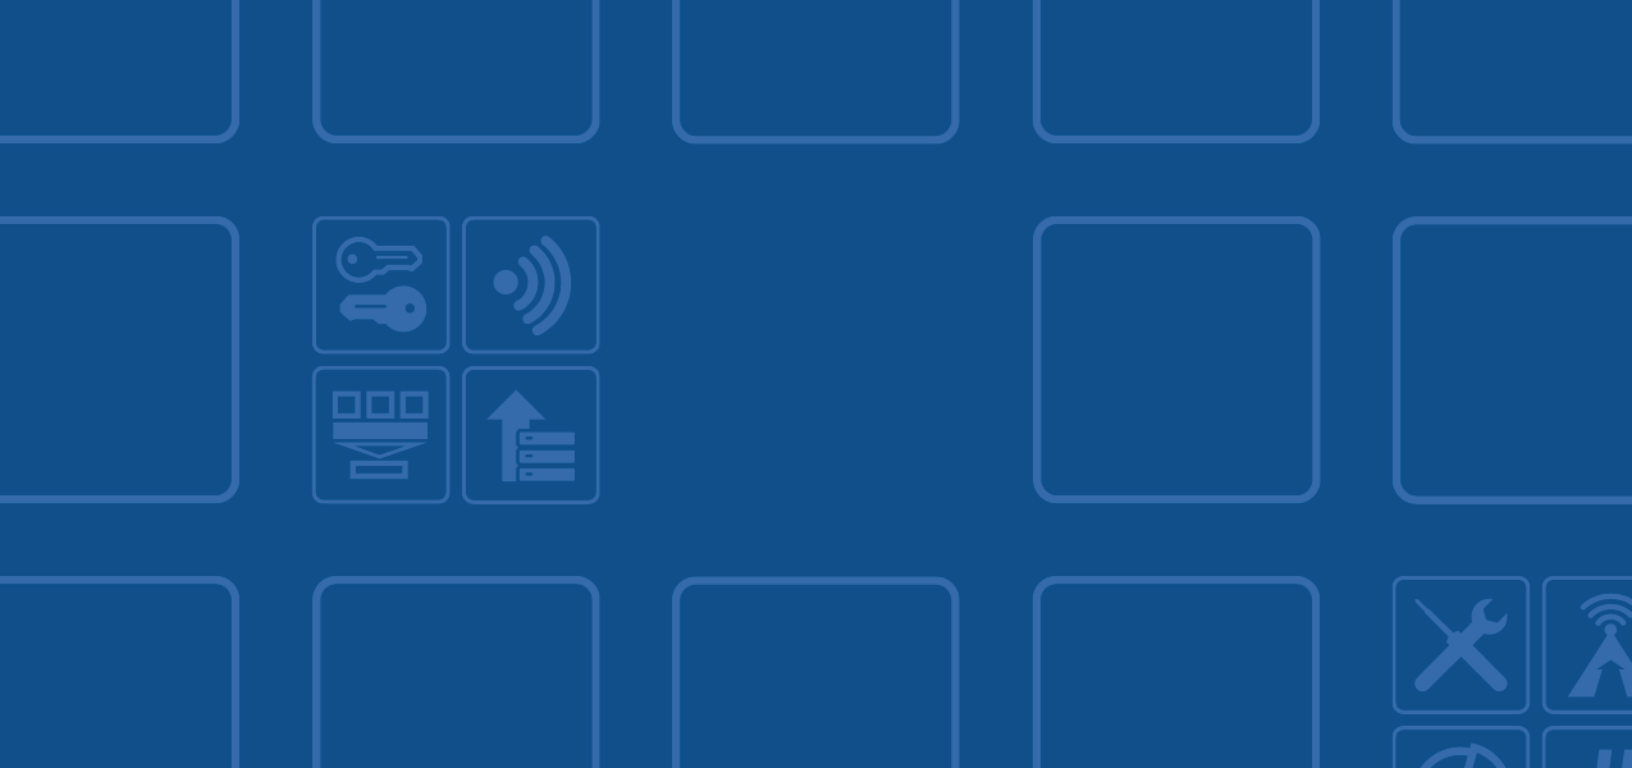

FortiCarrier Administration Guide

Version 5.4.0

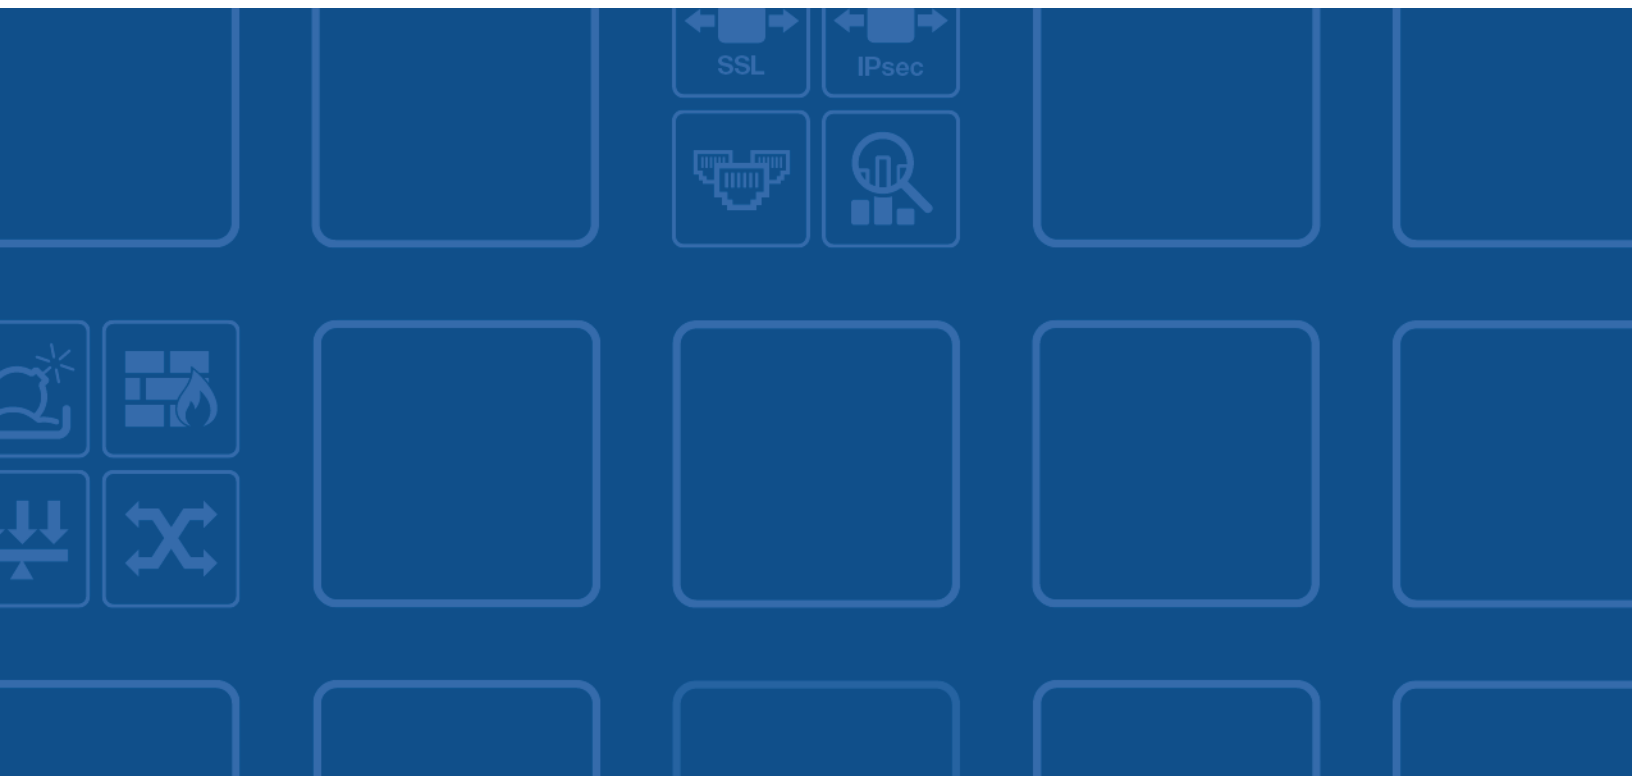

FORTINET DOCUMENT LIBRARY

https://docs.fortinet.com

FORTINET VIDEO GUIDE

https://video.fortinet.com

FORTINET KNOWLEDGE BASE

http://kb.fortinet.com

FORTINET BLOG

https://blog.fortinet.com

CUSTOMER SERVICE & SUPPORT

https://support.fortinet.com

FORTINET NSE INSTITUTE (TRAINING)

https://training.fortinet.com/

FORTIGUARD CENTER

https://fortiguard.com

FORTICAST

http://forticast.fortinet.com

END USER LICENSE AGREEMENT AND PRIVACY POLICY

https://www.fortinet.com/doc/legal/EULA.pdf

https://www.fortinet.com/corporate/about-us/privacy.html

FEEDBACK

Email: techdoc@fortinet.com

Tuesday, October 6, 2020

FortiCarrier - Administration Guide

Version 5.4.0

TABLE OF CONTENTS

Change Log

Date	Change Description
October 6, 2020	Misc minor fixes throughout.
September 23, 2020	Corrected timer values and added more information to the section: Default SSE timeout on page 45.
2017-05-19	FortiCarrier 5.4.0 Initial Release.
2017-07-27	Removed section on GTP.

Introduction

This guide provides information about the use and deployment of the FortiCarrier line of devices.These devices provide Carrier Grade NAT (CGN) also referred to as Large Scale NAT. The firmware for these devices is based largely on FortiOS and most of the detailed knowledge of the administration of a FortiCarrier unit can be found in the reference material for FortiOS. The primary difference between the FortiOS firmware and the firmware for FortiCarrier is that the latter has been optimized for its role in CGN and doesn't need many of the features found in FortiOS, especially those relating to UTM.

This document will therefore focus on the features that relate specifically to FortiCarrier rather than repeating information that relates to both firmwares.

Background and Compliance

- The growing lack of enough public IPv4 addresses to go around has forced ISPs to share the limited public IP addresses amongst their customers by using Network Address and Port Translation (NAPT).
  - RFC6888 recommends common requirements for carrier-grade NATs (CGN).
  - RFC4787/RFC5382/RFC5508 recommends specific CGN requirements on TCP/UDP/ICMP
- NAPT requires per connection and per NAT mapping state maintenance and CGN requires a high connection rate on top of this
- Logging is mandatory to record all NAT mappings or even all sessions
- The current software-based solution does not scale well to meet all the above requirements
- While IPv6 deployment is on the rise, it will need to coexist with IPv4 for many years
  - The use of IPV4 addresses and NAPT is still required in many scenarios to access IPv4 services

Similarities and differences

FortiCarrier is based on FortiOS firmware but focuses on and is optimized for functions like Carrier Grade NAT. FortiCarrier uses specialized FortiASIC-FP1 hardware to accelerate NAT performance and NAT policy handling using IP-pools for both TCP and UDP sessions.

Similarities to FortiOS

- The firmware is based on FortiOS with unneeded functions removed
- The CLI interface, policy control logging etc. are the same as FortiOS

Differences from FortiOS

Modified from FortiOS

- Many feature that do not relate to NAT, such as security profiles, are not available
- There are some additions to configuration options:

- In `firewall ippool`, to describe the IP-pool resource in more detail, new fields have been added , such as `session quota`, `port-range`, `single-port-allocation` mode, etc.
- In `firewall policy`, the options `endpoint-independent-mapping` and `endpoint-independent-filter` have been added for NAT mode

Specific to FortiCarrier

- There is specially designed hardware to achieve a significantly higher CPS than a normal FortiGate
- Due to a reliance on specialized hardware to perform it's function at the level that it does, there is no VM version of the appliance

Features

A large number of the features of FortiCarrier are dependent on the hardware, and while hardware can change from model to model, for the most part, any differences between models is a matter of degree rather than functionality.

The following is a list of features common to all of the FortiCarrier models:

In relation to FortiOS and FortiGate products

The FortiCarrier OS is based on FortiOS with many features removed. Removed features mainly include security features such as VPN and IPsec, and UTM features such as IPS and AV. Small modifications are added to FortiOS to include special features of FortiCarrier.

Session setup

The most important feature of FortiCarrier is that FortiASIC-FP1 hardware not only performs packet forwarding but also handles session setup. This is the reason for the system to achieve very high CPS and still leave host CPU plenty of room to handle other tasks. For hardware to perform the extra task and achieve high performance in session setup, FortiOS policies and routes will be converted to hardware policies and routes and there are several implications.

Session offloading

Comparing to other Fortinet NPU, the FortiCarrier FortiASIC-FP1 has the unique ability to setup sessions by itself and this ability also causes different behavior in packet handling in FortiCarrier system. Taking an example of NP6: NP6 will process packets if a session is found for the packet, otherwise the packet is forwarded to host, where policy/routing are checked, either the packet is denied or accepted, in the latter case, a hardware session(a pair) can be installed into NP6 for accelerated processing of subsequent packets.

This process is called "session offloading". FortiCarrier FortiASIC-FP1 also performs "session offloading" function as a NP6, i.e., accept session installation from host CPU. Below summarizes the processing of a packet as it passes through a FortiASIC-FP1:

1. IPv6 and IPv4 non-TCP/UDP packets are directly forwarded to host without hardware policy/routing check.

IPv6 hardware acceleration feature under in development now. When the feature is ready in future release, IPv6 forwarding can be offloaded by hardware as well.

2. For IPv4 TCP/UDP packets, if routing/policy lookup fail, the packets are dropped.
3. For IPv4 TCP/UDP packets, if the matched hardware policy is "drop", the packets are dropped.
4. For IPv4 TCP/UDP packets, if the matched hardware policy is "forward to host", the packets will be forwarded to host.

In this case, the FortiOS policy corresponding to this hardware policy may not be "forward" action, it can be any type of policy. This hardware policy is installed to bring the packets to host.

5. For IPv4 TCP/UDP packets, if the matched hardware policy is "processing", FortiASIC-FP1 will attempt to setup sessions based on the information obtained from the policy.

Session types

In FortiCarrier system, there are 3 types of sessions in terms of their residence:

1. In host only: e.g., an IPv6 session, or IPv4 non-TCP/UDP session.
2. In host and in FortiASIC-FP1: e.g., a session matching a policy with action of "forward to host". This is the common accelerated case for NP6.
3. In FortiASIC-FP1 only: those sessions are installed/torn down by hardware module in FortiASIC-FP1. Host can query or scan those sessions, but those sessions are not in host's session table.

The session statistics and session listing are also different for those 3 types.

FortiOS policies

Only certain configurations of the policies can enable FortiASIC-FP1 to perform sessions setup for matched traffic. In FortiOS policy configuration, in terms of NAT operation and IP pool settings, below summarize the action of the converted hardware policy:

- FOS policy No-NAT: forwarded to host. This will be improved in post-LPA release so that after session setup is done by host CPU, the sessions can be offloaded to FortiASIC-FP1for fast packet forwarding.
- FOS policy NAT with overload type: forwarded to host. (overload mode will be supported in post-LPA release.)
- FOS policy NAT with PBA or SPA type: FortiASIC-FP1 handle session setup and teardown.

In some cases, the host needs to install "trap" session (sessions to forward the matching traffic to host) of both directions once the initial packet reaches host by way of hardware policy, otherwise the host may not be able to receive the subsequent packets. And the subsequent packets may trigger host to further offload the session to hardware (NP6 behavior).

Pollicy look-up

TCAM device is used for policy look-up in FortiASIC-FP1subsystem. The look-up rate can be deterministic and very high, but some complex FortiOS policies defined in many independent dimensions (source/destination interfaces, source/destination IP addresses, services) may be converted to large number of TCAM rules.

