Fortinet Solutions for the MSSP

Contents

Introduction

There are a number of reasons why a growing number of subscribers are drawn into Managed Security Services (MSS), making this an attractive business for those service providers looking for new sources of revenue and to expand the customer base. To start with, subscribers benefit from gaining access to advanced technologies, sophisticated processes, and highly trained personnel that subscribers couldn't possibly replicate in house. MSS subscriptions also simplify budgeting and help subscribers keep expenses under control with predictable monthly or yearly subscriptions that, thanks to economies of scale, are offered at a lower cost. Additionally, subscribers benefit from faster service deployment, allowing them to quickly respond to market and regulatory changes.

While articulating a business case around a managed security service offering may seem to be an easy task to some, the Managed Security Service Provider (MSSP) faces the not so trivial challenges of designing the right service portfolio and identifying the best security solution partner. Fortinet has a complete portfolio of products and services that allow the provider to build modular and flexible service offerings based on Fortinet Unified Threat Management (UTM) platforms. The use of selected unified platforms greatly simplifies the service infrastructure as there are less moving parts. In addition, Fortinet platforms offer multiple deployment options, and allow service to grow over time to better align with the always changing business requirements. This translates into faster time-to-market and increased operational efficiency.

This document provides a technical overview to Fortinet's product and services portfolio for MSSPs. It discusses the most common MSSP service models, and it positions the products and features that can be leveraged to satisfy such models. To facilitate planning, multiple alternative designs are presented per service model, each one provided with a detailed analysis of their features and benefits. While the primary audience is the technical personal responsible for designing and implementing the solutions, content in this document may be useful to the business decision makers as well.

MSSP Common Requirements

Despite the fact every business environment is different due to industry, geographical and regulatory factors; most managed security service providers share some common objectives and requirements.

Reduced Total Cost of Ownership (TCO): This refers to costs associated with building and maintaining the MSSP infrastructure. It includes hardware and software costs, licensing costs, operational expenses, and other long term expenses such as product replacement and decommissioning. Building an MSSP service offering with unified platforms greatly simplifies the infrastructure, reducing ownership and operational costs. Fortinet UTM platforms allow growing services over time, extending the useful life of the equipment. Fortinet UTM platforms are an alternative to the deployment of multiple specialized devices, reducing power and cooling costs.

Centralized Management and Reporting: MSSP environments involve large amounts of devices and services that need to provisioned, monitored and administered. Centralized device management, logging and reporting become a fundamental necessity in such environments. Fortinet's FortiAnalyzer and FortiManager are carrier-grade solutions that allow the centralization of those functions, even in multi-tenant environments.

Service Level Agreement Compliance: MSSP services are often offered under service contracts that include a service level agreement (SLA). The SLA defines the terms under which service is provided. SLAs typically include measurable terms that define expected performance, availability, number of security incidents, etc. SLA may estipulate monetary or credit compensation to the subscriber in case of non compliance. Fortinet UTM solutions

provide a number of mechanisms to monitor and maintain SLA compliance. This helps the MSSP avoid penalties and protect customer satisfaction.

Regulatory Compliance: It is in the MSSP best interest to ensure continuous compliance to industry standards and government regulations governing subscribers and the MSSP itself. Fortinet UTM multi-layer protection helps ensure assets and data are continuously protected. Fortinet products also provide a wide set of tools to monitor compliance and to alert on any deviations or violations, so the necessary corrective actions can be taken on time.

High Availability and Resiliency: Customer assets and data must be protected at all times, without service interruption. To that end, the MSSP infrastructure must be designed with high availability and resiliency to ensure no disruption or service degradation occurs in the event of link, device and component failure. The infrastructure must also be able to sustain and recover from network misconfiguration and network attacks. Fortinet solutions deliver multiple-levels of redundancy, from link, module to device failover. Additionally, Fortinet UTM platforms are built-in with several mechanisms that deliver enhanced resiliency, including denial-of-service mitigation, statefull inspection and intrusion prevention.

Dependable Scalability: The MSSP infrastructure must be designed to accommodate new services and support user growth. Hardware acceleration and clustering technologies like those in Fortinet products give the MSSP room to grow in services and number of subscribers. Thanks to their feature modularity, the Fortinet UTM platforms allow the MSSP to enable new services as subscribers and the business demands it.

Ease of Integration: One of the biggest challenges in the deployment of security services is their integration into the carrier infrastructure. Fortinet UTM products support a number of features to ensure a seamless integration into the carrier routing and switching infrastructure. This includes support for a variety of routing protocols, Layer 2 technologies, and transparent deployment options to ensure insertion without a network redesign.

Multi-tenancy: MSSP environments are multi-tenant in nature, where parts of the infrastructure are shared across multiple subscribers or tenants. In these environments, it is crucial to ensure tenant resources and data are properly protected and isolated. Resources and data belonging to one tenant should not be accessible to unauthorized parties. Likewise, problems in one tenant environment should not degrade or disrupt the services of other tenants. Fortinet products offer a number of technical mechanisms to ensure adequate isolation between tenants. This includes the use of security controls, virtual domains, administrative domains, role-based access control, resource limiting, and standard based network isolation such as IPSec and VLANs.

