Firewall Performance

Marketing vs Reality

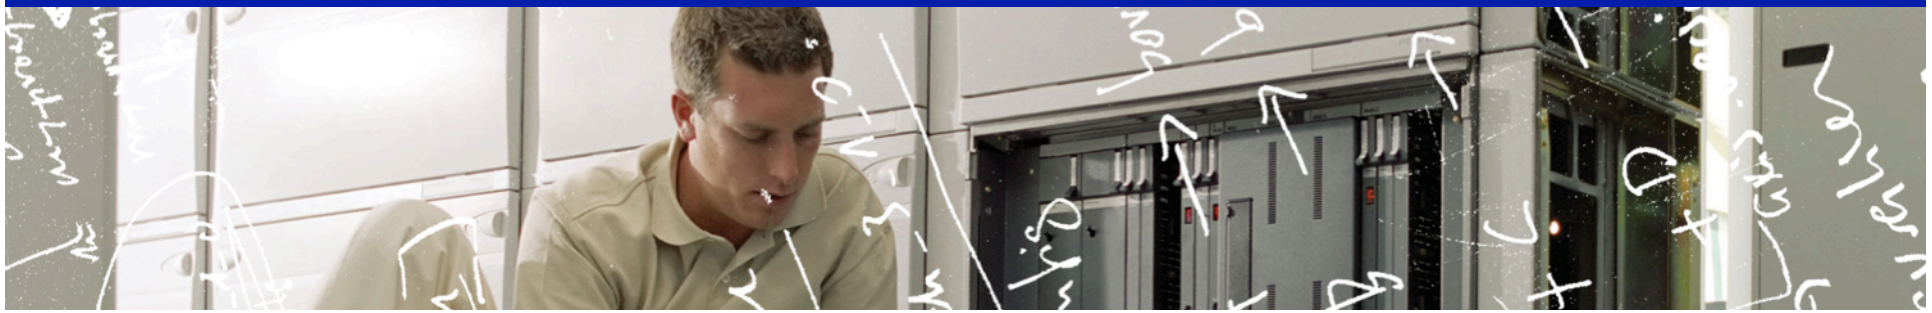

SWITCH

Serving Swiss Universities

Chris Welti, Alexander Gall
{chris.welti, gall}@switch.ch

Introduction

- NOC/PERT case we had with a customer (large university) March & October 2012
- connected redundantly with 10G
- using BGP for primary and secondary access
- Firewall caused dropped BGP sessions because of relatively small DOS-attack on host behind