For example:

- For protection purpose there is a restriction of how many TCAM rules a single FortiOS policy can create.
- For the CGN application, the FortiOS policies are usually simple and the number of rules in TCAM is not a problem.

Traffic links

As a high end network appliance, FortiCarrier has 4 FortiASIC-FP1s and each FortiASIC-FP1 has 4 10G links for traffic. All traffic links are connected to an internal switch (inside of FortiCarrier system), to which the front panel user ports are also connected. The user port traffic are distributed to FortiASIC-FP1s traffic links by the switch using certain hashing algorithm on incoming packets. To help accelerating NAT application, the hashing algorithm selection is different for traffic from private side (client side) and from public side(internet side). To achieve this asymmetrical hashing algorithm selection, an interface(physical, LAG or VLAN) must be marked as "private" or "public" side. This information is important for host and FortiASIC-FP1 to determine on which FortiASIC-FP1 the session will be installed and the packets will arrive.

NAT Features

NAT address mapping and filtering behavior:

Besides IP consistency, resource allocation performed by the FortiASIC-FP1 also maintains some desirable behavior for client address mapping and filtering as described in RFC4787.

- EIM: endpoint-independent-mapping, if a client uses an existing source port to connect to a different server, FortiASIC-FP1 reuses the existing mapping to create new sessions. This practice is more compatible for some applications to work with NAT devices, also it is more efficient way to use resource. Note, a new resource allocation counts towards the resource quota. If EIM is triggered, the new session does not cause new resource allocation and the new session only counts towards the session quota.

  Example:

  Client-A has an existing session, represented as 3-tuple of {A.a, B.b, S.s}, where A.a is the client IP and port, B.b is the mapped IP and port, and S.s is the server IP and port.

  When EIM is enabled, if the client uses A.a to connect to another server S1.s1, FortiASIC-FP1 will reuse the public IP and port at B.b, create session of {A.a, B.b, S1.s1}.

- EIF: endpoint-independent-filtering, if another server attempts to connect to a public IP and port which is used by some existing session, when EIF is enabled, FortiASIC-FP1 will create the session and reuse the mapping fro the existing session. When EIF is not enable, the server attempt to connect to the public IP and port will fail. This practice is recommended in RFC4787 for client applications which require such behavior.

  Example:

  Client-A has an existing session, {A.a, B.b, S.s}. When another server S1.s1 attempts to connect to public address and port B.b, when EIF is enabled, FortiCarrier will create the new session as {A.a., B.b, S1.s1}. When EIF is disabled, such connection will be checked in full-policy and it is likely dropped.

IP consistency (Fixed / Deterministic NAT)

- Short term IP consistency - a client always gets the same IP as its existing assignment as long as the client has existing sessions in the system. This is the desired behavior of NAT from a client's perspective and is maintained regardless of pool type.
- Long term IP consistency - • based on pre-determined (predictable through calculation) mapping of clients into available public resource. This is achieved by setting a special property of the pool in addition to pool type and by providing the range of client IP's with the pool. The pool sometimes is called fixed-NAT pool.
- Predictable private IP to public IP mapping
- Predictable private IP to public IP + port range mapping

Single-port-allocation mode (SPA)

- Better port resource sharing
- Limit of client session quota (maximum concurrent sessions)

Port Block Allocation (PBA)

- Lower logging overhead and higher cps
- Limit of client resource quota (number of port blocks)
- Limit of client session quota (maximum concurrent sessions)
- Port-block-size (8/16/32/64/128)
- Number of port blocks (<=8)

One-to-one mapping between client IP and public IP

- Static one to one mapping.
- Dynamic one to one mapping.
- DNAT policy (ingress policy support).

Full cone NAT support

- Endpoint-independent-mapping (EIM): same client's IP/port mapped to the same public IP/port.
- Endpoint-independent-filtering (EIF): allow different servers to use an existing public IP/port pair.

Support of port-preservation feature.

- Port-preserve (system-wide configuration to preserve client's incoming port, e.g., for certain application). Some "application" needs to have its source port preserved to be able to function properly. User can group such "application" (defined as a 2-tuple of L4 port and protocol) in a set, and new sessions matching such "application" will not translate the source port.

  Up to 15 such sets can be defined.

TCP/UDP are supported for NAT with hardware acceleration.

- Non-TCP/UDP are supported the same way as in FortiGate.
- IP-fragmentation is supported.

Other NAT features

- Essential-service pool (system-wide configuration to allow certain essential service traffic work when the client's quota is met.This enables better user experience)
- Hairpin support for traffic between two clients.
- User quota control on port resource usage
- User quota control on max concurrent session count

Public IP pool Management

IP pool based NAT resource management.

- Port-block-allocation (PBA): resources are assigned to a client in the unit of port-block from a public IP.

  The scheme has the advantage of higher throughput and reduced logging traffic.

- Single-port-allocation(SPA): a client is given a public IP and shares the ports with other clients.

  The scheme has the advantage of higher efficiency in resource utilization.

- Two types of quota can be enforced for a client.

1. Resource quota: the maximum number of port-blocks that can be assigned to a client. Only meaningful for pool type of port-block-allocation.
2. Session quota: the maximum number of concurrent sessions a user can have.

IP-pool configuration difference.

- FOS IP pool is converted to hardware resource pool. In hardware the resources are separately maintained for TCP and UDP protocols. So the configured FOS pool properties (number of IP's, pool type etc) are identical in the two hardware resource pools.
- Allow port range configure.
- Allow per-client resource quota (for PBA only) and per-client concurrent session quota.
- Allow multiple ranges of public IP to be configured. And IP pool IP addresses are not restricted to the interface's subnet.
- If a valid port-preservation-profile is configured, a client's source ports matching those defined in port-preservation-profile are preserved while NATing. In comparison, in FortiOS, the pool of the type "fixedport" will preserve any incoming source ports while NATing.
- If "fixed-allocation" property is selected for the pool and a client IP range is provided, each client will be statically assigned with a public IP.

Other IP Pool features

- Software assisted IP selection for dynamic NAT
  - Supports round robin
  - Other schemes available upon special request
- Hardware based allocation and de-allocation

Interface

- Currently the private and public side interfaces are hard-mapped as even and odd port index.
- Interface will be marked as "private" or "public" as an interface configuration field.
- Virtual interface (VLAN) will also be marked similarly.

Policy

- Most of the qualification fields are supported as in FortiGate.
- Support "hw-logging-mode" for hardware generated per-session/mapping logging. Need to configure "log syslogd setting" the same way as currently in FortiGate.

Hardware related

Hardware accelerated session setup for TCP or UDP traffic.

- TCAM based policy look-up with fast and deterministic performance
- TCAM based route look-up, including both LPM based route look-up and policy-based route look-up.
- Hardware based NAT IP/port resource allocation and deallocation.

Hardware based session operations.

- Forwarding or NATing.
- Time stamp update and aging based on timeout configurations.

Flexible per-application timeout configuration for hardware sessions.

- Default system-wide per-state timeout numbers can be tuned for TCP and UDP. Timeout number is used by hardware to signal events (e.g., session tear-down). Examples of states for TCP are: tcp-syn-sent, tcp-syn-wait tcp-established.
- User can also define an "application" in terms of TCP or UDP port, and define a timeout profile for the "application". Hardware sessions matching the "application" will follow the profile in timeout behavior.

Hardware based logging.

- per NAT mapping or per session (two logs per mapping/session)
- No impact on NAT and CPU performance
- Supporting per-session or per-mapping logging.
- Syslog over UDP
- Support up to 16 logging servers with hash based load balance: hash value computed from session information is used to select one of the syslog servers.
- Support 1 netflow-v9 logging server: binary format reducing logging traffic by a factor of about 3.
- NetFlow v9 or IPFIX over UDP (Customized template can be supported with advance notice)

Miscellaneous hardware features

- Performance in PBA mode: 5 Mcps
- Performance in SPA mode: 3 Mcps
- Route Look up
- Policy Look up
- Port Resource Management
- Acceleration

Application-Layer-Gateway (ALG) support

- Some "application" has the L3/L4 address information embedded in the upper layer payload. Such applications can not traverse a NAT device properly without CPU's help. User can define an "application" as a 2-tuple of L4 port and protocol, new sesson packets will be forwarded to host CPU (abbreviated as "host" from now on) for ALG handling.
- Supported the same way as in FortiGate, as long as the ALG can be identified by a pair of destination IP and port.
- FTP/TFTP/SIP/MGCP/ICMP Error/IP-options

HA

Support of session synchronization by hardware in high-availability mode.

- In HA mode, sessions installed by hardware itself are syncrhonized from primary and backup system, over a dedicated link. This greatly reduces the host's burden, as it relieves the host from synchronizing sessions and installing sessions on the primary and backup pair.

Miscellaneous HA features

- Active/Passive mode with HW based NAT/session sync
- Fail over time is similar to standard FortiGate

Other

- 5M/s NAT session setup (full TCP/UDP connections)
- 450M bi-directional concurrent sessions
- Per Application timeout control
- IP fragment translation
  - IP fragments are first de-fragmented and then translated
- ICMP translation
- GRE/PPTP/L2TP translation

Multiple FortiASIC-FP1s

There are four FortiASIC-FP1s in a FortiCarrier. Each FortiASIC-FP1 handles 40Gbps resulting in a total a throughput of 160 Gbps.

The FortiASIC-FP1s handle session setup and packet processing for those sessions. To be able to setup sessions, routes, policies, related ARP tables and IP-pool resources etc. are also loaded into the FortiASIC-FP1s.

Implications

- Packets hitting FortiASIC-FP1 policies (for example, NAT policy using IP-pool) will be completely handled by the FortiASIC-FP1s including resource-alloc/free, session setup/NAT/forwarding/teardown.
- Packets hitting other policies will be forwarded to the CPU, and normal processing will be applied.

Session/packet stats will be provided for those sessions completely handled by FortiASIC-FP1s (as they are not seen by the CPU).

PBA and SPA modes on Dynamic NAT

- For a PBA or SPA pool, SW programs a circular list of all the pool's public IPs to HW; the circular list has a pointer that walks the list in a round-robin fashion
- When a packet comes in, if HW NAT module finds it's the first packet coming from this private side IP, it will take the public IP (pointed by the pointer) from the circular list in step 1, and get a port or port block of this chosen public IP (the port or port block is randomized) to do NAT; the pointer will be advanced to the next public IP
- When a packet comes in and HW NAT module finds its private side IP is an "existing" one, meaning there are ongoing sessions alive, it will use the same public IP that already assigned to this public IP and choose a new port or port block for the current packet/flow.
- When all concurrent sessions of a private side IP has timed out, HW NAT module will remove this private to public IP mapping. So if after a while, the same private side IP comes in again, it will get a new public IP assigned, which may or may not be the same as previous mapping.
- The actual difference for PBA and SPA is the port allocation. In PBA mode, once a private IP got a port block, it'll use the ports inside this block until all ports are used up and then it'll try to get another port block; In SPA mode, each time a new port is needed, it'll randomly get one from all the available ports of this public IP. PBA has the benefits of higher performance.
- Fixed NAT feature is supported in both PBA and SPA modes. When Fixed NAT is enabled, in SPA case, the IP will be assigned by SW config and port is still randomly chosen; in PBA mode, not only the public IP, but the port block also needs to be assigned statically; the public IP assigned to a private IP won't change until the CLI config changes.

Deployment

As a NAT device, the system always has a private side (or inside) and a public side (outside). Currently private side must be using even-indexed physical port, public side must be using odd-indexed physical port. We will support interface role selection by configuration.

Graphic User Interface

Dashboard

The Dashboard of the FortiCarrier interface is similar to the one for FortiOS. It's main purpose being a quick summation of the information relevant to the appliance as a whole. Just like the one on FortiOS, it is made up of widgets to allow for customization.

The widgets applied to the dashboard can be made up of the following:

- System information
- License Information
- Unit Operation
- System Resources
- Alert Message Console
- CLI Console
- Interface History

Network

The **Network** section of the Interface helps with the intuitive configuration of settings that deal with networking. These settings are broken down into the following categories:

- **Interfaces**
- **DNS**
- **Static Routes**
- **Policy Routes**
- **RIP**
- **OSPF**
- **BGP**
- **Multicast**

Interfaces

Interfaces, both physical and virtual, enable traffic to flow to and from the internal network, and the Internet and between internal networks. The FortiCarrier unit has a number of options for setting up interfaces and groupings of interfaces.

System

The **System** section of the Interface helps with the intuitive configuration of settings that deal with appliance administration. These settings are broken down under the following headings:

- **Administrators**
- **Admin Profiles**
- **Settings**
- **HA**
- **SNMP**
- **Advanced**
- **Certificates**

Administrators

By default, the FortiGate has a super administrator account, called `admin`, which cannot be deleted. Additional administrators can be added for various functions, each with a unique user name, password, and set of access privileges.

Admin Profiles

Administrator profiles define what the administrator can do when logged into the FortiCarrier. When you set up an administrator account, you also assign an administrator profile dictating what the administrator will see. Depending on the nature of the administrator's work, access level or seniority, you can allow them to view and configure as much, or as little, as required.

super_admin profile

This profile has access to all components of FortiCarrier, including the ability to add and remove other system administrators. For some administrative functions, such as backing up and restoring the configuration using SCP, super_admin access is required. The super_admin profile cannot be deleted or modified, to ensure that there is always a method to administer the FortiCarrier.

The super_admin profile is used by the default `admin` account. It is recommended to add a password and rename this account once you have set up your FortiCarrier. In order to rename the default account, a second admin account is required. For more information, see "Admin Profiles" on page 17.

Creating profiles

To configure administrator profiles go to **System > Admin Profiles** and select **Create New**.

On the **New Administrator Profile** page, you define the components of FortiOS that will be available to view and/or edit. For example, you can configure a profile so that the administrator can only access the **Firewall Configuration**, which includes firewall policies, addresses, services, schedules, packet capture, and some other parts of the FortiGate configuration. Any other aspects of the FortiGate configuration, including VPNs and security profiles, will be hidden from this administrator.

Adding a local administrator

Only administrators with read-write for **Administrator Users** can create a new administrator account.

To add an administrator - GUI

1. Go to **System > Adminstrators**.
2. Select **Create New**.
3. Add a **Name** for the administrator.

The name of the administrator should not contain the characters <> () #"'. Using these characters in the administrator account name can result in a cross site scripting (XSS) vulnerability.

4. Enter the **Password** for the user. This may be a temporary password that the administrator can change later. Passwords can be up to 256 characters in length.
5. Enter the same password in the **Confirm Password** field.
6. Set **Type** to **Local User**.
7. Set the **Administrator Profile** from the drop down menu.
8. If you wish to **Restrict login to trusted hosts**, toggle the option, and fill in the IP addresses of the hosts.
9. If you wish to **Restrict admin to guest account provisioning only**, toggle the option, and select the **Guest Group**.
10. Select **OK**.

Policy & Objects

The Policy & Object section is for the configuration of the different policies used to manage the traffic going through the appliance and the objects that are used to identify attributes of the traffic or components used by the traffic. The headings that cover these areas of configurations are:

- **IPv4 Policy**
- **IPv6 Policy**
- **NAT64 Policy**
- **NAT46 Policy**
- **Addresses**
- **IP Pools**

IPv4 Policy

To configure a IPv4 policy in the GUI

1. Go to **Policy & Objects > IPv4 Policy**

The right side window will display a table of the existing IPv4 Policies.

- To edit an existing policy, double click on the policy you wish to edit
- To create a new policy, select the **Create New** icon in the top left side of the right window.

2. Make sure the policy has a name in the **Name** field

By default, a policy is required to have a name, but it is possible to toggle this requirement on or off in the CLI.