Managed Security Service Models

The actual implementation of managed security services may take many different forms but they often derivate from two common models or their combination, network based and customer premise equipment (CPE) based deployments. Figure 1 “MSSP Network Infrastructure” illustrates where in the provider infrastructure these managed security services are delivered.

Figure 1 – MSSP Network Infrastructure

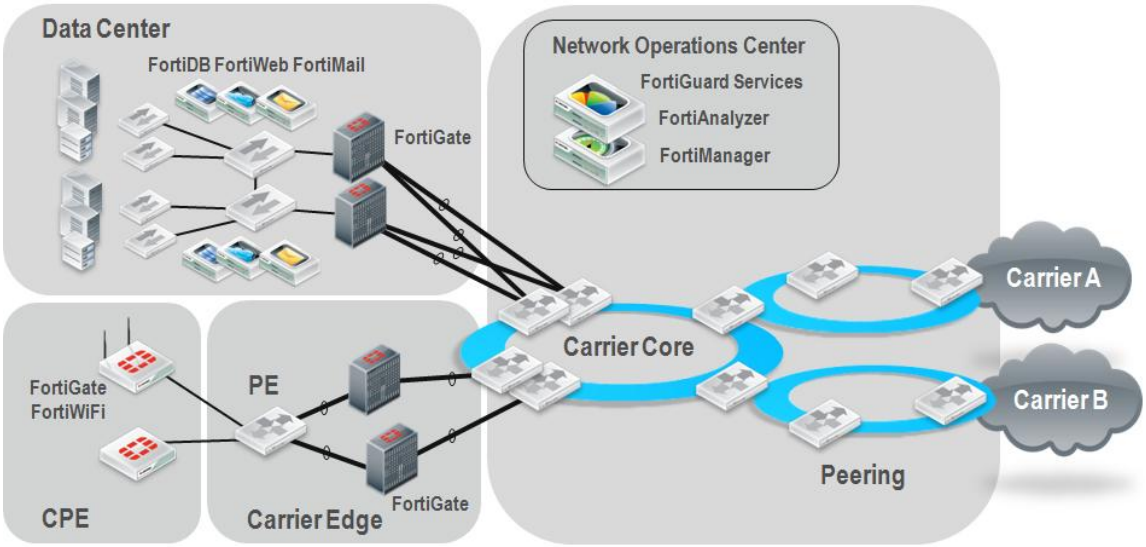

Network-based Services

The main feature of the network based model is that the equipment and services used to deliver the managed security services reside in the carrier's premises and remain under the provider's control. Depending of the type of service offered and the business model, these services may be deployed at various places in the carrier infrastructure such as data centers and the carrier subscriber edges.

Carriers providing hosting, private and public cloud services may offer managed security services at their data centers to protect their customers computing resources and data. Common security services offered include firewall, intrusion prevention, remote access VPN, and application security for web servers, email and databases. These data center services are traditionally delivered with the use of high performance appliances, and increasingly these days, with virtualized services that are part of a private or public cloud offering.

Carrier edge services typically consist of high performance security services deployed at the provider consumer edge with the objective to provide secure and clean access to large numbers of customers. Common services include firewall, web filtering, intrusion prevention, and denial of service (DoS) mitigation among others. These services are delivered by high performance appliances deployed at the consumer aggregation edge of the carrier Point of Presence (POP).

Figure 1 above gives an idea of the platforms and services implemented at both data center and carrier edge. Both network-based solutions are described in detail in this document.

CPE-based Services

In the CPE based model the service equipment is installed at the customer premises while it is monitored and managed by the provider. The objective of this service is to protect the computing resources and data residing at the customer site, and to provide secure connectivity between customer locations, and to external environments like the Internet. The customer premise equipment (CPE) may be owned or leased by the customer and consists of a physical appliance that delivers a number of services including routing, firewall, VPN, intrusion protection, and web filtering, to mention a few. The CPE devices may be deployed at several of the customer locations, such as main offices and branches. CPE equipment is chosen based on a number of factors, including the type of security services required, throughput and number of users.

Figure 1 “MSSP Network Infrastructure” illustrates how CPE services fit into the provider infrastructure and how they may relate to the carrier edge and data center based services.

Solutions for MSSPs

When it comes to deciding what managed security services to offer, there is a wide range of solutions providers may choose from. This section of the document provides a technical overview of some of the most common solutions, describing what Fortinet platforms and features are best for a given environment, and providing the design guidelines in accordance to best practices. To that end, data center and carrier edge network based solutions and CPE based solutions are discussed in detail.

Data Center Solutions

In terms of the services that can be offered, providers offering hosting services, private and public cloud services may enhance their service portfolios by adding high performance firewall, VPN and IPS, and by delivering application security services. Application security services may include messaging security with hosted antispam and antivirus, web and application firewall, and database security.

