FortiGate™ Transparent Mode

Technical Guide

-

FortiOS v4.0

UNIFIED THREAT MANAGEMENT SOLUTIONS

Contents

Index

Figures:

Additional Fortinet technical information is available from the Fortinet Knowledge Base http://kb.fortinet.com or Product documentation site http://docs.forticare.com.

Version 1.1 - June 2010
Version 1.2 April 2011
Comments and feedback : kb@fortinet.com

Trademarks

1. What is Transparent Mode?

1.1 Network operation in Transparent mode vs. NAT mode

NAT and Transparent (also called TP) mode are the two modes in which a Fortigate or a VDOM can operate.

NAT or Transparent mode will only depict the way that the Fortigate is making IP packets routing or Ethernet frames forwarding decision, while the UTM features will be available in both modes.

Any VDOM configured in a Fortigate can be set in any of those two modes.

For additional information about VDOM, please consult the Fortigate Administration Guide or VDOM and VLAN Guide at http://docs.forticare.com or the Knowledge Base at http://kb.fortinet.com

1.2 NAT mode network operation

In NAT mode, the FortiGate is processing IPv4 or IPv6 packets at the L3 OSI model, and from the network side, is acting as a router (or default-gateway) between the networks attached. Figure 1 shows an example of deployment in NAT mode.

The FortiGate in NAT mode supports the following routing protocols or routing capabilities: Static Routing, RIPv1/v2, OSPF, BGP, Multicast, or Policy Based Routing.

Figure 1: example of deployment in NAT mode

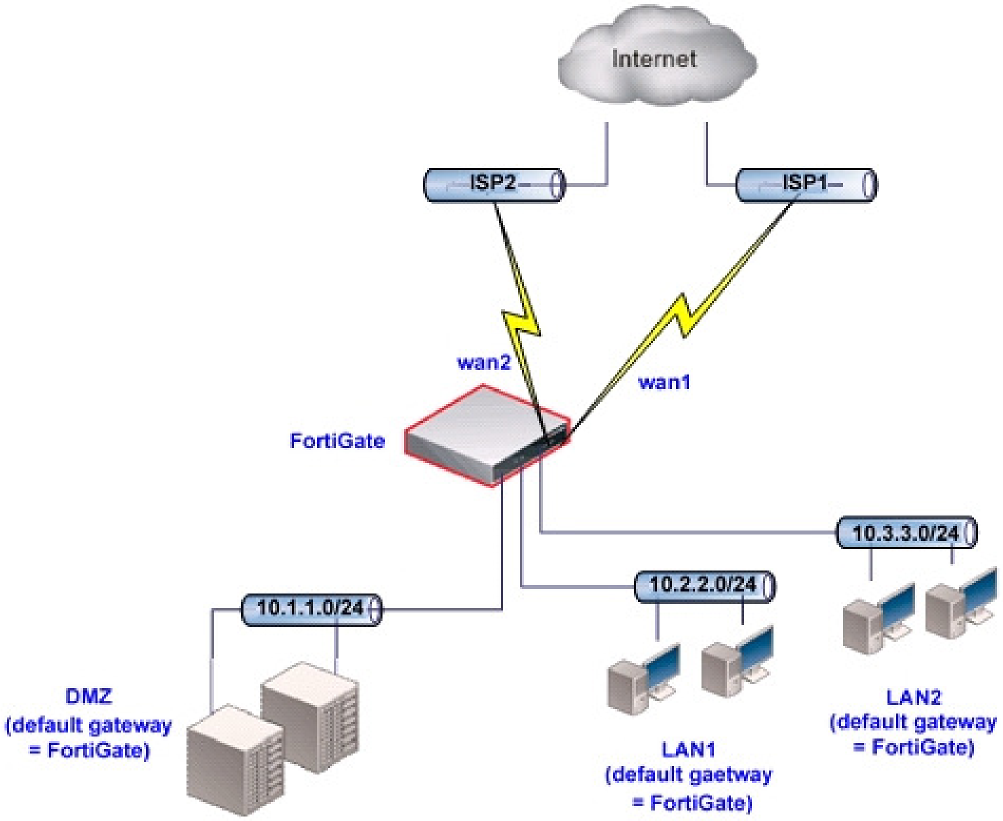

1.3 Transparent mode network operation

In Transparent mode, the FortiGate is processing packets at the L2 OSI network model, and from the network side is acting as L2 switch between different network elements (between a router and a mail server for example.)

A Fortigate in Transparent mode is deployed when, for example, the IP addressing cannot change on the existing network, or the routers in place must not be changed.

Placing a Fortigate in Transparent in a network does not require any changes on the existing devices as they do not know the presence of the Fortigate which is acting "transparently" on a L2 network level.

Figure 2 is showing some examples of deployment in Transparent mode.

Figure 2: examples of deployment in transparent mode.

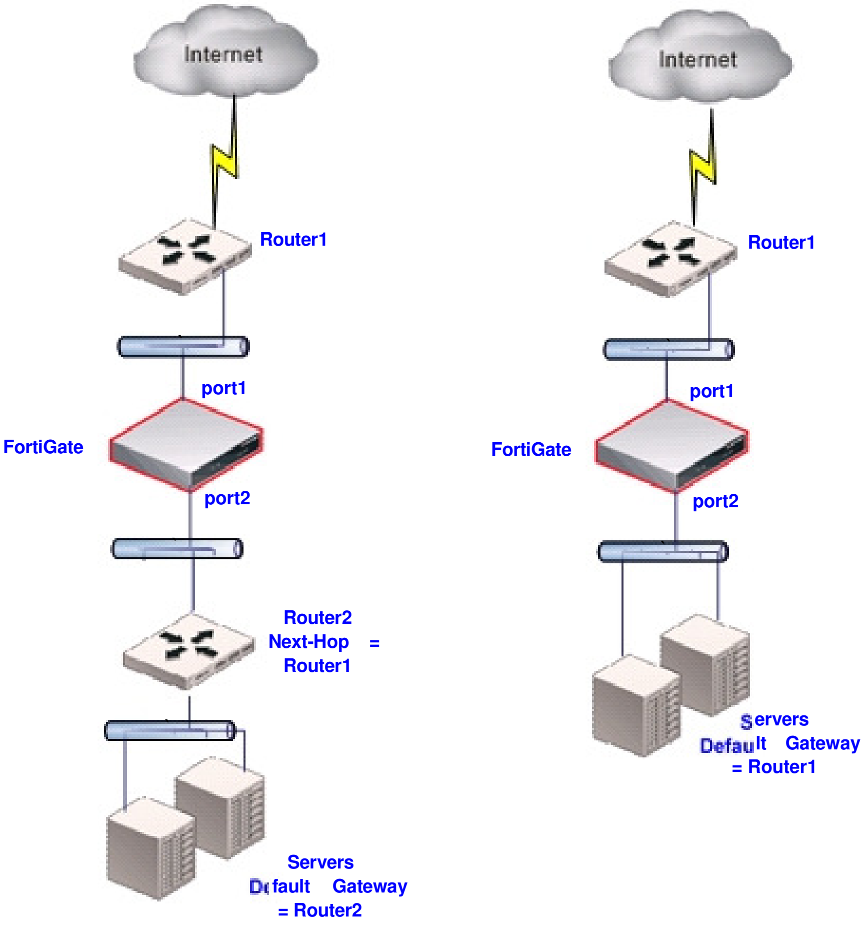

2. Feature Matrix - Maximum values

2.1 FortiGate features and capabilities matrix - NAT and Transparent mode

Feature/capability	NAT	Transparent	Comment
Unicast Routing / Policy Based routing	YES	NO	
VIP / IP pools / NAT	YES	YES	Configurable from CLI only in transparent mode
Multicast routing	YES	NO - Options available to forward multicast packets	
L2 forwarding	NO	YES	In Transparent mode, other frames than IP can be forwarded but with no UTM processing.
Firewall (packet filtering / NAT / Authentication)	YES	YES	
IPv6 capable	YES	YES	
Traffic shaping - TOS classification	YES	YES	
Hardware acceleration	YES	YES	
IPS	YES	YES	
Anti-Virus	YES	YES	
WEB filtering and Fortiguard WEB filtering	YES	YES	
Mail filtering and Anti-Spam	YES	YES	
Other UTM features (IM/P2P, DLP)	YES	YES	
IPSec gateway	YES	YES - Policy based mode only	
SSL gateway	YES	NO	
High-Availability (HA) - Virtual Cluster	YES	YES	
802.3ad (LACP / port aggregation)	YES	YES	
HA port redundancy	YES	YES	FortiGate hardware dependent

802.1q - VLAN trunking	YES	YES	
802.1d - Spanning Tree	NO	NO - option to forward BPDUs	
Logging and reporting (FAMS, syslog or FortiAnalyzer)	YES	YES	
Managed by FortiManager	YES	YES	

2.2 Maximum number of Interfaces in Transparent Mode

Maximum number of Interfaces (VLAN+ physical) per VDOM: 255 (all FortiGate models)

For any other scaling information please consult the FortiGate maximum values matrix available at http://docs.forticare.com

3. Transparent mode concept and terminologies

3.1 Setting Transparent mode

Setting Transparent mode is done as follows:

- If the Fortigate does NOT have VDOM enabled :

```
config system settings
    set opmode transparent
end
```

- If the Fortigate has got VDOM enabled :

```
config vdom
edit <vdom_name>
config system settings
    set  opmode  transparent
end
```

3.2 Network operation : MAC learning and L2 Forwarding Table

When operating in Transparent mode, the Fortigate behaves like a L2 switch in accordance with 802.1d principles:

- The FDB (forwarding Data Base) is populated with the network devices MAC addresses during a MAC learning process, based on the source MAC addresses seen in the Ethernet frames ingressing a FortiGate port. Static MAC entries can also be configured. See example in section Adding a static MAC entry
- Ethernet IP frames forwarding is based on known MAC address on each port
- As Spanning Tree in not running on the Fortigate, a port that comes up goes immediately into forwarding or flooding state. This last state will not occur once unicast MAC addresses are present in the FDB

Important note : If the FortiGate in Transparent mode bridges traffic to a router or host using a virtual MAC for one direction and a different physical MAC for the other direction (for example when VRRP or HSRP protocols are used), it is highly recommended to create a static MAC entry for the virtual MAC. This is to make sure that the virtual MAC address is present in the FDB.

3.3 Network operation : Broadcast, Multicast , Unicast forwarding

In Transparent like in NAT mode, no IPv4 packets are forwarded by the Fortigate from a port to another port unless a firewall policy is matched with action ACCEPT or IPSEC. Here below are some exceptions.

Note that only IP Ethernet frames are considered by the Fortigate. For other frame types, please see section Non IP frames forwarding

Here are some other consideration and exceptions.

Note : L2 (IP) means a L2 frame type 0x0800 (IP) or 0x0806 (ARP)

- **L2 (IP) Broadcast frames forwarding**:
  - **ARP** : by default, ARP broadcasts and ARP reply packets are flooded/forwarded on all ports or VLANs belonging to the same forwarding domain, without the need of firewall policies between the ports. This default behavior is necessary to allow the population of the FDB and allow further firewall policy lookup (see section Transparent mode Firewall processing for more details) . This option is configurable at the interface settings level with the parameter arpforward (enabled by default).

  - **Non ARP** : To forward other IP broadcasts see Non-ARP IP broadcast forwarding

L2 (IP) Multicast frames forwarding: the FortiGate does not forward frames with multicast destination MAC addresses by default. Multicast traffic such as one used by routing protocols or streaming media may need to traverse the FortiGate which should not interfere this communication.

Fortinet recommends that the FortiGate is set up using Multicast policies. This allows for greater control and predictability on traffic behavior. However Multicast traffic may be forwarded through a Transparent mode device using the *multicast-skip-policy* setting. This is detailed in the section Multicast processing and forwarding

- **L2 (IP) Unicast frames forwarding**: a frame with a unicast destination MAC address is subject to firewall processing before being forwarded. (see section Transparent mode Firewall processing for more details). This does not apply to ARP replies (see above).