3. Set the **Incoming Interface** parameter by selecting the field with the "**+**" next to the field label. Selecting the field will slide out a window from the right where you can select from the available interfaces. You can select one or more specific interfaces or you can select the **any** option. Choosing the any option will remove any other interfaces.
4. Set the **Outgoing Interface** parameter by selecting the field with the "**+**" next to the field label. (Same rules apply as with the above step.)
5. Set the **Source** parameter by selecting the field with the "**+**" next to the field label. The source in this case is either the source address, source user or source device of the initiating traffic. When the field is selected a window will slide out from the right. Tabs indicating **Address**, **User** or **Device** options are there to help categorize the options along with the option to search. In order to be able to select one of these options it needs to be configured as a firewall object before hand. The "**+**" icon next to the **Search** field is a shortcut for creating a new firewall object based on the tab that is currently selected. For the **Address** and **Device** tabs, single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
6. Set the **Destination Address** parameter by selecting the field with the "**+**" next to the field label. This field is similar to the **Source** field but address objects are the only available options to select. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
7. Set the **Schedule** parameter by using the drop down menu to select a preconfigured schedule. The "**+**" icon next to the **Search** field is a shortcut for creating a new schedule object.
8. Set the **Service** parameter by selecting the field with the "**+**" next to the field label. (Same mechanics for selection apply as with the other similar fields in this window.) Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option.
9. Set the **Action** parameter. Select one of the following options for the action:
   - **ACCEPT** - lets the traffic through to the next phase of analysis
   - **DENY** - drops the session

   Because the choice of **Action** determines the settings and options below this parameter in the window the rest of the step are associated with a specific **Action**.

Settings if the ACCEPT action is selected.

Network Options

10. Enable/Disable **NAT** by toggling the slider button.(gray means it is disabled)
11. Enable/Disable **Endpoint Independent Mapping** by toggling the slider button (gray means it is disabled).
12. Enable/Disable **Endpoint Independent Filtering** by toggling the slider button (gray means it is disabled).
13. Enable/Disable **Harware Session Statistics** by toggling the slider button (gray means it is disabled).
14. Select the **Hardware Logging Mode**:
    - **None**
    - **Per-Session**
    - **Per-Mapping**

15. Set the **IP Pool Configuration** by selection one of the options of:
    - **Use Outgoing Interface Address**
    - **Use Dynamic IP Pool**

    If the **Use Dynamic IP Pool** option is selected, an additional field will appear with the **+** icon. Selecting this field will slide out a window from the right where a preexisting IP Pool can be chosen. One or more IP Pools can be chosen and the "**+**" icon next to the **Search** field is a shortcut for creating a new IP Pool.

Settings if the DENY action is selected

10. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
11. Toggle whether or not to **Enable this policy**.The default is enabled.
12. Select the **OK** button to save the policy.

IPv6 Policy

To configure a IPv4 policy in the GUI

1. Go to **Policy & Objects > IPv4 Policy**

   The right side window will display a table of the existing IPv4 Policies.
   - To edit an existing policy, double click on the policy you wish to edit
   - To create a new policy, select the **Create New** icon in the top left side of the right window.
2. Make sure the policy has a name in the **Name** field

By default, a policy is required to have a name, but it is possible to toggle this requirement on or off in the CLI.

3. Set the **Incoming Interface** parameter by selecting the field with the "**+**" next to the field label. Selecting the field will slide out a window from the right where you can select from the available interfaces. You can select one or more specific interfaces or you can select the **any** option. Choosing the any option will remove any other interfaces.
4. Set the **Outgoing Interface** parameter by selecting the field with the "**+**" next to the field label. (Same rules apply as with the above step.)
5. Set the **Source** parameter by selecting the field with the "**+**" next to the field label. The source in this case is either the source address, source user or source device of the initiating traffic. When the field is selected a window will slide out from the right. Tabs indicating **Address**, **User** or **Device** options are there to help categorize the options along with the option to search. In order to be able to select one of these options it needs to be configured as a firewall object before hand. The "**+**" icon next to the **Search** field is a shortcut for creating a new firewall object based on the tab that is currently selected. For the **Address** and **Device** tabs, single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
6. Set the **Destination Address** parameter by selecting the field with the "**+**" next to the field label. This field is similar to the **Source** field but address objects are the only available options to select. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
7. Set the **Schedule** parameter by using the drop down menu to select a preconfigured schedule. The "**+**" icon next to the **Search** field is a shortcut for creating a new schedule object.

8. Set the **Service** parameter by selecting the field with the "**+**" next to the field label. (Same mechanics for selection apply as with the other similar fields in this window.) Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option.
9. Set the **Action** parameter. Select one of the following options for the action:
   - **ACCEPT** - lets the traffic through to the next phase of analysis
   - **DENY** - drops the session

   Because the choice of **Action** determines the settings and options below this parameter in the window the rest of the step are associated with a specific **Action**.

Settings if the ACCEPT action is selected.

Network Options

10. Enable/Disable **NAT** by toggling the slider button.(gray means it is disabled)
11. Set the **IP Pool Configuration** by selection one of the options of:
    - **Use Outgoing Interface Address**
    - **Use Dynamic IP Pool**

    If the **Use Dynamic IP Pool** option is selected, an additional field will appear with the **+** icon. Selecting this field will slide out a window from the right where a preexisting IP Pool can be chosen. One or more IP Pools can be chosen and the "**+**" icon next to the **Search** field is a shortcut for creating a new IP Pool.

Logging Options

12. Enable/Disable **Log Allowed Traffic** by toggling the slider button.(gray means it is disabled)

If Logging is enabled,you will have to chose between logging:

- **Security Events**
- **All Sessions**

13. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
14. Toggle whether or not to **Enable this policy**.The default is enabled.
15. Select the **OK** button to save the policy.

Settings if the DENY action is selected

10. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
11. Toggle whether or not to **Enable this policy**.The default is enabled.
12. Select the **OK** button to save the policy.

NAT64 Policy

1. Goto **Policy & Objects > NAT64 Policy**

   The right side window will display a table of the existing NAT64 Policies. In the menu bar at the top of the menu there is also a toggle to enable/disable **NAT64 Forwarding**.

   - To edit an existing policy, double click on the policy you wish to edit
   - To create a new policy, select the **Create New** icon in the top left side of the right window.
2. Set the **Incoming Interface** parameter by selecting the field with the "**+**" next to the field label. Selecting the field will slide out a window from the right where you can select from the available interfaces. You can select one or more specific interfaces or you can select the **any** option. Choosing the any option will remove any other

interfaces.

3. Set the **Outgoing Interface** parameter by selecting the field with the "**+**" next to the field label. (Same rules apply as with the above step.)
4. Set the **Source** parameter by selecting the field with the "**+**" next to the field label. The source in this case is either the source address, source user or source device of the initiating traffic. When the field is selected a window will slide out from the right. Tabs indcating **Address**, **User** or **Device** options are there to help categorize the options along with the option to search. In order to be able to select one of these options it needs to be configured as a firewall object before hand. The "**+**" icon next to the **Search** field is a shortcut for creating a new firewall object based on the tab that is currently selected. For the **Address** and **Device** tabs, single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
5. Set the **Destination Address** parameter by selecting the field with the "**+**" next to the field label. This field is similar to the **Source** field but address objects are the only available options to select. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
6. Set the **Schedule** parameter by using the drop down menu to select a preconfigured schedule. The "**+**" icon next to the **Search** field is a shortcut for creating a new schedule object.
7. Set the **Service** parameter by selecting the field with the "**+**" next to the field label. (Same mechanics for selection apply as with the other similar fields in this window.) Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option.
8. Set the **Action** parameter. Select one of the following options for the action:
   - **ACCEPT** - lets the traffic through to the next phase of analysis
   - **DENY** - drops the session

Settings if the ACCEPT action is selected.

Network Options

10. Skip the NAT setting.It is not configurable. This type of policy is intended only for traffic that is being NATed from IPv6 to IPv4, because without NATing the traffic couldn't reach its destination, so disabling NAT would be pointless.
11. Set the **IP Pool Configuration** by selection one of the options of:
    - **Use Outgoing Interface Address**
    - **Use Dynamic IP Pool**

    If the **Use Dynamic IP Pool** option is selected, an additional field will appear with the **+** icon. Selecting this field will slide out a window from the right where a preexisting IP Pool can be chosen. One or more IP Pools can be chosen and the "**+**" icon next to the **Search** field is a shortcut for creating a new IP Pool.

14. Enable/disable the **Log Allowed Traffic** parameter by toggling the slider button (gray means it is disabled).

    If the **Log Allowed Traffic** setting is enabled, choose whether to log just **Security Events** or **All Sessions** and determine whether or not to keep a record of the packets by toggling the **Capture Packets** setting on or off.

15. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
16. Toggle whether or not to **Enable this policy**.The default is enabled.
17. Select the **OK** button to save the policy.

Settings if the DENY action is selected

Enable the Log Violation Traffic setting by toggling the slider button.

10. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).

11. Toggle whether or not to **Enable this policy**.The default is enabled.
12. Select the **OK** button to save the policy.

NAT46 Policy

1. Go to **Policy & Objects > NAT46 Policy**

   The right side window will display a table of the existing NAT46 Policies.

   - To edit an existing policy, double click on the policy you wish to edit
   - To create a new policy, select the **Create New** icon in the top left side of the right window.
2. Set the **Incoming Interface** parameter by selecting the field with the "**+**" next to the field label. Selecting the field will slide out a window from the right where you can select from the available interfaces. You can select one or more specific interfaces or you can select the **any** option. Choosing the any option will remove any other interfaces.
3. Set the **Outgoing Interface** parameter by selecting the field with the "**+**" next to the field label. (Same rules apply as with the above step.)
4. Set the **Source** parameter by selecting the field with the "**+**" next to the field label. The source in this case is either the source address, source user or source device of the initiating traffic. When the field is selected a window will slide out from the right. Tabs indicating **Address**, **User** or **Device** options are there to help categorize the options along with the option to search. In order to be able to select one of these options it needs to be configured as a firewall object before hand. The "**+**" icon next to the **Search** field is a shortcut for creating a new firewall object based on the tab that is currently selected. For the **Address** and **Device** tabs, single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
5. Set the **Destination Address** parameter by selecting the field with the "**+**" next to the field label. This field is similar to the **Source** field but address objects are the only available options to select. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
6. Set the **Schedule** parameter by using the drop down menu to select a preconfigured schedule. The "**+**" icon next to the **Search** field is a shortcut for creating a new schedule object.
7. Set the **Service** parameter by selecting the field with the "**+**" next to the field label. (Same mechanics for selection apply as with the other similar fields in this window.) Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option.
8. Set the **Action** parameter. Select one of the following options for the action:
   - **ACCEPT** - lets the traffic through to the next phase of analysis
   - **DENY** - drops the session

Settings if the ACCEPT action is selected.

Network Options

10. Skip the NAT setting. This type of policy is intended only for traffic that is being NATed from IPv4 to IPv6, because without NATing the traffic couldn't reach its destination, so disabling NAT would be pointless.
11. Skip the **Use Outgoing Interface Address**. This is on by default and cannot be changed
12. Set the **Log Allowed Traffic** parameter by toggling the slider button (gray means it is disabled).

    If the **Log Allowed Traffic** setting is enabled, choose whether to log just **Security Events** or **All Sessions** and determine whether or not to keep a record of the packets by toggling the **Capture Packets** setting on or off.
13. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
14. Toggle whether or not to **Enable this policy**.The default is enabled.

15. Select the **OK** button to save the policy.

Settings if the DENY action is selected

Enable the Log Violation Traffic setting by toggling the slider button.

10. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
11. Toggle whether or not to **Enable this policy**.The default is enabled.
12. Select the **OK** button to save the policy.

Addresses

Addresses define sources and destinations of network traffic and are used when creating policies. When properly set up these objects can be used with great flexibility to make the configuration of policies simpler and more intuitive. The FortiCarrier compares the IP addresses contained in packet headers with a security policy's source and destination addresses to determine if the security policy matches the traffic.

The address categories and the types within those categories on the FortiCarrier can include:

- IPv4 addresses
  - IP address and Netmask
  - IP address range
  - Geography based address
  - Fully Qualified Domain Name (FQDN) address
  - Wildcard FQDN
  - IPv4 Address Group
- IPv6 addresses
  - Subnets
  - IP range
  - IPv6 Address Group

IP address and Netmask

The subnet type of address is expressed using a host address and a subnet mask. From a strictly mathematical stand point this is the most flexible of the types because the address can refer to as little one individual address or as many as all of the available addresses.

It is usually used when referring to your own internal addresses because you know what they are and they are usually administered in groups that are nicely differentiated along the lines of the old A, B, and C classes of IPv4 addresses. They are also addresses that are not likely to change with the changing of Internet Service Providers (ISP).

When representing hosts by an IP address with a netmask, the IP address can represent one or more hosts. For example, a firewall address can be:

- A single host such as a single computer with the address 192.45.46.45
- A range of hosts such as all of the hosts on the subnet 192.45.46.1 to 192.45.46.255
- All hosts, represented by 0.0.0.0 which matches any IP address

The netmask corresponds to the subnet class of the address being added, and can be represented in either dotted decimal or CIDR format. The FortiGate unit automatically converts CIDR formatted netmasks to dotted decimal format. Example formats:

- Netmask for a class A subnet of 16,777,214 usable addresses: 255.0.0.0, or /8
- Netmask for a class B subnet of 65,534 usable addresses: 255.255.0.0, or /16
- Netmask for a class C subnet of 254 usable addresses: 255.255.255.0, or /24
- Netmask for subnetted class C of 126 usable addresses: 255.255.255.128, or /25
- Netmask for subnetted class C of 62 usable addresses: 255.255.255.128, or /26
- Netmask for subnetted class C of 30 usable addresses: 255.255.255.128, or /27
- Netmask for subnetted class C of 14 usable addresses: 255.255.255.128, or /28
- Netmask for subnetted class C of 6 usable addresses: 255.255.255.128, or /29
- Netmask for subnetted class C of 2 usable addresses: 255.255.255.128, or /30
- Netmask for a single computer: 255.255.255.255, or /32
- Netmask used with 0.0.0.0 to include all IP addresses: 0.0.0.0, or /0

So for a single host or subnet the valid format of IP address and netmask could be either:

x.x.x.x/x.x.x.x, such as 192.168.1.0/255.255.255.0

or

x.x.x.x/x, such as 192.168.1.0/24

Static Route Configuration

A setting that is found in the IP/Netmask address type that is not found in the other address types is the enabling or disabling of **Static Route Configuration**. Enabling this feature includes the address in the listing of named addresses when setting up a static route.

To use in the GUI

1. Enable the **Static Route Configuration** in the address.
2. Go to **Network > Static Routes** and create a new route.
3. For a **Destination** type, choose **Named Address**.
4. Using the drop down menu, enter the name of the address object in the field just underneath the **Destination** type options.
5. Fill out the other information relevant to the route
6. Select the **OK** button

Creating an IP/Netmask address

1. Go to **Policy & Objects > Addresses.**
2. Select **Create New**. A drop down menu is displayed. Select **Address**.
3. In the **Category** field, chose **Address**.(This is for IPv4 addresses.)
4. Input a **Name** for the address object.
5. Use the **[Change]** link if you which to change the color of the icon when being viewed in the GUI.
6. In the **Type** field, select **IP/Netmask** from the drop down menu.
7. In the **Subnet/IP Range** field, enter the address and subnet mask according to the format x.x.x.x/x.x.x.x or the short hand format of x.x.x.x/x

8. In the **Interface** field, leave as the default **any** or select a specific interface from the drop down menu.
9. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
10. Select the desired on/off toggle setting for **Static Route Configuration**.
11. Input any additional information in the **Comments** field.
12. Press **OK**.