Customer setup

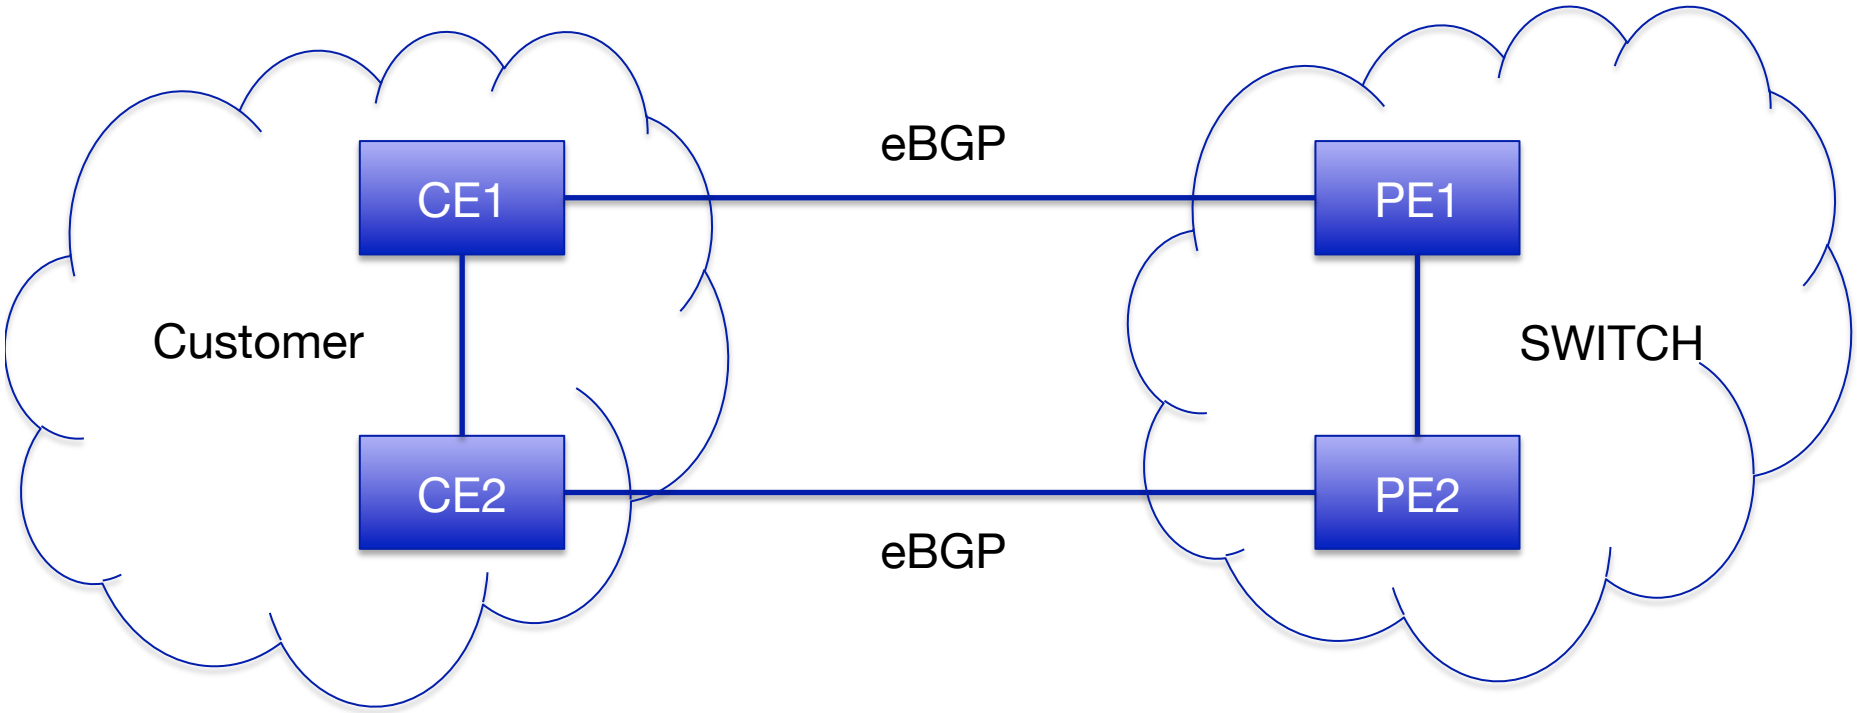

All Links 10GE

Symptoms

- customer contacts us that both BGP sessions were flapping
- primary session down -> failover to secondary
- primary sessions gets back up
- secondary starts to get down shortly after
- traffic always flapping between primary and secondary

View from SWITCH

- BGP flaps (neighbor down) on primary and secondary (not at the same time)
- physical interfaces were up
- no other links or BGP sessions affected
- router CPU utilization normal
- no input errors or output drops

View from customer

- HSRP flaps (route unreachable)
- BGP flaps primary and secondary (not at the same time)
- CPU utilization normal
- no input errors or output drops on router

Investigation

- bgp flaps only lasted about 30 mins
- long over when we investigated so looking at interface statistics
- Traffic stats show normal amount of traffic around 1Gbit/s