In addition to understanding the type of services that can be offered in a data center, there are other important aspects that should be considered like virtualization, multi-tenancy, centralized management and logging. Today most hosted services are provided in hybrid environments that combine physical appliances and virtualized services, therefore it is crucial to understand how to best integrate both physical and virtualized infrastructures. At the same time, as the provider infrastructure is likely shared across multiple customers, it is fundamental to ensure the appropriate segregation of resources and data between customers, making multi-tenancy a major aspect to be considered.

This section covers the following topics:

- High performance firewall, VPN and IPS
- Application security services
- Virtualization
- Multi-tenancy

High Performance Firewall, VPN and IPS

The high performance Firewall, VPN and IPS solution is designed for the most demanding environments, delivering the highest levels of throughput, number of concurrent connections, connections per second, and packets per second. The Firewall and IPS services deliver advanced protection to the customer servers and resources hosted by the provider, while VPN services can be implemented to provide secure access to both mobile users and selected customer locations.

The solution consists in the deployment of high performance FortiGate units at the distribution layer of the data center, creating a protection barrier between the carrier core and the access layer where the customer resources and servers reside. There are different ways the FortiGate unit can physically connect to the network; Figure 2 “High Performance Firewall, IPS and VPN” illustrates two common approaches, “on a stick” and “pass-through”. Multiple appliances can be implemented for high availability.

The “on a stick” approach, also known as “one arm”, consists in connecting multiple parallel links from the FortiGate unit to the aggregation switch. These physical links, often 1Gb or 10Gb Ethernet, may be bundled together in Link Aggregation Groups (LAGs). This provides great scalability and link redundancy. This approach

provides the best flexibility as services can be very easily inserted without the need to re-connect any of the physical links. It is just a matter of configuring VLANs and LAGs to manage how the traffic will flow in and out of the FortiGate unit. One important consideration in the "on a stick" design is that packets will likely traverse the switch multiple times as they flow between the external and internal networks. Consequently this design may introduce more latency and may be limited by the forwarding capacity of the aggregation switch, i.e. packets per second.

Figure 2 – High Performance Firewall, IPS and VPN

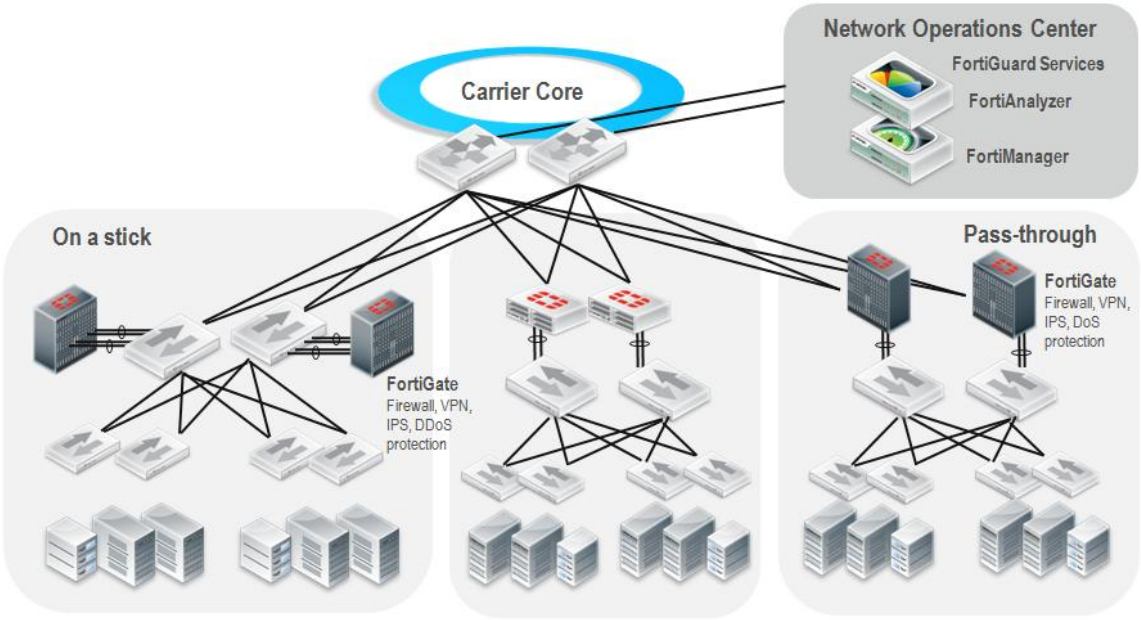

The "pass-through" design consists in placing the FortiGate unit inline between the core and distribution switches. Different physical links run between the FortiGate and the carrier core, and the access switches. The links, often 1Gb or 10Gb Ethernet, may also be bundled together in LAGs for better scalability and link redundancy. This design delivers better latency as packets for a given connection are processed once by the aggregation switches. In this design, the insertion of the FortiGate will more likely require adjusting the physical links between the core and aggregation layers.

