
CERN

HEPiX Fall 2017
Firewall Load Balancing Solution

Vincent DUCRET
vincent.ducret@cern.ch

Introduction

- Traffic between the CERN internal networks (LCG/GPN) and the external network (Internet) is filtered by firewalls
- Current setup has some limitations
- The Firewall Load-Balancing (FWLB) solution aims to bring scalability to the setup

Current firewall Setup

Limitations:

- Limited capacity
- Not possible to add a 3rd firewall
- Each FW is dedicated to a type of traffic
- Overall traffic is not equally distributed
- Routing constraints/complexity

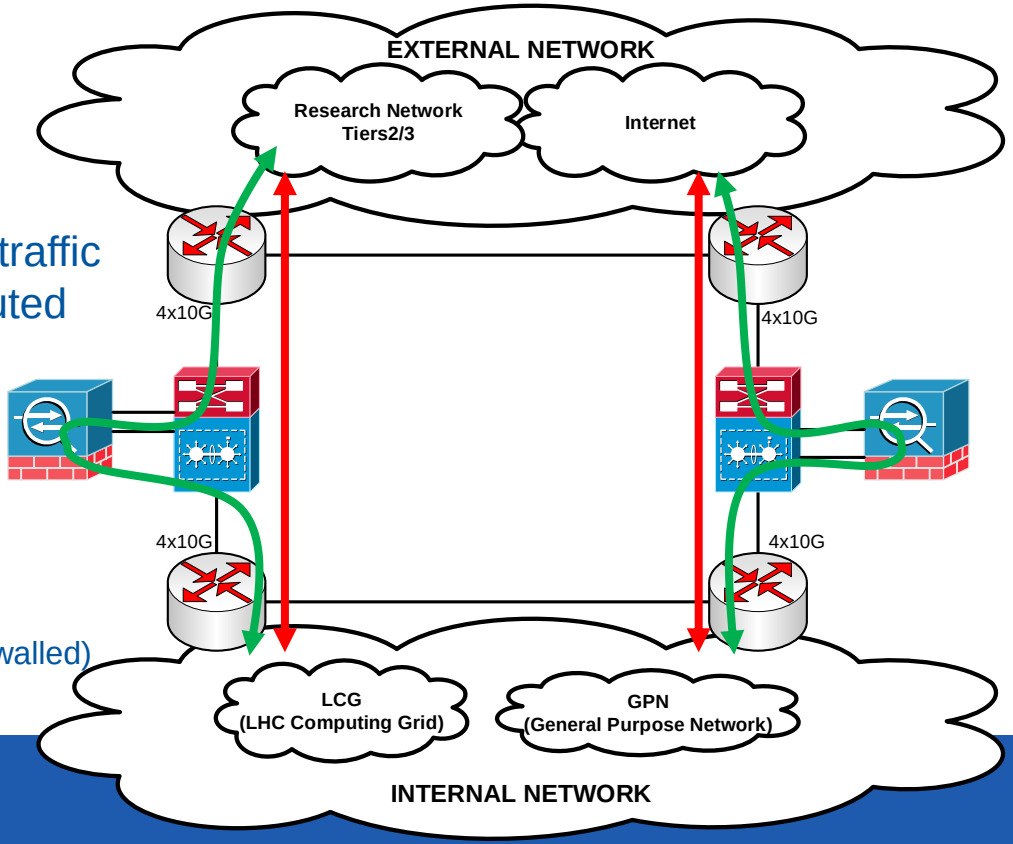

HTAR traffic (not firewalled)

Firewalled traffic

Four different options

1. Keep current setup but build a cluster of Firewalls

Pros: No major change on the design
Can use our spare firewall

Cons: Need exact same hardware (vendor lock)
Capacity increase is not linear (2x 10Gbps ➔ 14 Gbps)
Keep a separation by type of traffic (same routing constraints, unequal distribution of the traffic, etc.)