Traffic to customer

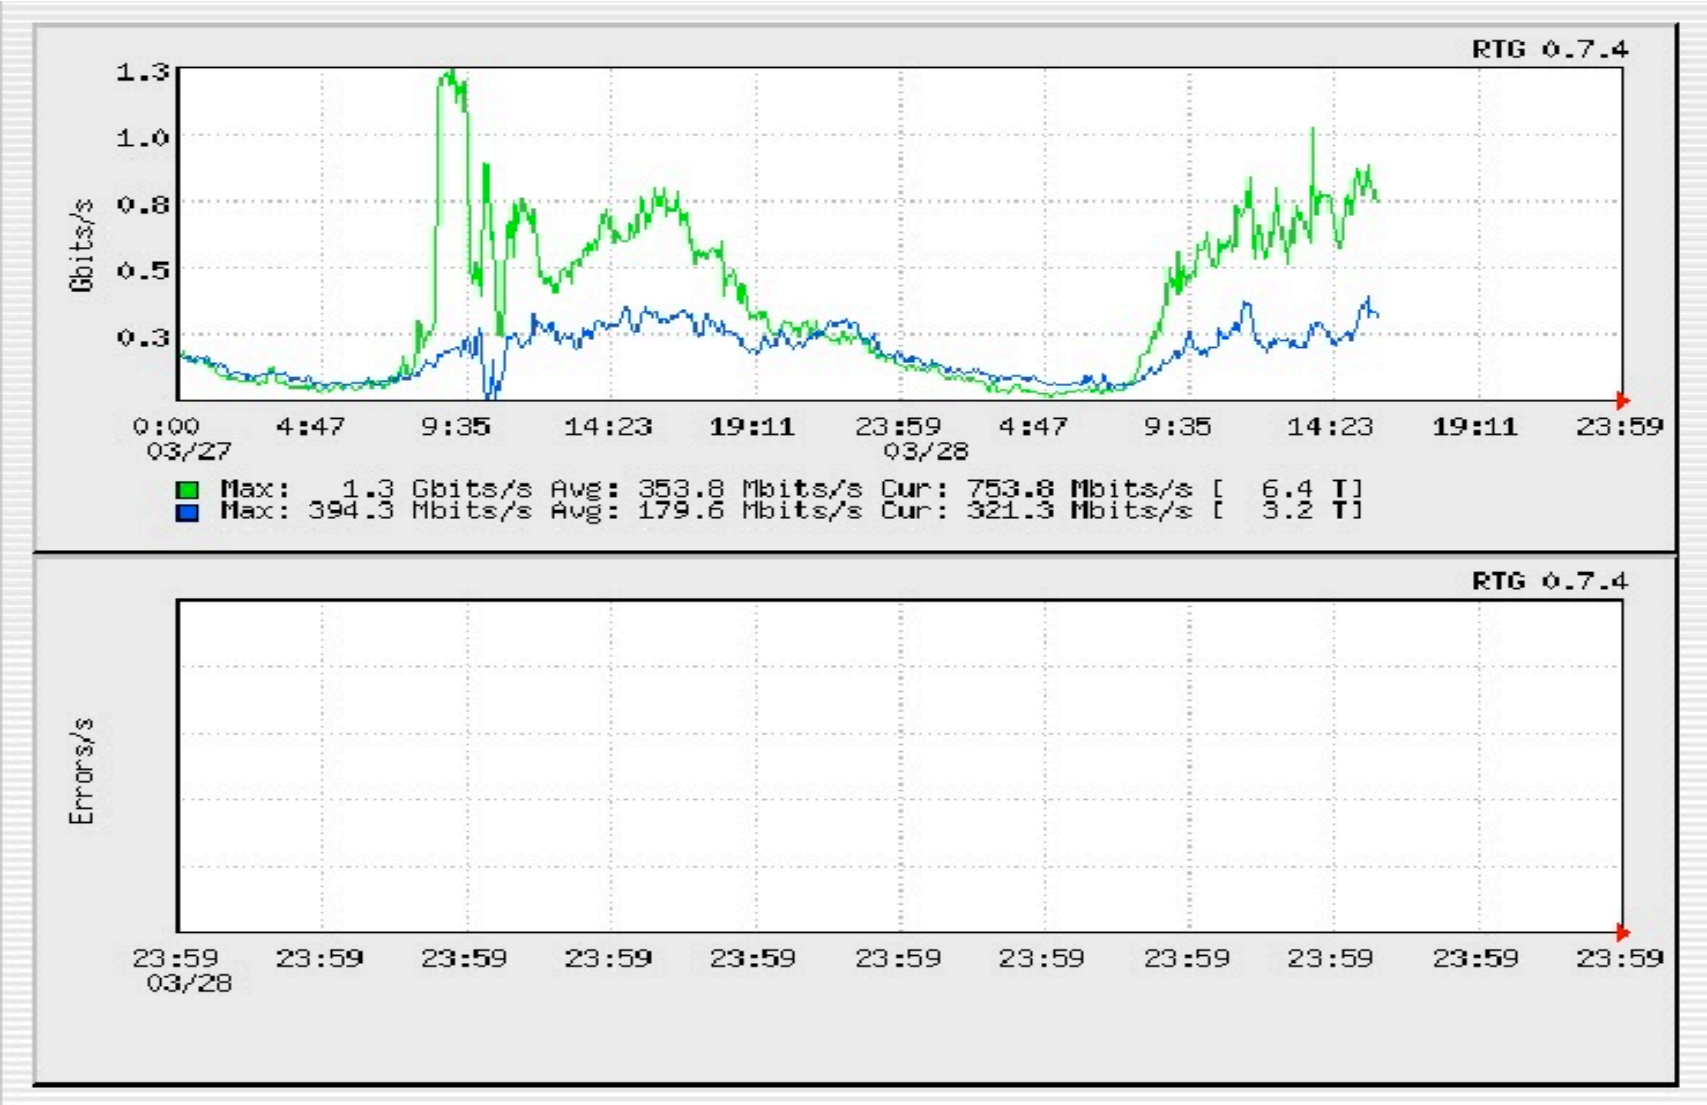

Investigation

- **but** packets/sec was high, around 1Mpps

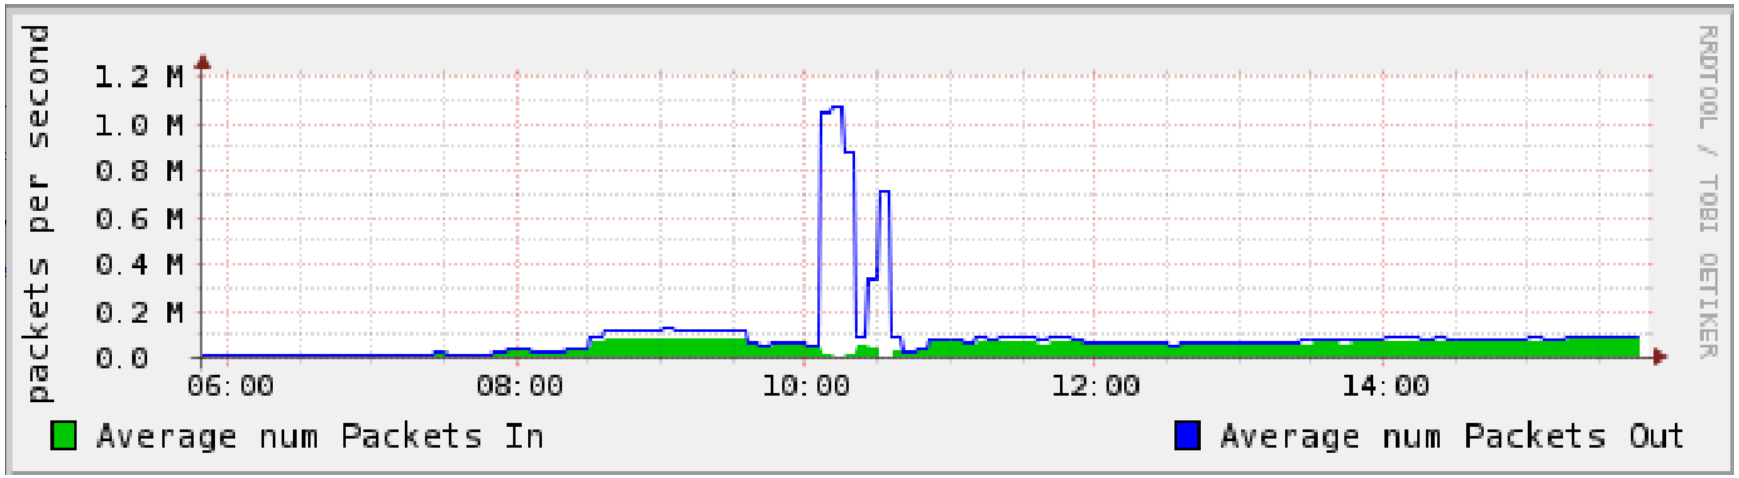

- from our nfsen netflow archive we could see it was a DDOS attack to a single host with small UDP/53 packets

