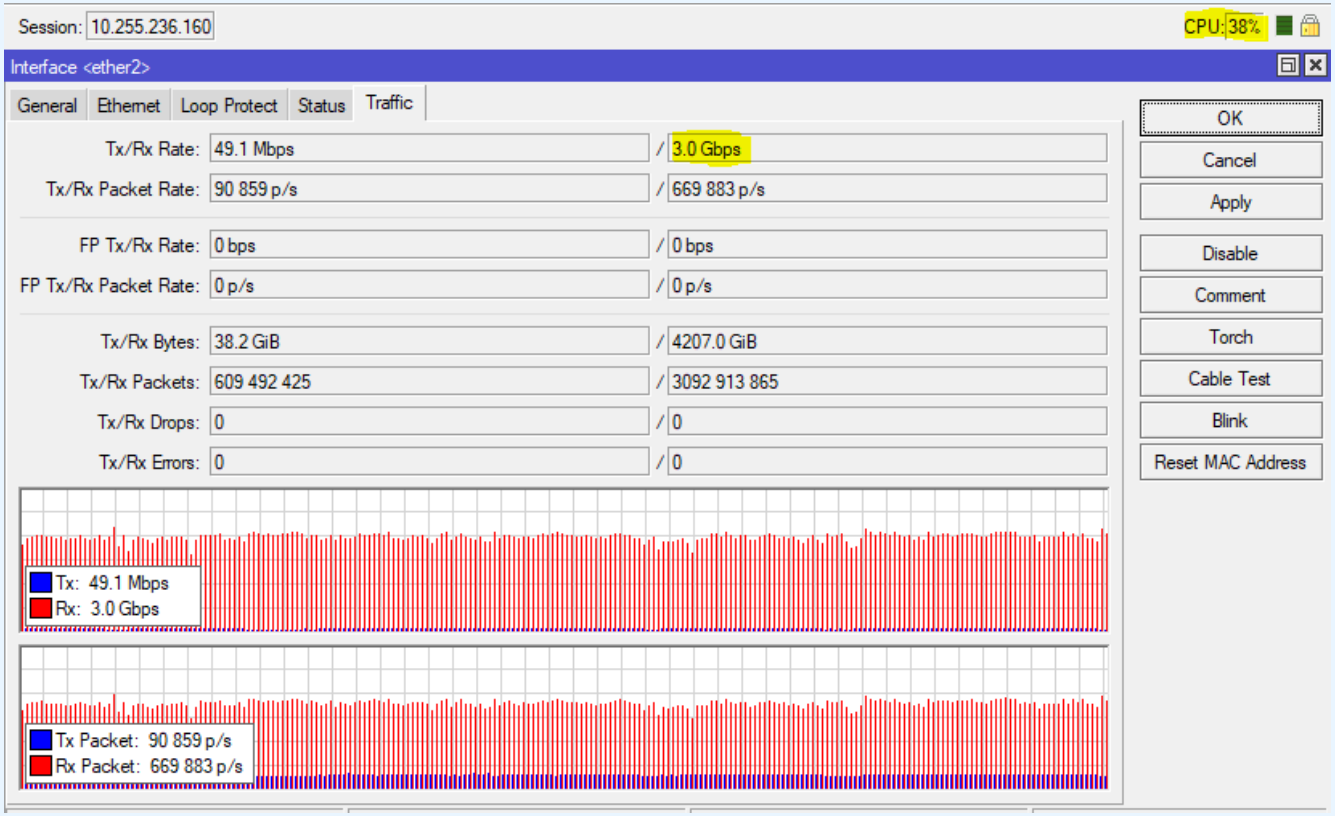

Design: CHR performance on VMWARE ESXi (96 byte MTU)

Platform	Hypervisor	CHR
Baltic Vengeance	VM Ware ESXi 6.5	6.41.3

Throughput: **1.1 Gbps** Peak VM CPU: **43%**

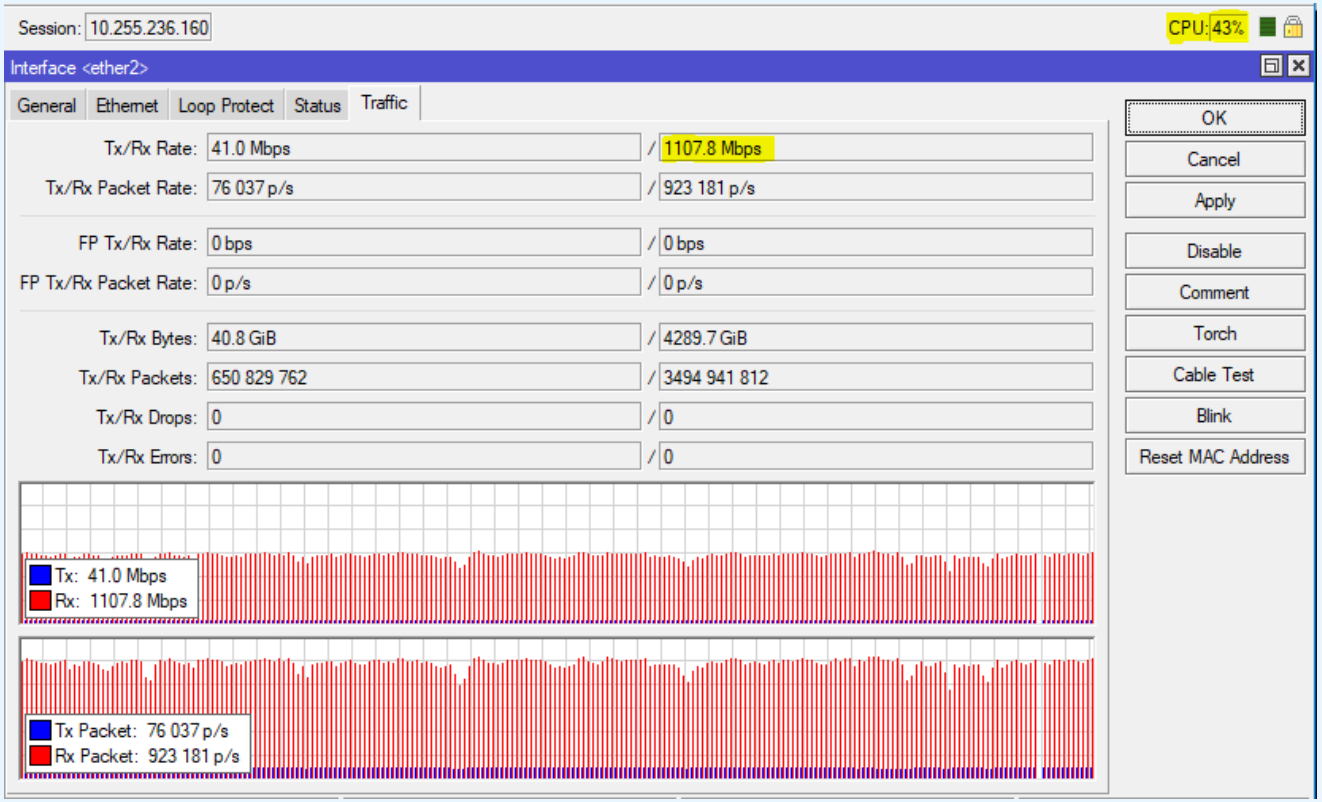

Design: CHR performance on VMWARE ESXi (full table + bw load)