3.4 Network operation : source MAC addresses in frames sent by or through the Fortigate

- Ethernet frames that are traversing and processed by the Fortigate and not altered from their original source and destination MAC addresses. End devices do not "see" the MAC address of the Fortigate. If NAT is enabled on a firewall policy, the source MAC address will be the Fortigate management MAC address.

- IP packets initiated by the Fortigate (remote management, access to FortiGuard server…) are sent in L2 Ethernet frames that have a source MAC address of the interface in the VDOM with the lowest MAC address.

See below an example with port2 and port3 in the same VDOM, remote access done via port2, but the sniffer trace showing MAC address of port2.

```
fgt300 (global) # diagnose hardware deviceinfo nic port2
[…]
Current_HWaddr                  00:09:0F:85:3F:C4
Permanent_HWaddr                00:09:0F:85:3F:C4

fgt300 (global) # diagnose hardware deviceinfo nic port3
[…]
Current_HWaddr                  00:09:0F:85:3F:C5
Permanent_HWaddr                00:09:0F:85:3F:C5


fgt300 (TP) # diagnose sniffer packet port3 "port 80" 6

3.774236 port3 -- 192.168.171.165.2619 -> 192.168.182.136.80: syn 3961770249
0x0000   0009 0f85 3fc4 0009 0f09 3204 0800 4500         ....?.... 2...E.
0x0010   0030 8071 4000 7e06 98d7 c0a8 aba5 c0a8         .0.q@.~........
0x0020   b688 0a3b 0050 ec23 d109 0000 0000 7002         ...;.P.#..... p.
0x0030   ffff d7e7 0000 0204 05b4 0101 0402
```

3.5 VLANs and Forwarding domain

Difference between a forwarding domain and a VLAN configured on a Fortigate:

- A forwarding domain is used to create separated broadcast domains between VLANs and allow independent VLAN learning - IVL (MAC addresses in the FDB). This would be equivalent to creating VLANs on a regular L2 switch.
  - Note: when VLANs are used in the network, configuring different forwarding domains is essential to avoid broadcast duplications. See also section Default VLAN forwarding behavior for additional information.

- VLANs configured on interfaces are only used for tagging packets egressing the port and classifying packets at ingress
- The packets processed by the direct interface (or port) itself are always sent untagged and must be received untagged.

3.6 Forwarding domain details and configuration

A forwarding domain is used to create separate broadcast domains and confine traffic across two or more ports. It also allows learning the same MAC in different VLANs (IVL). See section VLAN trunking and MAC address learning for more details.

A forwarding domain and its associated ID number are unique across one VDOM, or a Fortigate with VDOM disabled.

Notes:

- ○ Even though the forwarding domain ID is not in relation with the actual VLAN numbers, It recommended, for maintenance and troubleshooting purpose to configure of forwarding domain per VLAN and use the same forwarding domain ID as the VLANs ID.
- ○ Once forwarding domains are configured, it is possible to configure Firewall Policy(ies) only between ports or VLAN belonging to the same forwarding domain

Each new VDOM will create a new bridge instance in the Fortigate.

The figure 4 shows an example with 3 forwarding domains and VLANs configured (forwarding domain 0 is the default on a Fortigate or VDOM in TP mode).

In this example there are two VDOMs in TP mode: root and MGMT.

- Root VDOM has got :
  - ○ 3 forwarding domains, 0, 340, and 341.
  - ○ VLAN 340 configured on port1 ; packets will be tagged with ID 340 ○ VLAN 341 configured on port1 ; packets will be tagged with ID 341 ○ All other ports are untagged
- MGMT VDOM has got only the default forwarding domain 0

The expected behavior is the following:

- Packets untagged ingressing port1, port3 and port4 belong to the same broadcast domain in the root VDOM
- Packets tagged with VLAN 340 ingressing port1 and Packets untagged ingressing port2 belong to the same broadcast domain in the root VDOM
- Packets tagged with VLAN 341 ingressing port1 and Packets untagged ingressing port5 belong to the same broadcast domain in the root VDOM
- Packets untagged ingressing port6 belong to a different broadcast domain in the MGMT VDOM

Figure 4: forwarding domain

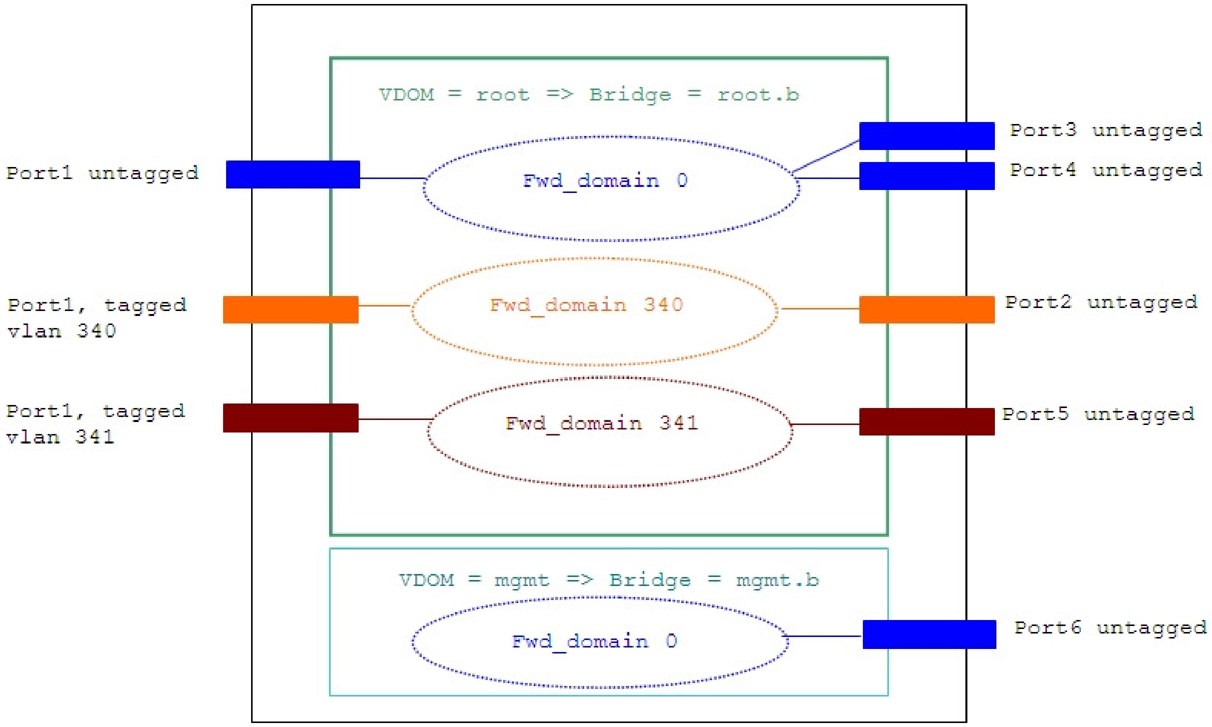

Configuration example for forwarding domain 340:

```
config system interface
    edit "VLAN340"
      set  forward-domain  340
      set interface "port1" set
      vlanid 340
next
    edit "port3"
      set forward-domain 340
next
end
```

3.7 VLAN details and configuration

A VLAN configured on a physical port is used to classify a packet in a broadcast domain in ingress, and to tag packet in egress.

VLAN on the Fortigate is conforming to the standard 802.1q. The

following rules apply to VLAN configuration:

- a VLAN ID can be used only once on the same physical port
- The same VLAN ID can be used on a different port
- The VLAN ID range is from 1 to 4094

3.8 Unknown VLAN processing - parameter vlanforward

When a Fortigate receives a tagged frame with an unknown VLAN ID, the processing is the following:

- By default, the ports are set with the parameter vlanforward = enable ; in this scenario, all tagged frames are forwarded to the ports belonging to the same forwarding domain. This allows inserting the Fortigate between two devices using trunk ports without any further configuration.
- When using forwarding domains in association with VLANs, the parameter vlanforward should be set to disable whenever applicable. In this scenario any new tagged frames with unknown VLAN IDs are dropped by the FortiGate. This is the solution recommended by Fortinet**.**

See also section Default VLAN forwarding behavior

VLAN configuration example: see section Configuration example: Forwarding domain, VLAN and trunk

3.9 VLAN trunking and MAC address learning

A Fortigate port becomes a trunk when 2 or more VLANs are configured on this port, in the same or different forwarding domains.

- ○ Important note: when trunks are configured on a Fortigate, it is essential, to avoid packets looping back on the VLANs of the trunk, to create forward-domain. This will confine all broadcasts and multicast traffic between the interfaces belonging to a same forward-domain.

In the case where a trunk port is configured with VLAN in different forwarding domains, the MAC address of the network device connected to this port is learned in the FDB of each forwarding domain. This is Independent VLAN Learning (IVL).

The figure 5 shows an example where the MAC address will be in the FBD of forwarding domain 400 and 401.

Figure 5 : IVL on a trunk

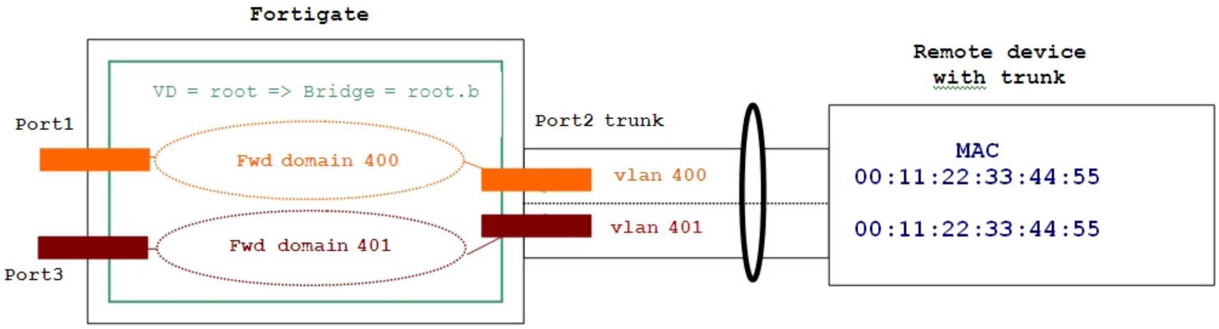

3.10 VLAN translation

A same forwarding domain can include several different VLANs. Therefore, a frame ingressing an interface with a certain VLAN ID can be forwarded to another port with another VLAN ID. This is sometimes referred as VLAN translation.

3.11 Verifiying the L2 FDB of a bridge instance

Viewing the FDB is done from the global mode. Type "config global" from the main prompt before executing the commands below.

- This command lists all bridge instances:

diagnose netlink brctl list

Example:

```
FGT # diag netlink brctl list
```

```
list bridge information
1. root.b           fdb: size=256   used=6     num=7     depth=2     simple=no
2. mgmt.b           fdb: size=256   used=5     num=4     depth=2     simple=no
Total 2 bridges
```

Here above we can see two bridge instances for 2 VDOMs in Transparent mode : `root` and `mgmt`

➢ This command will dump the L2 forwarding table for each VDOM bridge instance:

diagnose netlink brctl name host <VDOM_name>.b

Example for the root VDOM :

```
FGT# diag netlink brctl name host root.b
```

```
show bridge control interface root.b host.
fdb: size=256, used=6, num=7, depth=2, simple=no
Bridge root.b host table
port no device  devname mac addr                ttl     attributes
  2     7       wan2    02:09:0f:78:69:00       0       Local Static
  5     6       trunk_1 02:09:0f:78:69:01       0       Local Static
  3     8       dmz     02:09:0f:78:69:01       0       Local Static
  4     9       internal 02:09:0f:78:69:02      0       Local Static
  3     8       dmz     00:80:c8:39:87:5a       194
  4     9       internal 02:09:0f:78:67:68      8
  1     3       wan1    00:09:0f:78:69:fe       0       Local Static
```

3.12 ARP table in Transparent mode

In Transparent mode, the ARP table is used in the following situations:

- For IP traffic received or originated by the Fortigate itself, and in destination of the management device or next-hop.
- When IPSec is used, the Fortigate uses its ARP table to forward the traffic from the IPSec tunnel to the local destination host(s).