Assumed setup

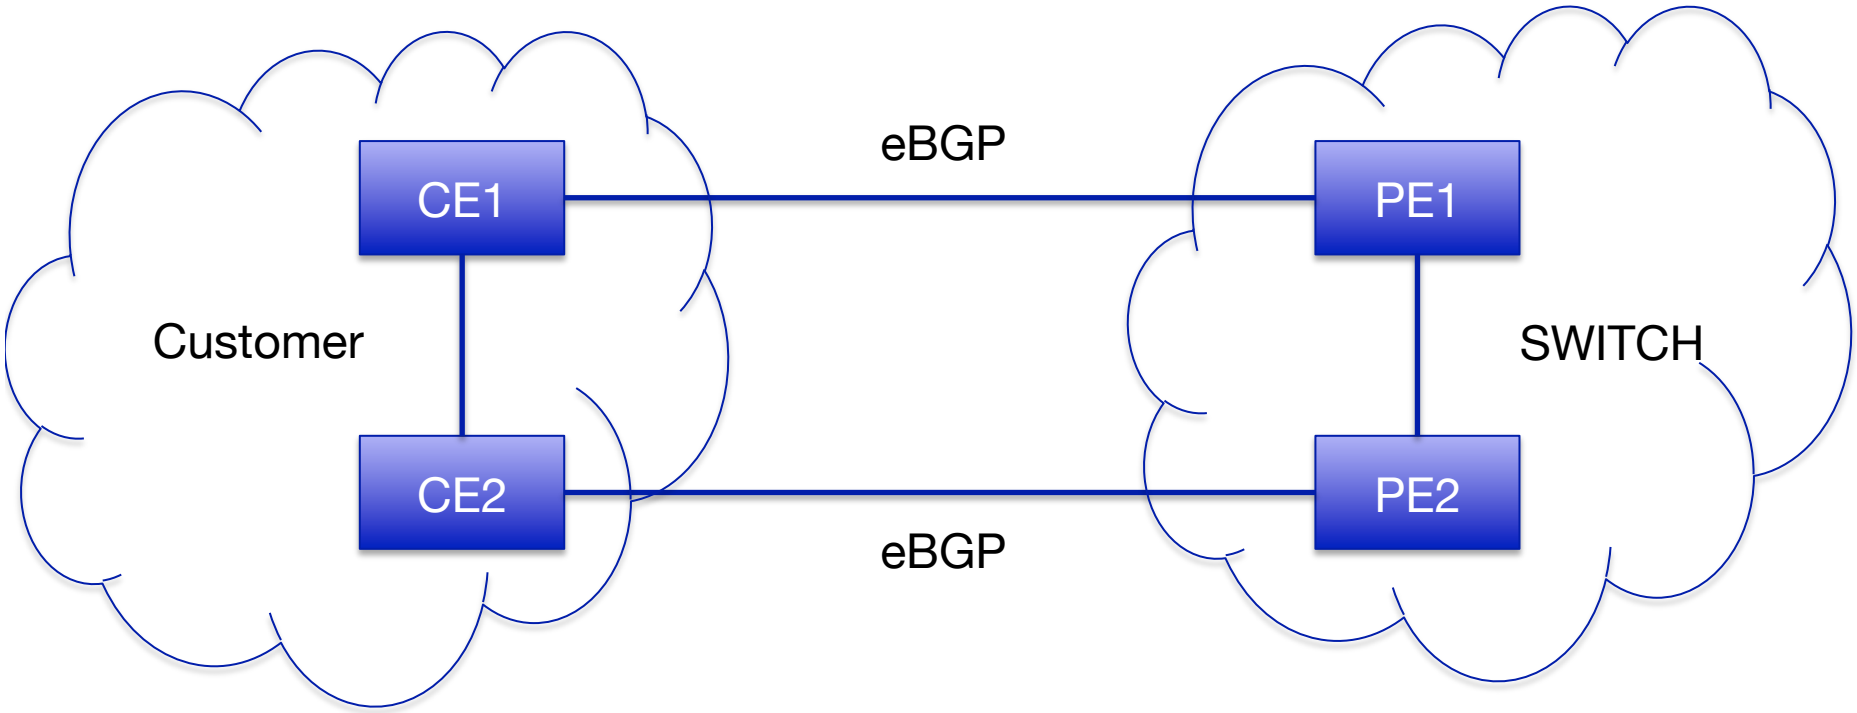

All Links 10GE

Investigation (2)

- Why is this a problem? Router can handle traffic easily...
- ... but customer has a **firewall** between their CE router and our PE router
- High-End Cisco ASA-5585-X-SSP40