Platform	Hypervisor	CHR
Baltic Vengeance	VM Ware ESXi 6.5	6.41.3

5 Gbps load plus reset of both upstreams: **Convergence Time – 11:05**

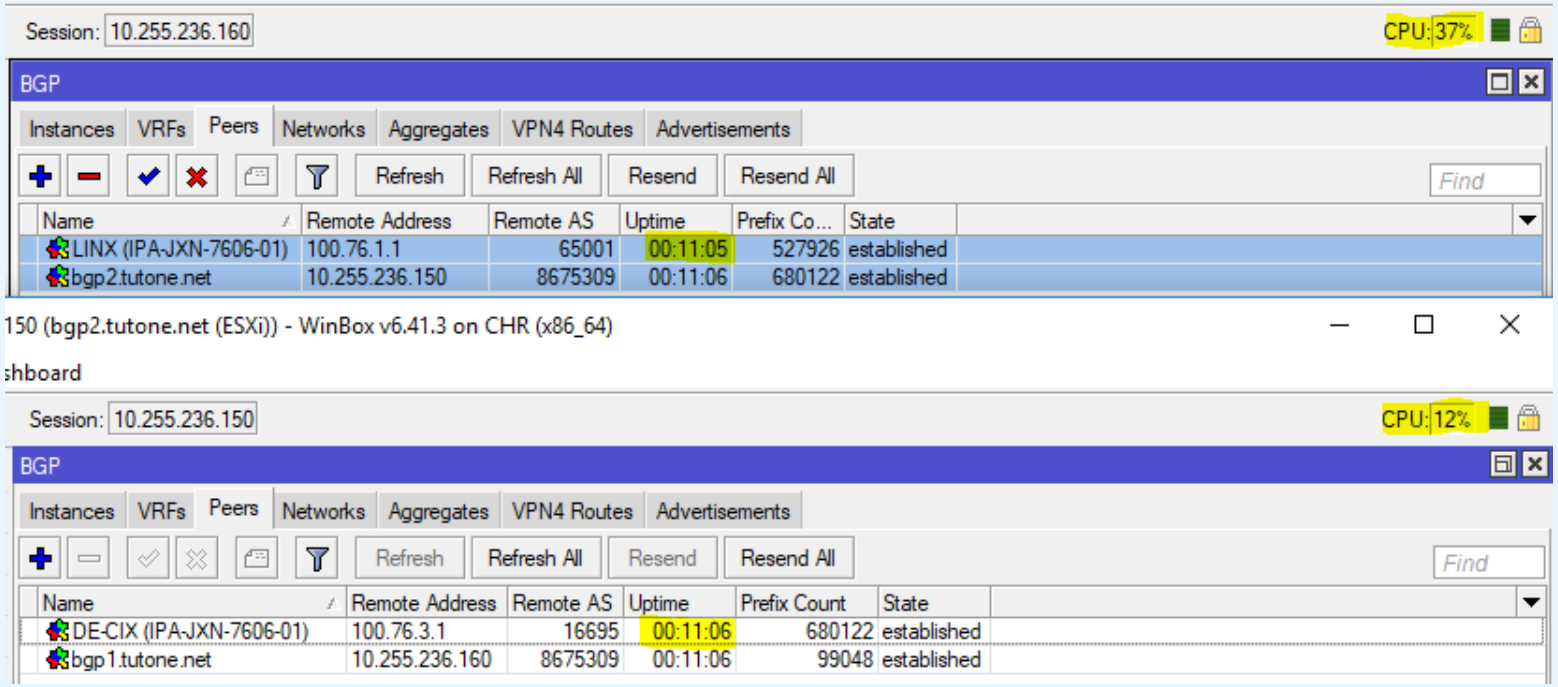

Design: CHR performance on Proxmox (KVM) (1 Full Table)

Platform	Hypervisor	CHR
Baltic Vengeance	**ProxMox (KVM) 5.1**	**6.41.3**

DE-CIX March 2018 Routing Table: **Time - 26 seconds**

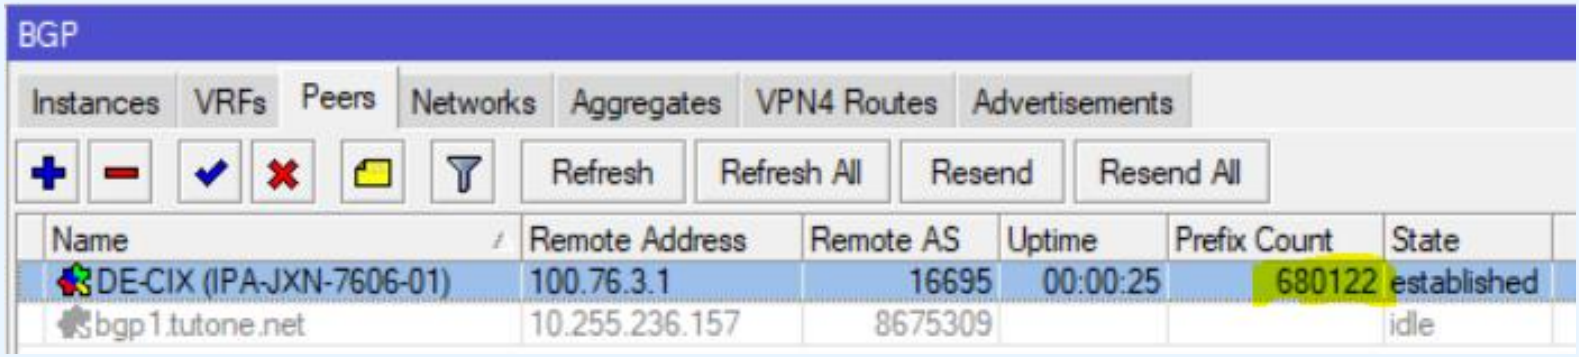

Design: CHR performance on Proxmox (KVM) (2 Full Tables)

Platform	Hypervisor	CHR
Baltic Vengeance	ProxMox (KVM) 5.1	6.41.3

DE-CIX and LINX March 2018 Routing Table: **Time – 1:34**

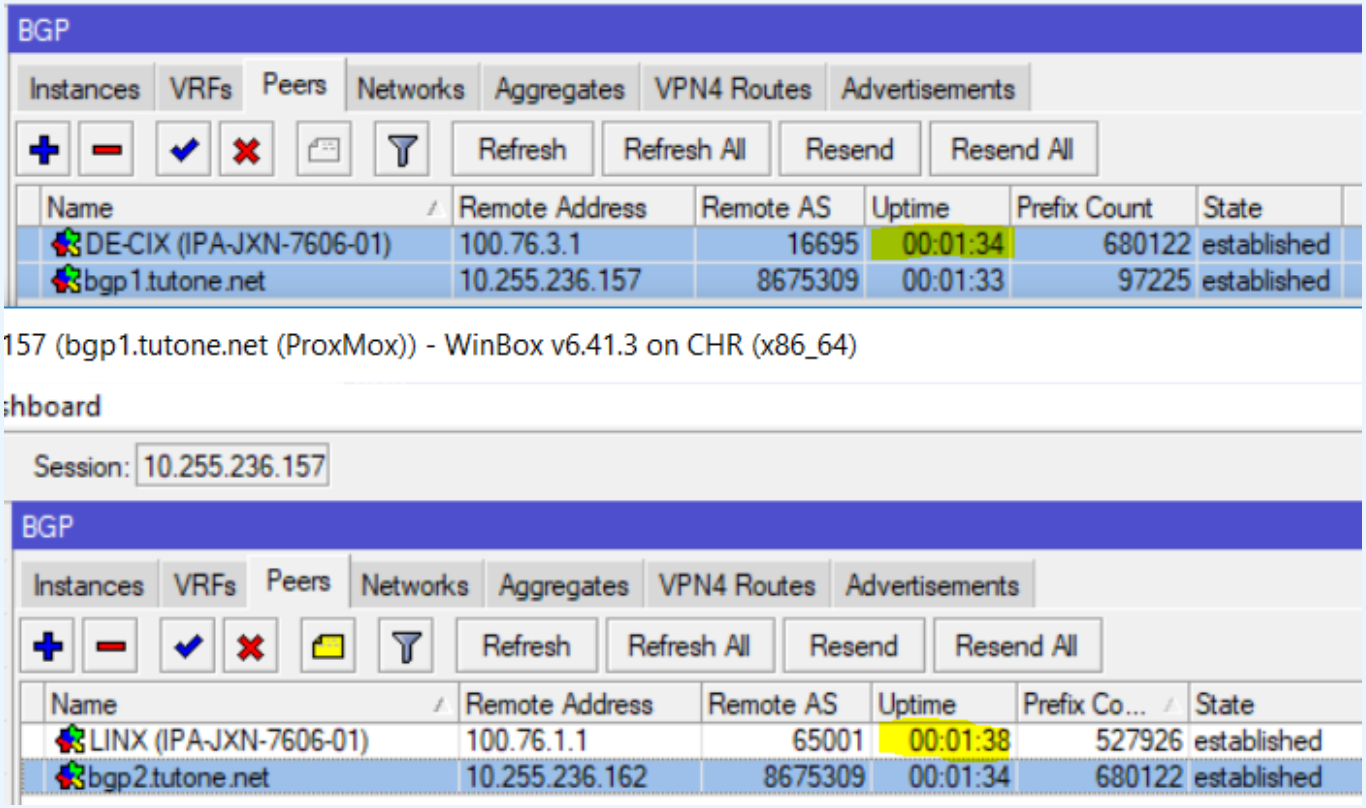

Design: CHR performance on Proxmox (KVM) (1500 byte MTU)