2. Keep current setup and use “bigger” Firewalls

Pros: no major change on the design
“easy” increase of the overall capacity

Cons: Price and time (tender required)
No use of our spare Firewall
Need router/switch upgrade (Price)
Keep a separation by type of traffic (same routing constraints, unfair distribution of the traffic, etc.)

3. Full review of the design with specific Load Balancers

Pros: “Long term" solution
Match all requirements

Cons: Unknown technology
Price and time (market study and tender required)
May need router/switch upgrade

4. Build an “intermediate” FW Load Balancer solution

Pros: Reuse existing and spare Firewalls
No need for switch/router upgrade
Allow us to get familiar with FW load-balancing
Total cost < 60% of current spare FW

Cons: May not match all requirements
May only be valid for 2/3 years prior to a global architecture review

Firewall Load Balancing Setup

More details about IDS in the presentation given by Adam Krajewski:
"Network Automation for Intrusion Detection System"

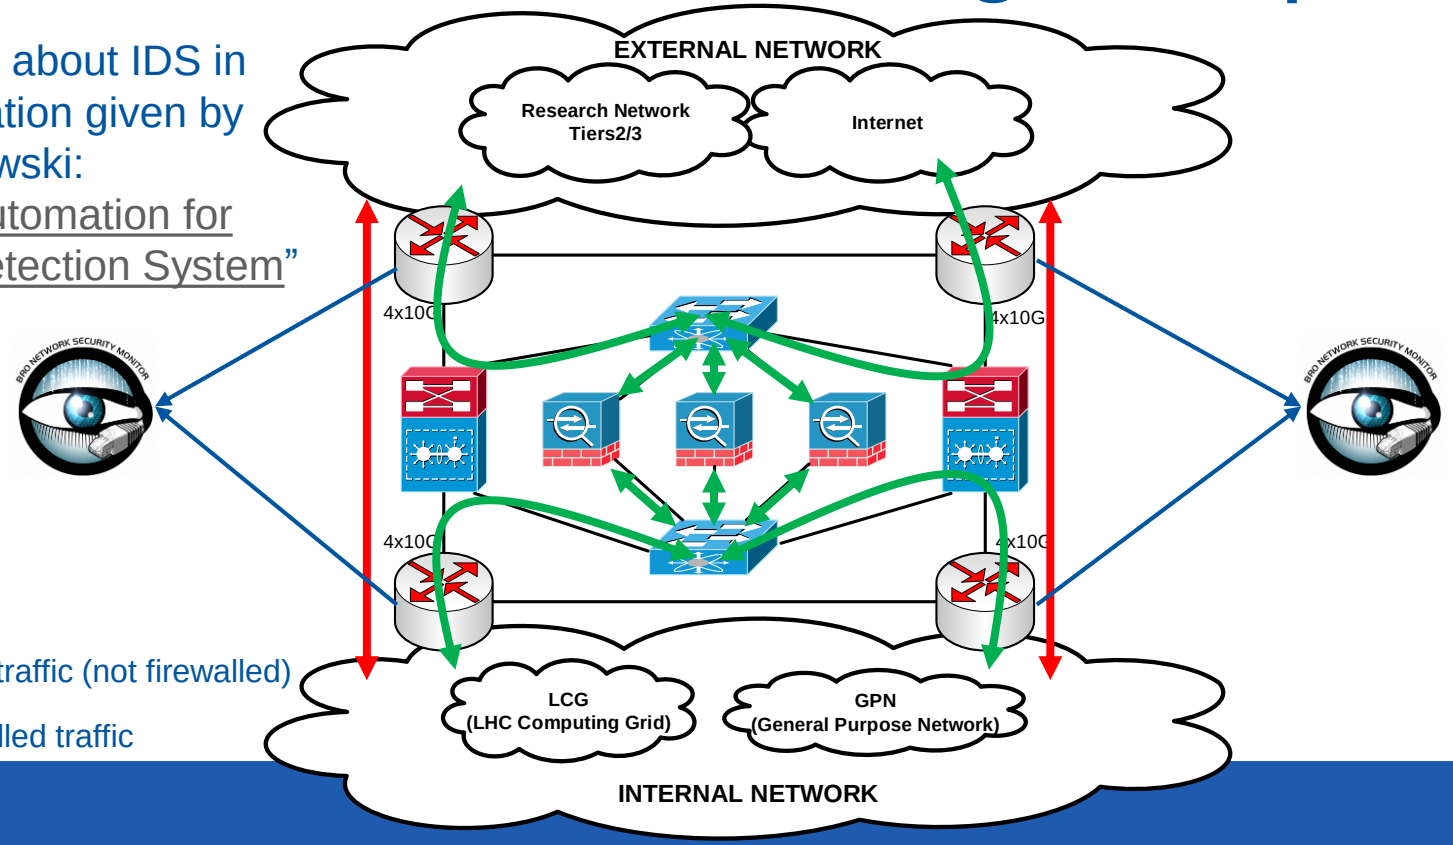

Firewall Load Balancing Setup

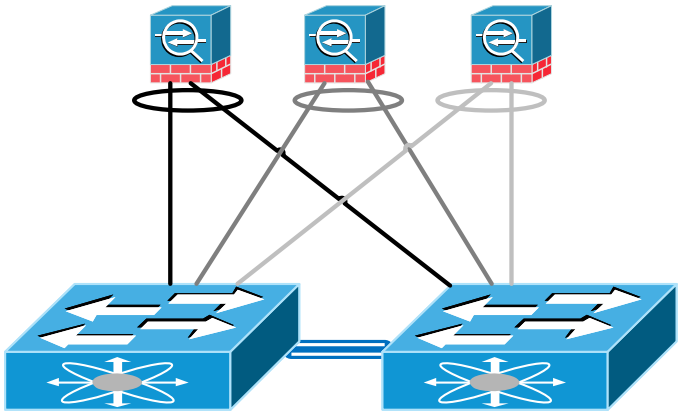

Physical Setup

Each Firewall have a 10Gbps link
to each Load Balancer
LAG make it a single logical connection

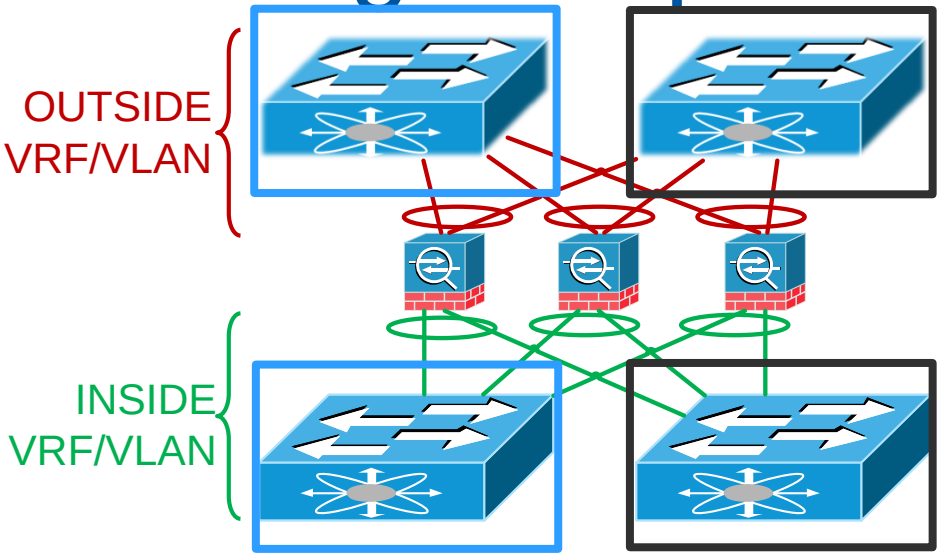

Logical Setup

Each Load Balancer is divided in two VRFs
Full redundancy for single device failure