Transparent firewall inbetween

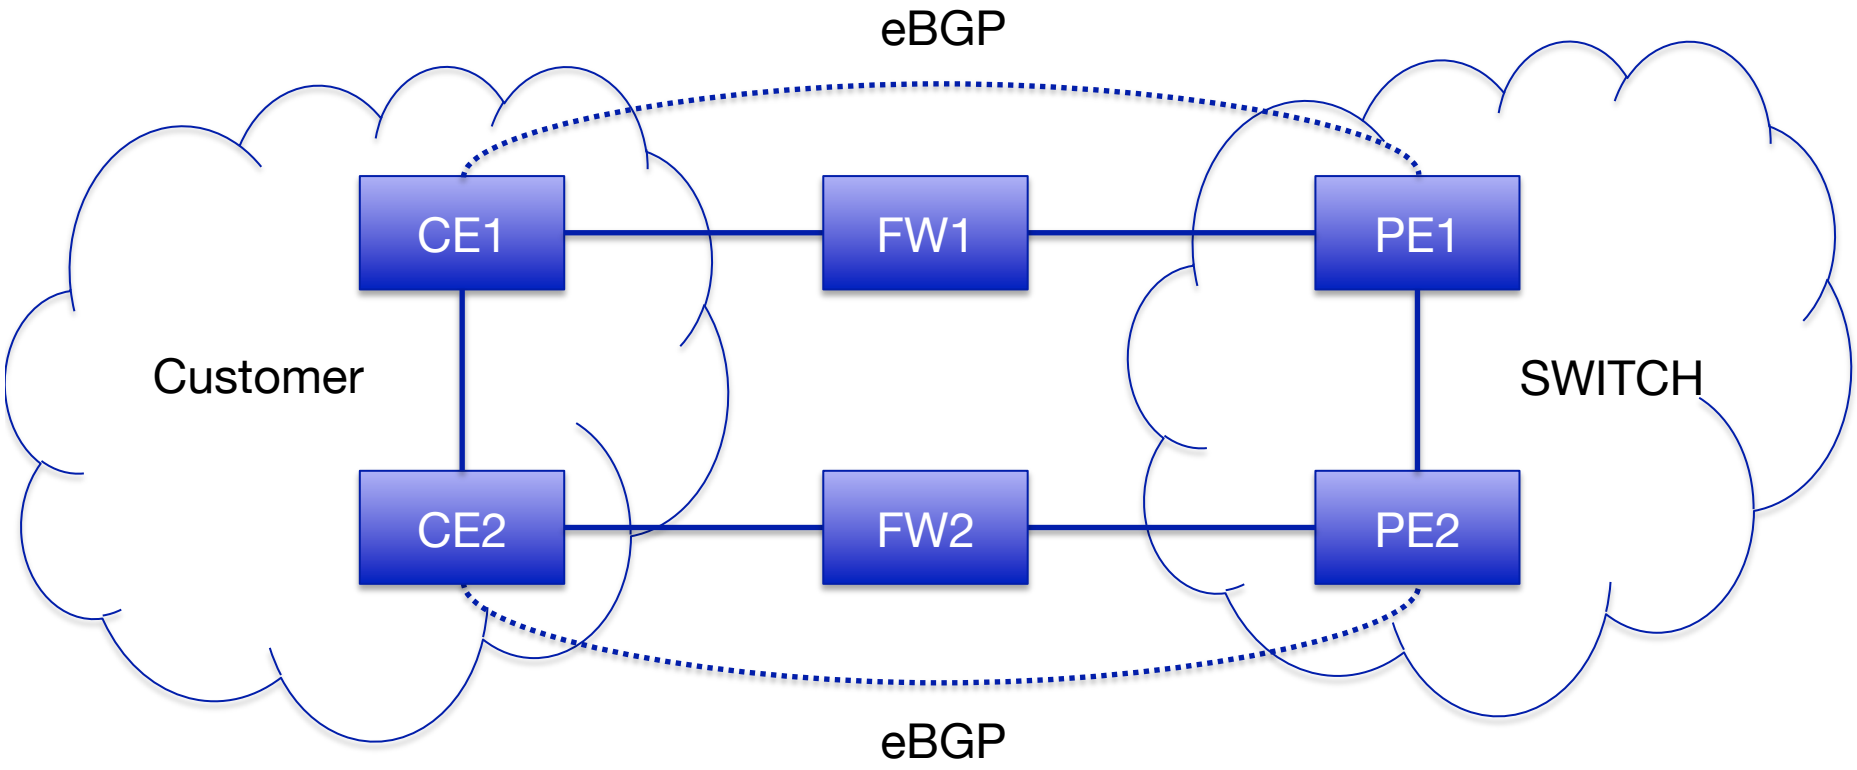

All Links 10GE

Reading Data Sheets

	ASA 5540	ASA5545-X	ASA5585 SSP40
Max Throughput	650Mbps	3Gbps	**20Gbps**
Real-World Throughput	-	1.5Gbps	**12Gbps**
Max VPN Throughput	325Mbps	400Mbps	**3Gbps**
64 Byte Packets/sec	-	900,000	**6,000,000**
Max Conns	400,000	750,000	**4,000,000**
Max Conns/sec	25,000	30,000	**240,000**
IPSEC VPN Peers	5000	2500	**10,000**
Max Interfaces	1xFE + 8x1GE	14x1GE	**12x1GE + 8x10GE**

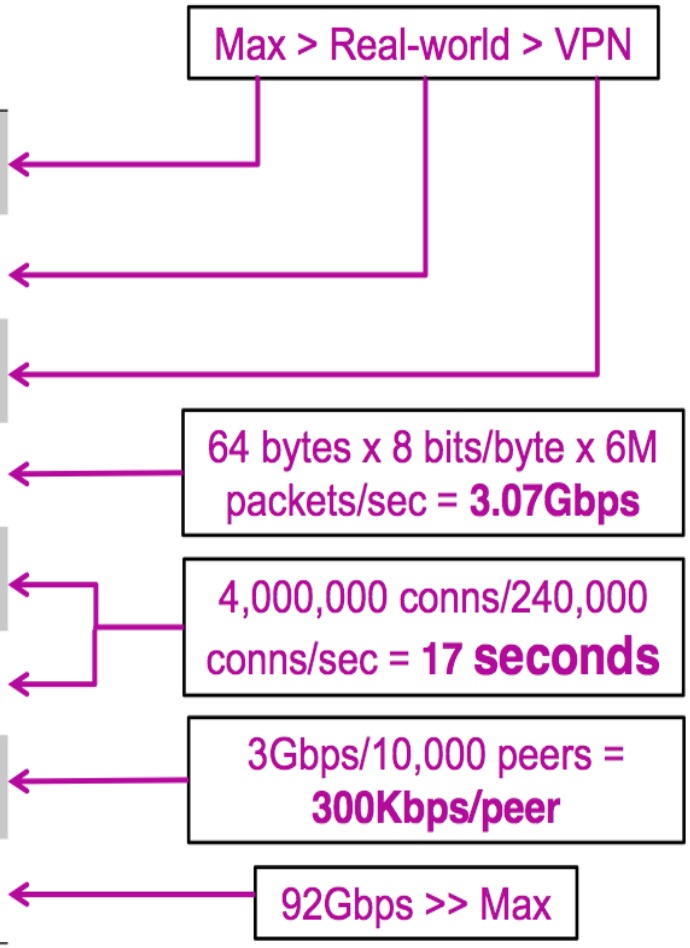

Investigation (3)