Platform	Hypervisor	CHR
Baltic Vengeance	ProxMox (KVM) 5.1	6.41.3

Throughput: **4.4 Gbps** Peak VM CPU: **28%**

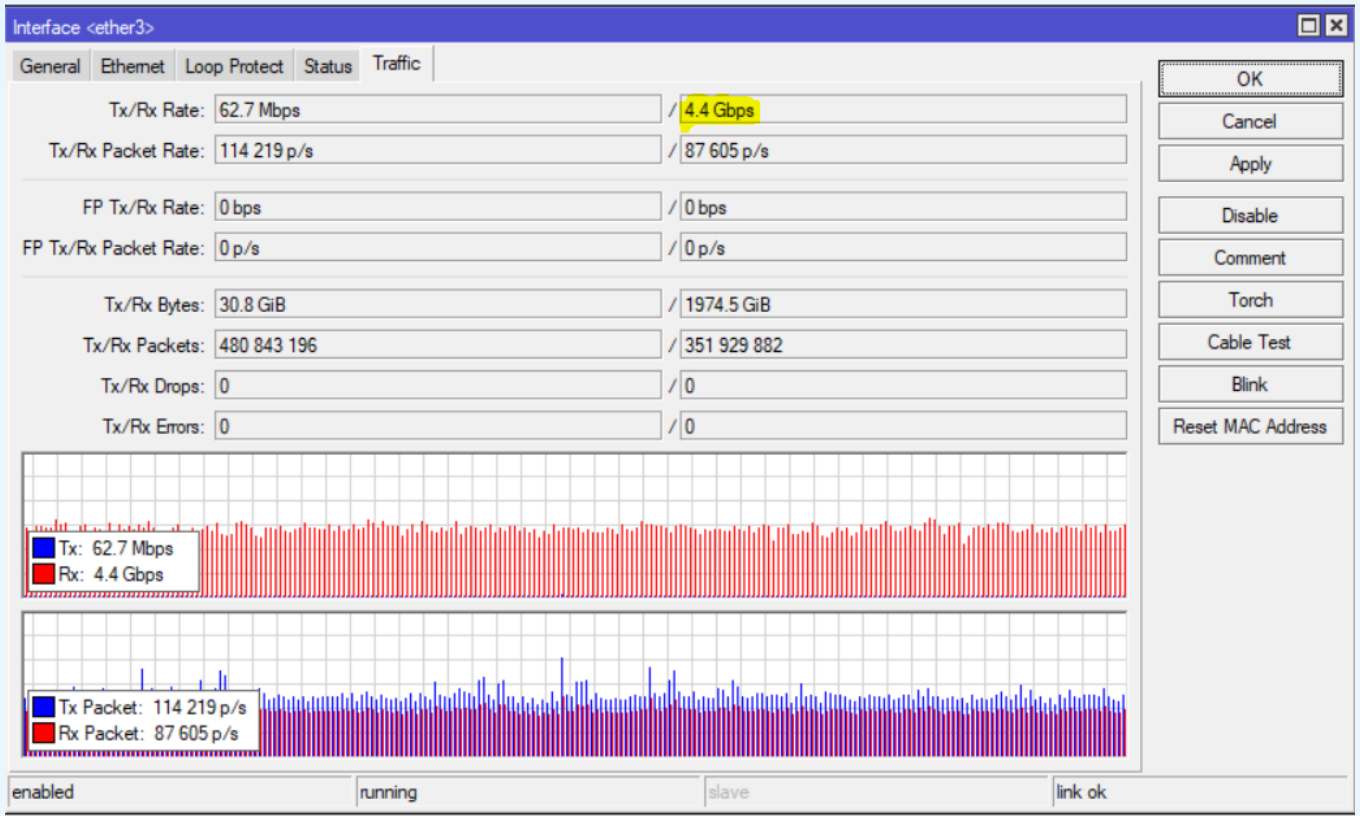

Design: CHR performance on Proxmox (KVM) (512 byte MTU)

Platform	Hypervisor	CHR
Baltic Vengeance	ProxMox (KVM) 5.1	6.41.3

Throughput: **1.8 Gbps** Peak VM CPU: **29%**

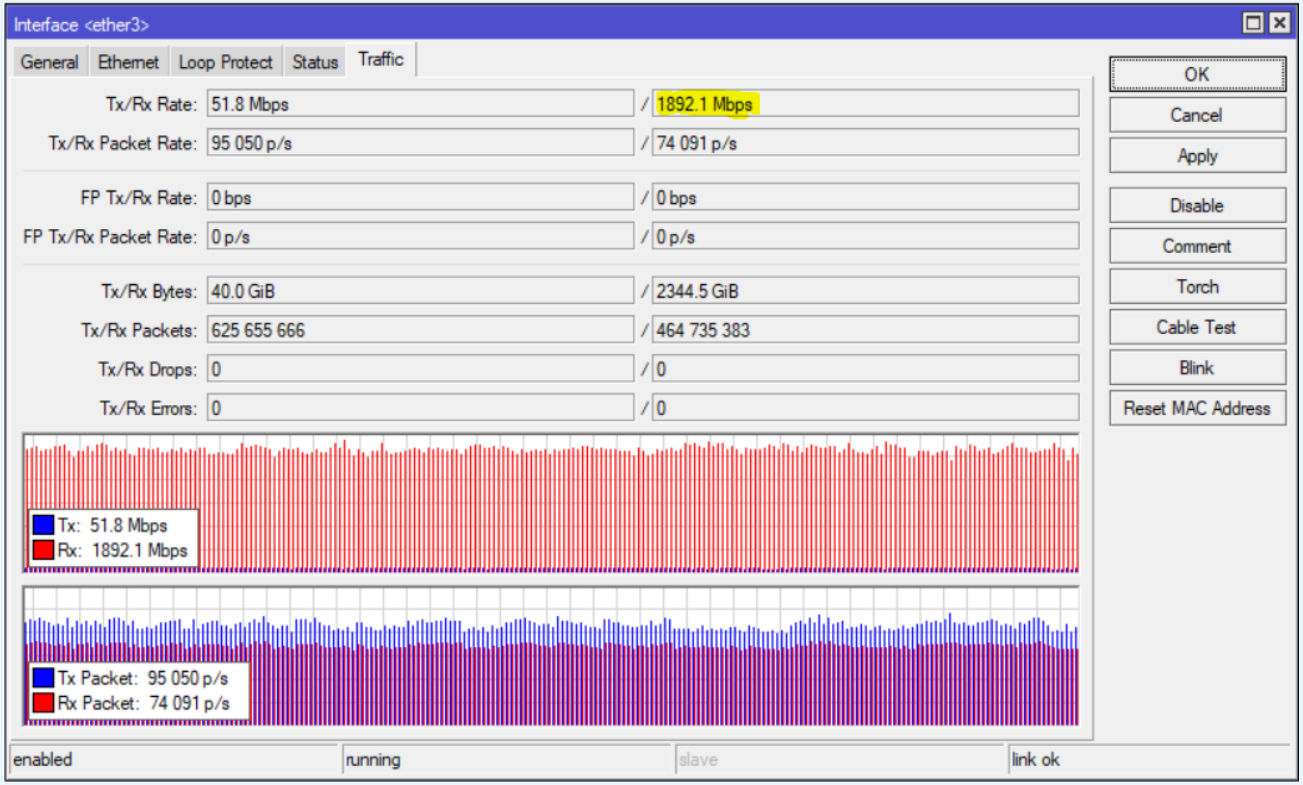

Design: CHR performance on Proxmox (KVM) (96 byte MTU)