Example

Example of a Subnet address for a database server on the DMZ:

Field	Value
Category	**Address**
Name	DB_server_1
Type	**IP/Netmask**
Subnet/IP Range	United States
Interface	**any**
Show in Address List	[on]
Static Route Configuration	[off]
Comments	

IP Range

Where the Subnet address is good a representing a standardized group of addresses that are subnets the IP Range type of address can describe a group of addresses while being specific and granular. It does this by specifying a continuous set of IP addresses between one specific IP address and another. While it is most common that this range is with a subnet it is not a requirement. For instance, 192.168.1.0/24 and 192.168.2.0/24 would be 2 separate subnets but if you wanted to describe the top half of one and the bottom half of the other you could describe the range of 192.168.1.128-192.168.2.127. It's also a lot easier that trying to calculate the correct subnet mask.

The format would be:

x.x.x.x-x.x.x.x, such as 192.168.110.100-192.168.110.120

There is a notation that is commonly used and accepted by some devices that follows the format:

x.x.x.[x-x], such as 192.168.110.[100-120]

This format is not recognized in FortiCarrier as a valid IP Range.

Creating a IP Range address

1. Go to **Policy & Objects > Addresses.**
2. Select **Create New**. A drop down menu is displayed. Select **Address**
3. In the **Category** field, chose **Address**(IPv4 addresses) or **IPv6 Address**.

4. Input a **Name** for the address object.
5. In the **Type** field, select **IP Range** from the drop down menu.
6. In the **Subnet / IP Range** field, enter the range of addresses in the following format: x.x.x.x-x.x.x.x (no spaces)
7. In the **Interface** field, leave as the default **any** or select a specific interface from the drop down menu. (This setting is not available for IPv6 addresses)
8. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example

Example of a IP Range address for a group of computers set aside for guests on the company network.

Field	Value
Category	**Address** or **IPv6 Address**
Name	Guest_users
Type	**IP Range**
Subnet / IP Range	192.168.100.200-192.168.100.240
Interface	**Port1**
Show in Address List	**[on]**
Comments	Computers on the 1st floor used by guests for Internet access.

IP Range addresses can be configured for both IPv4 and IPv6 addresses. The only differences in creating an IPv6 IP Range address is that you would choose IPv6 Address for the Category and the syntax of the address in the Subnet/IP Range field would be in the format of 2001:0db8:0000:0002:0:0:0:20-2001:0db8:0000:0004:0:0:0:20

Geography

Geography addresses are those determined by country of origin.

This type of address is only available in the IPv4 address category.

Creating a Geography address

1. Go to **Policy & Objects > Addresses.**
2. Select **Create New**. A drop down menu is displayed. Select **Address**.
3. In the **Category** field, chose **Address**. (This is for IPv4 addresses.)
4. Input a **Name**for the address object.
5. In the **Type** field, select **Geography** from the drop down menu.
6. In the **Country** field, select a single country from the drop down menu.
7. In the **Interface** field, leave as the default **any** or select a specific interface from the drop down menu.
8. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example: Geography-based Address

Your company is US based and has information on its web site that may be considered information that is not allowed to be sent to embargoed countries. In an effort to help reduce the possibility of sensitive information going to those countries you have be asked to set up addresses for those countries so that they can be block in the firewall policies.

- One of the countries you have been asked to block is Cuba
- You have been asked to comment the addresses so that other administrators will know why they have been created

1. Go to **Policy & Objects> Objects > Addresses** and select **Create New > Address.**
2. Fill out the fields with the following information

Category	Address
Name	Cuba
Type	Geography
Country	Cuba
Interface	any
Visibility	<enable>
Comments	Embargoed

3. Select **OK.**

FQDN Addresses

By using Fully Qualified Domain Name (FQDN) addressing you can take advantage of the dynamic ability of DNS to keep up with address changes without having to manually change the addresses on the FortiCarrier. FQDN addresses are most often used with external web sites but they can be used for internal web sites as well if there is a trusted DNS server that can be accessed. FQDN addressing also comes in handy for large web sites that may use multiple addresses and load balancers for their web sites. The FortiCarrier automatically maintains a cached record of all the addresses resolved by the DNS for the FQDN addresses used.

For example, if you were doing this manually and you wanted to have a security policy that involved Google you could track down all of the IP addresses that they use across multiple countries. Using the FQDN address is simpler and more convenient.

When representing hosts by an FQDN, the domain name can also be a subdomain, such as mail.example.com.

Valid FQDN formats include:

- `<host_name>.<top_level_domain_name>` such as example.com
- `<host_name>.<second_level_domain_name>.<top_level_domain_name>`, such as mail.example.com

When creating FQDN entries it is important to remember that:

- Wildcards are not supported in FQDN address objects
- While there is a level of convention that would imply it, “www.example.com” is not necessarily the same address of “example.com”. they will each have their own records on the DNS server.

The FortiCarrier keeps track of the DNS TTLs so as the entries change on the DNS servers the IP address will effectively be updated for the FortiGate. As long as the FQDN address is used in a security policy, it stores the address in the DNS cache.

There is a possible security downside to using FQDN addresses. Using a fully qualified domain name in a security policy means that your policies are relying on the DNS server to be accurate and correct. DNS servers in the past were not seen as potential targets because the thinking was that there was little of value on them and therefore are often not as well protected as some other network resources. People are becoming more aware that the value of the DNS server is that in many ways it controls where users and computers go on the Internet. Should the DNS server be compromised, security policies requiring domain name resolution may no longer function properly.

Creating a Fully Qualified Domain Name address

1. Go to **Policy & Objects > Addresses.**
2. Select **Create New**. A drop down menu is displayed. Select **Address**.
3. In the **Category** field, chose**Address**. (This is for IPv4 addresses.)
4. Input a **Name** for the address object.
5. In the **Type** field, select **FQDN** from the drop down menu.
6. Input the domain name in the **FQDN** field.
7. In the **Interface** field, leave as the default **any** or select a specific interface from the drop down menu.
8. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example: FQDN address

You have to great a policy that will govern traffic that goes to a site that has a number of servers on the Internet. Depending on the traffic or the possibility that one of the servers is down network traffic can go to any one of those sites. The consistent factor is that they all use the same Fully Qualified Domain Name.

- The FQDN of the web site: example.com
- The number of ISP connections off of the FortiGate firewall: 2

Configuring the address in the GUI

1. Go to **Policy & Objects> Objects > Addresses** and select **Create New > Address.**
2. Fill out the fields with the following information:

Category	Address
Name	BigWebsite.com
Type	FQDN
FQDN	bigwebsite.com
Interface	any
Show in Address List	<enable>
Comments	<Input into this field is optional>

3. Select **OK.**

Wildcard FQDN

There are a number of companies that use secondary and even tertiary domain names or FQDNs for their websites. Wildcard FQDN addresses are to ease the administrative overhead in cases where this occurs. Sometimes its as simple as sites that still use www. as a prefix for their domain name. If you don't know whether or not the www is being used it's simpler to use a wildcard and include all of the possibilities whether it be example.com, www.example.com or even ftp.example.com.

Wildcard FQDN addresses do not resolve to a specific set of IP addresses in the same way that a normal FQDN addresss does. They are intended for use in SSL exemptions and should not be used as source or destination addresses in policies.

Creating a Fully Qualified Domain Name address

1. Go to **Policy & Objects > Addresses.**
2. Select **Create New**. A drop down menu is displayed. Select **Address**
3. In the **Category** field**,** chose **Address.** (This is for IPv4 addresses.)
4. Input a **Name** for the address object.
5. In the **Type** fUncategorizedield, select **Wildcard FQDN**from the drop down menu.
6. Input the domain name in the **Wildcard FQDN** field.
7. In the **Interface** field, leave as the default **any** or select a specific interface from the drop down menu.
8. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
9. Input any additional information in the **Comments** field.
10. Press **OK.**

Example

Example of a FQDN address for a remote FTP server used by Accounting team:

Field	Value
Category	**Address**
Name	Example.com_servers
Type	**Wildcard FQDN**
Wildcard FQDN	*.example.com
Interface	**any**
Show in Address List	[on]
Comments	Secondary and tertiary domain names for example.com

Subnet Addresses

The Subnet Address type is one that is only used in reference to IPv6 addresses.It represents an IPv6 address subnet. This means that the address will likely be a series of hexadecimal characters followed by a double colon, followed by a "/", and then a number less than 128 to indicate the size of the subnet. An example would be:

```
fd5e:3c59:35ce:f67e::/64
```

- The hexidecimal characters represent the IPv6 subnet address.
- The "::" indicates 0's from that point to the left. In an actual address for a computer, the hexadecimal characters that would take the place of these zeros would represent the device address on the subnet.
- /xx, in this case /64 represents the number of bits in the subnet.This will make a range that can potentially include 18,446,744,073,709,551,616 addresses. For those wanting to use English rather than math, that is 18 Quintillion.

Creating a Subnet address

1. Go to **Policy & Objects > Addresses.**
2. Select **Create New**. A drop down menu is displayed. Select **Address**
3. In the **Category** field, chose **IPv6 Address**.
4. Input a **Name** for the address object.
5. In the **Type** field, select **Subnet** from the drop down menu.
6. In the **Subnet / IP Range** field, enter the range of addresses in IPv6 format (no spaces)
7. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
8. Input any additional information in the **Comments** field.
9. Press **OK**.

Example

Example of a IP Range address for a group of computers set aside for guests on the company network.

Field	Value
Category	**IPv6 Address**
Name	IPv6_Guest_user_range
Type	**Subnet**
Subnet / IP Range	fd5e:3c59:35ce:f67e::/64
Show in Address List	**[on]**
Comments	

Address Groups

Address groups are designed for ease of use in the administration of the device. If you have a number of addresses or address ranges that will commonly be treated the same or require the same security policies, you can put them into address groups, rather than entering multiple individual addresses in each policy refers to them.

The use of groups is not required. If you have a number of different addresses you could add them individually to a policy and the FortiCarrierl will process them just as quickly and efficiently as if they were in a group, but the chances are that if you have used a group once you could need to use it again and depending on the number of addresses involved entering them individually for each policy can become tedious and the likelihood of an address being missed becomes greater. If you have a number of policies using that combination of addresses it is much easier to add or subtract addresses from the group than to try and remember all of the firewall policies that combination of addresses was used in. With the group, you only have to make the one edit and it is used by any firewall policy using that address group.

Because security policies require addresses with homogenous network interfaces, address groups should contain only addresses bound to the same network interface, or to Any.

For example, if address 1.1.1.1 is associated with port1, and address 2.2.2.2 is associated with port2, they cannot be in the same group. However, if 1.1.1.1 and 2.2.2.2 are configured with an interface of Any, they can be grouped, even if the addresses involve different networks.

There are 2 Categories of Address groups to choose from:

- IPv4 Group
- IPv6 Group

You cannot mix different categories of addresses within a group, so whether or not it makes sense from an administrative purpose to group certain addresses together, if some are IPv4 and some are IPv6, it cannot be done.

Creating an Address Group

1. Go to **Policy & Objects > Addresses.**
2. Select the down arrow next to **Create New**, select **Address Group.**
3. Choose the **Category**, either **IPv4 Group** or **IPv6 Group**, that is applicable to the proposed selection of addresses.

4. Input a **Group Name**for the address object.
5. Use the **[Change]** link to change the color of the icon in the GUI. There are 32 different color options.

Depending on which **Category** has been chosen the configurations will differ slightly

IPv4 Group

1. Select the "+" in the **Members** field. You can select members of the group from the window that slides out from the left of the screen. It is possible to select more than 1 entry. Select the “X” icon in the field to remove an entry.
2. Select the desired on/off toggle setting for **Show in Address List**.
3. Select the desired on/off toggle setting for **Static Route Configuration** .

IPv6 Group

1. Select the "+" in the **Members** field. You can select members of the group from the window that slides out from the left of the screen. It is possible to select more than 1 entry. Select the “X” icon in the field to remove an entry.
2. Select the desired on/off toggle setting for **Show in Address List**.

Irrespective of the Category the groups all have the same final configuration options:

- Input any additional information in the **Comments** field.
- Press **OK.**

IP Pools

IP Pools are a mechanism that allow sessions leaving the FortiCarrier to use NAT. An IP pool defines a single IP address or a range of IP addresses to be used as the source address for the duration of the session. These assigned addresses will be used instead of the IP address assigned to that FortiCarrier interface.

When using IP pools for NATing, there is a limitation that must be taken into account when configuring the pool. If the IP address(es) within the pool are different from the IP address(es) that are assigned to the interface communications based on those IP addresses will fail. For example if the IP addresses assigned to an interface are 172.16.100.1 -172.16.100.14, you cannot choose 10.11.12.50 - 10.11.12.59 for the IP pool

There are 3 types of IP Pools that can be configured on the FortiGate firewall:

- **Port Block Allocation** - this setting is used to allocate a block of port numbers for IP pool users. Two variables will also have to be set. The block size can be set from 64 to 4096 and as the name implies describes the number of ports in one block of port numbers. The number of blocks per user determines how many of these blocks will be assigned. This number can range from 1 to 128.
- **Single Port Allocation** -
- **Overload** - this is the default setting. Internal addresses other than the one designated in the policy can use this address for the purposes of NAT.

Be careful when calculating the values of the variables. The maximum number of ports that are available on an address is 65,536. If you chose the maximum value for both variables you will get a number far in excess of the available port numbers.

4096 x 128 = 524,288

Source IP address and IP pool address matching when using a range

When the source addresses are translated to an IP pool that is a range of addresses, one of the following three cases may occur:

Scenario 1:

The number of source addresses equals that of IP pool addresses

In this case, the FortiGate unit always matches the IP addressed one to one.

If you enable fixed port in such a case, the FortiGate unit preserves the original source port. This may cause conflicts if more than one security policy uses the same IP pool, or the same IP addresses are used in more than one IP pool.

Scenario 2:

The number of source addresses is more than that of IP pool addresses

In this case, the FortiGate unit translates IP addresses using a wrap-around mechanism. If you enable fixed port in such a case, the FortiGate unit preserves the original source port. But conflicts may occur since users may have different sessions using the same TCP 5 tuples.

Scenario 3:

The number of source addresses is fewer than that of IP pool addresses

In this case, some of the IP pool addresses are used and the rest of them are not be used.

ARP Replies

If a FortiCarrierl interface IP address overlaps with one or more IP pool address ranges, the interface responds to ARP requests for all of the IP addresses in the overlapping IP pools. For example, consider a FortiCarrier unit with the following IP addresses for the port1 and port2 interfaces:

- port1 IP address: 1.1.1.1/255.255.255.0 (range is 1.1.1.0-1.1.1.255)
- port2 IP address: 2.2.2.2/255.255.255.0 (range is 2.2.2.0-2.2.2.255)

And the following IP pools:

- IP_pool_1: 1.1.1.10-1.1.1.20
- IP_pool_2: 2.2.2.10-2.2.2.20
- IP_pool_3: 2.2.2.30-2.2.2.40

The port1 interface overlap IP range with IP_pool_1 is:

(1.1.1.0-1.1.1.255) and (1.1.1.10-1.1.1.20) = 1.1.1.10-1.1.1.20

The port2 interface overlap IP range with IP_pool_2 is:

(2.2.2.0-2.2.2.255) & (2.2.2.10-2.2.2.20) = 2.2.2.10-2.2.2.20

The port2 interface overlap IP range with IP_pool_3 is:

(2.2.2.0-2.2.2.255) & (2.2.2.30-2.2.2.40) = 2.2.2.30-2.2.2.40

And the result is:

- The port1 interface answers ARP requests for 1.1.1.10-1.1.1.20
- The port2 interface answers ARP requests for 2.2.2.10-2.2.2.20 and for 2.2.2.30-2.2.2.40

Select Enable NAT in a security policy and then select Dynamic IP Pool. Select an IP pool to translate the source address of packets leaving the FortiCarrier unit to an address randomly selected from the IP pool. Whether or not the external address of an IP Pool will respond to an ARP request can be disabled. You might want to disable the ability to responded to ARP requests so that these address cannot be used as a way into your network or show up on a port scan.