All other forwarding decision is based on the FDB table or optional settings.

3.13 Configuration example: Forwarding domain, VLAN and trunk

The diagram in figure 6 illustrates this example:

Figure 6: network diagram with VLANs

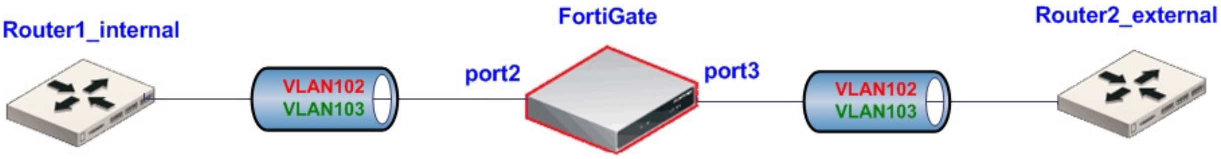

3.13.1 Step 1: Create VLANs and forwarding domains

```
config system interface
    edit    "vlan102_intern"
    set  forward-domain  102
    set   interface   "port2"
    set vlanid 102
next
    edit    "vlan102_extern"
    set  forward-domain  102
    set   interface   "port3"
    set vlanid 102
next
    edit    "vlan103_intern"
    set  forward-domain  103
    set   interface   "port2"
    set vlanid 103
next
    edit    "vlan103_extern"
    set  forward-domain  103
    set   interface   "port3"
    set vlanid 103
next
end
```

3.13.2 Step 2: Create the appropriate Firewall Policies

```
config  firewall  policy
    edit 1
        set srcintf "vlan102_extern"
        set dstintf "vlan102_intern"
            set srcaddr "all"
            set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ANY"  next
    edit 2
        set srcintf "vlan102_intern"
        set dstintf "vlan102_extern"
            set srcaddr "all"
            set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ANY"  next
    edit 3
        set srcintf "vlan103_intern"
        set dstintf "vlan103_extern"
            set srcaddr "all"
            set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ANY"  next
    edit 4
        set srcintf "vlan103_extern"
        set dstintf "vlan103_intern"
            set srcaddr "all"
            set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ANY"  next
end
```

4. FortiGate Remote management

4.1 Management IP configuration in Transparent mode

A Fortigate or a VDOM in Transparent can be assigned with a single IP address for remote access management, and multiple static routes can be configured. This can be used if in-band management wants to be applied.

When out-of-band management is desired (dedicated interface for remote management access), it is recommended to use a separate VDOM in NAT mode.

4.1.1 In-band management details and example

The management IP address of a VDOM in Transparent mode is bound to ALL ports or VLANs belonging to the same VDOM.

The remote access services such as Telnet, HTTPS, SNMP, etc… are subject to the same rules as in NAT mode, and must be enabled/disabled on each port.

Example of management IP configuration in Transparent mode:

```
config system settings
    set manageip 10.1.1.100/255.255.255.0
end

config router static
    edit 1
        set gateway 10.1.1.254
    next
end

config system interface
   edit port1
      set allowaccess ping ssh https snmp
end
```

If VDOM is enabled the above procedure must be preceded by:

```
config vdom
    edit <vdom_name>
```

It is also possible to add a second IP address for management and additional default routes:

```
config system settings
    set opmode transparent
    set manageip 192.168.182.136/255.255.254.0 10.1.1.1/255.255.255.0
end

config router static
    edit 1
        set gateway 192.168.183.254
    next
    edit 2
        set gateway 10.1.1.254
    next
end
```

Note : ping-server (dead gateway detection) is not supported in Transparent mode.

4.1.2 Out-of-band management details and example

When VDOM is enabled and the VDOMs are operating in Transparent mode, it is recommended, to avoid L2 loops and allow more routing flexibility, to keep one VDOM (generally the root VDOM) in NAT mode, with one or more VLAN or physical interface as out-of-band management.

Important note: the management VDOM must have IP connectivity to the Internet to allow communication with the FDS and retrieve services information (AV,IPS, Fortiguard filtering, Support…). All syslog and FortiManager communication also go through the management VDOM.

Figure 7, 8 and 9 show an example where the root VDOM is used as out-of-band management VDOM in NAT mode, with 2 physical ports in two different subnets.

Figure 7: FortiGate out-of-band remote management scenario with VDOM enabled

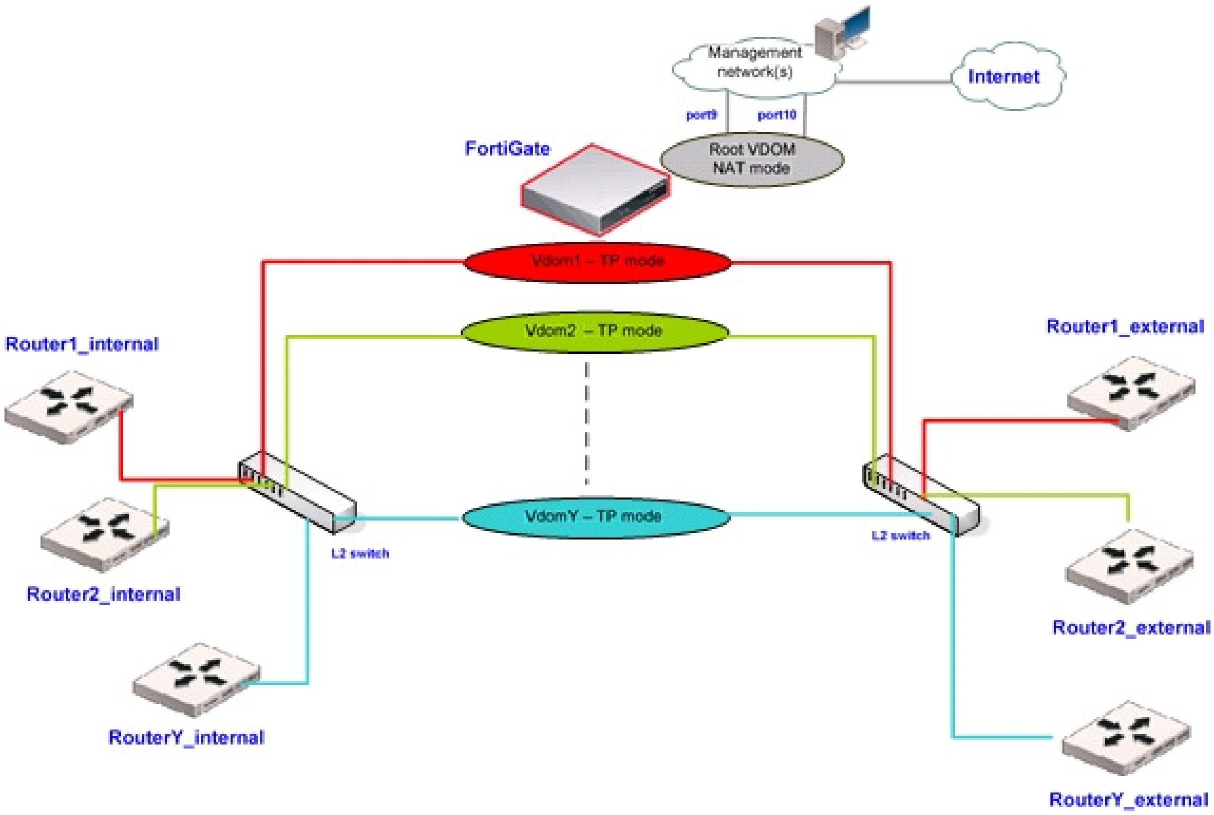

Figure 8: FortiGate Web based manager - Network window with VDOMs assignment

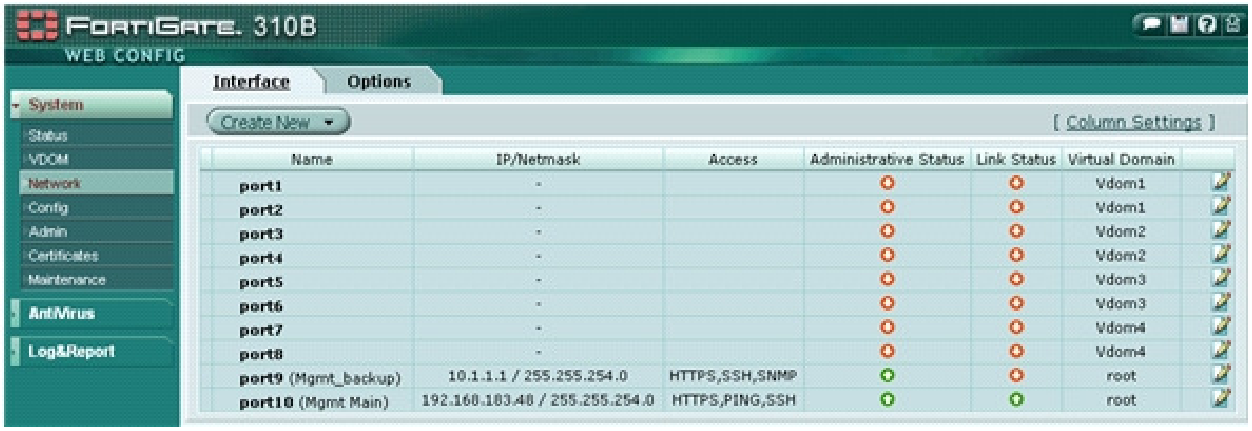

Name	IP/Netmask	Access	Administrative Status	Link Status	Virtual Domain
port1	-		🔴	🔴	Vdom1
port2	-		🔴	🔴	Vdom1
port3	-		🔴	🔴	Vdom2
port4	-		🔴	🔴	Vdom2
port5	-		🔴	🔴	Vdom3
port6	-		🔴	🔴	Vdom3
port7	-		🔴	🔴	Vdom4
port8	-		🔴	🔴	Vdom4
port9 (Mgmt_backup)	10.1.1.1 / 255.255.254.0	HTTPS,SSH,SNMP	🟢	🔴	root
port10 (Mgmt Main)	192.168.183.48 / 255.255.254.0	HTTPS,PING,SSH	🟢	🟢	root

Figure 9: FortiGate Web based manager - VDOM window with Management

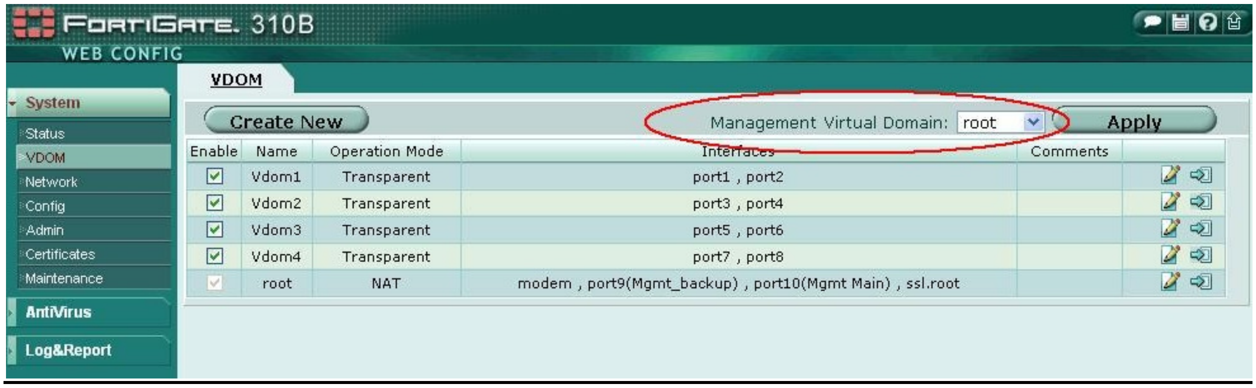

Management Virtual Domain: root

Enable	Name	Operation Mode	Interfaces	Comments
☑	Vdom1	Transparent	port1 , port2	
☑	Vdom2	Transparent	port3 , port4	
☑	Vdom3	Transparent	port5 , port6	
☑	Vdom4	Transparent	port7 , port8	
✓	root	NAT	modem , port9(Mgmt_backup) , port10(Mgmt Main) , ssl.root	

4.2 FortiManager, FortiAnalyzer

FortiManager, logging and reporting and FortiAnalyzer are supported similarly to NAT mode. For more information about this please consult the Fortinet documentation at http://docs.fortinet.com or the Knowledge base at http://kb.fortinet.com

Establishing a secured IPSec communication to a FortiAnalyzer is done as per the example hereafter (from global level if VDOM is enabled). This setting is independent from being in Transparent mode. However, as stated earlier in this section the management VDOM must have IP connectivity to the FortiAnalyzer.

```
FGT (global) # show system fortianalyzer
```

```
config system fortianalyzer
    set status enable
    set server 10.2.2.2
    set encrypt enable
    set  psksecret  fortinet
end
```

5. Transparent mode and HA

Note: for complete information about HA, please refer to the Fortigate Administration Guide or the HA Technical guides available at http://docs.fortinet.com or the Knowledge Base at http://kb.fortinet.com

Any other statement and feature description in this document apply to a Fortigate Cluster running in Active-Passive mode.