Platform	Hypervisor	CHR
Baltic Vengeance	ProxMox (KVM) 5.1	6.41.3

Throughput: **423 Mbps** Peak VM CPU: **28%**

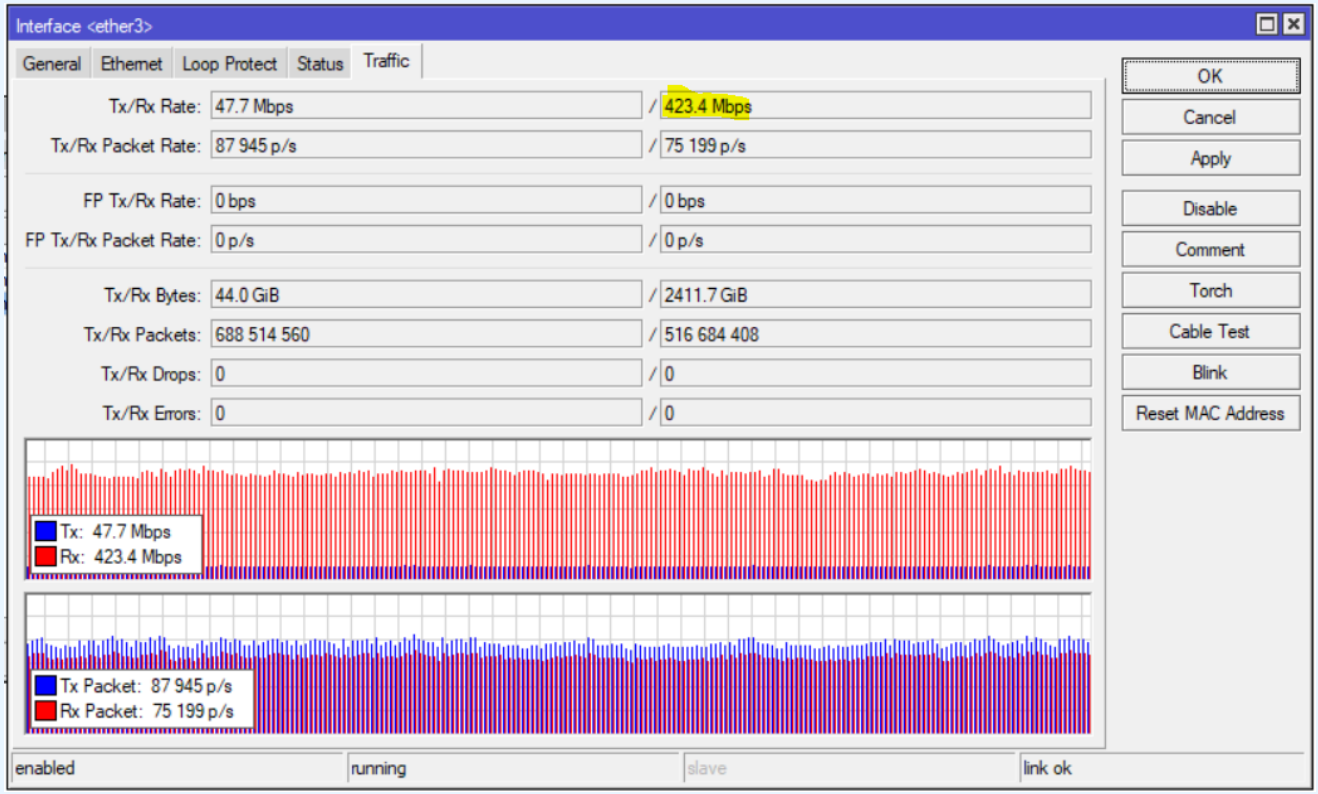

Design: CHR performance on PROXMOX (KVM)(full table + bw load)

Platform	Hypervisor	CHR
Baltic Vengeance	VM Ware ESXi 6.5	6.41.3

5 Gbps load plus reset of both upstreams: **Convergence Time – 9:03**

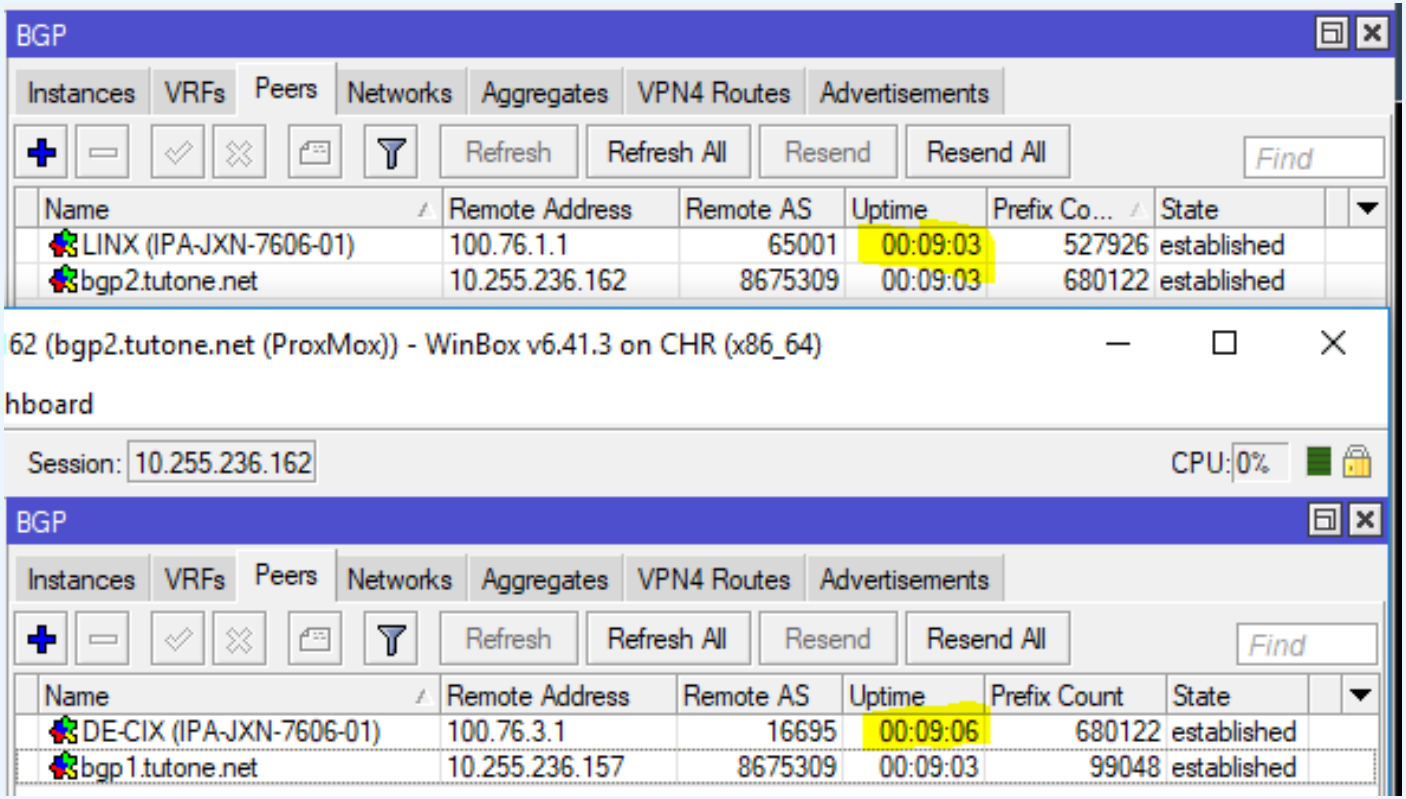

Design: CHR performance on Hyper V (Microsoft)(1 full table)