When configuring LAGs in platforms that offer hardware acceleration, it is important to understand what ports to use to preserve traffic acceleration. In certain platforms, traffic over a group of links served by different network processors (NPs) may be handled in software, instead of being accelerated in hardware. The general best practice is to bundle ports served by the same NP. Please consult the hardware documentation of the platform used for more information. Later this document presents several recommended deployment options for both "on a stick" and "pass-through" designs. The proposed data center designs deliver high availability by combining path redundancy, device failover, link backup, and other mechanisms available on FortiGate platforms. Multiple device failover options are discussed throughout the document including clustering, stateless failover, stateful active-passive and active-active failover. These failover implementations allow the provider to achieve N+1 redundancy by adding one or more FortiGate appliances or modules as backup, operating in either stand-by or active mode.

As illustrated in Figure 2, data center based services are centrally managed and monitor by the MSSP Network Operations Center (NOC) or Security Operations Center (SOC). The NOC/SOC is responsible for provisioning services, monitoring of UTM devices, and the proactive alerting on conditions that may require special attention,

such as link failures and security incidents. FortiGuard services, FortiManager and FortiAnalyzer are fundamental components that enhance the NOC/SOC operation.

The FortiGuard Subscription Services keep the UTM devices up to date with the latest antivirus, antispam, IPS, and Web Filtering definitions, protecting the MSSP subscribers against the latest content and network level threats. The FortiManager platform integrates seamlessly with the FortiAnalyzer products delivering a single point of command, control, analysis, and reporting. The FortiManager platform provides centralized policy-based provisioning, configuration, and update management for thousands of FortiGate, and FortiMail appliances, as well as FortiClient security agents. The FortiAnalyzer product family provides network-wide visibility by aggregating log and event data from large numbers of FortiGate and FortiMail security appliances. The "Centralized Management and Reporting" chapter provides more details on these important operational aspects.

Fortinet offers a number of FortiGate platforms that are a good fit for this type of data center based services. While other platforms may be a good alternative, this document describes the deployments using the FortiGate-5000 and the FortiGate-3950B product families.

There are different ways the FortiGate platforms can be implemented. This document describes the following deployment options:

- ELBC (FortiGate-5000 only)
- FortiSwitch-5203B (FortiGate-5000 only)
- FGCP
- External load balancer
- 802.3ad Link Aggregation

Enhanced Load Balance Cluster (ELBC)

ELBC is a load balancing feature available on the FortiGate-5140 Chassis equipped with a FortiSwitch-5003A or FortiSwitch-5003B and up to twelve (12) FortiGate-5001A or FortiGate-5001B modules acting as worker blades. ELBC balances the traffic across the FortiGate-5001A/B worker blades in the cluster, providing the highest levels of scalability for firewall, IPS, antivirus and web filtering.

ELBC is the solution that delivers the highest levels of throughput, concurrent connections, connections per second, and packets per second. As an example, a FortiGate-5140 Chassis equipped with FS-5003B and twelve FG-5001B blades may deliver 160Gbps of stateful firewall, 120 million concurrent connections, and 1.4 million of connections per second. ELBC is ideal for those providers looking for high performance firewall and IPS solution for hosting, public and private cloud offerings. It should be noted that some functions like VPN, multicast, certain complex NAT scenarios, and dynamic routing may not be supported depending on how ELBC and the FortiGate blades are configured.

With ELBC, all packets enter and leave the system throughout the FortiSwitch-5003A/B front panel ports, and not the ports of the FortiGate-5001A/B blades. As packets enter the system, the FortiSwitch blade applies a load balancing algorithm to determine to which FortiGate blade to forward the packet. The load balancing algorithm calculates a hash key value based on the source and destination addresses of the packet. Each FortiGate blade servers a hash key value assigned to it. The load balancing algorithm ensures that traffic for the same source and destination address pair is served by the same FortiGate blade in both directions. ELBC provides high availability by monitoring the health of the blades, and by detecting failures and automatically redistributing traffic to the remaining worker blades.

In a FortiGate-5140 Chassis configured with ELBC, the FG-5001A/B worker blades can operate in Transparent (Layer 2) or NAT/Route (Layer 3) mode, and a single chassis can support mixed modes across VDOMs. ELBC can also be deployed "on a stick" or "pass-through".

Figure 3 "ELBC Transparent Mode at the Data Center" shows one of the most recommended deployments for ELBC. If distribution and core switches are Layer 3 capable, the FortiGate blades in the ELBC cluster may be configured in transparent mode (Layer 2). Configuring the cluster in transparent mode has multiple benefits. To start with, the FortiGate blades integrate seamlessly into the network, without requiring a network redesign or change of IP addresses. Better path redundancy can also be achieved by implementing dynamic routing protocols between the core and distribution layers. Any routing protocol supported by the core and distribution platforms may be used with an ELBC layer 2 cluster, while a cluster in NAT/Route mode supports BGP and static routing. Additionally, a layer 2 cluster allows the forwarding of non-IP packets. As an ELBC cluster in NAT/Route mode acts as a layer 3 router, it does not forward non-IP packets.