5.1 HA MAC address assignment

If a cluster is operating in Transparent mode, the FortiGate Clustering Protocol (FGCP) assigns a virtual MAC address for the Master unit management IP address. Since you can connect to the management IP address from any interface, all of the FortiGate interfaces appear to have the same virtual MAC address.

5.2 Virtual Cluster

If VDOM (virtual domain) is enabled on a cluster operating Transparent Mode, HA Virtual Clustering can be configured in active-passive mode.

This will provide:

- Failover protection between two instances of a VDOM operating on two different FortiGate in the cluster.
- Load balancing between the FortiGate units on a per-VDOM basis

The figure 3 shows an example of HA Virtual Cluster design.

Figure 3: example of HA Virtual Cluster

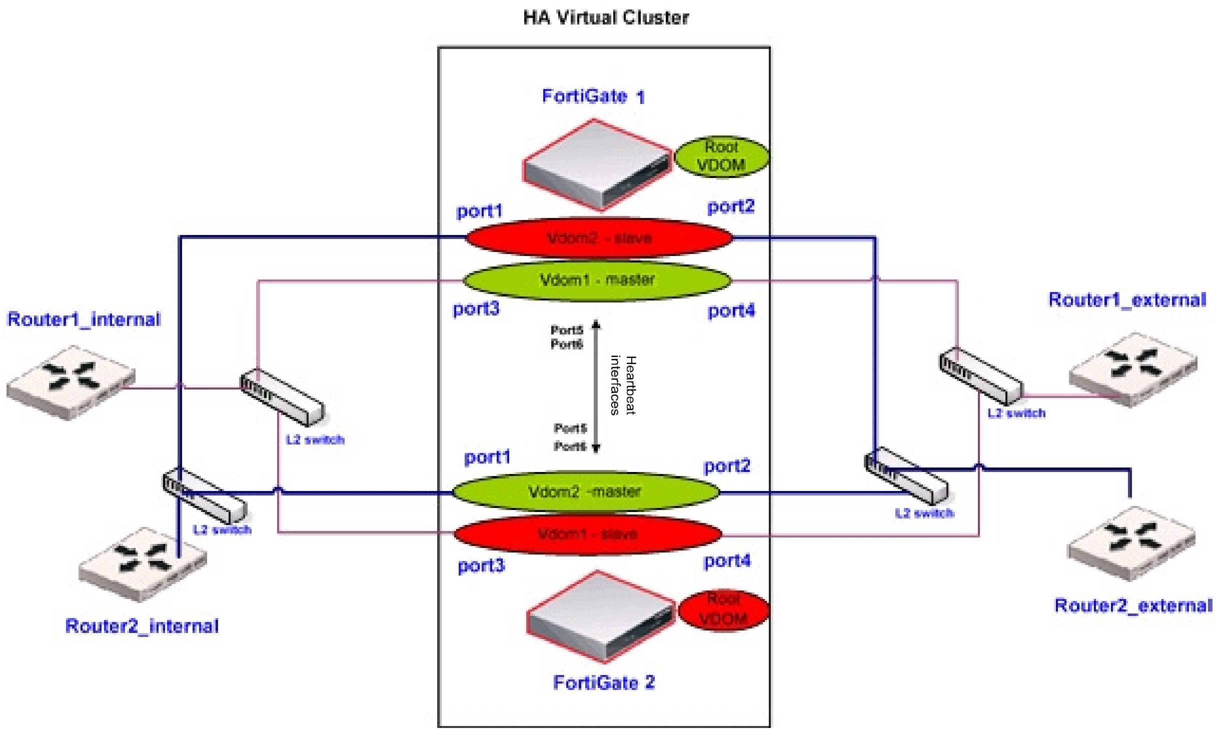

The roles have been defined such as, in normal operation:

- FortiGate1 is Master for Vdom1 and Slave for Vdom2
- FortiGate2 is Master for Vdom2 and Slave for Vdom1

In case of a failure or reboot of a FortiGate, the remaining unit will become Master for Vdom1 and Vdom2.

Note 1: the VDOMs given in this example are showing physical ports but a VDOM can also include VLAN interfaces.

Note 2: the L2 connectivity between the FortiGate is showing 4 separate L2 switches, but it could also be one single switch one each side configured with appropriate VLANs

5.2.1 HA Virtual Cluster CLI configuration example

- FortiGate 1:

```
FGT1 (global) # show  system ha
```

```
config system ha
    set mode a-p
    set hbdev "port5" 0 "port6" 0 set
    vcluster2 enable
    set  override  disable
    set priority 200
        config secondary-vcluster
            set override enable
            set priority 100
            set vdom "Vdom2"
         end
end
```

- FortiGate 2:

```
FGT2 (global) # show  system ha
```

```
config system ha
    set mode a-p
    set hbdev "port5" 0 "port6" 0 set
    vcluster2 enable
    set  override  disable
    set priority 100
        config secondary-vcluster
            set override enable
            set priority 200
            set vdom "Vdom2"
         end
end
```

6. Transparent mode Firewall processing

6.1 Firewall Policy lookup

In Transparent mode like in NAT mode, a firewall policy lookup is based on the couple < source interface + destination interface >.

The matching Firewall Policy will tell what processing and action to apply to the frames: Accept, Deny, IPSec, Protection Profile (content inspection), etc…

The Fortigate proceeds as follows to look for a matching Firewall Policy in Transparent mode:

- Step1: an Ethernet IP frame ingresses a port (or a VLAN on a port), corresponding to a specific bridge instance (from the port VDOM and Forwarding domain). This frame contains a destination MAC address that we will call MAC_D.
- Step2: The Fortigate is making a MAC_D address lookup in the bridge instance to determine the port where MAC_D has been learned. This will be the destination interface.
- Step3 : The Fortigate is then looking for a Firewall Policy corresponding to the couple < source interface + destination interface >. If multiple policies with the same couple < source interface + destination interface > exist, the Fortigate screens all of them from TOP to BOTTOM (as displayed in the configuration), until a match is found. Firewall Policy match lookup is based on “stop-on-match”, therefore the Firewall Policy ordering is important: from most specific in the first position to least specific in last position.

6.2 *Firewall Session list*

The flag **br** in the state line will indicate that this is a “bridged” session. See example below :

```
FGT# diagnose sys session list
```

```
session info: proto=17 proto_state=00 duration=59 expire=128 timeout=0 flags=000
00000 sockflag=00000000 sockport=0 av_idx=0 use=4
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 hakey=0
policy_dir=0    tunnel=/
state=may_dirty br rem
statistic(bytes/packets/allow_err): org=385/5/1 reply=0/0/0 tuples=2
orgin->sink: org pre->post, reply pre->post dev=3->4/4->3 gwy=0.0.0.0/0.0.0.0
hook=pre dir=org act=noop 192.168.182.93:1025->4.2.2.1:53(0.0.0.0:0)
hook=post dir=reply act=noop 4.2.2.1:53->192.168.182.93:1025(0.0.0.0:0)
misc=0 policy_id=1 auth_info=0 chk_client_info=0 vd=1 serial=000006d3 tos=ff/ff
imp2p=0 app=0
dd_type=0 dd_rule_id=0
```

7. Protection Profile in Transparent mode (AV, IPS, WEB filtering…)

The rules and concept for AV, IPS, and any other content inspection and protection profile settings are the same in NAT mode and Transparent mode.

When a protection profile is enabled on a firewall policy for content inspection, the FortiGate acts like a transparent proxy (*) for the protocols that need to be inspected (HTTP, SMTP…).

The FortiGate will therefore intercept the TCP sessions and create its own session from client to server and server to client. The source and destination MAC addresses of the original L2 frames are however not altered in this communication, as described in the section Network operation : source MAC addresses in frames sent by or through the Fortigate

(*) Devices in the network communicating through the FortiGate do not know the presence of the Fortigate.

Refer to the Fortigate Administration Guide for more information about UTM features and settings, at http://docs.forticare.com or the Knowledge Base at http://kb.fortinet.com

8. IPSec VPN in Transparent mode

8.1 Rules and details

IPSec can be used to tunnel network-layer (layer 3) traffic between two VPN peers or between a VPN server and its client. When an IPSec VPN tunnel is established between a FortiGate unit and a remote VPN peer or client, packets are transmitted using Encapsulated Security Payload (ESP) in tunnel mode (AH is not supported).

In Transparent mode, IPSec VPN is supported in Policy based configuration mode only.

IPSec VPN in Transparent mode can be used in those scenarios:

- Encrypt data over routed networks without changing anything on the routers. See example 1 later.
- Encrypt data over a non-routed transport network (extension of a LAN for example). See example 2 later.

The following rules apply to IPSec in Transparent mode:

- If both remote FortiGate IPSec gateways are not in the same broadcast domain (separated by routers) :
  - The hosts on each side must be on different subnets.
  - The FortiGate management IP addresses must be in the same subnet as the local hosts. **This is the preferred option.**
- If both remote FortiGate IPSec gateways are in the same broadcast domain (separated by optical switches for examples), the hosts on each side can be :
  - On the same subnet
  - On different subnet if the appropriate static route is configured on the remote Fortigate
  - The FortiGate management IP addresses can be in any different subnet than the local hosts
- A firewall Policy with the action IPSec is used to send traffic to the remote device into the tunnel. Therefore, it is important to place all remote devices on the appropriate ports of the Fortigate to allow a proper match < source interface + destination interface > . See section Transparent mode Firewall processing for more details.

Warning: This scenario requires that the remote hosts located on the remote FortiGate’s protected subnets have their MAC addresses hardcoded in FortiGate’s static MAC entry list. If this is not configured then it is expected to see outage in network communications.

8.2 *IPSec configuration example 1 : remote sites in different subnets*

This example provides a configuration example for IPSec VPN tunnels between three FortiGate in Transparent Mode (TP) in different subnets, as well as some troubleshooting steps.

The network scenario in figure 10 is used to illustrate this example.