Platform	Hypervisor	CHR
Baltic Vengeance	Hyper-V 2012	6.41.3

DE-CIX March 2018 Routing Table: **Time - 12 seconds**

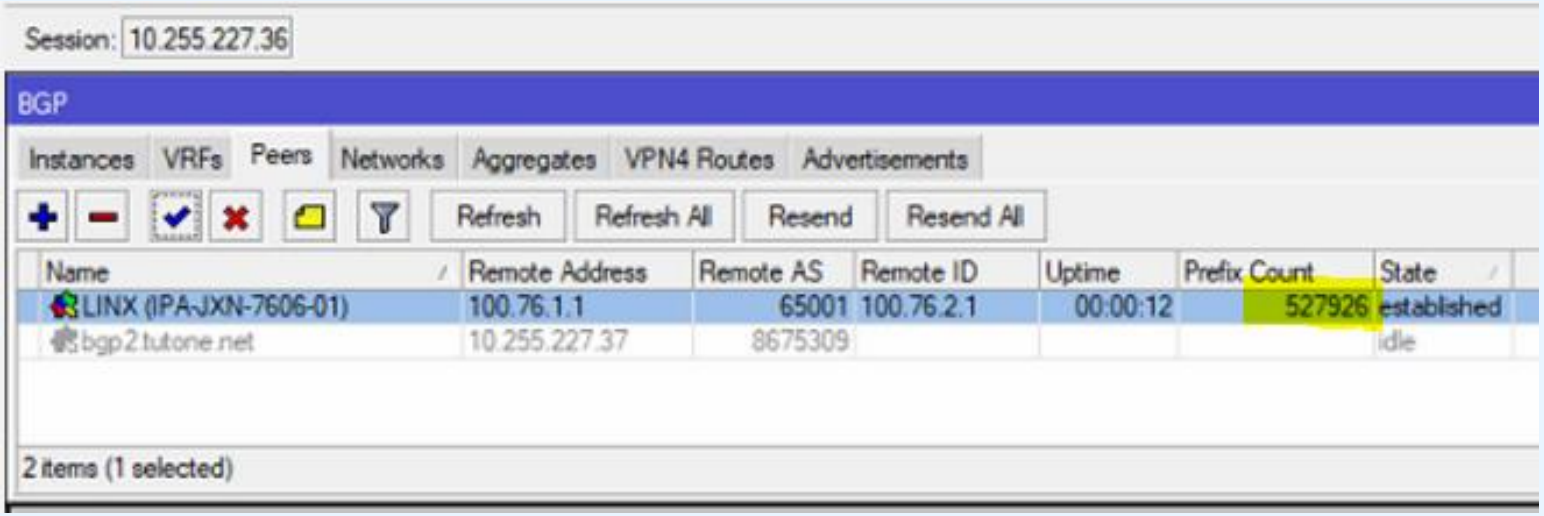

Design: CHR performance on Hyper V (Microsoft)(2 full tables)

Platform	Hypervisor	CHR
Baltic Vengeance	Hyper-V 2012	6.41.3

DE-CIX and LINX March 2018 Routing Table: **Time - 41 seconds**

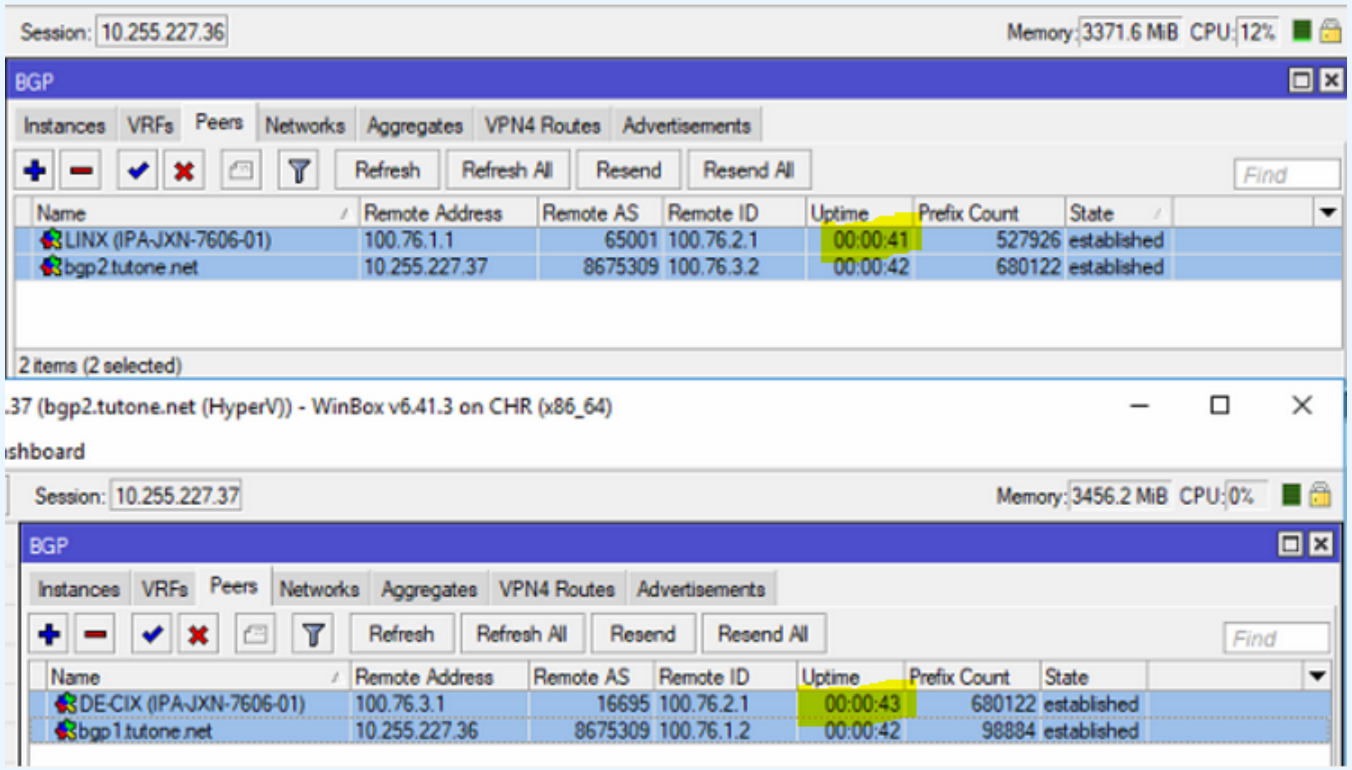

Design: CHR performance on Hyper-V (Microsoft)(1500 byte MTU)

Platform	Hypervisor	CHR
Baltic Vengeance	Hyper-V 2012	6.41.3

Throughput: **4.4 Gbps** Peak VM CPU: **32%**

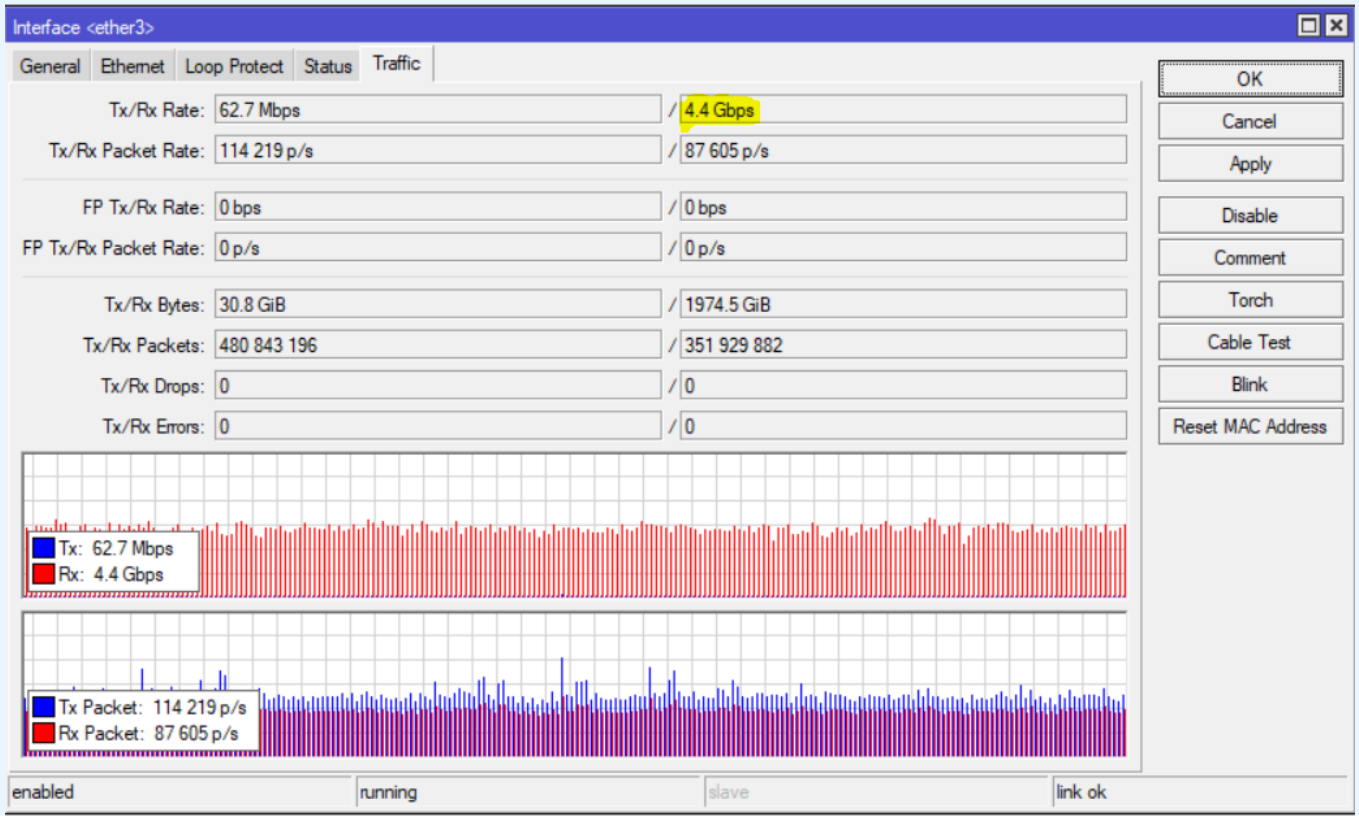

Design: CHR performance on Hyper-V (Microsoft)(512 byte MTU)