Figure 3 – ELBC Transparent Mode at the Data Center

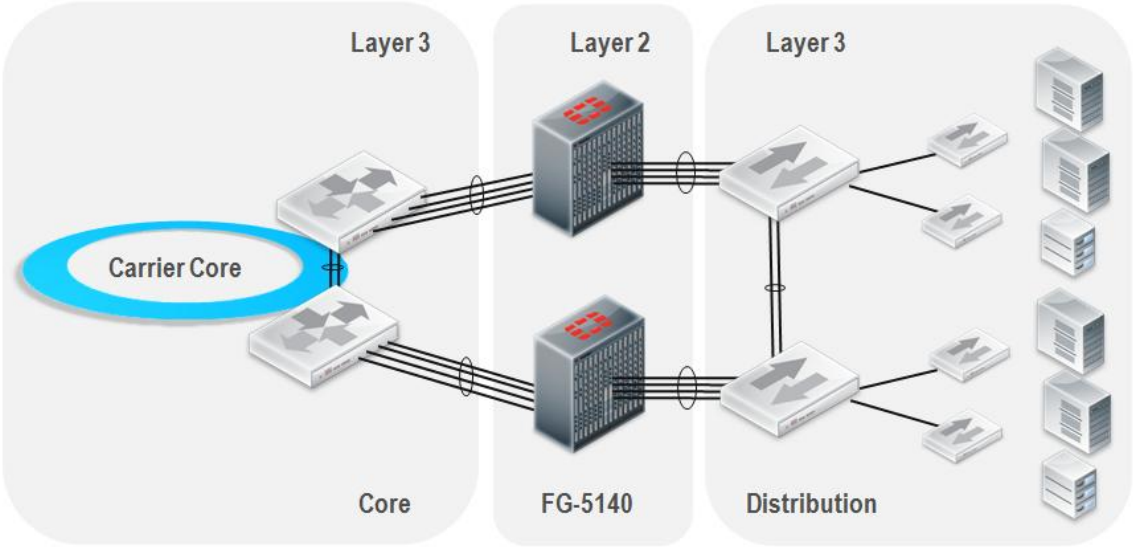

In this design redundancy may be implemented at multiple levels. Figure 3 shows two parallel paths, each with its own FG-5140 chassis. Path selection can be determined by the dynamic routing protocols running between the core and distribution layers. As the FortiGate cluster is configured in transparent mode, it does not participate in the routing decisions. The FG-5140 chassis can also be deployed with dual FortiSwitch-5003A/B blades, providing intra-chassis redundancy. Finally, the actual physical links can be grouped in Link Aggregation Groups (LAGs).

The example shown in Figure 3 uses a 4-Port-LAG, where four fabric ports are aggregated for the internal interface, and other four ports for the external interface, giving 40GB of bandwidth to each. The 4-Port-LAG configuration allows for a single internal/external interface pair, or service group. Other configuration options are available such as 2-Port-LAG and basic mode, the former allowing two service groups with 20GB per aggregated interface, and the later delivering four service groups with 10GB of bandwidth per interface.

FortiSwitch-5203B

The FortiSwitch-5203B module provides load balancing capabilities for the FortiGate-5140B and FortiGate-5060 Chassis equipped with one or more FortiGate-5001B modules acting as worker blades in an HA over chassis (HAOC) cluster. The FortiSwitch-5203B runs FortiOS code and is able to deliver firewall and VPN services, while load balancing the UTM functions across the FortiGate-5001B worker blades and the FortiSwitch-5203B itself. In

addition to the FortiGate-5140B chassis and FortiGate-5060 chassis, FortiSwitch-5203B boards may be installed in selected versions of the NEBS-compliant FortiGate-5140-R chassis.

The FortiSwitch-5203B is a good option for providers looking for a solution that delivers high performance for firewall, VPN and UTM functions, and that require dynamic routing, multicast, and complex NAT deployments not supported by ELBC. In terms of capacity the FortiSwitch-5203B solution performance is less than ELBC but more than FGCP (later discussed).

The deployment of the FortiSwitch-5203B requires the configuration of the HA over chassis mode on the board itself and the FortiGate-5001B blades. In this mode, the FortiSwitch-5203B operates as the primary unit, and it is the only responsible for load balancing the UTM sessions across the blades in the cluster. At this time, HA over chassis is supported for blades installed in the same chassis only.