Figure 10: IPSec example 1 - network diagram

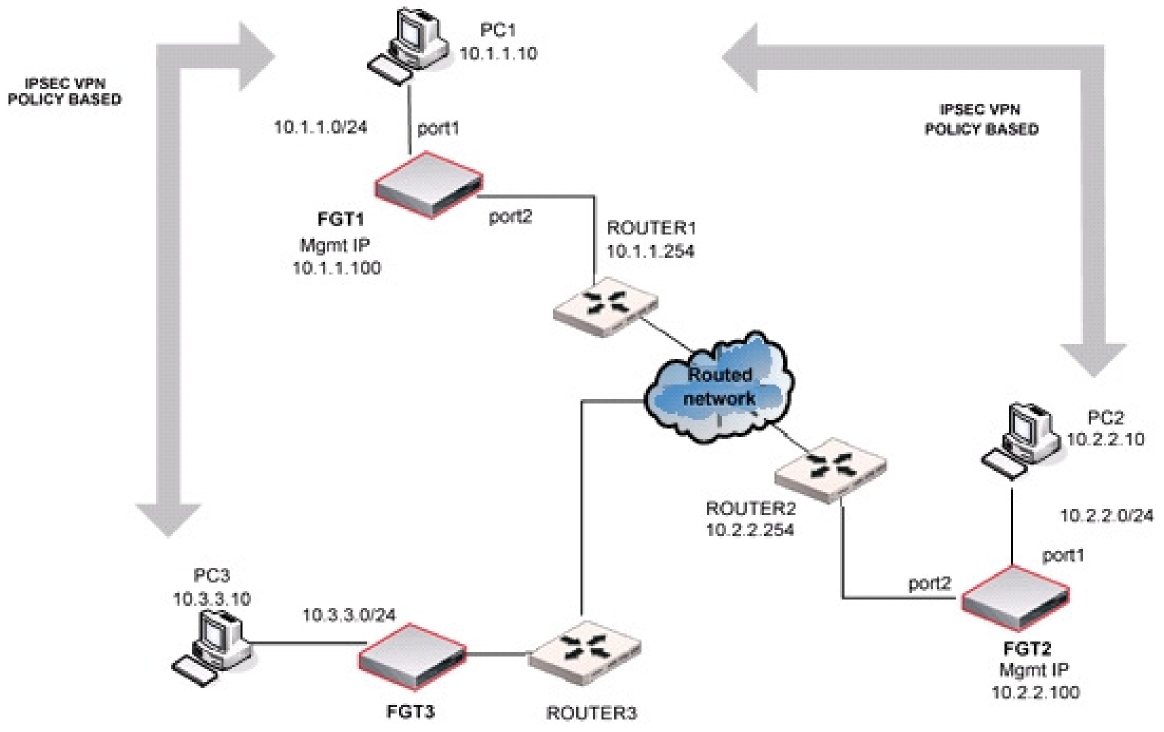

The expectation for this example is that PC1 will be able to communicate via the IPSec tunnels with PC2 and PC3, which are in different subnets.

The requirements for this example are:

- Because both FortiGate are not in the same broadcast domain (separated by routers), the hosts on each side must be on different subnets.
- FortiGate management IP addresses must be in the same subnet as the local hosts
- The default gateways (router1 ,router2, router3) for PC1 , PC2 and PC3 must be behind port2 in order for the FortiGate to match the appropriate Encrypt firewall policy (port1 --> port2)

8.2.1 Configuration of Fortigate 1 (FGT1)

Only relevant parts of configuration are provided

```
config system settings
    set opmode transparent
    set manageip 10.1.1.100/255.255.255.0
end

config router static
    edit 1
        set gateway 10.1.1.254
    next
end

config firewall address
    edit "10.1.1.0/24"
        set subnet 10.1.1.0 255.255.255.0
    next
    edit "10.2.2.0/24"
        set subnet 10.2.2.0 255.255.255.0
    next
    edit "10.3.3.0/24"
        set subnet 10.3.3.0 255.255.255.0
    next
end

config  vpn  ipsec  phase1
    edit "to_FGT2"
        set proposal 3des-sha1 aes128-sha1 des-md5 set
        remote-gw 10.2.2.100
        set psksecret fortinet
    next
    edit "to_FGT3"
        set proposal 3des-sha1 aes128-sha1 des-md5
        set remote-gw 10.3.3.100
        set psksecret fortinet
    next
end

config  vpn  ipsec  phase2
    edit "to_FGT2"
        set keepalive enable
        set phase1name "to_FGT2"
        set proposal 3des-sha1 aes128-sha1
        set dst-subnet 10.2.2.0 255.255.255.0
        set src-subnet 10.1.1.0 255.255.255.0
    next
    edit "to_FGT3"
        set keepalive enable
        set phase1name "to_FGT3"
        set proposal 3des-sha1 aes128-sha1
        set dst-subnet 10.3.3.0 255.255.255.0
        set src-subnet 10.1.1.0 255.255.255.0
    next
end

config  firewall  policy
    edit 1
        set srcintf "port1"
        set dstintf "port2"
            set srcaddr "10.1.1.0/24"
            set dstaddr "10.2.2.0/24"
        set action ipsec
        set  schedule  "always"
            set  service  "ANY"
        set inbound enable
        set outbound enable
        set  vpntunnel  "to_FGT2"
    next
    edit 3
        set srcintf "port1"
        set dstintf "port2"
            set srcaddr "10.1.1.0/24"
            set dstaddr "10.3.3.0/24"
        set action ipsec
        set  schedule  "always"
            set  service  "ANY"
        set inbound enable
        set outbound enable
        set  vpntunnel  "to_FGT3"
    next
end
```

8.2.2 Configuration of Fortigate 2 (FGT2)

Only relevant parts of configuration are provided

```
config system settings
    set opmode transparent
    set manageip 10.2.2.100/255.255.255.0
end
```

```
config  router  static
    edit 1
        set gateway 10.2.2.254
    next
end

config firewall address
    edit "10.1.1.0/24"
        set subnet 10.1.1.0 255.255.255.0
    next
    edit "10.2.2.0/24"
        set subnet 10.2.2.0 255.255.255.0
    next
end

config  vpn  ipsec  phase1
    edit "to_FGT1"
        set nattraversal disable
        set proposal 3des-sha1 aes128-sha1 des-md5 set
        remote-gw 10.1.1.100
        set psksecret fortinet
    next
end

config  vpn  ipsec  phase2
    edit "to_FGT1"
        set keepalive enable
        set phase1name "to_FGT1"
        set proposal 3des-sha1 aes128-sha1
        set dst-subnet 10.1.1.0 255.255.255.0
        set src-subnet 10.2.2.0 255.255.255.0
    next
end

config  firewall  policy
    edit 1
        set srcintf "port1"
        set dstintf "port2"
            set srcaddr "10.2.2.0/24"
            set dstaddr "10.1.1.0/24"
        set action ipsec
        set  schedule  "always"
            set  service  "ANY"
        set inbound enable
        set outbound enable
        set  vpntunnel  "to_FGT1"
    next
end
```

8.2.3 Troubleshooting procedure

All steps given when PC1 pings PC2

A- Verify if IPSec tunnels are up

```
FGT1 # diagnose vpn tunnel list
```

```
list all ipsec tunnel in vd 0
------------------------------------------------------
name=to_FGT2 ver=0 serial=1 10.1.1.100:0->10.2.2.100:0 lgwy=dyn tun=tunnel mode= auto
bound_if=0
proxyid_num=1  child_num=0  refcnt=7  ilast=0  olast=0
stat: rxp=0 txp=0 rxb=0 txb=0
dpd: mode=active on=1 idle=5000ms retry=3 count=0 seqno=1455
natt: mode=none draft=0 interval=0 remote_port=0
proxyid=to_FGT2 proto=0 sa=0 ref=1 auto_negotiate=0 serial=5
  src: 10.1.1.0/255.255.255.0:0
  dst: 10.2.2.0/255.255.255.0:0
```

The above tunnel is down (output given as example) !

```
FGT2 # diagnose vpn tunnel list
```

```
list all ipsec tunnel in vd 0
------------------------------------------------------
name=to_FGT1 10.2.2.100:0->10.1.1.100:0 lgwy=dyn tun=tunnel mode=auto bound_if=0
proxyid_num=1 child_num=0 refcnt=7 ilast=1 olast=1
stat: rxp=0 txp=0 rxb=0 txb=0
dpd: mode=active on=1 idle=5000ms retry=3 count=0 seqno=21 natt:
mode=none draft=0 interval=0 remote_port=0
proxyid=to_FGT1 proto=0 sa=1 ref=2 auto_negotiate=0 serial=1
  src: 10.2.2.0/255.255.255.0:0
  dst: 10.1.1.0/255.255.255.0:0
  SA: ref=3 options=00000009 type=00 soft=0 mtu=1436 expire=1771 replaywin=0 seq
no=1
  life: type=01 bytes=0/0 timeout=1773/1800
  dec: spi=c1a8e951 esp=3des key=24 9213fdf22b150e01abb3535d1a647044eebf772b92f2
f7ee
       ah=sha1 key=20 66a38bf99f0b2d234f64b5a05187995c4f56f6bb
  enc: spi=322067b4 esp=3des key=24 720e5680329937fb3630b7ed70bd41bb3114d3c269ae
8b61
       ah=sha1 key=20 e316113eb6ea03b014b3a5f9c1a3bd386637801a
```

The above tunnel is up !

B - Verify that destination local hosts are seen in the ARP table (necessary for IPSec despite being in TP mode)

```
FGT2 # get system arp
```

```
Address           Age(min)   Hardware Addr      Interface
10.2.2.10         2          00:50:56:00:76:04 root.b
10.2.2.254        0          00:09:0f:30:29:e4 root.b
```

C - Using the debug flow command on the initiator side (example on FortiGate1)

```
FGT1 # diagnose debug flow filter addr 10.1.1.10
FGT1 # diagnose debug flow show  console  enable FGT1
# diagnose debug enable
FGT1 # diagnose debug flow trace start 50
```

```
FGT1 # id=36870 trace_id=615 msg="vd-root received a packet(proto=1, 10.1.1.10:512-
>10.2.2.10:8) from port1."
id=36870 trace_id=615 msg="allocate a new session-00000636"
id=36870 trace_id=615 msg="Allowed by Policy-1: encrypt"
id=36870 trace_id=615 msg="enter IPsec tunnel-to_FGT2"
id=36870 trace_id=615 msg="SA is not ready yet, drop"
id=36870 trace_id=616 msg="vd-root received a packet(proto=1, 10.1.1.10:512->10.2.2.10:8) from
port1."
id=36870 trace_id=616 msg="Find an existing session, id-00000636, original direction"
id=36870 trace_id=616 msg="enter IPsec tunnel-to_FGT2"
id=36870 trace_id=616 msg="encrypted, and send to 10.2.2.100 with source 10.1.1.100"
id=36870 trace_id=616 msg="send out via dev-port2, dst-mac-00:09:0f:30:29:e0"
id=36870 trace_id=617 msg="vd-root received a packet(proto=1, 10.1.1.10:512->10.2.2.10:8) from
port1."
id=36870 trace_id=617 msg="Find an existing session, id-00000636, original direction"
id=36870 trace_id=617 msg="enter IPsec tunnel-to_FGT2"
id=36870 trace_id=617 msg="encrypted, and send to 10.2.2.100 with source 10.1.1.100"
id=36870 trace_id=617 msg="send out via dev-port2, dst-mac-00:09:0f:30:29:e0"
id=36870 trace_id=618 msg="vd-root received a packet(proto=1, 10.2.2.10:512->10.1.1.10:0) from
port2."
id=36870 trace_id=618 msg="Find an existing session, id-00000636, reply direction"
id=36870 trace_id=618 msg="send out via dev-port1, dst-mac-00:50:56:00:76:03"
id=36870 trace_id=619 msg="vd-root received a packet(proto=1, 10.1.1.10:512->10.2.2.10:8) from
port1."
id=36870 trace_id=619 msg="Find an existing session, id-00000636, original direction"
id=36870 trace_id=619 msg="enter IPsec tunnel-to_FGT2"
```

```
id=36870 trace_id=619 msg="encrypted, and send to 10.2.2.100 with source 10.1.1.100"
id=36870 trace_id=619 msg="send out via dev-port2, dst-mac-00:09:0f:30:29:e0"
id=36870 trace_id=620 msg="vd-root received a packet(proto=1, 10.2.2.10:512->10.1.1.10:0) from
port2."
id=36870 trace_id=620 msg="Find an existing session, id-00000636, reply direction"
id=36870 trace_id=620 msg="send out via dev-port1, dst-mac-00:50:56:00:76:03"
```

Note : the message "id=36870 trace_id=615 msg="SA is not ready yet, drop" simply means that the tunnel was not up yet.

