Virtual Servers and Services

https://campus.barracuda.com/doc/53248453/

The F-Series Firewall architecture is split into three layers:

- **Box Layer** – The box layer runs infrastructure services, responsible for logging, event, configuration, and control. The network subsystem is also part of the box layer, which creates some peculiarities with network configuration. Only the management IP address and the additional box layer IP addresses of the F-Series Firewall is allocated in the box layer. The box layer is always active.
- **Virtual Server Layer** – The virtual server layer builds on the box layer. It is a purely logical layer whose most important function is to make IP addresses available for the services (service layer) started on it. The virtual server layer introduces and activates all IP addresses that are needed for proper operation.
- **Service Layer** – The service layer introduces services such as the firewall, VPN, or DHCP. If the virtual server layer shuts down, all of its assigned services are also shut down and made unavailable.

Virtual servers represent the main operative instance on the NextGen Firewall F-Series next to global settings and box configuration objects. The virtual server layer manages all IP addresses that are required for the services running on the virtual servers. It introduces all IP addresses that are needed for proper operation except remote management and HA IP addresses. Even though virtual server contains the word "virtual", there is no virtualization layer between the box layer and the virtual server layer. The term “virtual” is used to describe the logically separated servers that are running on the system. You can create multiple virtual servers on each firewall.

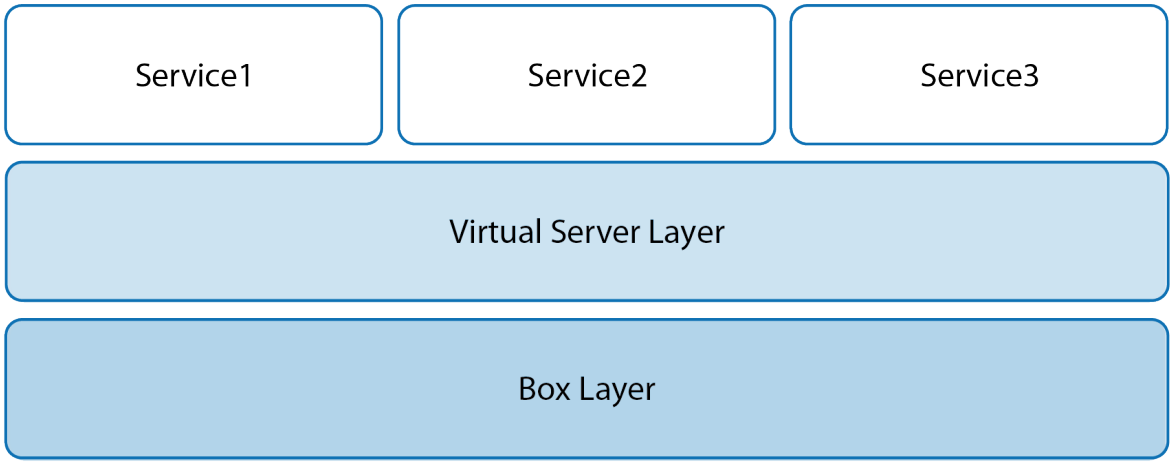

Virtual Servers

The virtual server layer runs on the box layer of the F-Series. It is a purely logical layer whose most important function is to make IP addresses available for the services (service layer). Introduce all IP addresses on the virtual server that will be used for the services running in the virtual server. These IP addresses must be in one of the networks for which a directly attached network route exists on box level. It is not recommended to use the management or additional box level IP addresses because this could causes problems in HA cluster configurations. When a virtual server is started, it assigns IP addresses to its services, causing the box layer to automatically activate pending routes of directly attached network routes. Once created virtual servers cannot be renamed.

Preconfigured Virtual Server S1

By default, the virtual server S1 is already created on every F-Series Firewall except the larger hardware models. To avoid duplicated IP addresses within networks, this virtual server listens on the loopback IP address 127.0.0.9. On stand-alone F-Series Firewall systems, you can keep the default virtual server S1. On firewalls you want to manage with a Control Center, create a new virtual server because virtual server names must be unique in the cluster.

Virtual Servers in the Barracuda NextGen Control Center

On the Barracuda NextGen Control Center, virtual servers are created in the Control Center cluster. The setup procedure is very similar to the procedure on a stand-alone firewall, which means that you can create a server and assign the network IP addresses and services. Virtual servers act as separate configuration entities, so you can copy them from one cluster to another.

For more information, see How to Configure Virtual Servers.

HA Monitoring and Transparent Failover

A virtual server is transferable between members of a high availability cluster. If the primary unit fails, the virtual server, including its assigned IP addresses and all services, is instantly transferred to the secondary unit. You can also create virtual servers with services to run only on a secondary unit that, in case of a failover, are transferred to the primary unit and vice versa.

For HA failover, the management IP address and the 1st virtual server IP address are monitored by default. To configure transparent monitoring for HA clusters, create monitoring policies for interfaces and IP addresses. The virtual server stays up as long as these health check targets are reachable.

For more information, see Virtual Server Monitoring and High Availability.

Services

The service layer runs on the virtual server layer of the firewall. It introduces the services such as firewall, HTTP proxy, VPN, and DHCP. The services use the configured IP addresses of the virtual server on which they are running. If the virtual server shuts down, all of the assigned services and IP addresses are also shut down and made unavailable. If the Firewall F-Series Firewall is deployed in a high availability cluster, the services and necessary IP addresses transparently failover to the other HA unit.

Service Limitations

Some services can only be introduced once on an F-Series Firewall:

- **Forwarding Firewall** – Because the Firewall module is based on the kernel, you can only have one firewall service per firewall.
- **VPN** – Because the VPN service is based on the kernel, it can only be introduced once on a firewall. The forwarding firewall and VPN service must be in the same virtual server.
- **Access Control** and **Mail Gateway** – These services provide a user interface that is always bound to the first introduced service of their type.
- **HTTP Proxy** – The HTTP proxy service can be introduced multiple times, but the HTTP proxy fail cache interface can only be used by one service.
- **HTTP Proxy** and **Web Filter** – You must also configure the HTTP proxy service and the web filter service on the same system and assign them to the same virtual server.

For more information, see How to Configure Services, NextGen F-Series Services or Shared Services.

Figures

1. layers1-01.png