Like ELBC, in an HAOC cluster all packets enter and leave the system throughout the FortiSwitch-5203B front panel ports, and not the ports of the FortiGate-5001B blades. Unlike ELBC, session setup, firewall, stateful inspection, VPN, and session helpers are handled by the FortiSwitch-5203B itself, and are not load balanced across the worker blades. UTM functions like IPS and antivirus are balanced across the various FortiGate-5001B blades and the FortiSwitch-5203B itself. The FortiSwitch-5203B uses a weighted load balancing algorithm to decide where to forward an UTM session. It is possible to adjust weights to influence load balancing, for instance, setting the weight of the primary unit (always the FortiSwitch-5203B) to 0 will ensure all UTM sessions are load balanced to the worker blades, leaving the FortiSwitch-5203B dedicated to session setup and traffic load balancing. High availability is provided by monitoring the health of the blades in the cluster, and by detecting failures and automatically redistributing traffic to the remaining worker blades. In an HAOC cluster, only a FS-5203B blade can become the primary unit.

As illustrated in Figure 4 – “FS-5203B Deployment at the Data Center”, in the FortiSwitch-5203B solution the FG-5001B worker blades can operate in Transparent (Layer 2) or NAT/Route (Layer 3) mode, and a single chassis can support mixed modes across VDOMs. This solution can also be deployed “on a stick” or “pass-through”.

Figure 4 – FS-5203B Deployment at the Data Center

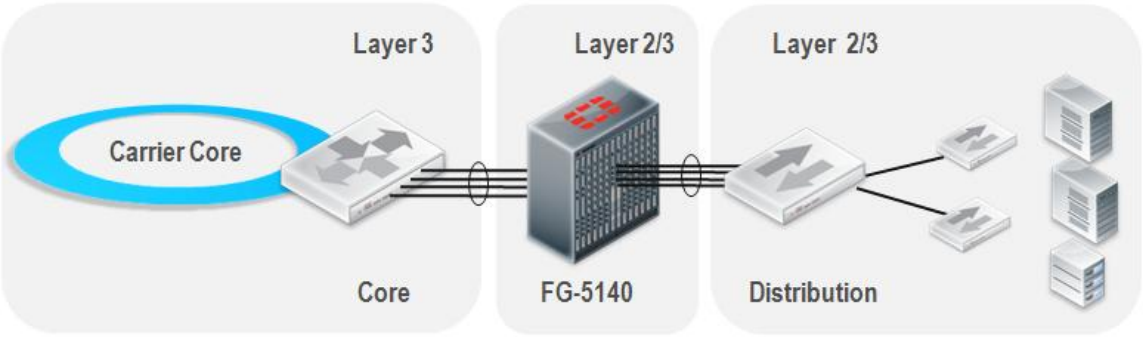

The FortiSwitch-5203B solution provides greater deployment flexibility compared with ELBC. In case the distribution switches are Layer 2 capable only, the FG-5001B worker blades can be configured in NAT/Route mode, allowing them to participate in dynamic routing. The FortiSwitch-5203B solution also supports VRRP, facilitating multipath redundancy in that scenario. The FG-5001B worker blades can also be deployed in transparent mode, similarly to the example provided for ELBC, which may facilitate deployment in the case the distribution switches are Layer 3 capable.

When deciding between NAT/Route and Transparent, it is important to understand that while both modes have essentially the same feature set, due to its nature some features are not available in transparent mode. DHCP, dynamic routing, SSL VPN, virtual IP are some of the features not available in transparent mode.

Multiple levels of redundancy may be implemented. Two parallel paths can be implemented, each one with its own FG-5000 chassis. In that scenario path selection can be determined by a dynamic routing protocol. Optionally, two FS-5203B blades and multiple FortiGate worker blades may be deployed in the same FG-5000 chassis for intra-chassis redundancy. In this case, the primary FS-5203B blade can also load balance sessions to the backup FS-5203B blade as well as the worker blades. Finally, an additional level of redundancy may be achieved by grouping physical links in Link Aggregation Groups (LAGs).

FortiGate Clustering Protocol (FGCP)

FGCP is a Fortinet proprietary solution for high availability and UTM load balancing. With FGCP, two or more identical FortiGate units are configured as members of a High Availability (HA) cluster. All FortiGate units in the cluster must be the same model with the same modules installed and must run the same FortiOS firmware build. The FGCP cluster behaves as a single FortiGate unit operating in NAT/Route or Transparent mode. The units in the cluster share state and configuration information using an HA heartbeat mechanism. In the event one unit fails, the other units in the cluster automatically detect the failure and take over the functions handled by the failed unit without service disruption.

FGCP is ideal for providers looking for a failover solution available to all FortiGate models that integrate seamlessly into the network; that provides enhanced scalability for resource-intensive UTM processing and without sacrificing any functionality available in FortiOS. FGCP is available to both FortiGate blades and appliances.

FGCP can be configured in active-passive (A-P) or active-active (A-A) HA mode, the later providing increased performance with UTM load balancing.

FGCP active-passive HA provides host standby failover protection and link failover. An active-passive HA cluster consists of a primary unit that processes communication sessions, and one or more subordinate units in standby. At any given time, only the primary unit processes communication sessions. The subordinate units are responsible for monitoring the health of the primary unit and making sure their configurations are synchronized with the one of the primary unit. In case the primary unit fails, FGCP automatically elects a new primary unit from the pool of subordinate units taking over all the functions handled by the failed unit. In terms of performance, an active-passive HA cluster can deliver up to the capacity of the primary unit.

