Group Based Policy

How to integrate security functions in EVPN/VxLAN

ITNOG-10 21/4/2026

Nicola Modena - CCIE #19119 / JNCIE-SP #986
nicola@modena.to - http://tierzero.it

What I will talk about

Group-Based Policy incorporates security functionality into EVPN/VxLAN environments, offering a scalable approach for distributing and integrating security functions across network infrastructures.

This presentation leverages this technology to design a highly scalable security solution that overcomes the limitations of the classical monolithic firewall-centric design.

About me

Nicola Modena - CCIE #19119 / JNCIE-SP #986 Emeritus
Independent Network Architect

More than 25 years experience designing and implementing
service provider and large enterprise networks.
https://tierzero.it | nicola@modena.to

Introducing Group Based Policy for EVPN/VxLAN

It’s this technology more usable than for micro-segmentation ?

Group Based Policy – IETF Draft

DATA PLANE

CONTROL PLANE

UNIFIED DEFINITION

draft-smith-vxlan-group-policy

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: April 25, 2019

M. Smith
Cisco Systems, Inc.
L. Kreeger
Arrcus, Inc.
October 22, 2018

VXLAN Group Policy Option
draft-smith-vxlan-group-policy-05

Abstract

This document defines a backward compatible extension to Virtual eXtensible Local Area Network (VXLAN) that allows a Tenant System Interface (TSI) Group Identifier to be carried for the purposes of policy enforcement.

draft-wlin-bess-group-policy-id-extended-community-03

Workgroup: bess
Internet-Draft:
draft-wlin-bess-group-policy-id-extended-community-03
Published: 20 October 2023
Intended Status: Standards Track
Expires: 22 April 2024

W. Lin
Juniper Networks
J. Drake
Individual
D. Rao
Cisco Systems

Group Policy ID BGP Extended Community

Abstract

Group Based Policy can be used to achieve micro or macro segmentation of user traffic. For Group Based Policy, a Group Policy ID, also known as Group Policy Tag, is used to represent a logical group that shares the same policy and access privilege. This specification defines a new BGP extended community that can be used to propagate Group Policy ID through a BGP route advertisement in the control plane. This is to facilitate policy enforcement at the ingress node when the optimization of network bandwidth is desired.

draft-lrss-bess-evpn-group-policy-02

BESS WorkGroup
Internet-Draft
Intended status: Standards Track
Expires: 8 December 2025

W. Lin
Juniper
D. Rao
A. Sajassi
Cisco
L. Kreeger
Arrcus
J. Rabadan
Nokia
6 June 2025

EVPN Group Policy
draft-lrss-bess-evpn-group-policy-02

Abstract

Group Based Policy can be used to achieve micro or macro segmentation of user traffic. For Group Based Policy, a Group Policy ID, also known as Group Policy Tag, is used to represent a logical group that shares the same policy and access privilege. This document defines a backward compatible extension to Virtual eXtensible Local Area Network (VXLAN) that allows a Group Policy ID to be carried for the purposes of policy enforcement at the egress Network Virtualization Edge (NVE). It also defines a new BGP Extended Community that can be used to propagate Group Policy ID through a BGP route advertisement in the control plane. This is to facilitate policy enforcement at the ingress NVE when feasible.

2014 — 2022 — 2024

POLICY APPLIED BIT

16 BIT = 0-65535

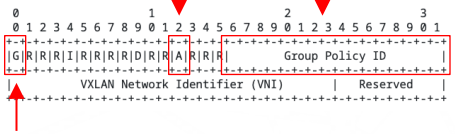

GROUP POLICY TAG PRESENT

ADDED TO EVPN TYPE 2,3,5

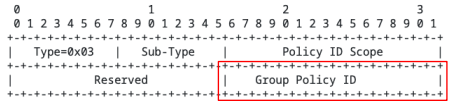

16 BIT = 0-65535

WOW! Multivendor standard! Is it interoperable ? **NO COMMENT**

What problems do we want to solve ?

move beyond architecture with monolithic firewall to **create distributed and scalable security solutions**

Monolithic “firewall centric” architecture

Still the most used architecture in enterprise networks

Limits:

- Firewall performs all the functions
  (not just security)
- Limited performance
- Limited flexibility
- Scale-up is the only possible upgrade
- High costs
- Maintenance is critical for all the services

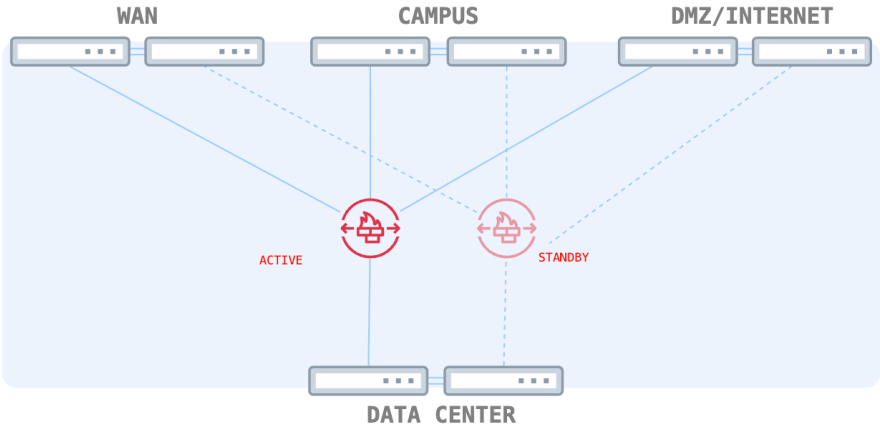

And even with more firewalls, traffic is inspected more than once !!

Modern firewall architecture

Fortigate 3700F

IPS	NGFW	Threat Protection	Interfaces
86 Gbps	80 Gbps	75 Gbps	Multiple 400 GE QSFP-DD, 200 GE QSFP56, 100 GE QSFP28, 50 GE SFP56, 40 GE QSFP+,

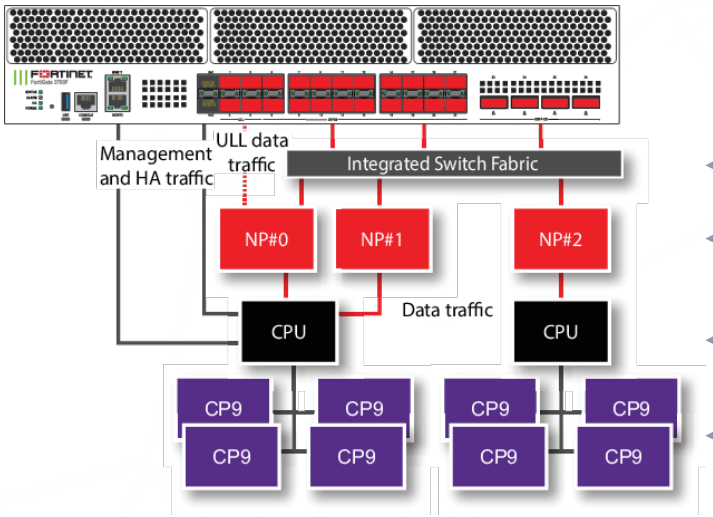

ASIC - L2/L3 Functions

NPU- Network Processor – L4 Functions

Control Plane - Management

CP - Content Processor – L7 Functions

Modern firewalls use different internal components to hardware-accelerate the different functions

https://docs.fortinet.com/document/fortigate/7.6.6/hardware-acceleration/753806/fortigate-3700f-and-3701f-fast-path-architecture