D - Using the debug flow command on the receiver side (example on FortiGate2)

```
FGT2 # diagnose debug flow filter addr 10.1.1.10
FGT2 # diagnose debug flow show  console  enable FGT2
# diagnose debug enable
FGT2 # diagnose debug flow trace start 50
```

```
FGT2 # id=36870 trace_id=51 msg="vd-root received a packet(proto=1, 10.1.1.10:512-
>10.2.2.10:8) from port2."
id=36870 trace_id=51 msg="allocate a new session-00000435"
id=36870 trace_id=51 msg="Allowed by Policy-1:"
id=36870 trace_id=51 msg="send out via dev-port1, dst-mac-00:50:56:00:76:04"
id=36870 trace_id=52 msg="vd-root received a packet(proto=1, 10.2.2.10:512->10.1.1.10:0) from
port1."
id=36870 trace_id=52 msg="Find an existing session, id-00000435, reply direction"
id=36870 trace_id=52 msg="enter IPsec tunnel-to_FGT1"
id=36870 trace_id=52 msg="encrypted, and send to 10.1.1.100 with source 10.2.2.100"
id=36870 trace_id=52 msg="send out via dev-port2, dst-mac-00:09:0f:30:29:e4"
```

E - Using the sniffer trace (example on FortiGate2)

```
FGT2 # diagnose sniffer packet any "host 10.2.2.10" 4
```

```
9.460021 root.b out arp who-has 10.2.2.10 tell 10.2.2.100
9.460028 port2 out arp who-has 10.2.2.10 tell 10.2.2.100
9.460034 port1 out arp who-has 10.2.2.10 tell 10.2.2.100
9.460462 port1 in arp reply 10.2.2.10 is-at 0:50:56:0:76:4
9.460462 root.b in arp reply 10.2.2.10 is-at 0:50:56:0:76:4
[...]
49.477368 port2 in 10.1.1.10 -> 10.2.2.10: icmp: echo request
```

```
49.477444 port1 out 10.1.1.10 -> 10.2.2.10: icmp: echo request
49.477898 port1 in 10.2.2.10 -> 10.1.1.10: icmp: echo reply
50.510023 port2 in 10.1.1.10 -> 10.2.2.10: icmp: echo request
50.510079 port1 out 10.1.1.10 -> 10.2.2.10: icmp: echo request
50.510524 port1 in 10.2.2.10 -> 10.1.1.10: icmp: echo reply
```

Note : the above ARP process in Transparent mode with IPSec is allowing to Fortigate to :
- Identify the MAC address of the destination device 10.2.2.10
- Populate the MAC table (see below), which in turn will give a destination interface and allow a Firewall policy look-up

F - Check the FDB entries for the destination

```
FGT2 # diagnose netlink brctl name host root.b
```

```
port no  device  devname   mac addr                ttl
[..]
  1       2       port1     00:50:56:00:76:04       0
```

8.3 *IPSec configuration example 2 : remote sites in the same subnet and one remote subnet*

This example provides a configuration example for IPSec VPN tunnels between two FortiGate in Transparent Mode (TP) in the same subnet separated by a L2 transparent network, and one remote subnet on the second site.

The network scenario in figure 11 is used to illustrate this example.

Warning: This scenario requires that PC1's MAC address is added to the FortiGate's static MAC table. The preferred scenario would be to have a router installed between the 2 FortiGate's.

Figure 11: IPSec example 2 - network diagram

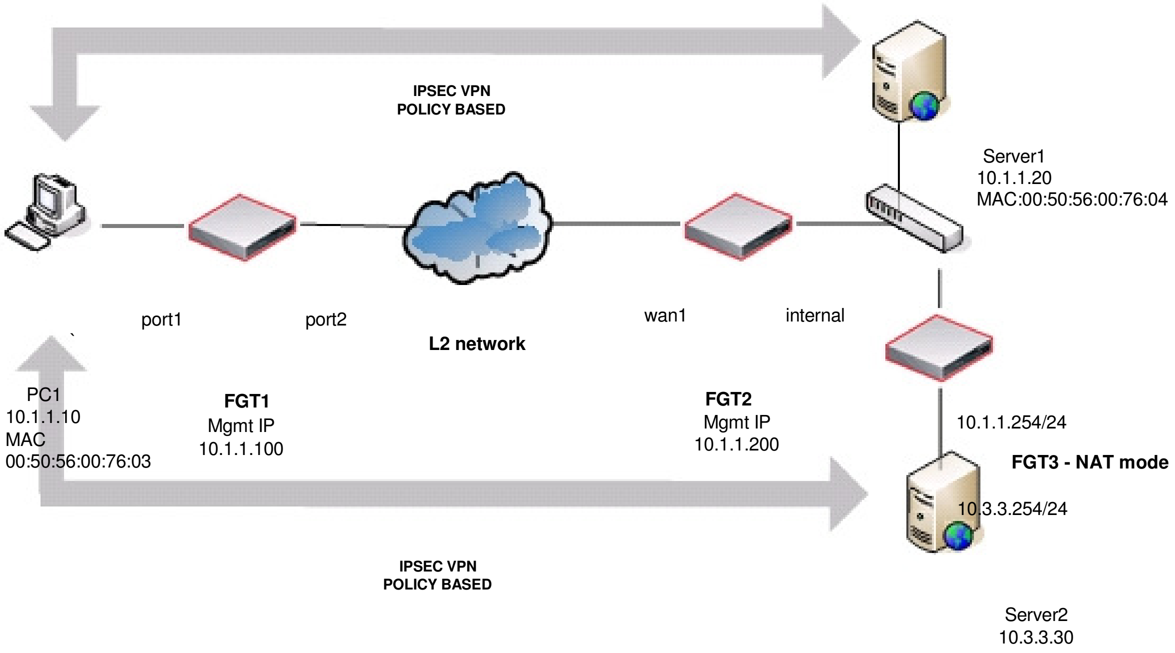

The expectation for this example is that PC1 will be able to communicate via the IPSec tunnel with Server1 in the same subnet, and Server2 in a different subnet.

The requirements for this example are:

- The default gateway (FGT3) for PC1 and all remote device must be behind port2 of FGT1, in order for this FortiGate to match the appropriate Encrypt firewall policy (port1 --> port2)
- Despite being in Transparent mode, **FGT2 must have a valid route to Server2**
- FGT3 is used as a router between subnet 10.1.1.0/24 and 10.3.3.0/24.

PC1 MAC address added to FGT2 static MAC entries
Server1 MAC address added to FGT1 static MAC entries

8.3.1 Configuration of Fortigate 1 (FGT1)

Only relevant parts of configuration are provided

```
config system settings
    set opmode transparent
    set manageip 10.1.1.100/255.255.255.0
end

config router static
    edit 1
        set gateway 10.1.1.252
    next
end

config system mac-address-table
    edit 00:50:56:00:76:04 ==>Server1
        set interface port2
       next
end

config  firewall  address
    edit "all"
    next
    edit "Server1"
        set subnet 10.1.1.20 255.255.255.255
    next
    edit "Server2"
        set subnet 10.3.3.30 255.255.255.255
    next
    edit "10.1.1.0/24"
        set subnet 10.1.1.0 255.255.255.0
    next
    edit "gateway"
        set subnet 10.1.1.254 255.255.255.255
    next
end

config  vpn  ipsec  phase1
    edit "to_FGT2"
        set proposal 3des-sha1 aes128-sha1 des-md5 set
        remote-gw 10.1.1.200
        set psksecret fortinet
    next
end
```

```
 config vpn ipsec phase2
 edit "to_FGT2"
  set keepalive enable
  set phase1name "to_FGT2"
  set proposal 3des-sha1 aes128-sha1
  set src-subnet 10.1.1.0 255.255.255.0
 next
end

 config firewall policy
   edit 1
    set srcintf "port1"
    set dstintf "port2"
    set srcaddr "10.1.1.0/24"
     set dstaddr "Server1"
     set action ipsec
     set schedule "always"
     set service "ANY"
     set inbound enable
     set outbound enable
     set vpntunnel "to_FGT2"
     next
 edit 2
    set srcintf "port1"
    set dstintf "port2"
    set srcaddr "10.1.1.0/24"
     set dstaddr "Server2"
     set action ipsec
     set schedule "always"
     set service "ANY"
     set inbound enable
     set outbound enable
     set vpntunnel "to_FGT2"
     next
 edit 3
    set srcintf "port1"
    set dstintf "port2"
    set srcaddr "10.1.1.0/24"
     set dstaddr "gateway"
     set action ipsec
     set schedule "always"
     set service "ANY"
     set inbound enable
     set outbound enable
     set vpntunnel "to_FGT2"
    next

 end
```

Note : Firewall Policy 3 is not mandatory and is only used to allow PC1 to test a ping reachability to its default gateway 10.1.1.254

8.3.2 Configuration of Fortigate 2 (FGT2):

Only relevant parts of configuration are provided

```
config system settings
set opmode transparent
set manageip 10.1.1.200/255.255.255.0
end

config router static
edit 1
set gateway 10.1.1.252
next

edit 2
set dst 10.3.3.0 255.255.255.0 set gateway 10.1.1.254
next
end

config system mac-address-table
edit 00:50:56:00:76:03
set interface wan1
next
end

config firewall address
 edit "all"
next
edit "PC1"
set subnet 10.1.1.10 255.255.255.255
next

edit "10.1.1.0/24"
set subnet 10.1.1.0 255.255.255.0
next

edit "10.3.3.0/24"
set subnet 10.3.3.0 255.255.255.0
next
end

 config vpn ipsec phase1
  edit "to_FGT1"
 set proposal 3des-sha1 aes128-sha1 des-md5 set remote-gw 10.1.1.100
 set psksecret fortinet
 next
 end

 config vpn ipsec phase2
  edit "to_FGT1"
 set keepalive enable
 set phase1name "to_FGT1"
 set proposal 3des-sha1 aes128-sha1
 set dst-subnet 10.1.1.0 255.255.255.0
 next
 end
```

```
config firewall policy
 edit 1
set srcintf "internal"
set dstintf "wan1"
set srcaddr "10.1.1.0/24" set dstaddr "PC1"
set action ipsec
set schedule "always"
 set service "ANY" set inbound enable
set outbound enable
set vpntunnel "to_FGT1"
next

edit 2
set srcintf "internal"
set dstintf "wan1"
set srcaddr "10.3.3.0/24" set dstaddr "PC1"
set action ipsec
set schedule "always"
 set service "ANY" set inbound enable
set outbound enable
set vpntunnel "to_FGT1"
next
end
```

8.3.3 Troubleshooting procedure

8.3.3.1 Check the ARP entries of PC1

```
C:\ arp -a
```

```
Interface: 10.1.1.10 --- 0x20003
  Internet Address      Physical Address      Type
  10.1.1.20             00-50-56-00-76-04     dynamic
  10.1.1.254            00-09-0f-85-3f-c8     dynamic
```

Note : MAC address **00-09-0f-85-3f-c8** is the FGT3 interface in subnet 10.1.1.0/24

8.3.3.2 FDB entries of FGT1

```
FGT1 (global) # diagnose netlink brctl name host Vdom1.b
```

```
show bridge control interface Vdom1.b host. fdb:
size=256, used=6, num=6, depth=1
Bridge Vdom1.b host table
port no device  devname mac addr                ttl     attributes
  1     10      port1   00:50:56:00:76:03       0
  2     9       port2   00:50:56:00:76:04       44      static
  2     9       port2   00:09:0f:85:3f:c8       13
  1     10      port1   00:09:0f:88:2f:69       0       Local Static
  2     9       port2   00:09:0f:88:2f:68       0       Local Static
  2     9       port2   00:09:0f:23:01:d6       0
```

Note 1: MAC address **00:09:0f:23:01:d6** is "internal" port MAC address of FGT2 00:09:0F:23:01:D6. This is the MAC address used for management in the Transparent mode VDOM of FGT2, chosen between the lowest MAC address between wan1 (00:09:0F:78:00:74) and internal (00:09:0F:23:01:D6)

8.3.3.3 ARP entries of FGT2

```
FGT2 (TP) # get system arp
```

```
Address           Age(min)   Hardware Addr      Interface
10.1.1.20         82         00:50:56:00:76:04 TP.b
10.1.1.100        13         00:09:0f:88:2f:68 TP.b
10.1.1.254        76         00:09:0f:85:3f:c8 TP.b
```