IP pools and zones

Because IP pools are associated with individual interfaces IP pools cannot be set up for a zone. IP pools are connected to individual interfaces.

Fixed Port

Some network configurations do not operate correctly if a NAT policy translates the source port of packets used by the connection. NAT translates source ports to keep track of connections for a particular service.

However, enabling the use of a fixed port means that only one connection can be supported through the firewall for this service. To be able to support multiple connections, add an IP pool, and then select Dynamic IP pool in the policy. The firewall randomly selects an IP address from the IP pool and assigns it to each connection. In this case, the number of connections that the firewall can support is limited by the number of IP addresses in the IP pool.

Configuring IP Pools

To configure a new Dynamic IP Pool:

1. Go to **Policy & Objects > IP Pools**
2. Select the Create New button
3. Select from the **IP Pool Type** options, **IPv4 Pool** or **IPv6 Pool**. The available fields will defer between the two IP Pool types.

IPv4 Pool

4. Input a string value for the **Name** field for the pool
5. Input detailed descriptions for the IP Pool in the **Comments** field.
6. Select the Type from the options:

Port Block Allocation

The following options are available to the Port Block Allocation type:

- Set the **External IP Range** parameter by selecting the field with the "**+**" next to the field label. A window will pull out from the side of the screen displaying the available options. Single or multiple options can be selected.
- Set the **Block Size**. The value here is based on power of 2 and up to 128.
- Set the Blocks per User . Use integer values.
- Enable or disable **Fixed Allocation.**
  - If used, a Client IP Range needs to be chosen by selecting the field with the "**+**" next to the field label.

Single Port Allocation

The following options are available to the Single Port Allocation type:

- Set the **External IP Range** parameter by selecting the field with the "**+**" next to the field label. A window will pull out from the side of the screen displaying the available options. Single or multiple options can be selected.
- Enable or disable **Fixed Allocation.**
  - If used, a Client IP Range needs to be chosen by selecting the field with the "**+**" next to the field label.

Overload

The following options are available to the Overload type:

- Set the **External IP Range** parameter by selecting the field with the "**+**" next to the field label. A window will pull out from the side of the screen displaying the available options. Single or multiple options can be selected.

The following settings are available to all three types:

7. Enable/disble **ARP Reply** using the check box.
8. Enter an integer value for the **TCP Session Quota**
9. Enter an integer value for the **UDP Session Quota**
10. Select the **OK** button

IPv6 Pool

4. Input a string value for the **Name** field for the pool
5. Input detailed descriptions for the IP Pool in the **Comments** field.
6. Describe the **External Range** but entering the start and end IP addresses in the two fields.
7. Select the **OK** button

Log & Reports

The Log & Report section contains the following sections for configuring their related options:

- **System Events**
- **Log Settings**
- **Alert Email**

Monitor

The Monitor section does not have sections for configuring settings of the appliance but does contain the following monitors to observe the activities of the appliance:

- **Routing Monitor**
- **DHCP Monitor**

- **WAN Link Monitor**

Configuration

Initial access

Basic configuration

The basic FortiCarrier configuration involves creating policies to allow sessions from internal networks to reach the Internet or other external networks. You can create these policies from the GUI by going to **Policy & Objects > IPv4 Policy**. In addition to IPv4 policies you can create IPv6, NAT 64 and NAT 46 policies.

Just like any firewall policy, FortiCarrier policies define incoming and outgoing interfaces, source and destination addresses, schedules and services that are allowed.

FortiCarrier policies also include advanced source NAT options that are enabled by selecting **NAT** and configuring the various NAT options. These advanced options include endpoint independent mapping and filtering, hardware-based session statistics and hardware-based logging.

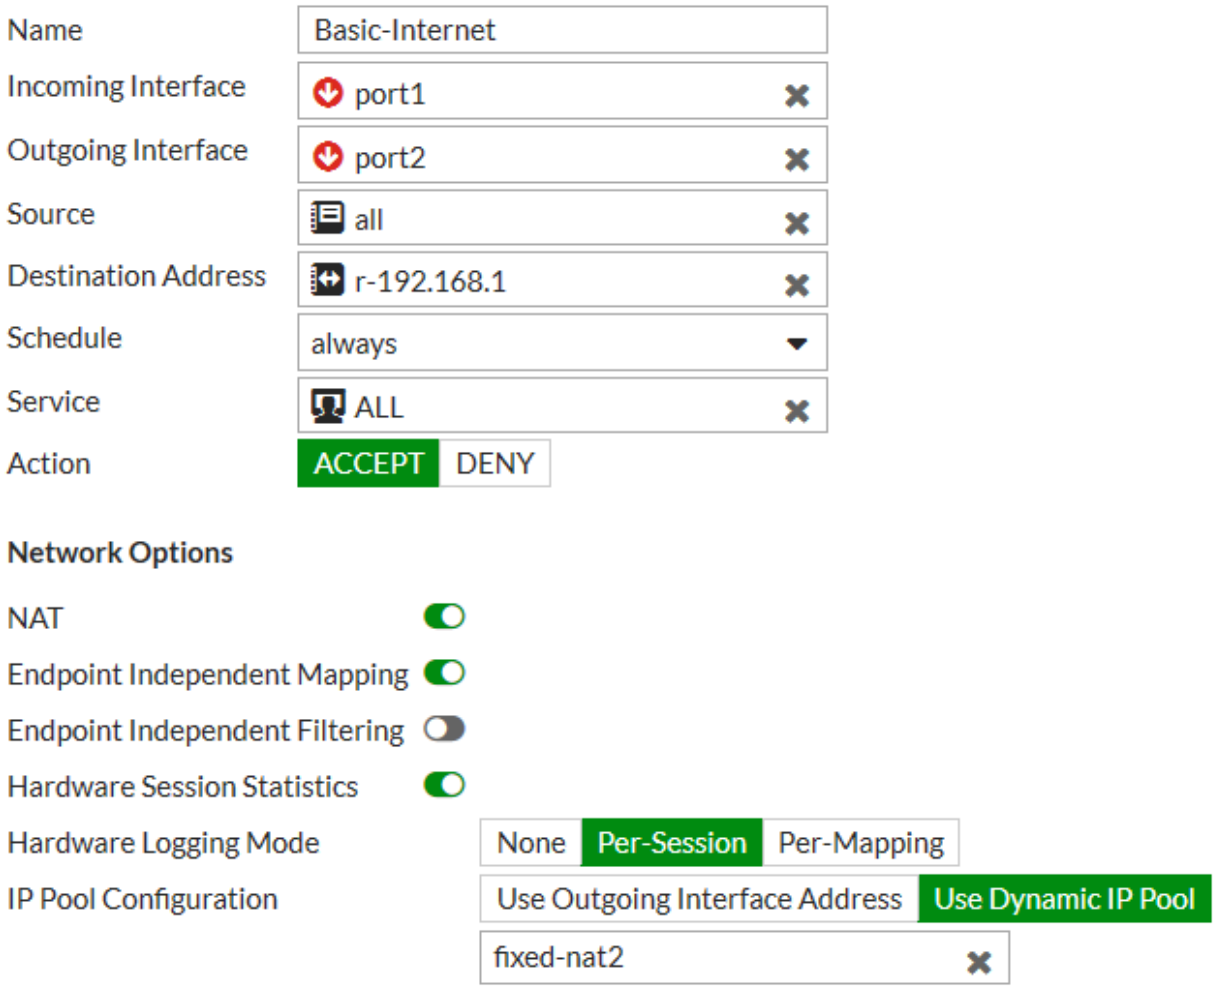

In addition to the NAT options in the policy, another important way to apply FortiCarrier NAT features is by configuring IP Pools. Go to **Policy & Objects > IP Pools** to configure IP pools. Using IP Pools you can control and apply port block allocation, single port allocation, and overload settings.

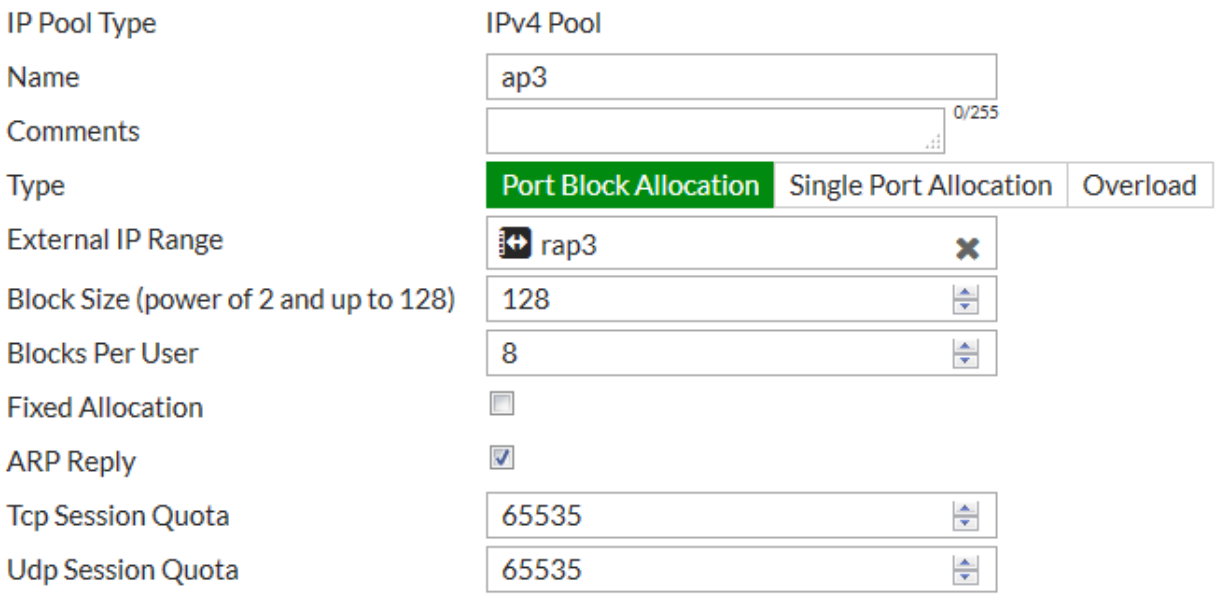

Finally, FortiCarrier also supports sessions that do not require NAT. In this case, you still go to **Policy & Objects > IPv4 Policy**, configure a policy and disable NAT.

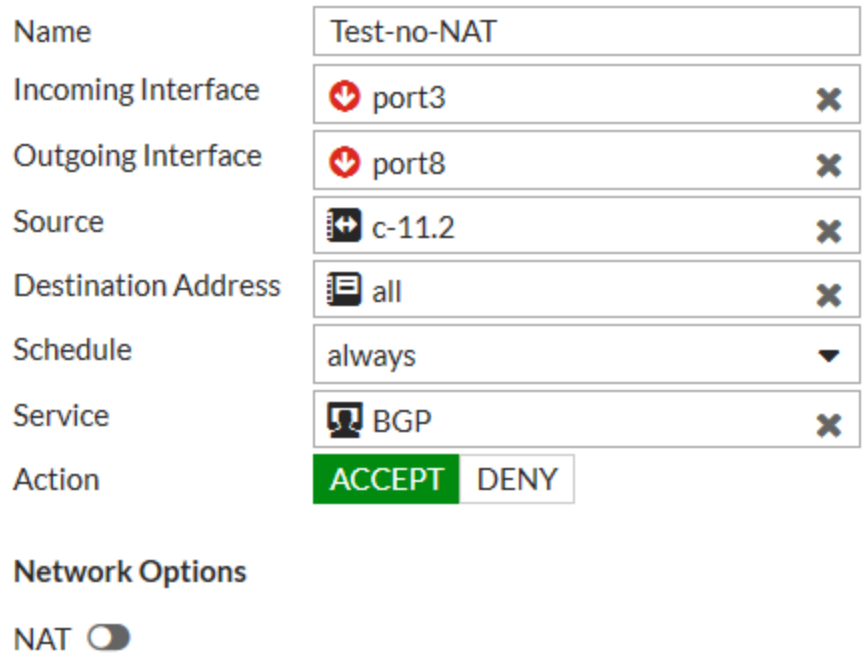

Interface configuration

As in FortiOS, an interface can be configured as a physical interface, a VLAN interface, a link-aggregation interface (LAG), or a VLAN interface on top of a LAG. There is a new flag to indicate whether the interface is used in NAT device at "client" side(or private side) or at "internet" side(or public side). The flag must be set when configuring an interface, and the default value is "private". For policies using FortiASIC-FP1 accelerated session

setup, the source interfaces must all be marked as "private", and the destination interfaces must all be marked as "public".

For policies using FortiASIC-FP1 to perform "offload session" only, the restriction does not apply.

Syntax:

```
set cgn-nat-role private
set cgn-nat-role public
```

IP-pool configuration.

Type.

IP pool configuration supports 3 types of IP pool.

1. **port-block-allocation**: default type, support FortiASIC-FP1 accelerated session setup, resource are assigned in the unit of a block of ports. This type has the advantage of higher CPS throughput and reduced logging traffic.
   Additional 2 parameters must be set: "block-size" and "num-blocks-per-user". Note, "block-size"x"num-blocks-per-suer" is the effective resource quota of the client of this pool.
2. **single-port-allocation**: support FortiASIC-FP1 accelerated session setup, resource are assigned to client in the unit of single port, from the entire range of available ports. This type has the advantage of higher efficiency in resource utilization.
3. **overload**: this type has the same meaning as in FortiOS, it does not support FortiASIC-FP1 accelerated session setup.

When this type of pool is attached to a FOS policy, the corresponding hardware policy will forward the matching packets host, and the host will allocate resource and install FortiASIC-FP1 sessions as in a NP6 system.

For this type, host does not need to keep track of resource for its simplicity. The overload mode is conflicting with EIM/EIF properties.

Syntax:

```
set type port-block-allocation
set type single-port-allocation
set type overload
```

IP range.

One IP pool can support multiple IP ranges (post-LPA release, for LPA only one range). There is a limit of 64K on the number of IP's in single pool, and a limit of 256K on the total number of IP's of all pools. As in FortiOS, IP pool ranges can not overlap.

Syntax:

For LPA,

```
set startip
set endip
```

For post-LPA,

```
set iprangename
```

Port range.

The port range is used for both TCP and UDP protocols. There is only one port range.

- If left unspecified, system default port range is used 5117-65535
- If specified, the size of the range must be larger than 4096 ports.

Syntax:

```
set port-start
set port-end
```

Block-size and number of block per user.

When port-block-allocation type is selected, the block-size and the number of block per user must be configured.

- The default value is: block-size = 128, num-blocks-per-user = 8.
- The block-size can be selected from 16, 32, 64 and 128.
- The num-blocks-per-user is in the range of 1 to 8.
- The product of block-size and num-blocks-per-user is the effective resource quota per user.
- There is no resource quota for SPA pool.

Syntax:

```
set block-size
set num-blocks-per-user
```

Session quota.

A client is also limited by the number of concurrent sessions, regardless of the pool type. The limit is specified separately for TCP and UDP.

Syntax:

```
set tcp-session-quota
set udp-session-quota
```

Fixed-allocation.

FortiASIC-FP1 supports statical mapping of client to available public IP's. This mode can be used together with PBA or SPA type. For SPA type, the IP assignment to a client is fixed, the port is allocated dynamically as the port resource are shared among all clients mapped to this public IP. For PBA type, not only the IP but also the port-blocks are statically allocated for the client. When this mode is set, a range of client IP addresses need to be provided. The assignment of clients to public resource is of round-robin nature and can be predicted given the client IP. There is a limit of the number of clients in each pool at 1,000,000.