Distribute functions into the EVPN fabric

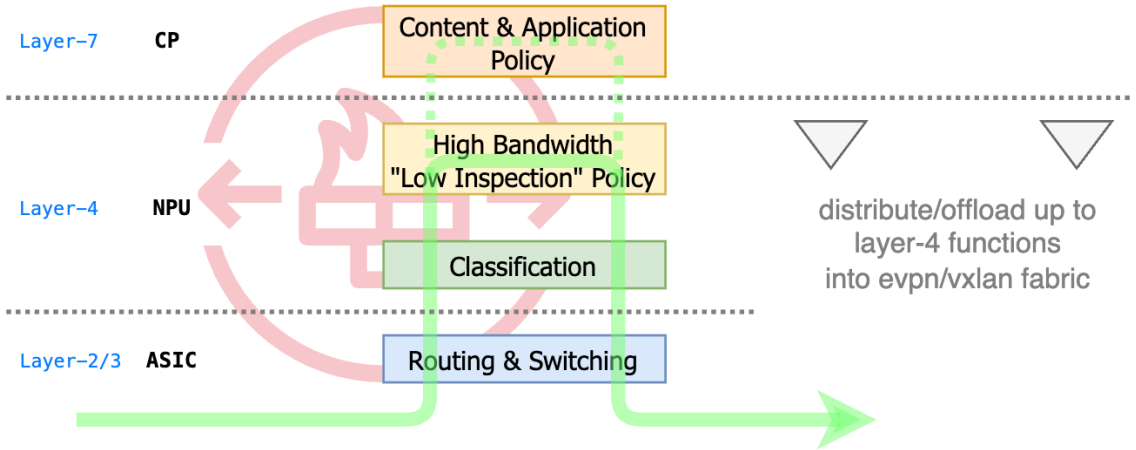

- The different firewall functions are performed by specialized internal components.
- delegate/offload the functions down to layer 4 (including HA and LB) in the EVPN/VxLAN fabric
- And use the firewall only where it adds value, creating a distributed and scalable solution

Step 1 - Distribute L2/L3 functions

Transform the infrastructure into an EVPN/VxLAN fabric

Migrate to EVPN/VxLAN fabric

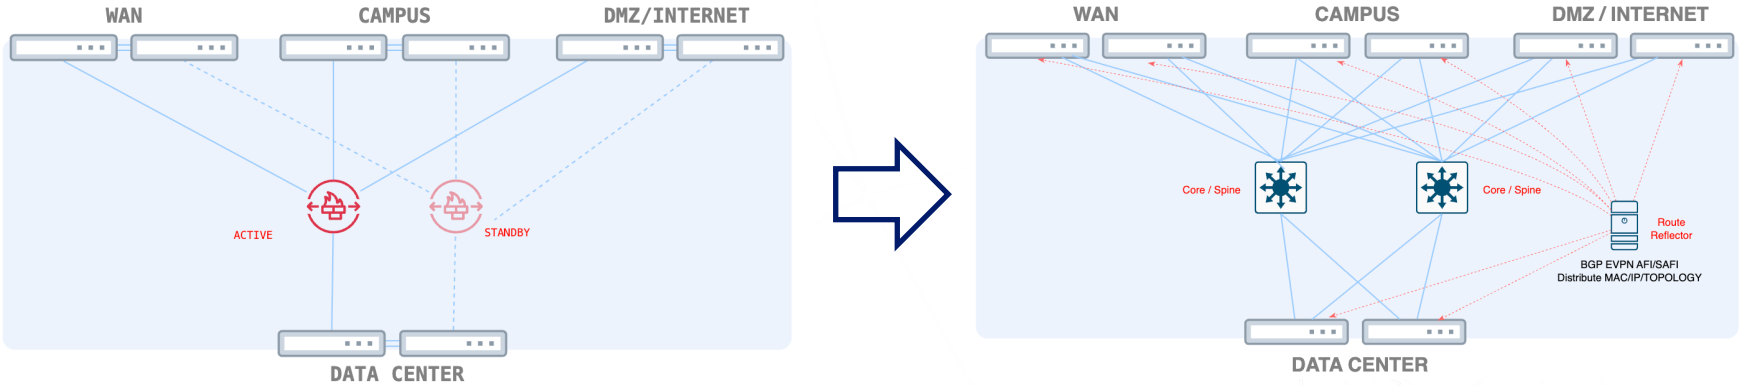

- EVPN/VxLAN on IP Fabric, widely supported in switches, routers, hypervisors and even firewalls
- BGP is used to advertise the topology, mac, ipv4 / ipv6, prefixes and multicast….

No more firewalls as default-gateway!!