Note : it is important to have the entry for 10.1.1.254 which is the route to 10.3.3.0/24

8.3.3.4 IPSec Tunnel verification on FGT1

```
FGT1 (Vdom1) # diagnose vpn tunnel list
```

```
list all ipsec tunnel in vd 3
------------------------------------------------------
name=to_FGT2 10.1.1.100:0->10.1.1.200:0 lgwy=dyn tun=tunnel mode=auto bound_if=0
proxyid_num=1 child_num=0 refcnt=10 ilast=0 olast=0
stat: rxp=2754 txp=2945 rxb=308448 txb=176700
dpd: mode=active on=1 idle=5000ms retry=3 count=0 seqno=166 natt:
mode=none draft=0 interval=0 remote_port=0
proxyid=to_FGT2 proto=0 sa=1 ref=2 auto_negotiate=0 serial=1
  src: 10.1.1.0/255.255.255.0:0
  dst: 0.0.0.0/0.0.0.0:0
  SA: ref=3 options=00000009 type=00 soft=0 mtu=1436 expire=1271 replaywin=0 seqno=1e1 life:
  type=01 bytes=0/0 timeout=1750/1800
  dec:  spi=3f148cb7  esp=3des  key=24  834832201a0dbbf60b0098106f08380538dbd94cacd1ad31
       ah=sha1 key=20 b0257a135cba745b956bef3d4b8a6e65934c074b
  enc:  spi=1895305e  esp=3des  key=24  4d3092f0b3f84184d4779f85a9953230bf9bc28bd93c0afa
       ah=sha1 key=20 0c70acf6ad2193ec5934e2a4332fd09f32016e60
  npu_flag=00 npu_rgwy=10.1.1.200 npu_lgwy=10.1.1.100 npu_selid=0
```

8.3.3.5 Sniffer trace on FGT1 when PC1 pings all 3 remote destinations

```
FGT1 (Vdom1) # diagnose sniffer packet any "icmp" 4
```

```
interfaces=[any]
filters=[icmp]
0.342268 port1 in 10.1.1.10 -> 10.3.3.30: icmp: echo request
0.342844 port2 in 10.3.3.30 -> 10.1.1.10: icmp: echo reply
0.342884 port1 out 10.3.3.30 -> 10.1.1.10: icmp: echo reply
0.771700 port1 in 10.1.1.10 -> 10.1.1.20: icmp: echo request
0.772504 port2 in 10.1.1.20 -> 10.1.1.10: icmp: echo reply
0.772539 port1 out 10.1.1.20 -> 10.1.1.10: icmp: echo reply
0.907377 port1 in 10.1.1.10 -> 10.1.1.254: icmp: echo request
0.907850 port2 in 10.1.1.254 -> 10.1.1.10: icmp: echo reply
0.907883 port1 out 10.1.1.254 -> 10.1.1.10: icmp: echo reply
```

8.3.3.6 Sniffer trace on FGT1 filtered on IPSec protocol

```
FGT1 (Vdom1) # diagnose sniffer packet port2 "proto 50" 6
```

```
interfaces=[port2]
filters=[proto 50]
pcap_lookupnet: port2: no IPv4 address assigned

1.249003 port2 -- 10.1.1.100 -> 10.1.1.200:  ip-proto-50 92
0x0000   0009 0f23 01d6 0009 0f88 2f68 0800 4500         ...#..... /h..E.
0x0010   0070 c9e6 0000 3f32 9a48 0a01 0164 0a01         .p... ?2.H...d..
0x0020   01c8 1895 305f 0000 01e2 02b6 37b6 8b2c         ....0_..... 7..,

1.249478 port2 -- 10.1.1.200 -> 10.1.1.100:  ip-proto-50 92
0x0000   0009 0f88 2f68 0009 0f23 01d6 0800 4500         ..../h...#... E.
0x0010   0070 2e31 0000 3f32 35fe 0a01 01c8 0a01         .p.1..?25......
0x0020   0164 3f14 8cb8 0000 01e2 324d 66e2 9236         .d?...... 2Mf..6
```

Note : from the above trace, the MAC address **0009 0f88 2f68** is the MAC address of FGT1 port2 . This is the MAC address used for management in the Transparent mode VDOM of FGT1, chosen between the lowest MAC address between port1 (00:09:0F:88:2F:69) and port2 ((00:09:0F:88:2F:68)

8.3.3.7 Debug flow on FGT1 filtered on Server3

```
FGT1 (Vdom1) # diagnose debug flow filter addr 10.3.3.30 FGT1
(Vdom1) # diagnose debug flow show console enable FGT1 (Vdom1)
# diagnose debug enable
FGT1 (Vdom1) # diagnose debug flow trace start 10
```

```
id=20085 trace_id=11 msg="vd-Vdom1 received a packet(proto=1, 10.1.1.10:512->10.3.3.30:8) from
port1."
id=20085 trace_id=11 msg="Find an existing session, id-00004e85, original direction"
id=20085 trace_id=11 msg="enter IPsec tunnel-to_FGT2"
id=20085 trace_id=11 msg="encrypted, and send to 10.1.1.200 with source 10.1.1.100"
id=20085 trace_id=11 msg="send out via dev-port2, dst-mac-00:09:0f:23:01:d6"
id=20085 trace_id=12 msg="vd-Vdom1 received a packet(proto=1, 10.3.3.30:512->10.1.1.10:0) from
port2."
id=20085 trace_id=12 msg="Find an existing session, id-00004e85, reply direction"
id=20085 trace_id=12 msg="send out via dev-port1, dst-mac-00:50:56:00:76:03"
```

Note : From the trace above, **dst-mac-00:09:0f:23:01:d6** is "internal" port MAC address of FGT2 00:09:0F:23:01:D6. This is the MAC address used for management in the Transparent mode VDOM of FGT2, chosen between the lowest MAC address between wan1 (00:09:0F:78:00:74) and internal (00:09:0F:23:01:D6)

9. Replay traffic scenario

Situations can arise where an identical TCP packet enters twice the Fortigate via 2 different ports. This can be due to a firewall or other network device redirecting packets out on the same port it has received it.

The FortiGate will in this condition detect a replay packet and drop it.

If the network topology or culprit devices cannot be changed to avoid this, the workaround on the Fortigate can be to disable TCP replay verification packets.

```
config system global
    set anti-replay | loose | strict | disable |
end
```

Note : in v3.0 this command was :

```
config system global
    set conn-tracking disable
end
```

The debug flow diagnosis output (*) hereafter shows the message indicating this condition:

```
id=20085 trace_id=179 msg="vd-VDOM_VLAN1 received a packet(proto=6, 10.10.253.9:10709
>10.10.248.5:25) from TO_EXTERNAL ."
id=20085 trace_id=179 msg="Find an existing session, id-00041475, original direction"
id=20085 trace_id=179 msg="replay packet, drop"
```

(*) For additional diagnosis and troubleshooting procedures, please consult the Knowledge Base at http://kb.fortinet.com

10. Other options in Transparent mode

10.1 Adding a static MAC entry in the FDB (forwarding Database)

Example on how to add a static MAC entry (this is a per VDOM command):

```
config system mac-address-table
    edit 00:01:02:03:04:05
        set interface "port3"
    next
end
```

Verifying the FDB table:

```
fgt300 (global) # diagnose netlink brctl name host TP.b
```

```
port no device  devname mac addr                 ttl     attributes
[…]
  2     4       port3   00:01:02:03:04:05        0       Static
```

10.2 Non-ARP IP broadcast forwarding

To forward other IP broadcasts than ARP, use the parameter broadcast-forward :

```
      config system interface
          edit "port2"
              set broadcast-forward enable
          next
      end
```

10.3 Multicast processing

10.3.1 Multicast IP and MAC addresses

Multicast MAC addresses are destined for an identified group of devices in the broadcast domain. Broadcasting video might actually use multicast MAC address to address only the devices that need a specific stream. Routing protocols such as RIP or OSPF also use multicast frames.

A Multicast MAC address is identified by the low order bit (0x01) in the first octet which indicates that this packet is a Layer 2 multicast packet. There are several well-known and standard formats of a multicast frames. The 01:00:5E prefix has been reserved by the IANA for use in mapping L3 IP Multicast addresses into L2 MAC addresses (*).

When mapping L3 to L2 addresses, the low order 23 bits of the L3 IP multicast address is mapped into the low order 23 bits of the MAC address.

Example:

- IP 224.0.0.2 maps to MAC 01-00-5E-00-00-02
- IP 224.10.8.5 maps to MAC 01-00-5E-0A-08-05

Some example of multicast MAC addresses:

- Spanning Tree protocol uses the multicast MAC 01-80-c2-00-00-00 for its standard BPDU.
- VRRP and HSRP protocols use a multicast MAC that start with 01-00-5e-xx-xx-xx
- RIP2 : 01-00-5E-00-00-09

Multicast IP addresses use the Class D address space. The **224.0.0.0 to 239.255.255.255** IP address range is reserved for multicast groups. The multicast address range applies to multicast groups, not to the originators of multicast packets.

10.3.2 Multicast packets forwarding

In Transparent mode, a FortiGate does not forward, by default, frames with multicast destination MAC addresses. Multicast traffic such as one used by routing protocols or streaming media may need to traverse the FortiGate which should not interfere in this communication.

Fortinet recommends that the FortiGate is set up using Multicast policies. This allows a finer control.

Figure 12: L2 Multicast frames forwarding

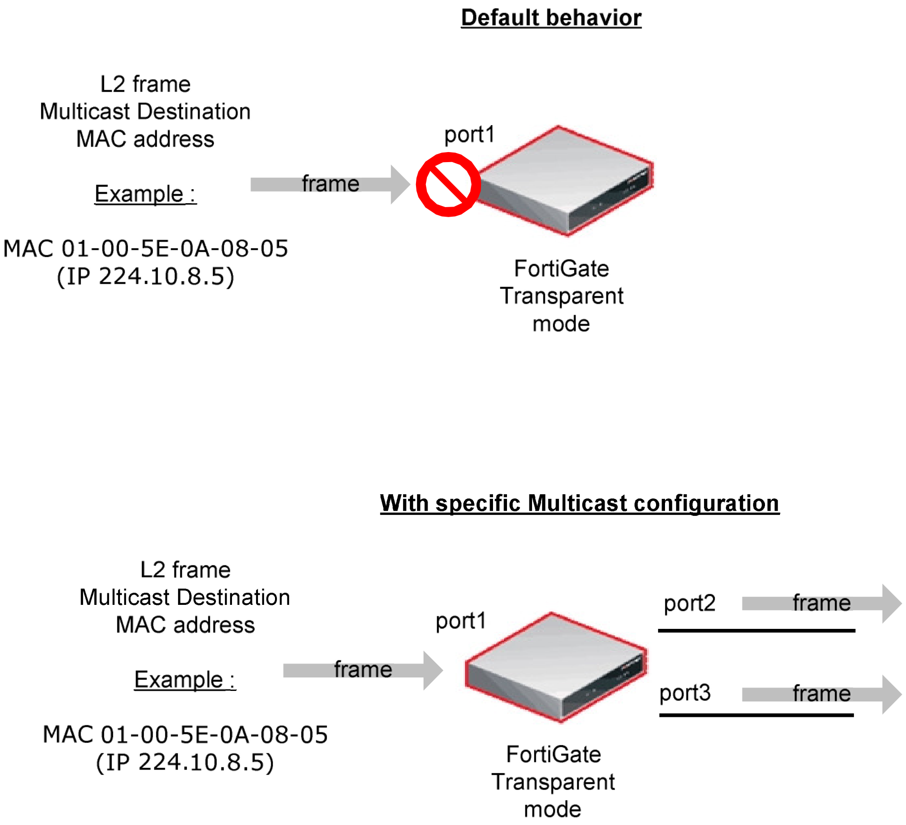

10.3.3 Forwarding all multicast traffic with policy

Multicast traffic may have to be forwarded through a Transparent mode device using the *multicast-skip-policy* setting. This is the configuration for this solution:

```
config system settings
    set multicast-skip-policy enable
end
```

If VDOM is enabled the above procedure must be preceded by:

```
config vdom
    edit <vdom_name>
```

In that case, no check is performed on sources/destinations/interfaces. A multicast packet received on an interface is flooded unconditionally to all interfaces (except the incoming interface) belonging to the same forwarding domain.