Syntax:

```
set fixed-allocation
set client-address-start
set client-address-end
```

Static route configuration

LPM route lookup is supported by hardware by using TCAM rules. Simple configuration of static routes are supported, the "dst"/"gateway"/"device" can be configured to form a static route. Options such as "dynamic-gateway" and "internet-service" are not supported.

Policy route configuration

Policy routing is supported by hardware by using TCAM rules. The "key" part of the policy route supports "input-device", "src", "dst" fields. The next hop part of the configurations supports "gateway" and "output-device". Only single path route is supported.

Logging configuration

When enabled in policy, FortiASIC-FP1 sends logging messages to configured servers. To reduce burden of logging servers, there can be up to 16 syslog servers. FortiASIC-FP1 calculates hashing value over session tuple information to select logging server. There can be only one netflow-v9 logging server. Either syslog or netflow-v9 logging mode can be selected but not both. Also note, the hardware session logging configurations are separated from FortiOS syslog configurations, i.e., host still uses its own logging mechanism to do event and host session logging.

To configure hardware session logging, the steps are:

1. Define logging servers. Logging server is defined as server IP address (FQDN not supported) and destination port.
2. Select logging mode: netflow-v9 or syslogd mode can be selected.

syntax:

```
config log hw_logging servers
config log hw_logging settings
```

ALG configuration

Application-layer-gateway is the necessary mechanism to help some applicatons pass NAT device properly. FortiCarrier supports all ALG types supported by FortiOS. And the configurations are also the same as in FortiOS by configuring "sys session-helper". ALG is defined as a 2-tuple of destination port and protocol. Once configured, the traffic matching the 2-tuple will be forwarded to host.

ALG traffic is forwarded to host, the application-sse-timeout option is used only by hardware installed sessions. So the two use cases of 2-tuple are not compatible, hence can't be used at the same time, i.e., a 2-tuple is either for ALG use or for application-sse-timeout use.

HA configuration

FortiCarrier runs HA mode in the same way as FortiOS. In HA mode, host is responsible to synchronize configurations between primary and backup systems, and to synchronize host sessions when session synchronization is enabled. The sessions setup by FortiASIC-FP1 can be synchronized over a dedicated connection between two systems. This session synchronization is done by hardware.

Without hardware session synchronization, in a fail over, backup can still take over the traffic, and FortiASIC-FP1 can setup sessions for new sessions, but existing sessions from the primary are lost. Enabling hardware session synchronization can prevent existing session loss to achieve better high availability.

One of the physical front-panel ports can be selected as the dedicated connection for hardware session synchronization. This port once selected for this purpose cannot be used for other task, e.g., heartbeat, or normal traffic ports, or in LAG.

An new field is added in existing HA configuration. Check is performed to make sure only unused physical port may be selected for this purpose.

syntax:

```
set hw-session-sync-dev
```

Advanced configurations

Port preservation

Port-preservation feature is to allow certain source ports of the client be preserved while NATing. This behavior is necessary for some applications to function properly. Port-preservation is an option in IP pool configuration. Once a valid port-preservation profile is attached to the IP pool, traffic hitting policies using this pool will have the port-preservation property as defined by the port-preservation profile.

The following restrictions apply:

- There can be maximum 15(0~14) port-preservation profiles (hardware limit).
- Each profile can have up to 32 ports to be preserved (software limit).
- The preserved ports apply to both TCP and UDP of the pool (software limit).

The steps to configure port-preservation are as below:

1. Configure preserved ports.
2. Configure port-preservation-profile by adding preserved ports.
3. Attach port-preservation-profile to IP pool.

syntax:

```
config firewall cgn_hw_cfg preserved-port
config firewall cgn_hw_cfg port-preservation-profile
```

In IP pool configuration:

```
set profile-id
```

Virtual IP

Virtual IP (VIP) is a FortiOS supported feature to create a static one to one mapping between a client and a public IP. In FortiCarrier FortiASIC-FP1 also supports this feature by using ingress policy. The configuration steps are the same as in FortiOS.

1. Create "firewall VIP" object to define the mapping between client IP(mappedip) and public IP (extip).
2. Create policy to point the "dstaddr" to the VIP object.

Syntax:

Create VIP:

```
set extip
set extintf
set mappedip
```

Default SSE timeout

SSE timeouts are hardware tear down sessions with timeout (age-out) events based on the state of the session. These timeouts are set using the following command, shown with the default times in seconds:

```
config system cgn_hw_cfg sse-tmo
   set tcp-ttl 3600
   set tcp-halfclose 120
   set tcp-halfopen 10
   set tcp-timewait 10
   set tcp-rst 60
   set udp-ttl 30
end
```

Where:

`tcp-ttl` is the TCP timeout, default value 3600 seconds.

`tcp-halfclose` is the timer for half closed TCP sessions, default value 120 seconds. This timer closes TCP sessions that are in the FIN or RST state.

`tcp-halfopen` is the timer for half open TCP sessions, default value 10 seconds. This timer closes sessions are in the SYN sent or SYN wait state.

`tcp-timewait` is the timer for TCP time wait sessions, default value 10 seconds. This timer closes sessions that are in the FIN state.

`tcp-rst` is the timer for TCP reset sessions, default value 60 seconds.

`udp-ttl` is the timer for UDP sessions, default value 30 seconds.

The above timeout times only apply to hardware sessions. For host (or CPU) sessions, the timeouts are controlled by the host session table and you can use the following command to adjust timers for host sessions (shown with default values):

```
config system global
   set tcp-halfclose-timer 30
   set tcp-halfopen-timer 30
   set udp-idle-timer 60
end
```

Application SSE timeout

Hardware also supports customized timeout profile for certain "applications". An application is defined as a 2-tuple of destination port and protocol (TCP or UDP). When a new connection's destination port and protocol matches such 2-tuple, hardware will use the customized timeout profile instead of the default one.

The useable TCP timeout profile is about 11, and useable UDP timeout profile is about 88. The reason of the difference is that TCP profile uses more table entries than UDP profile.

Below is the command to set the application SSE timeout.

Syntax:

```
config firewall cgn_hw_cfg app-tmo
```

Example configuration 1

Configuration examples with comments in line(#):

Configure firewall policy

```
config firewall policy
   edit 1
      set srcintf "port2" #can be multiple interface just as in normal FortiGate
      set dstintf "port1"
      set srcaddr "all" #can be defined address object as normal FortiGate
      set dstaddr "all"
      set action accept
      set schedule "always"
      set service "ALL"
      set nat enable #NAT must be enabled.
      set endpoint-independent-mapping enable #option: EIM
      set endpoint-independent-filtering enable #option: EIF
      set hw-logging-mode session #option: enable per-session or per-mapping hardware
          logging.
      set ippool enable #ippool must be used.
      set poolname "pool1" #Can use multiple pools, but all pools of one policy must
      #have the same properites, e.g., pool type, port-block size etc.
      next
   end
```

Configure the IP Pool

```
config firewall ippool
   edit "pool1"
      set type port-block-allocation #select port-block-allocation or single-port-
          allocation type.
      set startip 101.1.1.4 #range of IP
      set endip 101.1.1.255
      set block-size 64 #option to change block-size(PBA only)
      set num-blocks-per-user 4 #option to set number of blocks per user(PBA only)
      set tcp-session-quota 65535 #session quota
      set udp-session-quota 65535
      set port-start 12000 #range of port(applied to both TCP and UDP)
      set port-end 50000
      next
```

```
   end
```

Routing setup is the same as in FortiGate.

```
config router static
   edit 2
      set dst 6.6.1.0 255.255.255.0
      set gateway 7.7.11.1
      set device "port1"
      next
   end
```

Logging setup is the same as in FortiGate.

```
config log syslogd setting
   set status enable
   set server "6.6.1.3"
   end
```

High Availability (HA)

FortiCarrier uses the FortiGate Clustering Protocol (FGCP) for Active-Passive HA with session pickup and HA reserved management interfaces. A FortiCarrier HA cluster must include two, and only two, FortiCarrier units. The FortiCarrier units must be the same model with the same hardware configuration.

When session pickup is disabled, HA configuration is the same as FortiOS, and any interfaces can be used for the HA heartbeat.

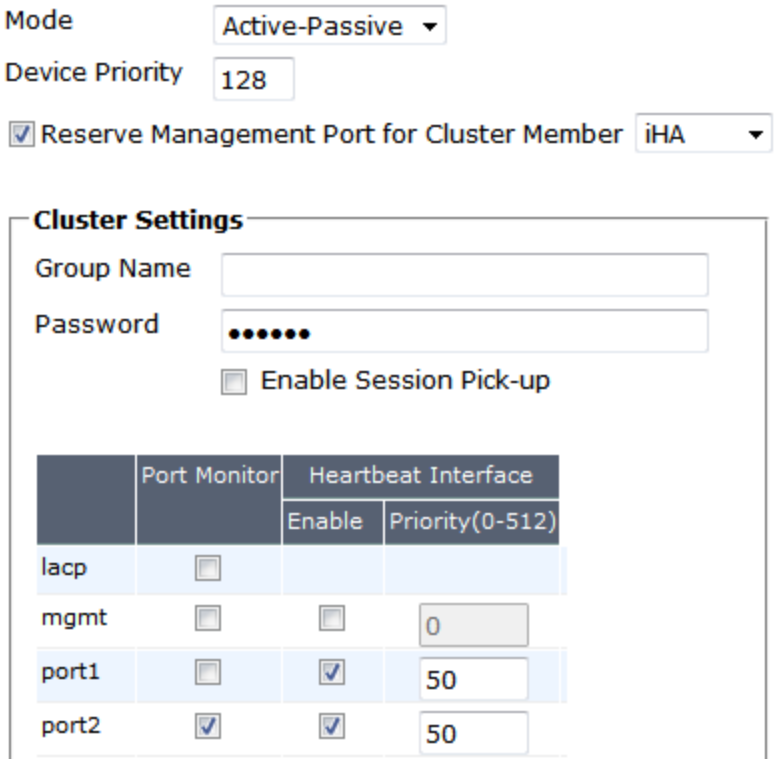

When session pickup is enabled, not only software states, but also hardware states are stored in the FortiASIC-FP1 processors. These sessions are synchronized among HA cluster members.

Therefore, besides HA heartbeat link, another 4 dedicated network links are required for hardware to synchronize session information between HA primary and backup. To match the throughput of four FortiASIC-FP1 processors in each system, 4 dedicated network links are needed. They can be specified by the following CLI commands:

```
conf sys ha
   set hw_sess_sync_dev ìport5 port6 port 7 port8î
   end
```

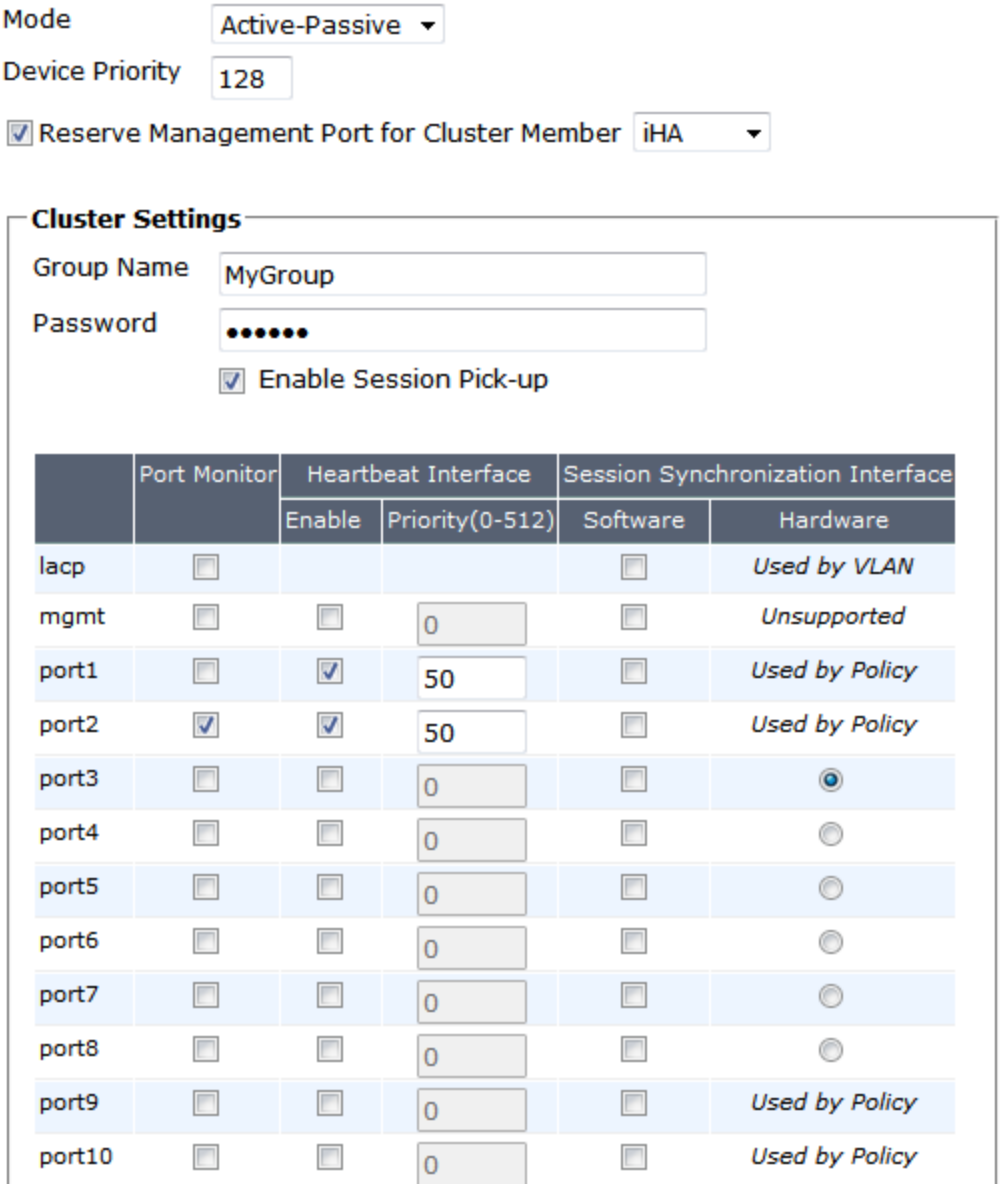
Mode
Active-Passive
Device Priority
128
Reserve Management Port for Cluster Member
iHA
Cluster Settings
Group Name
MyGroup
Password
Enable Session Pick-up
Port Monitor
Heartbeat Interface
Session Synchronization Interface
Enable
Priority(0-512)
Software
Hardware
lacp
Used by VLAN
mgmt
Unsupported
port1
50
Used by Policy
port2
50
Used by Policy
port3
port4
port5
port6
port7
port8
port9
Used by Policy
port10
Used by Policy

FortiCarrier specific CLI guide

IP pool

Syntax

```
config firewall ippool
   edit <pool name>
      set type <one-to-one|port-block-allocation|single-port-allocation>
      set tcp-user-quota <integer>
      set udp-user-quota <integer>
      set extra-tcp-user-quota <integer>
      set extra-udp-user-quota <integer>
      set tcp-session-quota <integer>
      set udp-session-quota <integer>
      set port-start <integer>
      set port-end <integer>
      set fixed-allocation <enable|disable>
      end
```

Options

type

`set type <one-to-one|port-block-allocation|single-port-allocation>`

- Set the type of IP pool
- The `fixed-port-range` and `overload` options have been removed

Options	Description
one-to-one	One to one mapping
port-block-allocation	PBA mode
single-port-allocation	SPA mode

tcp-user-quota

`set tcp-user-quota <integer>`

- Set the maximum number of TCP ports allowed per user

integer	1 - 32768

udp-user-quota

`set udp-user-quota <integer>`

- Set the maximum number of UDP ports allowed per user

integer	1 - 32768

extra-tcp-user-quota

```
set extra-tcp-user-quota <integer>
```

- When regular TCP user quota is used up, this option sets the limit of how many extra TCP quota can be allocated for essential services

integer	0 - 16

extra-udp-user-quota

```
set extra-udp-user-quota <integer>
```

- When regular UDP user quota is used up, this option sets the limit of how many extra extra UDP quota can be allocated for essential services

integer	0 - 16

tcp-session-quota

```
set tcp-session-quota <integer>
```

- Set the maximum number of TCP sessions allowed per user

integer	128 - 65535

udp-session-quota

```
set udp-session-quota <integer>
```

- Set the maximum number of UDP sessions allowed per user

integer	1 - 32768

port-start

```
set port-start <integer>
```

- Set the start of the port range

integer	1024 - 65535

port-end

```
set port-end <integer>
```

- Set the end of the port range

integer	1024 - 65535

fixed-allocation

```
set fixed-allocation <enable|disable>
```

- Toggles the fixed-allocation feature.
- This option is only available when port type is PBA or SPA

Options	Description
enable	Enables the Fixed Allocation feature
disable	Disables the Fixed Allocation feature

Central NAT

Central NAT, while it does relate to NAT, is not supported in FortiCarrier.