Deploy Anycast-Gateway and offload all Switching, Routing, LB & HA functions to the fabric

Most scalable solutions where migration can be incremental in any component

Step 2 - Classification

identifies traffic closest to the source

Create the object database

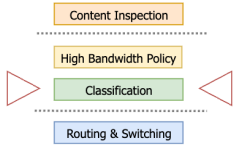

JUNOS

```
nmodena@leaf-02> show configuration firewall family any
filter GBP-CLASSIFY-MAC {
  micro-segmentation;
   term MAC-04-30 {
     from {
       mac-address {
         aa:bb:cc:00:04:30/48;
     }}
     then gbp-tag 102;
}}
filter GBP-CLASSIFY-VLAN {
  micro-segmentation;
  term VLAN-24 {
    from {
      vlan-id 24;
    }
    then gbp-tag 100;
}}
```

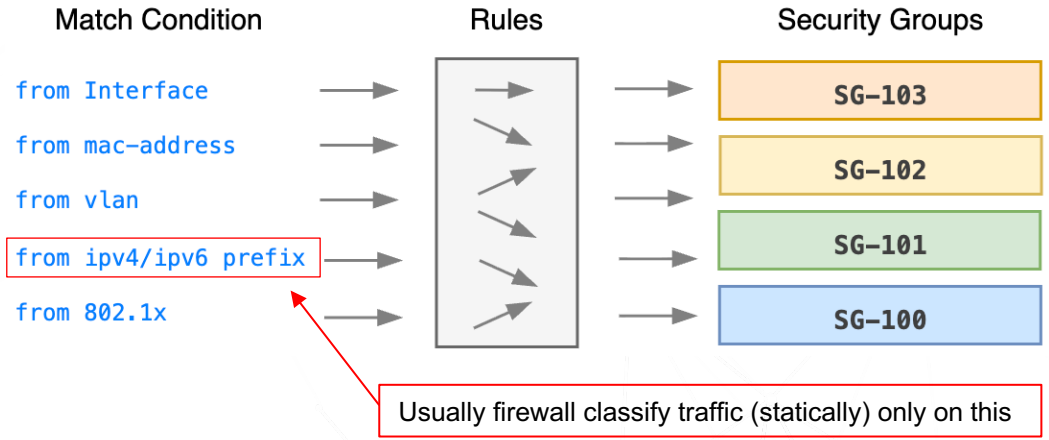

Group Policy ID (0–65535) represents a Security Group of hosts/prefixes **that share the same policy.**
Mapping could be DYNAMIC, based on connected interface, authentication, location, etc or fixed on MAC/IP

MAC and IP/IPv6 addresses are organized “Heterogeneous Group” and tagged with the Security Group

- It’s a **BGP extended community** added to EVPN type-2,3,5 (control-plane)
- It’s a metadata value that is transmitted in the **header of VxLAN packets** (data plane)

Populate the enriched FIB

JUNOS

```
nmodena@leaf-03> show ethernet-switching table vlan-id 24

Ethernet switching table : 8 entries, 8 learned
   Vlan                MAC                MAC         GBP    Logical                SVLBNH/      Active
   name                address            flags       tag    interface              VENH Index   source
   BD24                00:00:cc:1e:00:01  DRP                esi.652                             05:00:00:fd:e8:00:00:4e:38:00
   BD24                2c:6b:f5:ec:72:f0  DRP         100    vtep.32769                          10.100.0.12
   BD24                aa:bb:cc:00:02:30  DR          100    vtep.32769                          10.100.0.12
   BD24                aa:bb:cc:00:03:30  DR          101    vtep.32769                          10.100.0.12
   BD24                aa:bb:cc:00:04:30  DR          102    vtep.32769                          10.100.0.12
   BD24                aa:bb:cc:00:05:30  D           100    ge-0/0/4.0
   BD24                aa:bb:cc:00:06:30  D           101    ge-0/0/5.0
   BD24                aa:bb:cc:00:07:30  D           102    ge-0/0/6.0
```

Group Policy ID as part of type-2 EVPN signaling

```
nmodena@leaf-03> show route table bgp.evpn.0 match-prefix "2:*aa:bb:cc:00:04:30" detail | match comm
                Communities: target:65000:20024 encapsulation:vxlan(0x8) gbp-tag:0L:102
```

FIB & RIB also contains the associated Security Group for each entry

- Security Group for local entries from local classification rules
- Security Group for remote entry from EVPN signaling