Firewall Load Balancing Setup

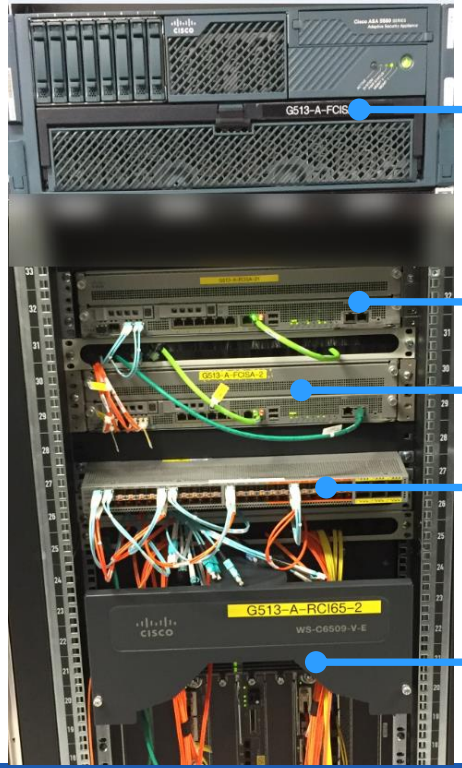

Load balanced FW2

Load balanced FW1

Production FW

Cisco Nexus 5K for FWLB

Router in charge of PBR for HTAR/Firewalled traffic

- Based on two Cisco Nexus 5672UP (each with 32x 10Gbps + 4x 40Gbps)

- Firewall Load Balancing is done using the ITD feature (Intelligent Traffic Director)
  = PBR with some automated features

FW LB setup details

- Load balancing is based on source IP (relies on PBR)
- Traffic to/from a single server will always cross the same FW
- Fully redundant setup
  - If one FW fails, traffic is redirected to the remaining ones (30-60 seconds detection; not transparent for active sessions)
  - If one load balancer fails, all firewalls and routers still have a connection to the second load balancer
  - If one router fails, traffic is rerouted to the other one

Configuration details

Manual configuration

```
itd device-group FIREWALLS-IN
  node ip X.Y.Z.1 weight 12
  node ip X.Y.Z.2 weight 4
  probe icmp frequency 3 timeout 2 retry-down-count 3 retry-up-count 5

itd ITD-FW-IN
  device-group FIREWALLS-IN
  ingress interface Vlan100
  failaction node reassign
  load-balance method src ip buckets 16
  no shutdown
```

Configuration details

Auto-generated ACLs

```
ip access-list ITD-FW-IN_itd_bucket_1
  10 permit ip 1.1.1.0 255.255.255.15 any
ip access-list ITD-FW-IN_itd_bucket_2
  10 permit ip 1.1.1.16 255.255.255.15 any
(…)
ip access-list ITD-FW-IN_itd_bucket_15
  10 permit ip 1.1.1.224 255.255.255.15 any
ip access-list ITD-FW-IN_itd_bucket_16
  10 permit ip 1.1.1.240 255.255.255.15 any
```

Configuration details

Auto-generated PBR rules

```
route-map ITD-FW-IN_itd_pool permit 0
  description auto generated route-map for ITD service ITD-FW-IN
  match ip address ITD-FW-IN_itd_bucket_1
  set ip next-hop X.Y.Z.1
route-map ITD-FW-IN_itd_pool permit 1
  description auto generated route-map for ITD service ITD-FW-IN
  match ip address ITD-FW-IN_itd_bucket_2
  set ip next-hop X.Y.Z.2
(…)
route-map ITD-FW-IN_itd_pool permit 15
  description auto generated route-map for ITD service ITD-FW-IN
  match ip address ITD-FW-IN_itd_bucket_16
  set ip next-hop X.Y.Z.1
```