10.3.4 Configuring firewall multicast-policy

The use of firewall multicast-policy allows a finer control over the multicast packets. Hereafter are some commented examples. Note that the parameter multicast-skip-policy mentioned above must be left to disabled.

Those policies can only be configured from the CLI.

1- Simple policy

```
config  firewall  multicast-policy
   edit 1
      set  action  accept
   next
end
```

In that case, no check is performed on sources/destinations/interfaces. A multicast packet received on an interface is flooded unconditionally to all interfaces (except the incoming interface) belonging to the same forwarding domain.

2- To restrict incoming and outgoing interfaces:

```
config  firewall  multicast-policy
edit 1
      set srcintf "port1"
      set dstintf "port2"
      set action accept
next
end
```

3- To be more restrictive (example to allow RIP2 packets from port1 to port2 and sourced by 10.10.0.10) :

```
config  firewall  multicast-policy
edit 1
      set srcintf "port1"
      set  srcaddr  10.10.0.10  255.255.255.255
      set dstintf "port2"
      set  dstaddr  224.0.0.9  255.255.255.255
      set action accept
next
end
```

4- This policy will allow all 224.0.0.0/255 range (OSPF, RIPv2, DVMPR…) from port1 to port2

```
config firewall multicast-policy
edit 1
      set  srcintf  "port1"
      set dstintf "port2"
      set  dstaddr  224.0.0.0  255.255.255
      set action accept
   next
end
```

10.4 Configuring VIP (DNAT) in Transparent mode

The following example shows how to configure a VIP and a sniffer trace with the result:

```
config firewall vip
    edit "vip1"
        set extip 192.168.183.48
        set extintf "vlan160_p2"
        set  mappedip  192.168.182.78
    next
end

config  firewall  policy
    edit 4
        set srcintf "vlan160_p2"
        set dstintf "vlan18_p3"
            set srcaddr "all"
            set dstaddr "vip1"
        set action accept
        set schedule "always"
        set service "ANY"  next
end
```

Note : if the mappedip is on a different subnet than the management IP, the Fortigate must have a valid route to this destination

The sniffer trace below shows the destination IP 192.168.183.48 being NATted to 192.168.182.78:

```
fgt300 (TP) # diagnose  sniffer packet any "icmp" 4
```

```
interfaces=[any]
filters=[icmp]
4.126138 vlan160_p2 in 192.168.182.93 -> 192.168.183.48: icmp: echo request
4.126190 vlan18_p3 out 192.168.182.93 -> 192.168.182.78: icmp: echo request
4.126196 port3 out 192.168.182.93 -> 192.168.182.78: icmp: echo request
4.126628 vlan18_p3 in 192.168.182.78 -> 192.168.182.93: icmp: echo reply
4.126661 vlan160_p2 out 192.168.183.48 -> 192.168.182.93: icmp: echo reply
4.126667 port2 out 192.168.183.48 -> 192.168.182.93: icmp: echo reply
```

10.5 Configuring NAT (SNAT) in Transparent mode

Source NAT is an option available in Transparent mode and configurable in CLI only. The following example shows how to configure it and a sniffer trace with the result:

```
config firewall ippool
    edit "nat-out"
        set endip 192.168.183.48
        set startip 192.168.183.48
        set interface vlan18_p3  next
end

config  firewall  policy
    edit 3
        set srcintf "vlan160_p2"
        set dstintf "vlan18_p3"
            set srcaddr "all"
            set dstaddr "all"
        set action accept
        set ippool enable
            set poolname "nat-out"
        set schedule "always"
        set service "ANY"
        set nat enable
     next
end
```

The sniffer trace below shows the source IP 192.168.182.93 being source NATted to 192.168.183.48:

```
fgt300 (TP) # diagnose sniffer packet any "host 10.2.2.1" 4
```

```
interfaces=[any]
filters=[host 10.2.2.1]
4.891970 vlan160_p2 in 192.168.182.93 -> 10.2.2.1: icmp: echo request
4.892003 vlan18_p3 out 192.168.183.48 -> 10.2.2.1: icmp: echo request
4.892007 port3 out 192.168.183.48 -> 10.2.2.1: icmp: echo request
4.933216 vlan18_p3 in 10.2.2.1 -> 192.168.183.48: icmp: echo reply
4.933249 vlan160_p2 out 10.2.2.1 -> 192.168.182.93: icmp: echo reply
4.933253 port2 out 10.2.2.1 -> 192.168.182.93: icmp: echo reply
```

10.6 Non-IPv4 Ethernet frames forwarding

In the situation where non IP frames (or non Ethernet II) frames need to be accepted on a port, the parameter l2forward can be enabled (disabled by default).

This can be used to forward frames such as PPPoE PADI, Appletalk, on other ports belonging to the same forwarding domain.

The procedure is the following:

```
config  system  interface
  edit port1
      set l2forward enable
  next
  edit port2
      set l2forward enable
  next
end
```

10.7 CDP and VTP packets processing

In order to pass CDP(*) or VTP(*) packets through a FortiGate in Transparent mode, the parameter stpforward must be applied on the port configuration.
VTP and CDP packets are sent to the destination MAC address 01-00-0C-CC-CC-CC

(*) VTP : Cisco VLAN Trunk Protocol - CDP : Cisco Discovery Protocol

The example below will allow CDP and VTP packets to be sent from port3 up to the Remote unit, through two VDOMs, via one physical port and three port aggregations.

The following diagram illustrates an example:

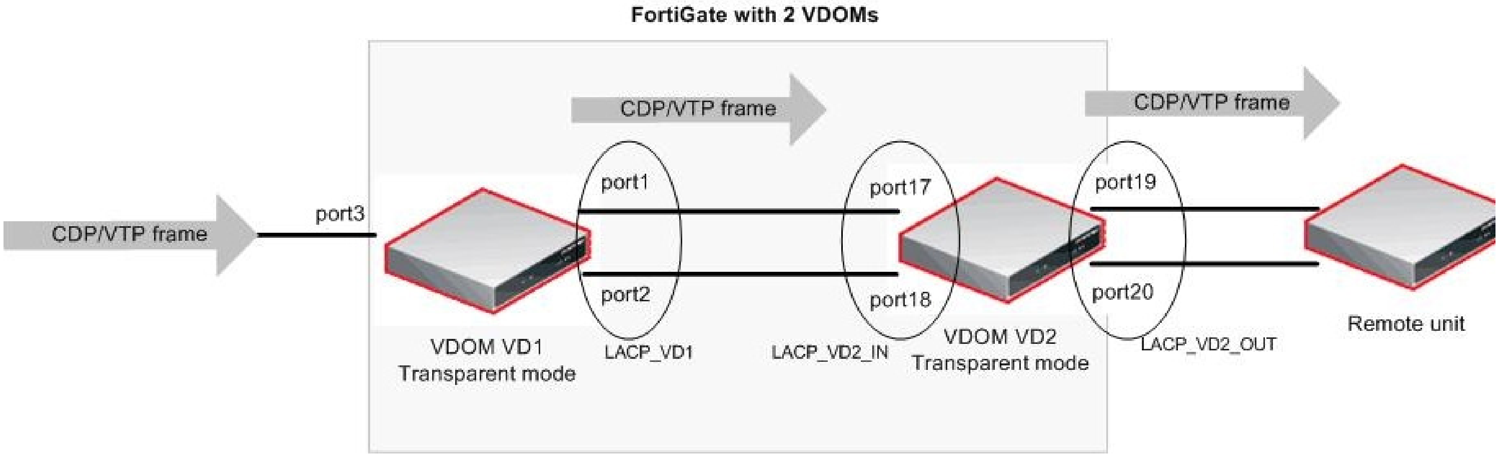

Port and Port aggregation configuration:

```
config system interface
    edit "port1"
        set vdom "VD1"
    next
    edit "port2"
        set  vdom  "VD1"
    next
    edit "port3"
        set vdom "VD1"
        set stpforward enable
    next
    edit "port5"
        set  vdom  "VD3"
    next
    edit "port6"
        set  vdom  "VD3"
    next
    edit "port17"
        set  vdom  "VD2"
    next
    edit "port18"
        set  vdom  "VD2"
    next
    edit "port19"
        set  vdom  "VD2"
    next
    edit "port20"
        set  vdom  "VD2"
    next
    edit   "LACP_VD2_IN"
        set vdom "VD2"
        set  stpforward  enable
        set type aggregate
            set member "port17" "port18"
    next
    edit   "LACP_VD2_OUT"
        set vdom "VD2"
        set  stpforward  enable
        set type aggregate
            set member "port19" "port20"
    next
    edit "LACP_VD1"
        set vdom "VD1"
        set  stpforward  enable
        set type aggregate
            set member "port1" "port2"
    next
end
```

Notes:

- When using aggregation, the stpforward setting needs to be applied only on the port aggregation level, not on the physical port
- This will also forward regular Spanning Tree BPDUs

Verification with a sniffer trace:

```
FGT# diagnose sniffer packet any "" 4
```

```
41.365434 port3 in llc unnumbered, ui, flags [command], length 72
41.365437 LACP_VD1 out llc unnumbered, ui, flags [command], length 72
41.365439 port2 out llc unnumbered, ui, flags [command], length 72
41.365479 LACP_VD2_IN in llc unnumbered, ui, flags [command], length 72
41.365482 LACP_VD2_OUT out llc unnumbered, ui, flags [command], length 72
41.365484 port19 out llc unnumbered, ui, flags [command], length 72
```

See above the CDP packet flow from port3, LACP_VD1 (port2), LACP_VD2_IN, LACP_VD2_OUT (port19)

Note: the following sniffer trace command will filter only CDP or VTP packets :

```
FGT# diagnose sniffer packet port_name "ether host 01-00-0C-CC-CC-CC"
```

Note: Cisco NATIVE VLAN is carrying CDP/VTP frames. The frames of this VLAN must be received on the Fortigate physical interfaces (not VLAN sub-interface). Physical interfaces are the only ones that can send/accept non-tagged packets.

10.8 Default VLAN forwarding behavior

By default, the parameter **vlanforward** is enabled on each physical interface of a Fortigate or VDOM in Transparent mode.

This allows to forward all VLANs traffic of a trunk that was connecting two network devices and where the Fortigate has been introduced, without having to perform any further configuration.

It is recommended to configure forwarding domains for each VLAN and disable this parameter in order to avoid packet from looping into the trunk from one VLAN to another.

Configuration example:

```
config  system  interface
  edit port1
      set vlanforward disable
  next
  edit port2
      set vlanforward disable
  next
end
```

10.9 Spanning Tree BPDUs forwarding

Spanning tree (spt, rstp, pvst, pvst+) BPDUs are not forwarded by default in Transparent Mode.

Introducing a FortiGate in the network must be done with case as L2 loops might be introduced or Spanning Tree broken.
To forward spanning tree BPDUs, the following setting can be applied on each interface where this is required:

```
config  system  interface
  edit port1
      set stpforward enable
  next
  edit port2
      set stpforward enable
  next
end
```

11. Transparent mode reminder and best practices

1) Create forwarding domains when VLANs are used and set vlanforward to disable on all relevant physical interface.

2) The forward-domain ID can be different to the VLAN ID, but it is recommended for troubleshooting and readability to keep them the same.

3) Only interfaces from the same forwarding domains can have Firewall Policies between each others.

4) In order to allow IVL (independent VLAN learning), the VLANs must be placed in separate forwarding domains

5) If an out-of-band management is desired, use if possible a VDOM in NAT mode as management VDOM and create (an) other Transparent mode VDOM(s) for the user traffic.

6) As Spanning Tree BPDUs are not forwarded by default, insert the Fortigate with caution to avoid L2 loops

7) Multicast packets are not forwarded by default; this might cause routing protocols (RIP2,OSPF) disruption.

8) When using HSRP or VRRP configure static MAC entries for the Virtual MAC addresses.