Security Functions for L2 and L3

NOKIA
SR Linux

```
A:nmodena@srleaf2# info from state network-instance DATA24 route-table ... group-based-policy-tag | as table
+---------------------+----+------------------+----------------------+-----------------+---------+
|     Ipv4-prefix     | Id |    Route-type    |    Route-owner       | network-instance | GBP-tag |
+=====================+====+==================+======================+=================+=========+
| 0.0.0.0/0           |  0 | bgp-evpn         | bgp_evpn_mgr         | DATA24          |     123 |
| 192.0.2.1/32        |  0 | bgp-evpn         | bgp_evpn_mgr         | DATA24          |     124 |
| 192.0.2.2/32        |  0 | bgp-evpn         | bgp_evpn_mgr         | DATA24          |     124 |
| 10.0.24.0/24        |  0 | bgp-evpn         | bgp_evpn_mgr         | DATA24          |         |
| 10.0.24.0/24        | 10 | local            | net_inst_mgr         | DATA24          |         |
| 10.0.24.1/32        | 10 | host             | net_inst_mgr         | DATA24          |         |
| 10.0.24.11/32       |  0 | bgp-evpn-ifl-host | bgp_evpn_ifl_host_mgr | DATA24          |     100 |
| 10.0.24.12/32       |  0 | bgp-evpn-ifl-host | bgp_evpn_ifl_host_mgr | DATA24          |     101 |
| 10.0.24.21/32       |  0 | bgp-evpn-ifl-host | bgp_evpn_ifl_host_mgr | DATA24          |     100 |
| 10.0.24.22/32       |  0 | bgp-evpn-ifl-host | bgp_evpn_ifl_host_mgr | DATA24          |     101 |
| 10.0.24.255/32      | 10 | host             | net_inst_mgr         | DATA24          |         |
+---------------------+----+------------------+----------------------+-----------------+---------+

A:nmodena@srleaf2# show network-instance default protocols bgp routes evpn route-type 5 detail
...
ip-prefix          : 0.0.0.0/0
Communities        : [target:65000:3024, bgp-tunnel-encap:VXLAN, gbp-tag:0:123, mac-nh:1a:cb:18:ff:00:00]
```

Group Policy ID as part of type-5 EVPN signaling

All Switching and Routing operations involve a source and a destination Security Group

An Accept/Deny or custom security policy must be associated to the forwarding operations

Step 3 – Infrastructure Policy

Micro Segmentation and High Bandwidth Layer-4 policy

Define visibility matrix

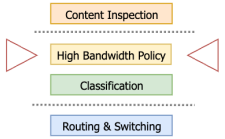

NOKIA
SR Linux

```
--{ running }--[ network-instance SW24 group-based-policy acl ]--
A:nmodena@srleaf2# info
    entry 10 {
        description " L2 and L3 host isolation inside SG "
        match {
            source-group [ SG-100 ]
            destination-group [ SG-100 ]
        }
        action {
            drop {
    }}}
    entry 20 {
        description " https access to public web server "
        match {
            protocol tcp
            destination-group [ SG-200 ]
            destination-port {
                value 443
        }}
        action {
            accept {
    }}}
```

all the host also within the same security group are isolated

any source can communicate with https server in SG-200

Visibility map between Security Groups

SRC / DST	SG-100	SG-101	SG-102	SG-200	SG-300	SG-400
SG-100	drop	P-101-R	P-102-R	P-200-R	accept	accept
SG-101	P-101	accept	drop	P-200-R	accept	accept
SG-102	P-!02	drop	accept	P-200-R	accept	accept
SG-200	P-200	P-200	P-200	drop	accept	accept
SG-300	accept	accept	accept	accept	accept	accept
SG-400	accept	accept	accept	accept	accept	accept

- Accept / Deny or Security Policy (up to Layer 4) can be enforced within/between Security Groups
- “TCP Stateful” are controlled just by checking SYN/ACK flags –> no state distribution <-
- Some platforms (and orchestrators) automatically create the “reverse” ACL for return traffic