Firewall policy

Syntax

```
config firewall policy
   edit <policy name>
      set endpoint-independent-mapping <enable|disable>
      set endpoint-independent-filtering <enable/disable>
      set hw-logging-mode <none|session|mapping>
      end
```

Options

endpoint-independent-mapping

```
set endpoint-independent-mapping <enable|disable>
```

- Toggle Endpoint Independent Mapping

Options	Description
enable	Enables the Fixed Allocation feature
disable	Disables the Fixed Allocation feature

endpoint-independent-filtering

```
set endpoint-independent-filtering <enable/disable>
```

- Toggle Endpoint Independent Filtering

Options	Description
enable	Enables the Endpoint Independent Filtering feature
disable	Disables the Endpoint Independent Filtering feature

hw-logging-mode

```
set hw-logging-mode <none|session|mapping>
```

- Set the type of logging to be carried out

Options	Description
none	no hardware logging
session	Enables hardware logging for session setup/teardown
mapping	Enables hardware logging for NAT mapping setup/teardown

Firewall Policy Limitations

Firewall parameter	Maximum number per policy
Source interfaces	16
Destination interfaces	16
IP ranges	16
Services	16
IP pools	8

Firewall address

Syntax

```
conf firewall address
   edit <address name>
      set type <ipmask|iprange|wildcard>
      end
```

Options

type

- Sets the type of address object

Options	Description
ip mask	Subnet method of describing an address group
ip range	Method that uses a starting and ending address and included all the addresses between the two (inclusively)
wildcard	Uses wildcards in place of one or more octets

- Firewall address types that are not supported by FortiCarrier
  - FQDN
  - Geography
  - Wildcard-FQDN

Network interface

Syntax

```
conf system interface
   edit <interface name>
      set cgn-nat-role <private|public>
      end
```

Options

cgn-nat-role

```
set cgn-nat-role <private|public>
```

- Determine whether a port is on NAT private side or public side

Options	Description
private	This interface is connected to a private subnet with non routable IP addresses.
public	This interface is connected to a publicly veiwable subnet

Routing

Syntax

```
conf router static
   edit <name>
```

Options

Routing Limitations

- Settings in the for Routing that are not supported in FortiCarrier:
  - `internet-service`
  - `internet-service-custom`
  - `weight`
  - `priority`
  - `distance`
  - `blackhole`
  - `virtual-wan-link`
  - `dynamic-gateway`

Policy based routing

Syntax

```
conf router policy
   edit <name>
```

Options

Policy Routing limitations

- gateway and output-device must both be specified.
- `src-negate` and `dst-negate` are not supported.

Logging

The syslog filter feature is not supported in FortiCarrier, so `config log syslogd filter` command is not available.

VDOM

Multi-vdom is not supported in FortiCarrier

Troubleshooting

This short writing summarizes some diagnose commands added on FortiCarrier CLI to help troubleshoot issues may be encountered in the field. The CLI commands are described first, followed by some use cases.

Diagnose commands

The newly added diagnose commands are all in the form of “diagnose npu cgn1 <command>”, below each command is explained and example outputs are given. As there are 4 FortiASIC-FP1 devices in the system, some commands can be called per-device by providing a device ID. Most of such commands can also be accessed from visiting the FortiOS /proc file system.

dce

Checks FortiASIC-FP1 packet drop counters. The counters are clear-on-read. The non-zero counters are shown with a name (cause) which may be analyzed to understand the packet drop condition in hardware.

Category: Per-device

Syntax:

```
diag npu cgn1 dce <device-id>
```

Example:

```
FCR38E-198-modi # diag npu cgn1 dce 0
RPE0_LSM_FUL_POL:0000000000000596 [dc] RPE1_LSM_FUL_POL:0000000000000720 [dd]
RPE2_LSM_FUL_POL:0000000000000752 [de] RPE3_LSM_FUL_POL:0000000000000824 [df]
```

sse-hw

Check FortiASIC-FP1 SSE sessions stats. Show detailed stats of hardware sessions.

Category: Per-device

Syntax

diag npu cgn1 sse-hw 1

Example

```
FCR38E-198-modi # diag npu cgn1 sse-hw <dev-id>
llmgr-base = ffffc90002486180, 1
tx-cmd:        00000000  tx-cmd_err:    00000000
rx-rsp:        00000000  rx-rsp_err:    00000000
rx-unsolicited:00000000
Counters        SSE0            SSE1            SSE2            SSE3            Total
```

```
-------------- -------------- -------------- -------------- -------------- --------------
SRHSUCC        2983076742     2984767807     2983870940     2983329559     3345110456
SRHFAIL        2245670669     2245864690     2245692525     2245663771     392957063
UPDSUCC        894822133      895118200      895008929      894911060      3579860322
AGEDOUT        741393717      741422702      741438900      741419402      2965674721
INSSUCC        1482831830     1482936531     1482857609     1482838039     1636496753
INSFAIL        0              0              0              0              0
DELSUCC        741317704      741392726      741298386      741298727      2965307543
DELFAIL        0              0              0              0              0
DEPFAIL        0              0              0              0              0
ENTCNT         120411         121115         120334         119921         481781
SRH_CNT        2987684547     2989435769     2988429366     2987898559     3363513649
SRH_MATCH      2241622926     2243284816     2242371660     2241850036     379194846
SRH_MISS       746061626      746150956      746057710      746048526      2984318829
INS_MATCH      0              0              0              0              0
```

pdq

Check FortiASIC-FP1 packet descriptor queue status.

Category: Per-device

Syntax

```
diag npu cgn1 pdq <dev-id>
```

Example

```
FCR38E-198-modi # diag npu cgn1 pdq 1
pba_num = 00000bee
----------- ----------- ----------- ----------- -----------
PDQ: isw-ipdq, isw-i, 6 banks
93/93   1a/1a   87/87   b6/b6   39/39   0/0
----------- ----------- ----------- ----------- -----------
PDQ: isw-ppdq, isw-p, 6 banks
94/94   21/21   8c/8c   bc/bc   39/39   0/0
----------- ----------- ----------- ----------- -----------
PDQ: dsw=>sse, sse, 4 banks
d6/d6   b9/b9   c0/c0   cf/cf
----------- ----------- ----------- ----------- -----------
PDQ: rpe:sse=>rpe, rpe, 4 banks
d9/d9   bc/bc   c0/c0   d1/d1
----------- ----------- ----------- ----------- -----------
PDQ: rpe:fp-pipe, rpe, 4 banks
db/db   be/be   c0/c0   d3/d3
```

```
---------- ---------- ---------- ---------- ----------
PDQ: rpe:fp-vpdq, rpe, 4 banks
fe/fe   f9/f9   de/de   8d/8d
---------- ---------- ---------- ---------- ----------
PDQ: rpe:sp-pipe, rpe, 4 banks
e1/e1   cf/cf   e4/e4   4c/4c
---------- ---------- ---------- ---------- ----------
PDQ: rpe:sp-vpdq, rpe, 4 banks
e1/e1   cf/cf   e4/e4   4d/4d
---------- ---------- ---------- ---------- ----------
PDQ: ehp: qosw, ehp, 4 banks
51/51   28/28   9a/9a   9e/9e
---------- ---------- ---------- ---------- ----------
PDQ: nat: merge-0/1/SDR/RX, nat, 1 banks
c4/c4   1a/1a   e1/e1   1d/1d
---------- ---------- ---------- ---------- ----------
PDQ: fwd: DSW2HRX, HRX, 1 banks
3a/3a
---------- ---------- ---------- ---------- ----------
PDQ: tunpdq: HRX, HRX, 1 banks
3a/3a
---------- ---------- ---------- ---------- ----------
PDQ: rxpdq-0~7: HRX, HRX, 1 banks
```

NOTE: 3a/3a 0/0 0/0 0/0 0/0 0/0 0/0 0/0

prp

Check driver prp-manager status.

Category: Per-device

Syntax

```
diag npu cgn1 prp <dev-id>
```

Example

```
FCR3805E_192 # diag npu cgn1 prp 1
PRP-stats:(dbg: 00000000, dctrl: 00000000, tcnt: 000027a9)
CMDQ-stats:
tx-cmd:        000ff613  tx-cmd_err:    00000000
rx-rsp:        00000000  rx-rsp_err:    00000000
rx-unsolicited:000804a2
FIFO stats
CLI-FIFO(E/EF/D/DF): 00000021 00000000 00000021 00000000
```

```
add_pool(a/f):    00000007/00000000
add_res(a/f):     0000000e/00000000
refill(a/f):      00000012/00000000
mgr-busy(a/f):    00000005/00000000
add_entry(a/f):    0016b2e2/00000000
add_lsv(a/f):    001fde80/00000000
write_ilt(a/f):    000926f0/00000000
misc     (a/f):    00004060/00000000
nodealloc : 00000651/00000000 nodefree  : 00000651/00000000 lsv_response: 0007f800/00000000
```

xgmac-stats

Show traffic MIBS stats for 4 traffic links of each FortiASIC-FP1. Each FortiASIC-FP1 has 7x10G links, 4 of them are for user traffic, 2 for inter-FortiASIC-FP1 or inter-system, 1 for session logging.

Category: Per-device

Syntax

```
diag npu cgn1 xgmac-stats <dev-id>
```

Example

```
FCR38E-198-modi # diag npu cgn1 xgmac-stats 1
Counters        XE0             XE1             XE2             XE3
--------------- --------------- --------------- --------------- ---------------
RX_BCAST        0               23              1               0
RX_MCAST        11317           0               2268            16340
RX_UCAST        17475637571     17475049058     17474463599     17475967360
RX_PAUSEFRM     0               0               0               0
RX_UNDERSIZE    0               0               0               0
RX_OVERSIZEP    43              94              33              67
RX_FRAG         0               0               0               0
RX_JAB          0               0               0               0
RX_FCS          0               0               0               0
RX_WFULL        0               0               0               0
RX_GOODOCTET    9531287858930   9531464326278   9531178563660   9531660023763
RX_OCTET        9531287860796   9531464327104   9531178565668   9531660024377
--------------- --------------- --------------- --------------- ---------------
TX_BCAST        0               0               0               0
TX_MCAST        0               0               0               0
TX_UCAST        17473632163     17473044073     17472472637     17473971769
TX_COL          0               0               0               0
TX_LATECOL      0               0               0               0
TX_EXCESSCOL    0               0               0               0
```

```
TX_UNDERRUN     0               0               0               0
TX_XPX_QFULL    0               0               0               0
TX_GOODOCTET    9530365222398   9530543864118   9530274251689   9530747459028
TX_OCTET        9530365222472   9530543868833   9530274251689   9530747459254
--------------- --------------- --------------- --------------- ---------------
PKT1024TOMAX    11444702782     11445035248     11444714708     11445201120
PKT512TO1023    6269308         6260442         6262252         6267484
PKT256TO511     757136021       757063904       757004737       757138252
PKT128TO255     11317           0               2268            0
PKT65TO127      22741161753     22739733732     22738954659     22741348755
PKT64           0               0               0               0
--------------- --------------- --------------- --------------- ---------------
```

nat-stats

Show or clear accumulated resource usage in various types.

Category: Per-device

Syntax

```
diag npu cgn1 nat-stats <dev-id> <op: 0: read, 1: clear>
```

Example

```
FCR38E-198-modi # diag npu cgn1 nat-stats  0 0
nat stats for dev-(0):
NAT_MAX_SESS_REACHED_TCP: 524026
NAT_PORTS_FREED_TCP: 3040327890
NAT_PORTS_ALLOCATED_TCP: 3040569504
FCR38E-198-modi # diag npu cgn1 nat-stats  0 1
```

user-port

Show packet stats of each physical interface and VLAN interface.

Category: Sys-Stats

Syntax

```
diag npu cgn1 user-port
```

Example

```
FCR3805E_192 # diag npu cgn1 user-port
LIF PORT                  RX-PKT                          RX-BYTE                          TX-PKT
                    TX-BYTE
134 port1               24244594/+24238758              1959433119/+1958102511
24276178/+24276171              28785114721/+28785114157
```

```
SW                    24244636                          2153384899                          24276178
              28882219433

135 port2             24283220/+24277384             28689617737/+28688287129
24237538/+24237536              2054759821/+2054759671

SW                    24283262                          28883893349                         24237538
               2151709973

140 port7                    0/+0                              0/+0                               3/+0
                      222/+0

SW                           0                                 0                                  3
                      234

141 port8                    0/+0                              0/+0                               4/+0
                      316/+0

SW                           0                                 0                                  4
                      332

OID VID                   RX-PKT                          RX-BYTE                            TX-PKT
                    TX-BYTE

134 1010                     0/+0                              0/+0                               2/+0
                      158/+0

134 1020                     0/+0                              0/+0                               2/+0
                      158/+0

140 305                      0/+0                              0/+0                               2/+0
                      158/+0

140 405                      0/+0                              0/+0                               1/+0
                       64/+0

141 306                      0/+0                              0/+0                               2/+0
                      158/+0

141 406                      0/+0                              0/+0                               2/+0
                      158/+0
```

session

Show concurrent and accumulated SSE sessions of each FortiASIC-FP1 and overall. Starting from 2nd column are: concurrent sessions, accumulated/incremental session-inserts, rate of session-inserts, accumulated/incremental session-deletes, rate of session-deletes.

Category: Sys-stats

Syntax

```
diag npu cgn1 session
```

Example

```
FCR38E-198-modi # diag npu cgn1 session
sess-ent              sess-ins                                  sess-del
FPGA-0:          479164,       6216412273/177672  ,   178175/s          3108452627/89097   ,    89117/s
FPGA-1:          481195,       6216869537/176686  ,   176370/s          3108007757/88191   ,    88128/s
FPGA-2:          481126,       6220680451/176367  ,   176522/s          3110474573/88208   ,    88253/s
FPGA-3:          480893,       6220217680/177063  ,   176585/s          3109193483/88670   ,    88280/s
Total:          1922378,      24874179941/707788  ,   707652/s         12436128440/354166  ,   353778/s
```

sw-pat

Show the relation between FortiOS IP-pool and the hardware pool. Note, hardware pool is always paired (TCP and UDP). Any FortiOS IP-pool can only belong to one hardware pool (pair). For now, there is a software limit of total number of public IP's in the system (256K).