Configuration details

Auto-generated PBR rules

```
interface Vlan100
  description GAR1 IN 0513-A-FI11
  no shutdown
  mtu 9000
  no ip redirects
  ip address A.B.C.D/X
  ip policy route-map ITD-FW-IN_itd_pool
```

ITD will adapt the PBR rules in case of FW failure

Configuration details

Python scripts to ease operation

```
cli alias name addfw21
    source itd-fw.py -fwin X.Y.Z.1 -fwout A.B.C.1 -weight 12 -add

cli alias name removefw21
    source itd-fw.py -fwin X.Y.Z.1 -fwout A.B.C.1 -rem

cli alias name addfw22
    source itd-fw.py -fwin X.Y.Z.2 -fwout A.B.C.2 -weight 4 -add

cli alias name removefw22
    source itd-fw.py -fwin X.Y.Z.2 -fwout A.B.C.2 -rem

cli alias name itdreset source itd-clear.py
```

FWLB setup limitations

- Current platform supports IPv4 load-balancing only
  - ➔ IPv6 traffic is not concerned by this FWLB solution
    IPv6 flow unchanged (dedicated firewalls, no HTAR)
- Current platform supports 32x10Gbps ports and 6x 40Gbps
- Not possible to transparently remove a Firewall (let current sessions “die”)

But this is considered sufficient for the next 2/3 years.

It lets us get familiar with FW LB solution prior to a global FW design review in the future (new/bigger firewalls, larger LB devices with IPv6 support, etc.)

FW LB PILOT Setup (active since mid-July 2017)

Pilot FWLB traffic
Wired connection Bldgs.28/31/600

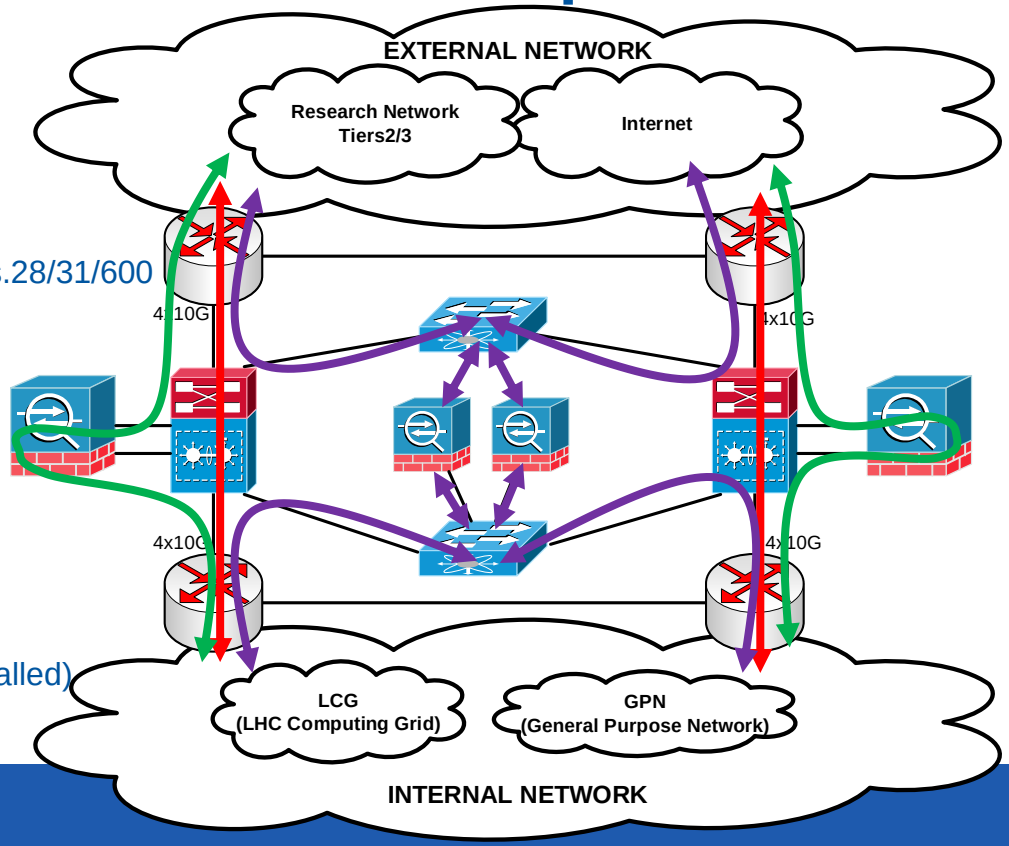

HTAR traffic (not firewalled)