The “state” VxLAN gbp flag

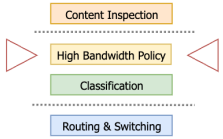

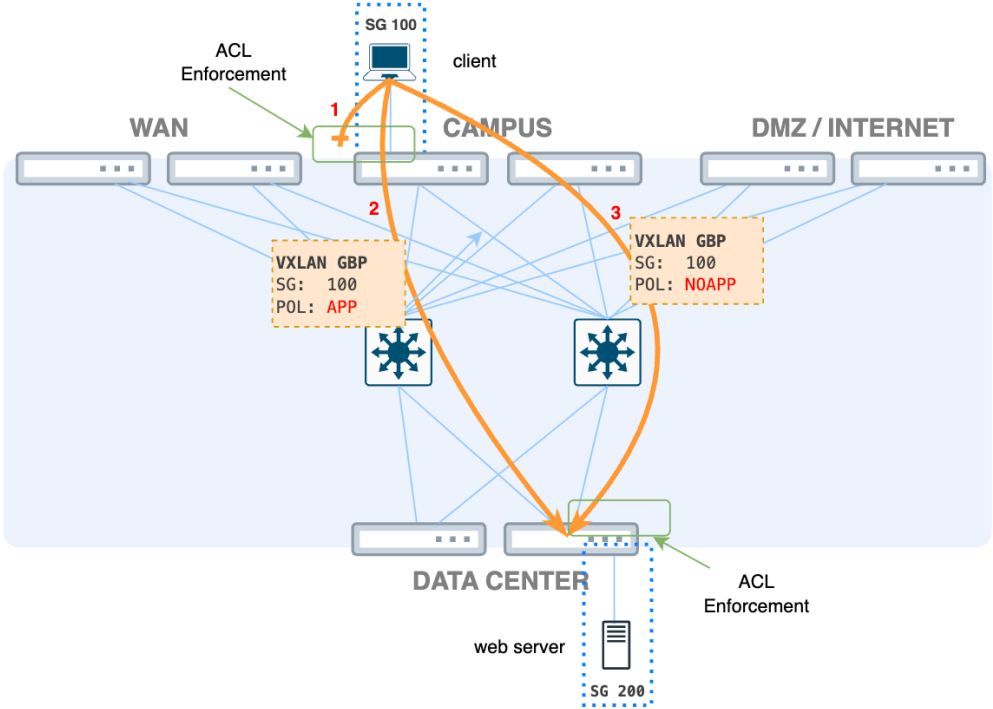

1) Local policy with deny (es not tcp/443)
   -> DROP

2) Local policy with accept (es tcp/443)
   -> VxLAN flag Policy “Applied”
   The destination switch accepts the packet

3) Delegate check to the destination
   -> VxLAN flag policy “not Applied”
   The destination switch must apply the policy

- Security policies are usually enforced locally but may be delegated to the destination switch (VTEP)
- Policy application is tracked in any packet **VxLAN/gbp header state flag**

Why does this solution scale?

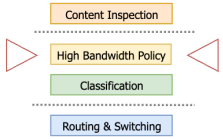

PRO

- Group IDs are easy to look up in ASICs and reduce the size of policy tables (like MPLS)
- The state of policy application is stored in the packet; there is no state distribution (like SR)
- The policies are performed in hardware for maximum performance
- Policy can be enforced close to the source (no wasted bandwidth)
- Only involved traffic classification and policy enforcement into different devices (smaller rule bases)

CONS

- Not a real stateful inspection for TCP
- Not stateful for UDP or other protocols
- Not all the session control that is usually performed by the Firewall NPU

EXTRA

- New DPU-equipped “Smart Switch” can perform more security functions in hardware