Category: Sys-stats

Syntax

```
diag npu cgn1 sw-pat
```

Example

```
FCR3600E-TOP # diag npu cgn1 sw-pat
PRP HA mode: master
prp-mgr-cfg: tbl-entry=12582912, shared=1, ip-max-chain=8
SPA     :  3 pools, 65369 IP's, 3949398873 resource
pool-(24/25): refcnt=1, sz=1, ffff880fa5bc7300, {2:pool14 ,}
pool-(28/29): refcnt=2, sz=1, ffff880fa4846ee0, {3:ap9 ,}
pool-(30/31): refcnt=3, sz=1, ffff880fa643c820, {4:fixed-nat ,}
```

pool

Show detailed stats of resource in each pool. For each pool, 1st line shows pool type, protocol, session quota etc. Stats are for each FortiASIC-FP1, "shared" and total. Currently all pools are operating in "shared" mode – i.e., resource of one IP can be shared among all FortiASIC-FP1's. If the pool is not being used by any sessions, per-FortiASIC-FP1 stats are not shown.

Starting from the 2nd column, they are

1. Number of IP addresses in active and idle list;
2. Total number of resource;
3. Number of resource in prp-mgr (software);
4. Number of resource in prp (being buffered by hardware);
5. Number of resource in NAT/SSE (being used by hardware);
6. Number of resource in a cache (software buffer of resource before pushing down to hardware).

Category: Sys-Stats

Syntax

```
diag npu cgn1 pool
```

Example

```
FCR3600E-TOP # diag npu cgn1 pool
----------PRP_MGR-stats:----------
ip(a/i)     total   in-mgr   in-prp   in-nat   in-cache
Pool[24], SPA, TCP, 0/64/65535/4, q=1944, hr/pr=0/0:
Shared:     0/0       5d1c65   5c71f5        0        0        0 e:        0        0
```

```
Total :   194/0     005d1c65 005c71f5 0000aa70 00000000 00000000 e: 00000000 00000000
Pool[25], SPA, UDP, 1/64/65535/4, q=1944, hr/pr=0/0:
Shared:     0/0       5d1c65   5c71f5        0        0        0 e:        0        0
Total :   194/0     005d1c65 005c71f5 0000aa70 00000000 00000000 e: 00000000 00000000
Pool[28], SPA, TCP, 0/64/65535/4, q=1944, hr/pr=0/0:
Shared:     0/0     75fccffc 7524d6bc        0        0        0 e:        0        0
Total : 1fff0/0     75fccffc 7524d6bc 00d7f940 00000000 00000000 e: 00000000 00000000
Pool[29], SPA, UDP, 1/64/65535/4, q=1944, hr/pr=0/0:
Shared:     0/0     75fccffc 7524d6bc        0        0        0 e:        0        0
Total : 1fff0/0     75fccffc 7524d6bc 00d7f940 00000000 00000000 e: 00000000 00000000
Pool[30], SPA, TCP, 0/64/65535/4, q=1944, hr/pr=0/0:
FPGA-0:  7ef8/0            0        0  11f8bea    9f208     3284 e:        0        0
FPGA-1:  7ef8/0            0        0  11f9a0c    9fb47     2b45 e:        0        0
FPGA-2:  7ef8/0            0        0  11fb728    a1186     2e2b e:        0        0
FPGA-3:  7ef8/0            0        0  11f8e72    9f572     2fd3 e:        0        0
Shared:     0/0     750d1ef8 7066035a        0        0        0 e:        0        0
Total : 1fbe0/0     750d1ef8 7066035a 047e6b90 0027f447 0000bbc7 e: 00000000 00000000
Pool[31], SPA, UDP, 1/64/65535/4, q=1944, hr/pr=0/0:
Shared:     0/0     750d1ef8 7436dc78        0        0        0 e:        0        0
Total : 1fbe0/0     750d1ef8 7436dc78 00d64280 00000000 00000000 e: 00000000 00000000
```

ha

Show HA status from driver and hardware's point of view. Driver maintains 4 HA state (S1, S2, M1 and M2), and some status related to hardware session sync (HA-RLT status, NAT-auto-drain status). Also it maintains the status of "driver step synchronization" between primary and backup.

Category: Sys-stats

Syntax

```
diag npu cgn1 ha
```

Example

```
FCR3600E-TOP # diag npu cgn1 ha
ha-status:
state = M2(2), rlt = 0, last rlt-status = 00000000
BOOTUP          :        1             0
SWITCH-OVER     :        2             0
SLAVE-JOIN      :        2             0
SLAVE-DEPART    :        1             0
RLT             :        9             9
NAT             :        0             0
TIMER/RLT-START :     1019           171
```

```
BUSY            :        0
--drv-step--
LS: 6, PS: 4, SEQ:63, ACK:63
RX-MSG-CNT      : 00000143      00000000
RX-CMD-CNT      : 00000080      00000000
RX-RSP-CNT      : 00000063      00000000
TX-CMD-CNT      : 00000063      00000000
TX-RSP-CNT      : 00000080      00000000
LOCAL-STEP-CNT  : 00000099      00000000
IDLE            : 00000006      00000000
DEL-RSC         : 00000006      00000000
ADD-RSC         : 00000010      00000000
DONE            : 00000021      00000000
```

sw-pat

Show the relation between FortiOS IP-pool and the hardware pool. Note, hardware pool is always paired (TCP and UDP). Any FortiOS IP-pool can only belong to one hardware pool (pair). For now, there is a software limit of total number of public IP addresses in the system.

Category: Sys-stats

Syntax

```
diag npu cgn1 sw-pat
```

Example

```
FCR3600E-TOP # diag npu cgn1 sw-pat
PRP HA mode: master
prp-mgr-cfg: tbl-entry=12582912, shared=1, ip-max-chain=8
SPA     :  3 pools, 65369 IP's, 3949398873 resource
pool-(24/25): refcnt=1, sz=1, ffff880fa5bc7300, {2:pool14 ,}
pool-(28/29): refcnt=2, sz=1, ffff880fa4846ee0, {3:ap9 ,}
pool-(30/31): refcnt=3, sz=1, ffff880fa643c820, {4:fixed-nat ,}
```

igr-fpga

Set any one of the 4 FortiASIC-FP1's or all 4 FortiASIC-FP1's to process traffic.

Category: Debug

Syntax

```
diag npu cgn1 igr_fpga <dev-id or -1 >
```

Example

```
FCR3805E_192 # diag npu cgn1 igr_fpga  -1
```

fp_enable

Disable or enable FortiASIC-FP1 fast-path processing. FortiASIC-FP1 fast-path processing, if enabled, will process TCP/UDP traffic (including policy/route lookup, session setup and SNAT/DNAT based on sessions). If disabled, FortiASIC-FP1 will forward all packets to host.

Category: Debug

Syntax

```
diag npu cgn1 fp_enable <dev-id> <op: 0: disable, 1: enable>
```

Example

diag npu cgn1 fp_enable -1 0

ruby

Run ruby shell to check FortiASIC-FP1 internal status.

Category: Debug

Syntax

```
diag npu cgn1 ruby
```

Example

```
diag npu cgn1 ruby
```

bcm-shell

Run bcm-shell to check Broadcom switch status. Note, it can only be run from console. To exit, type “Ctrl-C”.

Category: Debug

Syntax

```
diag npu cgn1 bcm-shell
```

Example

```
diag npu cgn1 bcm-shell
```

tcam-NHI

Check NHI usage in TCAM database.

Category: Debug

Syntax

```
diag npu cgn1 tcam nhi
```

Example

```
FCR3805E_192 # diag npu cgn1 tcam nhi
70 NH:
NHI  type  gw:oif        lpm    pbr   ref
0x  1   0  00000000:1023    193     0    193
0x  2   1  00000000:   0    104     0    104
0x 20   2  00000000: 135     1     0      1
0x 21   2  00000000: 134     1     0      1
0x 22   2  00000000: 134     1     0      1
0x 23   2  00000000: 134     1     0      1
0x 24   2  00000000: 189     1     0      1
0x 25   2  00000000: 189     1     0      1
0x 26   2  00000000: 189     1     0      1
```

tcam-route

Check LPM route usage in TCAM database.

Category: Debug

Syntax

```
diag npu cgn1 tcam route
```

Example

```
FCR3600E-TOP # diag npu cgn1 tcam route
98 LPM:
prefix/pl  type   nhi      gw:oif    rule#
ffffffff/32,  0, 0x  1, 00000000:1023, 0x0000
01030000/32,  0, 0x  1, 00000000:1023, 0x0000
01030001/32,  0, 0x  1, 00000000:1023, 0x0001
010300ff/32,  0, 0x  1, 00000000:1023, 0x0002
010b0000/32,  0, 0x  1, 00000000:1023, 0x0003
40097ff0/29,  0, 0x  2, 00000000:   0, 0x0023
40097ff8/30,  0, 0x  2, 00000000:   0, 0x0022
40097ffc/31,  0, 0x  2, 00000000:   0, 0x0021
40097ffe/32,  0, 0x  2, 00000000:   0, 0x0020
01030000/24,  0, 0x 20, 00000000: 188, 0x0018
010b0000/24,  0, 0x 21, 00000000: 134, 0x0019
010c0000/24,  0, 0x 22, 00000000: 135, 0x001a
010d0000/24,  0, 0x 23, 00000000: 144, 0x001b
010e0000/24,  0, 0x 24, 00000000: 145, 0x001c
010f0000/16,  0, 0x 25, 00000000: 136, 0x001d
```

```
01100000/16,  0, 0x 26, 00000000: 137, 0x001e
05030200/24,  0, 0x 27, 00000000: 165, 0x001f
020b0000/16,  0, 0x 2e, 02000b01: 134, 0x005e
020c0000/16,  0, 0x 2f, 02000c01: 135, 0x005d
020d0000/16,  0, 0x 30, 02000d01: 144, 0x005c
020e0000/16,  0, 0x 31, 02000e01: 145, 0x005b
020f0000/16,  0, 0x 32, 02000f01: 136, 0x005a
02100000/16,  0, 0x 33, 02001001: 137, 0x0059
ca100000/16,  0, 0x 33, 02001001: 137, 0x0002
ca110000/16,  0, 0x 33, 02001001: 137, 0x0001
```

tcam-full-policy

Check full-policy table in TCAM. 1st column is FortiOS policy-ID, 2nd column is hardware policy-ID.

Category: Debug

Syntax

```
diag npu cgn1 tcam fulpol
```

Example

```
FCR3805E_192 # diag npu cgn1 tcam fulpol
37 policies:
fos-pol   hw-pol   num-tcam   1st-rule
0xffff    0x0002   0x0001     0x0000
0x013a    0x0012   0x0001     0x0001
0x013a    0x0013   0x0001     0x0003
0x0133    0x0020   0x0001     0x0013
0x0133    0x0021   0x0001     0x0014
0x0132    0x0022   0x0001     0x0015
0x0132    0x0023   0x0001     0x0016
0x0131    0x0024   0x0001     0x0017
0x0131    0x0025   0x0001     0x0018
0x0130    0x0026   0x0001     0x0019
0x0130    0x0027   0x0001     0x001b
0x0065    0x0034   0x0001     0x002f
0x0065    0x0035   0x0001     0x0030
```

tcam-ingress-policy

Check ingress-policy table in TCAM.

Category: Debug

Syntax

```
diag npu cgn1 tcam igrpol
```

Example

```
FCR3805E_192 # diag npu cgn1 tcam igrpol
2 policies:
fos-pol   hw-pol   num-tcam   1st-rule
0xffff    0x0010   0x0001     0x0000
0xffff    0x0011   0x0001     0x0001
```

tcam-daemon-exit

Restart tcam-daemon.

Category: Debug

Syntax

```
diag npu cgn1 tcam exit
```

Example

```
FCR3805E_192 # diag npu cgn1 tcam exit
TCAMD has been restarted!
```

Besides the above “diagnose npu cgn1 <cmd>” commands, the “diagnose system session” commands can be used to list and clear sessions in hardware.

Hardware session list

The command lists/clears sessions installed by hardware. Note, when there are a lot of SSE sessions, doing a list can cause SSE to burst large amount of messages to host. To list/clear sessions installed by host, don’t use “diag sys session file hwsel 0xF”.

Category: Debug

Syntax

```
diag sys session filter hwsel 0xF
diag sys session list
diag sys session clear
```

Examples

```
FCR3805E_192 # diag npu cgn1 session
sess-ent              sess-ins                              sess-del
FPGA-0:                4,          1123072/0       ,     8140/s              561400/0       ,     4650/s
FPGA-1:                5,          1122592/0       ,     7647/s              561137/0       ,     4365/s
```

```
FPGA-2:                6,          2057621/0       ,     7696/s              1028835/0       ,     4396/s
FPGA-3:                3,          1123571/0       ,     7751/s               562047/0       ,     4437/s
Total:                18,          5426856/0       ,    31234/s              2713419/0       ,    17848/s

FCR3805E_192 # diag sys session filter  hwsel  0xF

FCR3805E_192 # diag sys session list
session info: proto=6 proto_state=00 duration=0 expire=0 timeout=0 flags=00000000 sockflag=00000000 sock-
port=0 av_idx=0 use=0
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 policy_dir=0 tunnel=/ vlan_cos=0/0
state=
statistic(bytes/packets/allow_err): org=0/0/0 reply=0/0/0 tuples=2
speed(Bps/kbps): 0/0
orgin->sink: org pre->pre, reply pre->pre dev=0->0/0->0 gwy=0.0.0.0/0.0.0.0
hook=pre dir=org act=snat 3.16.38.85:11869->3.8.69.147:80(64.2.10.132:5219)
hook=pre dir=reply act=dnat 3.8.69.147:80->64.2.10.132:5219(3.16.38.85:11869)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=0 auth_info=0 chk_client_info=0 vd=0
serial=00000000 tos=00/00 app_list=0 app=0 url_cat=0
dd_type=0 dd_mode=0
npu_state=00000000
no_ofld_reason:  npu-flag-off
SSE-session-0:
…
FCR3805E_192 # diag sys session clear

FCR3805E_192 # diag npu cgn1 session
sess-ent              sess-ins                                sess-del
FPGA-0:                0,          1123072/0       ,     2366/s               561402/2       ,     1339/s
FPGA-1:                0,          1122592/0       ,     2373/s               561140/1       ,     1341/s
FPGA-2:                0,          2057621/0       ,     2388/s              1028837/1       ,     1357/s
FPGA-3:                0,          1123571/0       ,     2405/s               562048/0       ,     1370/s
Total:                 0,          5426856/0       ,     9532/s              2713427/4       ,     5407/s
```

Typical use cases

No traffic passing:

1. Check stats of user ports (“diag npu cgn1 user-port”);
2. Check stats of FortiASIC-FP1 10G traffic links;
3. Check DCE counters to see whether packets are dropped for any reasons;
4. If the dropping is due to MAC filter, check whether the traffic come in with wrong destination MAC address for any reasons;
5. If the dropping is due to route lookup, check whether the routes are added correctly by verifying FortiOS setting and TCAM configurations (“diag npu cgn1 tcam route/nhi”);
6. If the dropping is due to full-policy lookup, check whether full-policies are added correctly by cross-checking FortiOS policy setup and TCAM configurations (“diag npu cgn1 tcam fulpol”).
7. If the dropping is due to “SSE PDQ”, check PDQ status to see whether the queue is stuck.
8. If the dropping is due to lack of PRP resource, check whether resource pool is configured correctly by verifying FortiOS ip-pool setup and PRP-manager configurations (“diag npu cgn1 pool”).
9. If none of the above work, disable FortiASIC-FP1 fast-path and see if the packets can reach host.

Portion of the traffic not passing: Besides the above 5, 6, 8 items, also check whether any DCE drop counters are related to session quota being exceeded, e.g., a client opens too many connections, and is exceeding its session quota or resource quota(for port-block-allocation case).

Supported RFCs