Platform	Hypervisor	CHR
Baltic Vengeance	Hyper-V 2012	6.41.3

Throughput: **1.6 Gbps** Peak VM CPU: **33%**

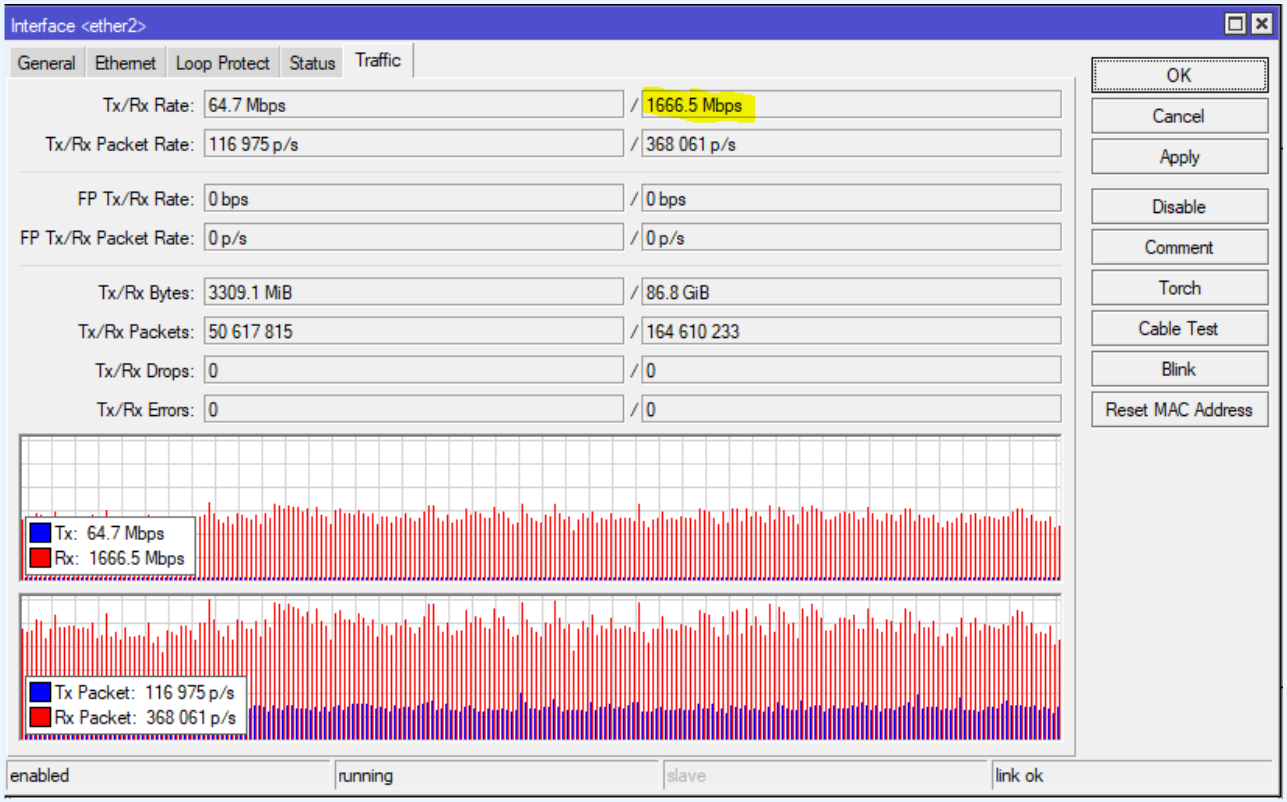

Design: CHR performance on Hyper-V (Microsoft)(96 byte MTU)

Platform	Hypervisor	CHR
Baltic Vengeance	Hyper-V 2012	6.41.3

Throughput: **646 Mbps** Peak VM CPU: **37%**

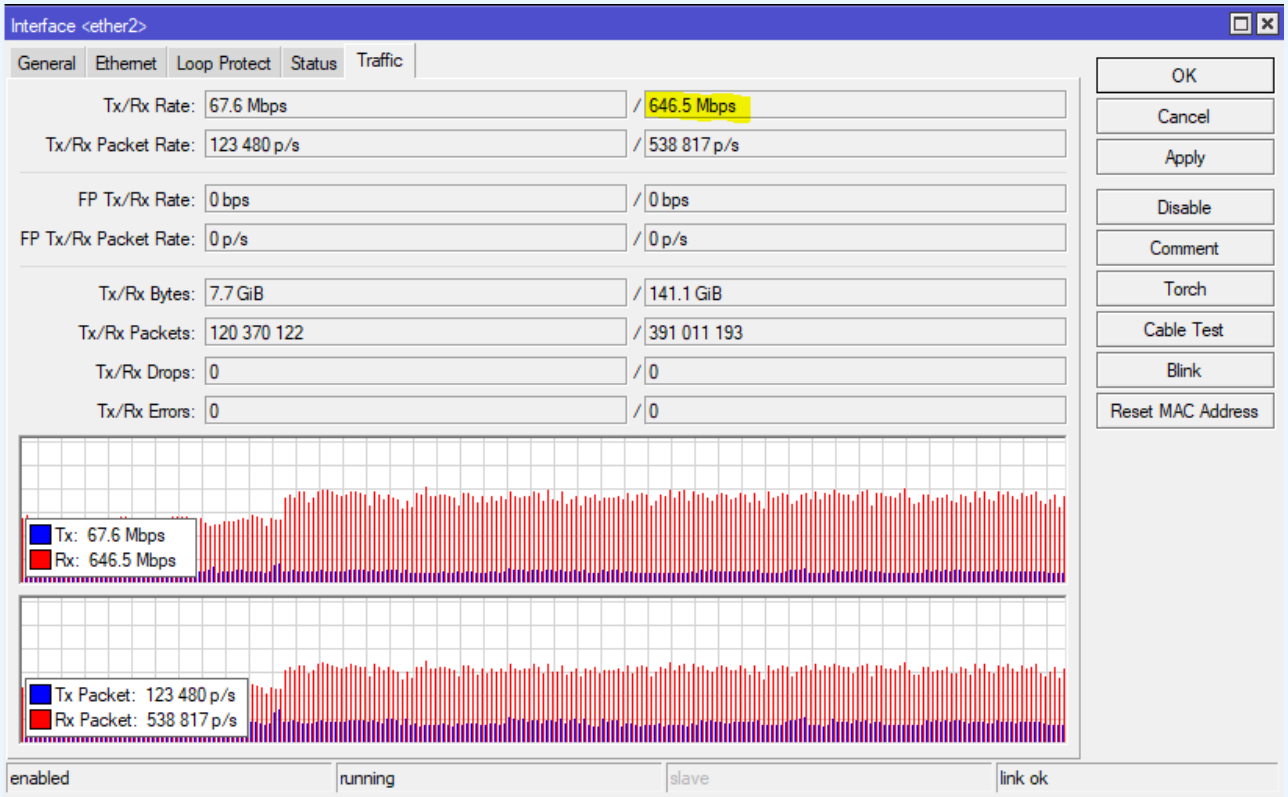

Design: CHR performance on Hyper-V (Microsoft)(96 byte MTU)