- So from the spec sheets should not be an issue that leads to packet loss
- No problems seen in FW CPU/memory logs either, no interface errors on the external 10GE
- trying to confirm it was indeed the DOS traffic that caused this
- -> build test scenario

Confirm issue with high pps

- generated 4 UDP flows (100 byte packets) from a single GE connected host in our office towards a /32 announced from their backup router
- approx. 1.3M pps at 1Gbit/s total
- enough to cause BGP sessions to flap
- customer opens case with Cisco

Cisco Feedback

- TAC engineer acknowledged problem, issue with small packets and high rates
- drops on internal 10GE data channels
- cause is internal interface use DMA without interrupts, CPU periodically polling packets from RX ring not often enough?
- static fixed buffer size for packets (1550 b)
- too many packets towards RX ring -> drops at interface FIFO queue, even backpressure, will cause more losses than necessary (packets for others RX rings also impacted)

Cisco Feedback (2)

- max. packet rate per interface at about 20-60000 at 1GE and about 8-10 times more for 10GE according to FW engineer
- massive gap to their marketing numbers of 6Mpps

Simplified ASA5585 Block Diagram

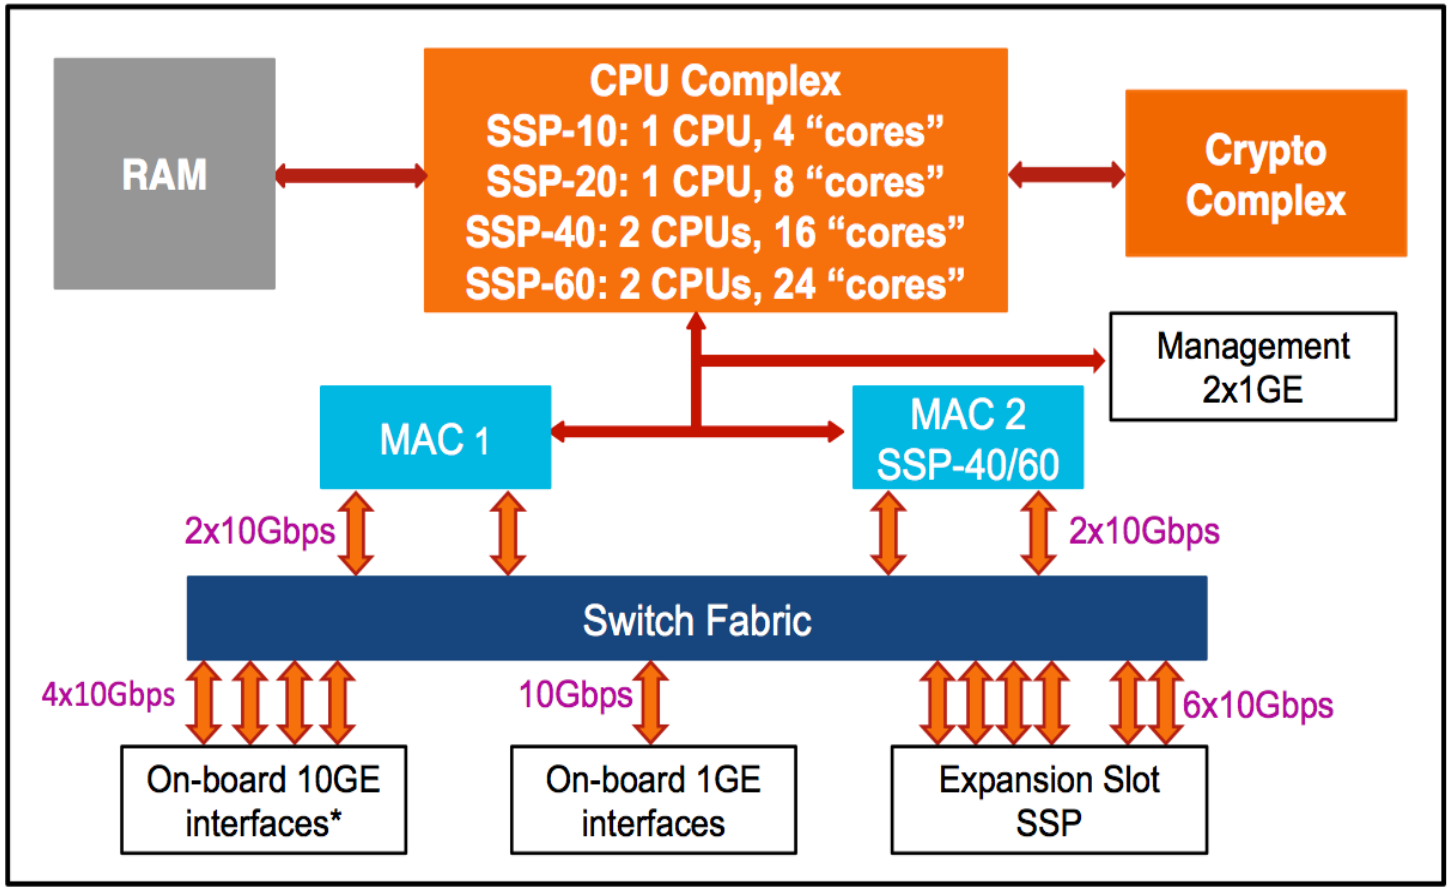