FGCP active-active HA provides the same failover capabilities as active-passive, plus the ability to load balance resource-intensive UTM functions across the units in the cluster. Protocol recognition, virus scanning, IPS, web filtering, email filtering, data leak prevention (DLP), application control, and VoIP content scanning and protection are examples of UTM processing load balanced. By load balancing the resource-intensive UTM processing among all cluster units, an active-active HA cluster may provide better UTM performance than a standalone FortiGate unit or than an active-passive HA cluster. All non-UTM sessions are not load balanced and are processed by the primary unit. UDP, ICMP, multicast, and broadcast sessions are never load balanced and are always processed by the primary unit. VoIP, IM, P2P, IPSec VPN, HTTPS, SSL VPN, HTTP multiplexing, SSL offloading, WAN optimization, explicit web proxy, and WCCP sessions are also always processed only by the primary unit. An active-active HA cluster consists of a primary unit that receives all communication sessions and load balances them among the primary unit and all of the subordinate units. In an active-active cluster the subordinate units are also considered active since they also process UTM sessions. In terms of performance, an active-active HA

cluster can deliver up to the capacity of the primary unit for firewall and non-UTM traffic, but it may provide better performance than a standalone unit for UTM traffic.

By default, FGCP does not maintain an HA session table and after a failover most TCP sessions do not resume and must be restarted as the cluster renegotiates. If session failover is required, FGCP may be configured with session pick-up. Session pickup ensures that non-UTM sessions are picked up by the new primary unit after a failover. It should be noted however that enabling session pick-up generates more heartbeat traffic as a larger portion of the session table must be synchronized. Consider enabling session pickup delay to improve performance by reducing the number of sessions that are synchronized.

Figure 5 – "FGCP Active-Active HA Deployment at the Data Center" illustrates two active-active HA scenarios, the one on the left forms an HA cluster with four FortiGate blades installed in two FG-5000 chassis, and the one of the right with an HA cluster of four FortiGate appliances like the FG-3950B.

In addition to the FGCP requirement of same hardware and software, members of a cluster must also have identical connections to the network. This means, units in the cluster must have the same interfaces connected to the same network segments, making them Layer 2 adjacent. In Figure 5 the internal and the external interfaces of each unit in the two clusters converge into two separate switches, making sure every unit is Layer 2 adjacent to the other cluster members.

Figure 5 – FGCP Active-Active HA Deployment at the Data Center

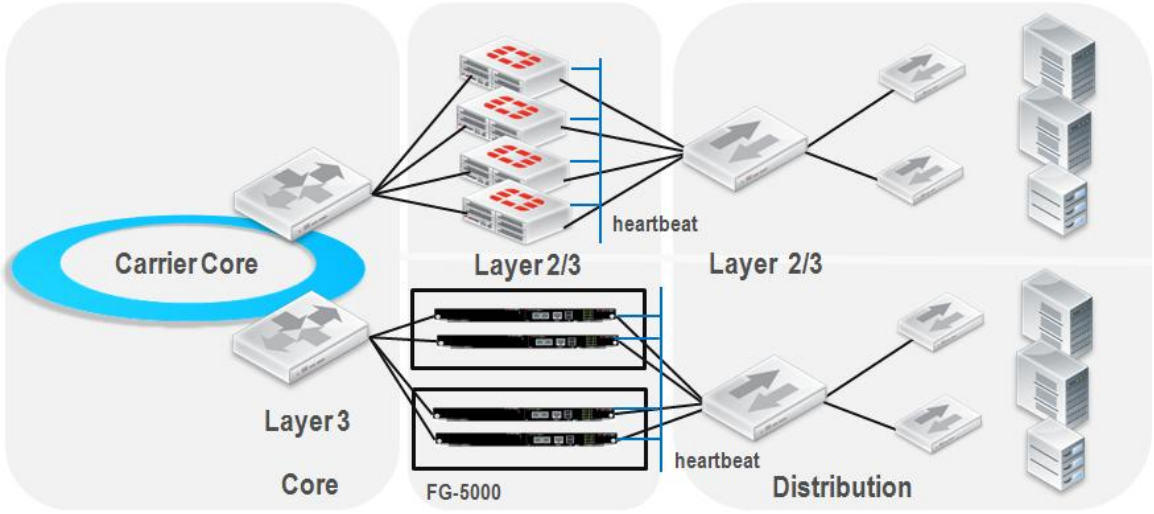

Another best practice is to implement multiple dedicated heartbeat interfaces. HA heartbeat communication is fundamental to the operation of the cluster; it is the primary mechanism used to monitor the health and status of cluster and to ensure configuration synchronization. If heartbeat communication fails, units in the cluster become unaware of each other, resulting in a Split Brain condition where all units take the primary role and disrupting communication. In addition, heartbeat packets contain sensitive cluster configuration information and can consume a considerable amount of network bandwidth. Thus it is highly recommended to isolate the heartbeat interfaces from the user data interfaces, and to use multiple heartbeat interfaces for backup. This can be done by using two or more dedicated interfaces and connecting them to a separate switch, or using a dedicated heartbeat VLAN. If using blades like the FortiGate-5001B, use the base backplane interfaces to handle the HA heartbeat communication through the chassis backplane, leaving the front ports of the blades available for network connections. It is also highly recommended to ensure links used for heartbeat communication do not exceed 10ms