Platform	Hypervisor	CHR
Baltic Vengeance	Hyper-V 2012	6.41.3

5 Gbps load plus reset of both upstreams: **Convergence Time – :43**

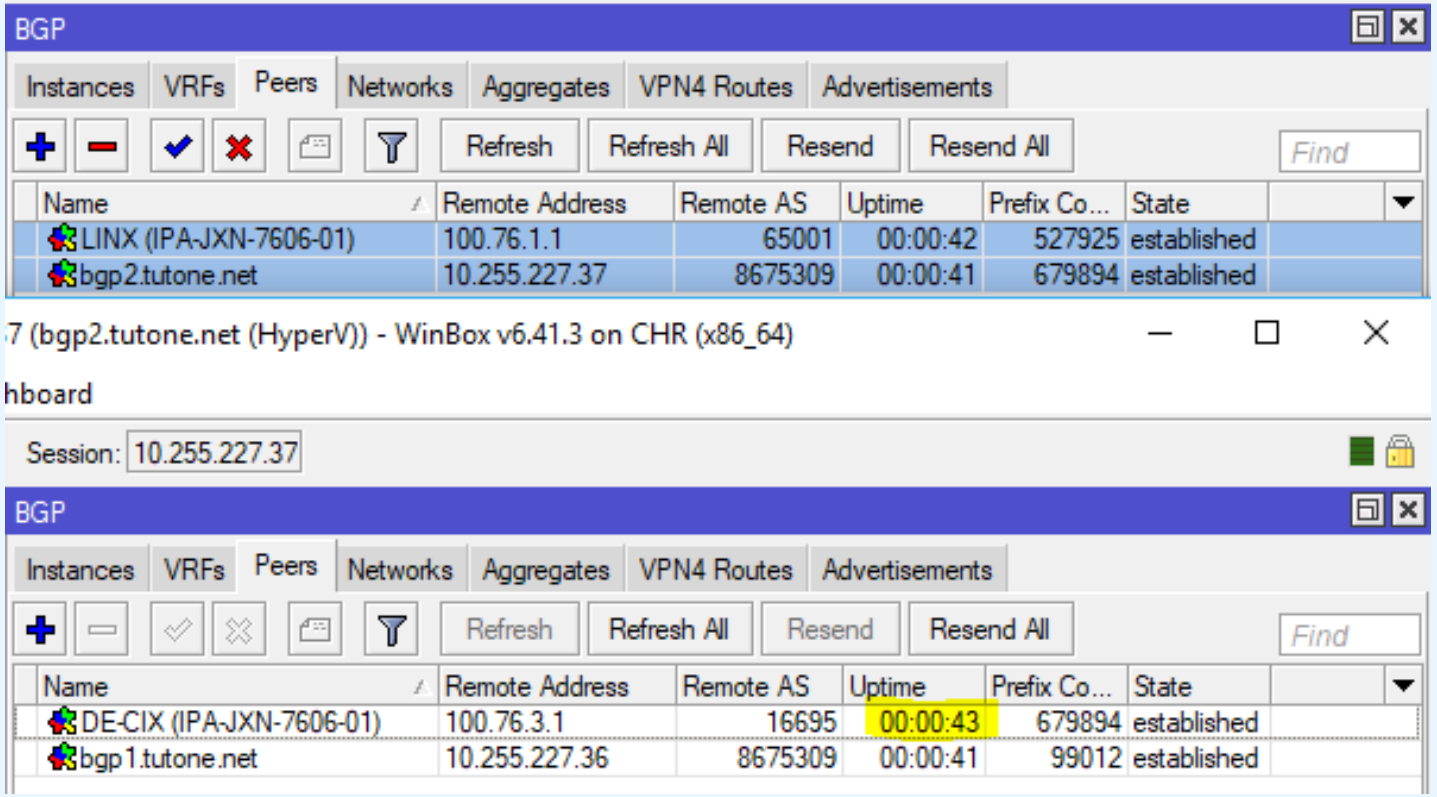

Design: CHR real world performance example - PogoZone

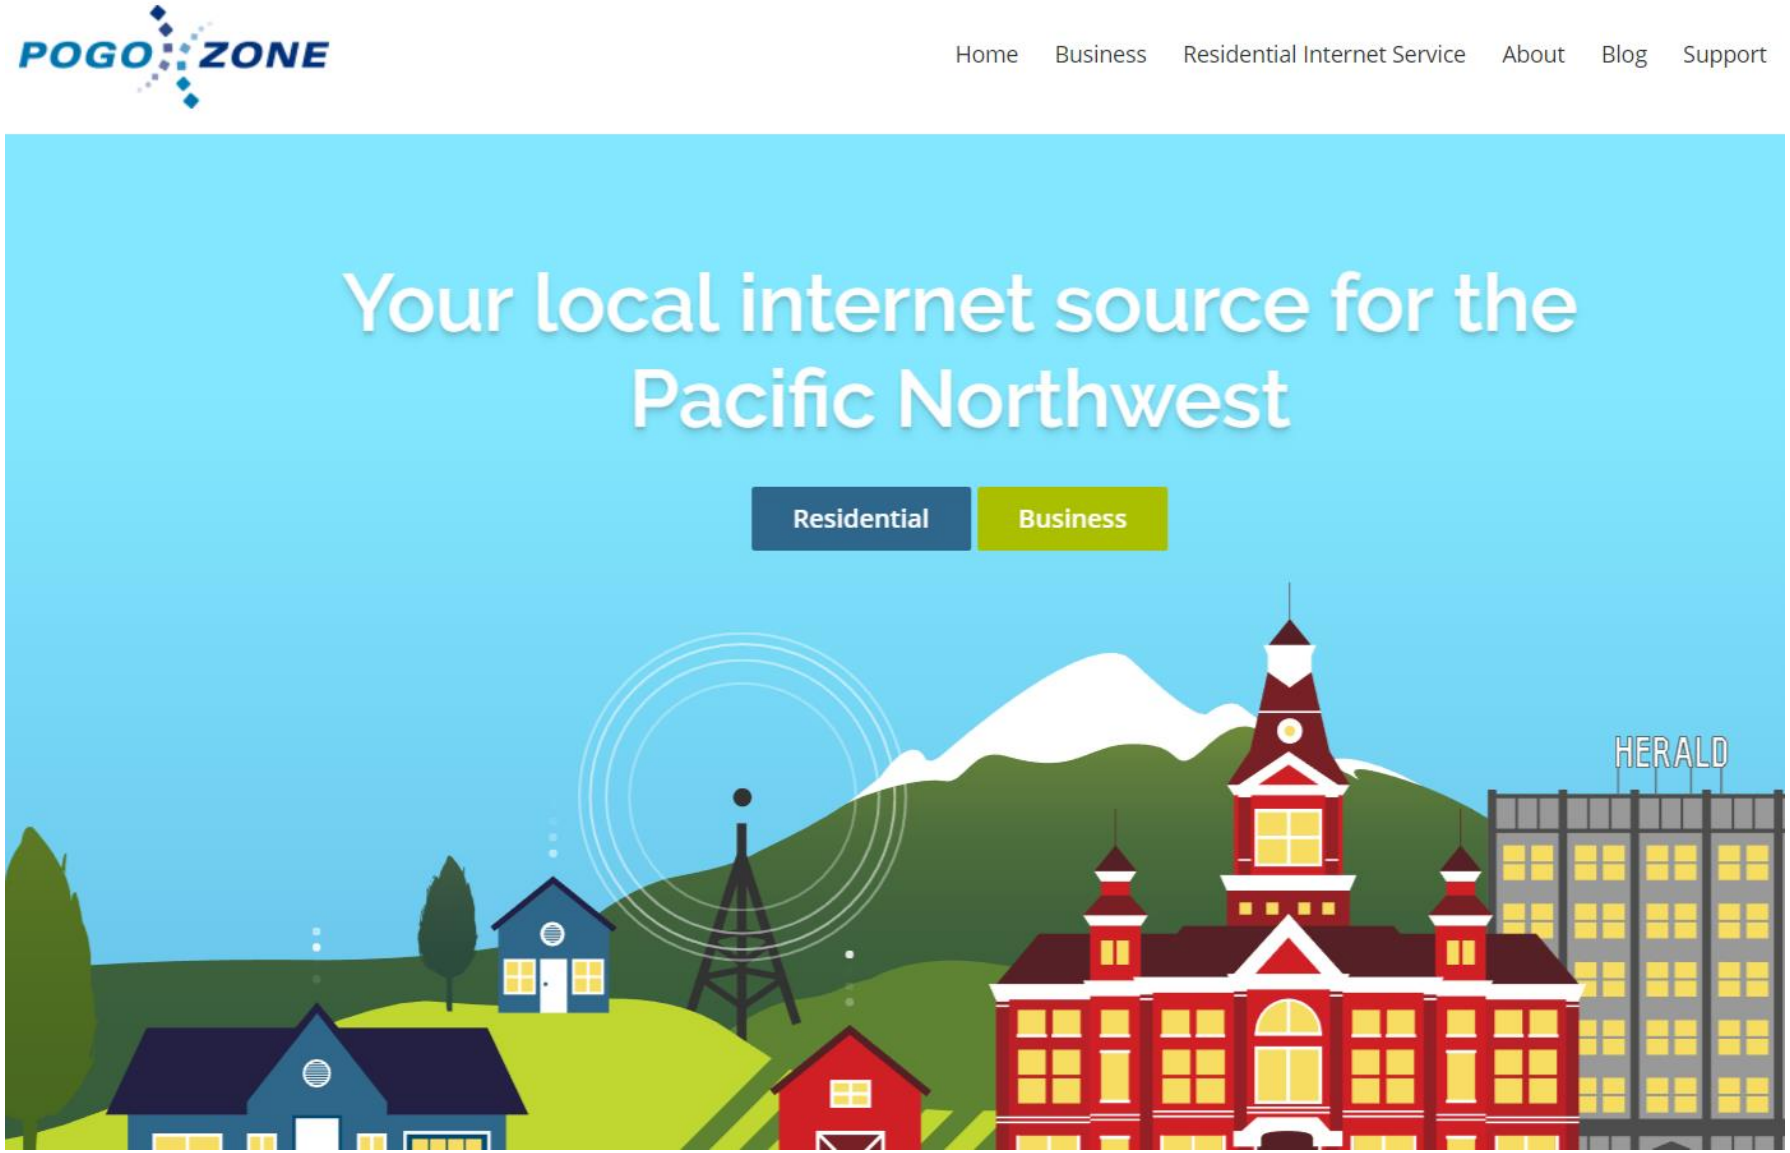

Design: **CHR real world performance example - PogoZone**

US WISP - Throughput: **1.6 Gbps** Peak VM CPU: **32%**

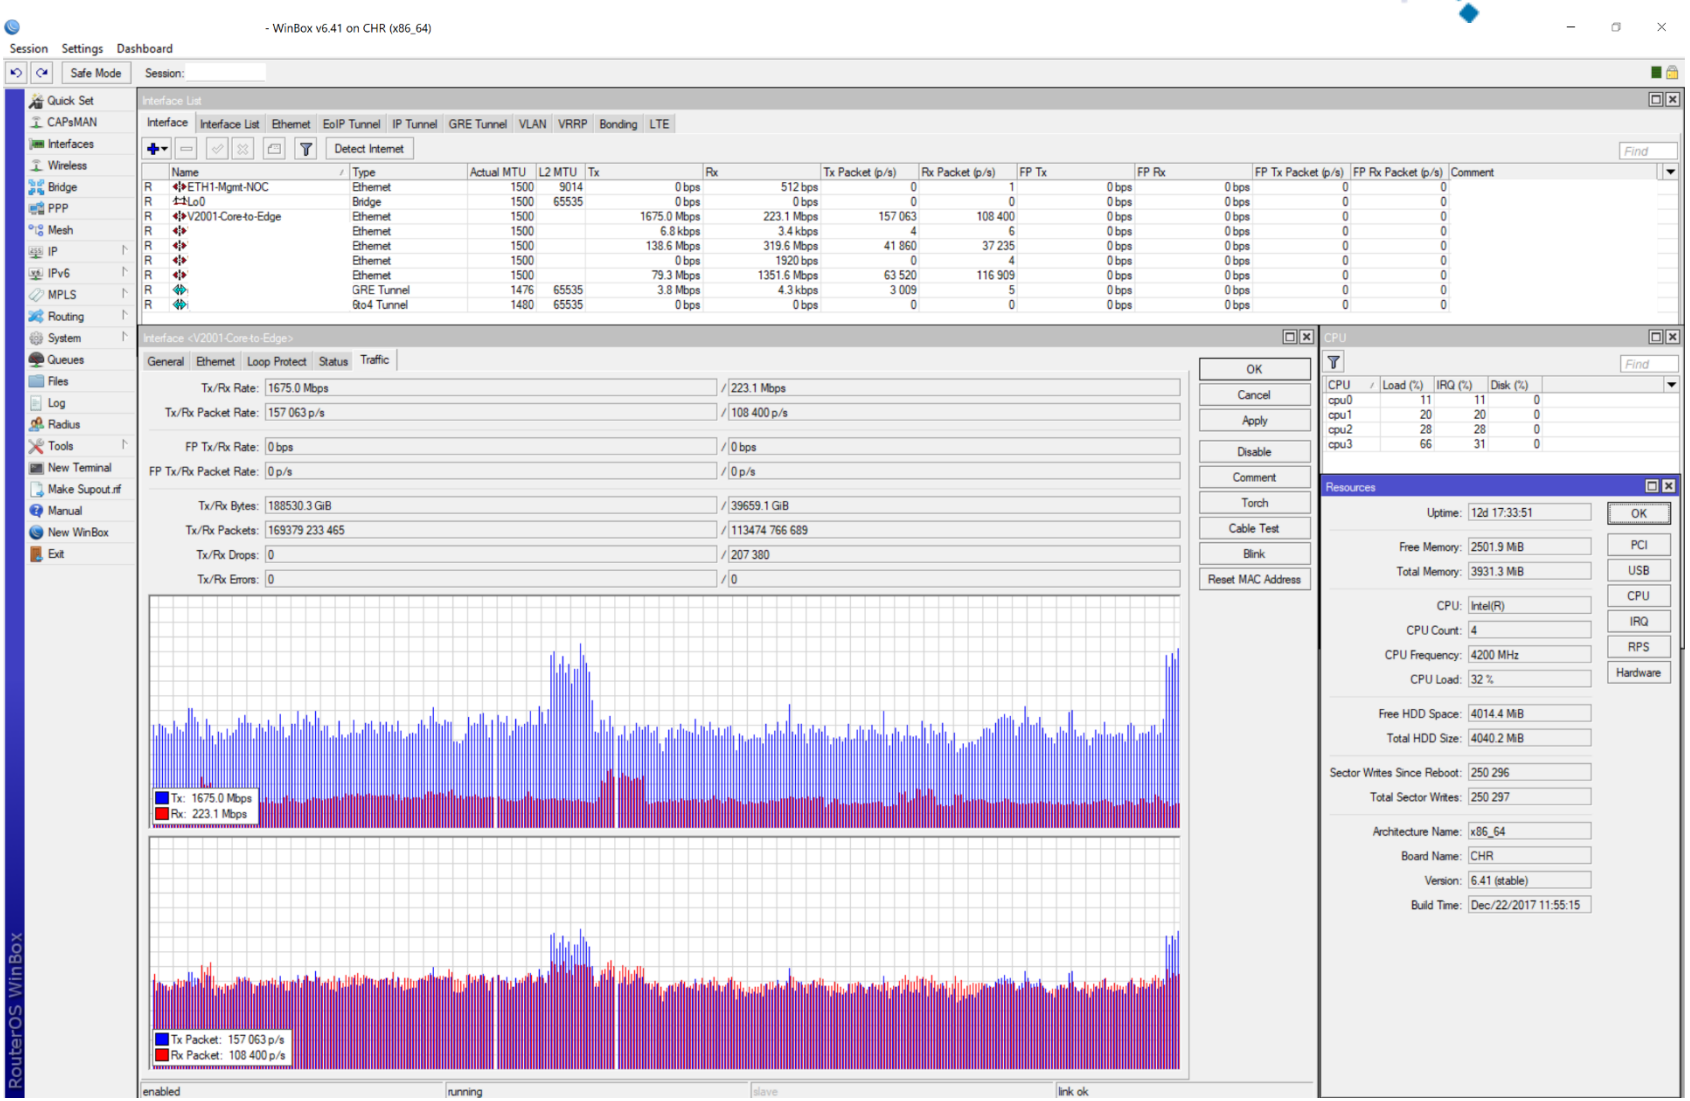

Design: **CHR results – which hypervisor ???**

- **Which Hypervisor is the best?**
  - #1 **Hyper-V** was faster than ESXi and ProxMox all the way around and the most consistent with results in testing. It was by far the clear winner in routing convergence times
  - #2 **ProxMox KVM** delivered much better convergence speeds than ESXi but not quite as good as Hyper-V
  - #3 **VM Ware ESXi** was the slowest in everything but raw throughput and it only has a marginal edge over Hyper-V in that area.

Questions??