*SSP-20/40/60

ASA NIC Architecture

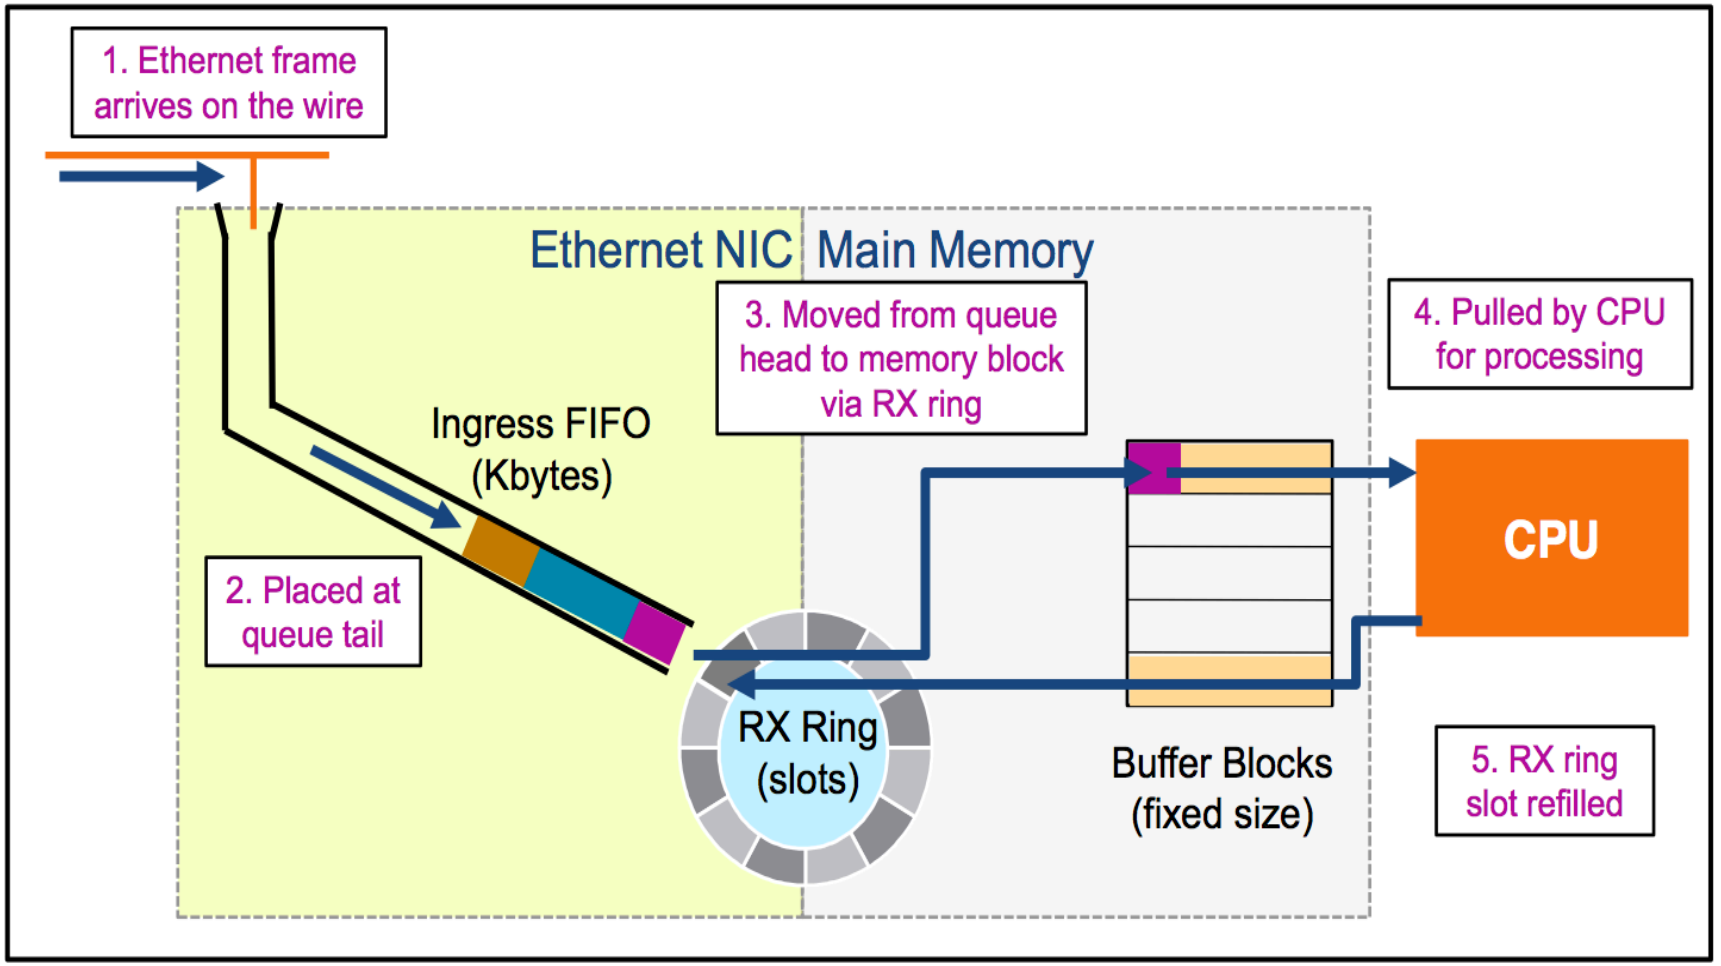

Ingress Load-Balancing on 10GE and MAC

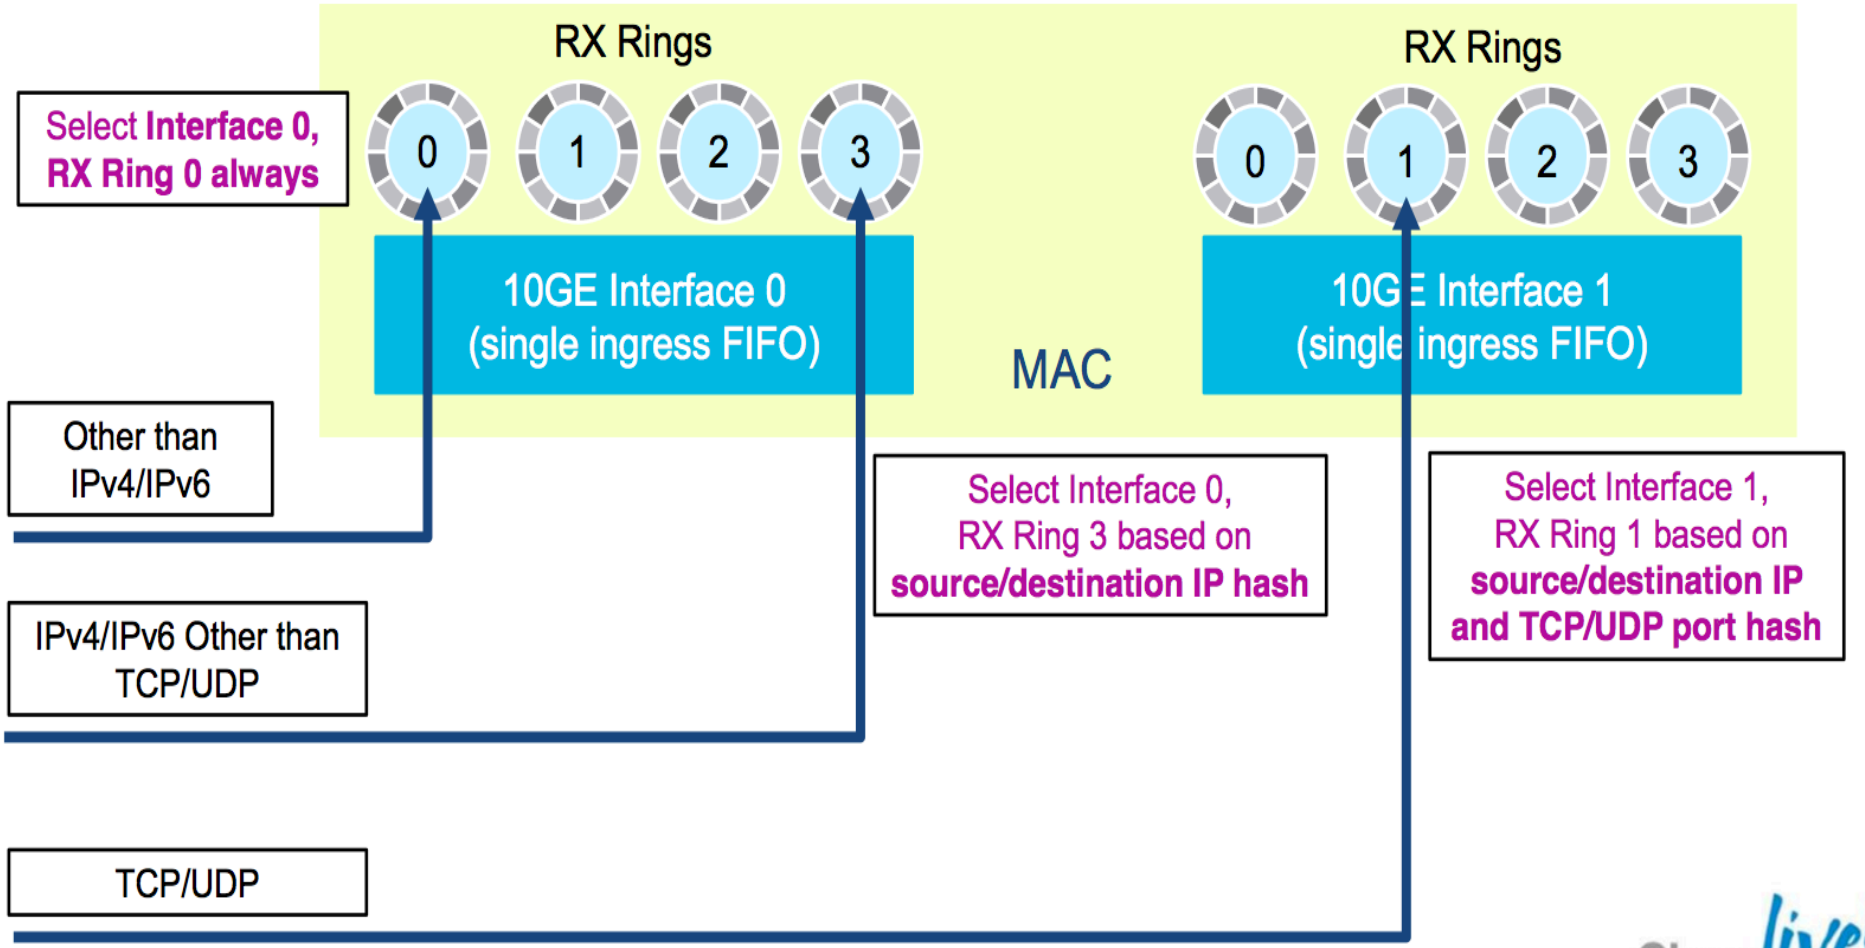

Conclusions

- Be aware that numbers from data sheets might not apply for your use cases
- Don't trust numbers you have not confirmed yourself
- When you put a single big perimeter firewall at the entry edge of your network, you can **lose *all*** your traffic when something goes wrong
- The smart thing would be to use distributed firewalls (preferably on hosts themselves)

Reading recommendations

- www.ciscolive365.com
  - BRKSEC-3020
  - BRKSEC-3021