Firewalled traffic

FW LB PILOT Setup (active since mid-July 2017)

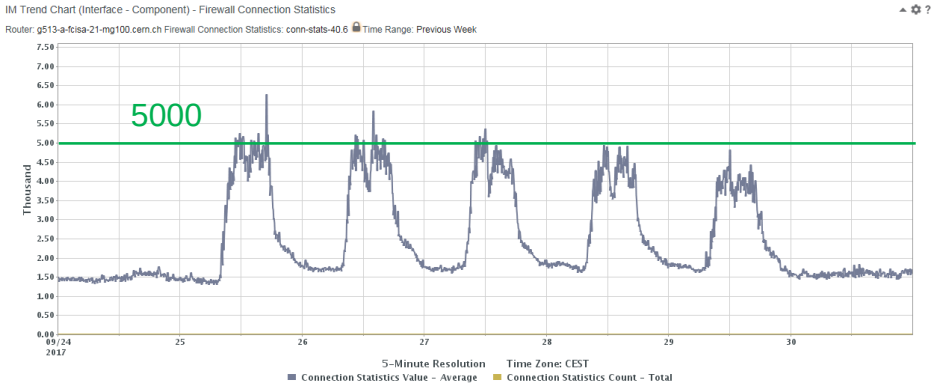

1st Firewall (newer model) manages ~5000 connections

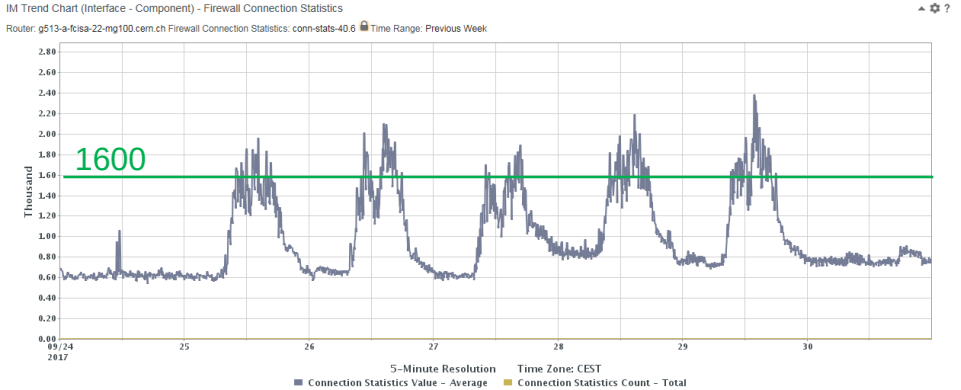

2nd Firewall (older model) manages ~1600 connections

Pilot load balancing configuration is ¾ to 1st FW and ¼ to 2nd FW

NOTE: as a comparison, at the same time, production firewalls manage between 350K and 750K connections

Planning overview

- Summer 2017: Pilot (traffic from IT buildings)
- Q4-2017: move half of the devices to 2nd network hub
- Q4-2017/Q1-2018: migrate all GPN/LCG traffic to FWLB solution (several steps)
- 2018/2019 global review of the firewall setup

- Are you doing Firewall Load Balancing?
- Which solution are you using?
- How scalable is it?
- Does it support dual stack IPv4/IPv6?

If you don’t know, please forward this presentation (or at least this slide) to your “network” colleagues. They can contact me via vincent.ducret@cern.ch

Any questions?


CERN