Step 4 – Firewall insertion

Service chaining for firewall and advanced security inspection

Introducing firewall inspection

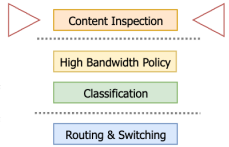

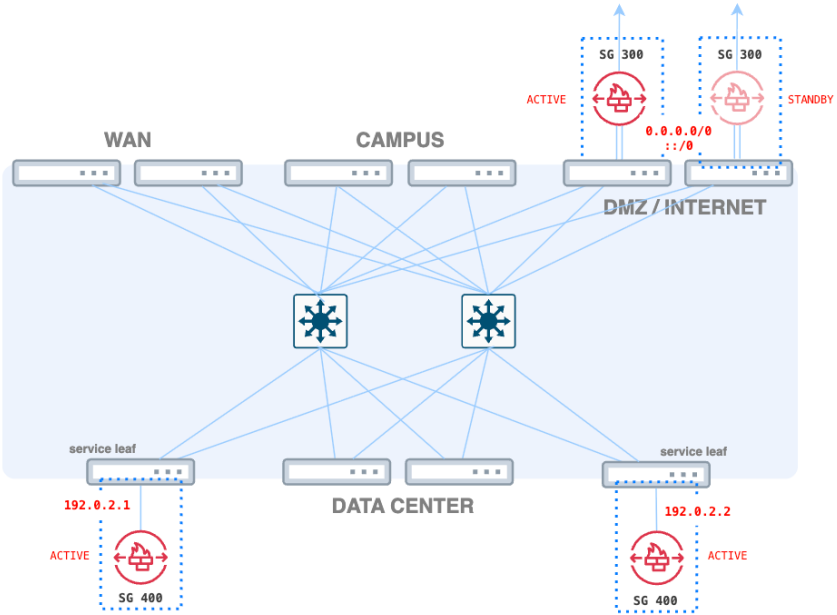

- Multiple firewalls can be connected anywhere into the fabric, also with specialized roles
- Multi-homed and single-arm firewall are supported
- High Availability and Load Balancing on scale-out firewalls groups can be performed by the fabric
- A firewall cluster must be identified or associated to a Security Group

Delegating inspection to a firewall

CISCO NX-OS

Internet traffic is identified by the default route

```
security-group 300 name SG-300
  match external-subnets vrf DATA24 ipv4 0.0.0.0/0
  match external-subnets vrf DATA24 ipv6 0::/0
!
class-map type security match-any CLASS-IP46
  match ipv4
  match ipv6
!
policy-map type security PERMIT-IP46
  class CLASS-IP46
    permit
!
vrf context DATA24
  security contract source 100 destination 300 policy PERMIT-IP46 bidir
```

automatically create policy for return traffic

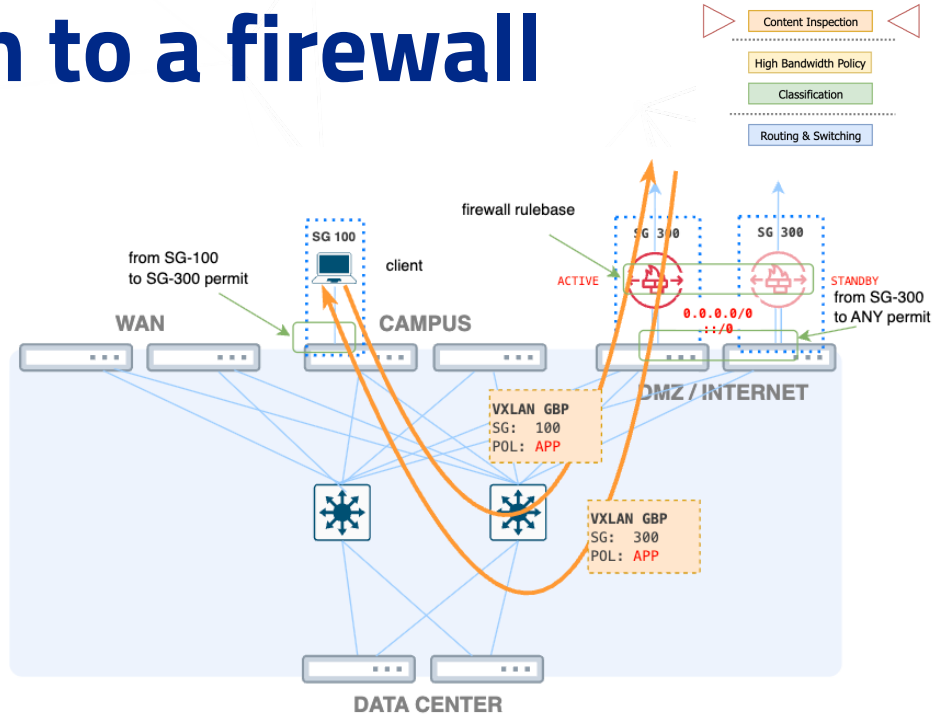

- Delegate internet access control to a border firewall by creating an SG from the default-route
- Create a "permit" policy to transport any traffic sourced in SG100 to the SG300 of the firewall cluster
- The fabric handles high availability and optimal traffic forwarding tracking 0.0.0.0/0 and ::/0 prefixes.
- The firewall enforces all the security policies.

Firewall service insertion

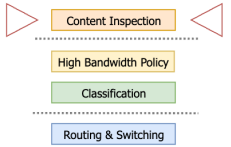

check firewall availability (HA)

CISCO NX-OS

```
epbr service FW-CLUSTER-01
  vrf DATA24
  security-group 400
  probe icmp frequency 1 timeout 1 source-interface loopback1
  service-end-point ip 192.0.2.1
  service-end-point ip 192.0.2.2
!
epbr service-chain HTTP-INSPECTION
  load-balance method src-dst-ipprotocol
  10 set service FW-CLUSTER-01 fail-action drop
!
class-map type security match-any CLASS-HTTP
  match ip tcp stateful dport 80
!
class-map type security match-any CLASS-HTTPS
  match ip tcp stateful dport 443
!
policy-map type security WEB-SERVICES
  class CLASS-HTTPS
    permit
  class CLASS-HTTP
    service-chain HTTP-INSPECTION
!
vrf context DATA24
  security contract source 100 destination 200 policy WEB-SERVICES bidir
```

Active / Active firewall cluster
Symmetric HASH load balancing

Check flags for “stateful” tcp

automatically create policy for return traffic

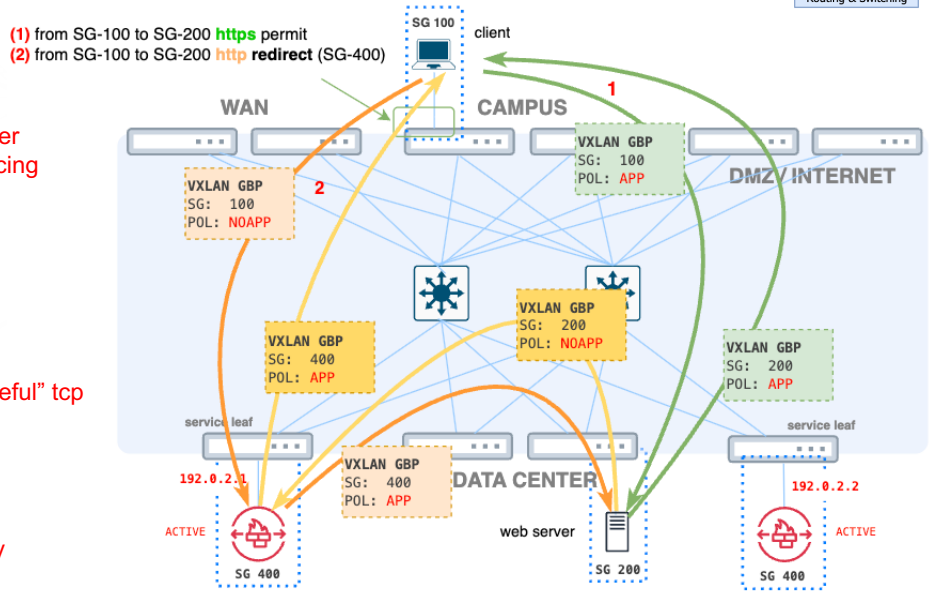

- A firewall inspection service is configured into the fabric and assigned to SG-400
- https traffic is directly permitted by the fabric
- http traffic is redirected to the firewall cluster for inspection (transiting through SG-400)
- HA probes and symmetric load balancing for active/active operations are performed by the fabric

Conclusion

let's try to recap

Conclusion

“Monolithic Firewall” architecture are over

EVPN/VxLAN IP fabric enables high scalability

Group Based Policy integrate security functions directly into Control and Data planes

Group Based Policy leverages micro and macro segmentation

It’s possible to offload switching, routing, HA, LB and up to layer-4 inspection to the fabric

Any questions ?

you can find me:

nicola@modena.to

linkedin.com/in/nmodena

Itnog telegram group

This presentation (and future updates) at https://github.com/nmodena/blog

Special thank to Ivan Pepelnjak, Christian Biasibetti e Alessandro De Prato for the revision