of latency; higher latency may impede synchronization. HA heartbeat packets are non-TCP packets that use Ethertype values 0x8890, 0x8891, and 0x8890. The FGCP uses link-local IP4 addresses in the 169.254.0.x range for HA heartbeat interface IP addresses.

With respect to the FGCP integration into the network, the fact a FGCP cluster functions as a single appliance greatly simplifying its deployment, independently the distribution switches are configured as Layer 3 or Layer 2 devices. In addition, whether the FGCP cluster is configured in active-passive or active-active, NAT/Route or Transparent mode, under normal operation the FGCP cluster always provides a single path. This means a FGCP cluster in NAT/Route mode provides a single Layer 3 path for static or dynamic routing. Similarly, when configured in transparent mode, the FGCP cluster provides a single Layer 2 path eliminating the possibility of Layer 2 loops.

When deciding between NAT/Route and Transparent, it is important to understand that while both modes have essentially the same feature set, due to its nature some features are not available in transparent mode. DHCP, dynamic routing, SSL VPN, virtual IP are some of the features not available in transparent mode.

While implementing FGCP is quite transparent to most network environments, some configuration changes may be required to the surrounding switches. The following are some of the well known situations:

- FGCP relies on gratuitous ARP packets to refresh the MAC forwarding tables of adjacent switches after a failover. Switches configured to not update their forwarding tables in response to gratuitous ARP packets may not properly re-direct traffic to the new primary unit.
- Configuration changes may be required when connecting an active-active HA cluster to switches configured with the spanning tree protocol. For an active-active HA cluster to be compatible with the spanning tree algorithm, the FGCP requires that the sum of maximum age and forward delay should be less than 20 seconds. The maximum age and forward delay settings are designed to prevent layer 2 loops. If there is no possibility of layer 2 loops in the network, you could reduce the forward delay to the minimum value.
- In an active-active HA cluster the load balancing of UTM traffic requires the primary unit in the cluster to re-directs traffic into the subordinate units. As a result, data packets with the same source and destination IP addresses may be received by different ports in the same switch. Switches configured with anti-spoofing mechanisms may drop these packets.

As with the previous designs presented, multiple levels of redundancy may be implemented to extend the high availability of an environment using FGCP. Two parallel paths can be implemented, each one with its own FGCP cluster. In that scenario path selection can be determined by a dynamic routing protocol. FGCP is also compatible with LACP (802.3ad), therefore the actual physical links connecting the cluster units can be grouped in Link Aggregation Groups (LAGs) for higher performance and better reliability.

Finally, FGCP can also be configured with VDOMs, which is called Virtual Clustering. This feature is an extension of the FGCP for a cluster of two FortiGate units operating with multiple VDOMS enabled. Virtual clustering operates in active-passive mode to provide failover protection between two instances of a VDOM operating on two different cluster units. You can also operate virtual clustering in active-active mode to use HA load balancing to load balance sessions between cluster units. Alternatively, by distributing VDOM processing between the two cluster units you can also configure virtual clustering to provide load balancing by distributing sessions for different VDOMs to each cluster unit.

For more information on the FGCP please see the High Availability chapter of the FortiOS handbook: http://docs.fortinet.com/fgt.html

External Load Balancer

A group of FortiGate devices can also be load balanced by an external load balancing device like the FortiBalancer 2000, delivering better scalability and redundancy than a standalone FortiGate appliance. Using an external load balancer provides linear scalability as higher performance can be achieved by adding more FortiGate units to the group. However, it should be noted that the overall scalability of the environment is limited by the capacity of the load balancer in terms of factors such as maximum throughput, maximum number of concurrent connections, session rate, and packets per second. External load balancers can be used with both FortiGate blades and appliances.

Figure 6- "FortiGate appliances with an external load balancer" illustrates the scenario where multiple FortiGate appliances are load balanced by a pair of FortiBalancer 2000 appliances. The FortiGate appliances can be deployed in either NAT/Routed or Transparent mode, and depending on the load balancer design this solution can follow the "on a stick" or "pass-through" models.

When deciding between NAT/Route and Transparent, it is important to understand that while both modes have essentially the same feature set, due to its nature some features are not available in transparent mode. DHCP, dynamic routing, SSL VPN, virtual IP are some of the features not available in transparent mode.

Figure 6 – FortiGate appliances with an external load balancer

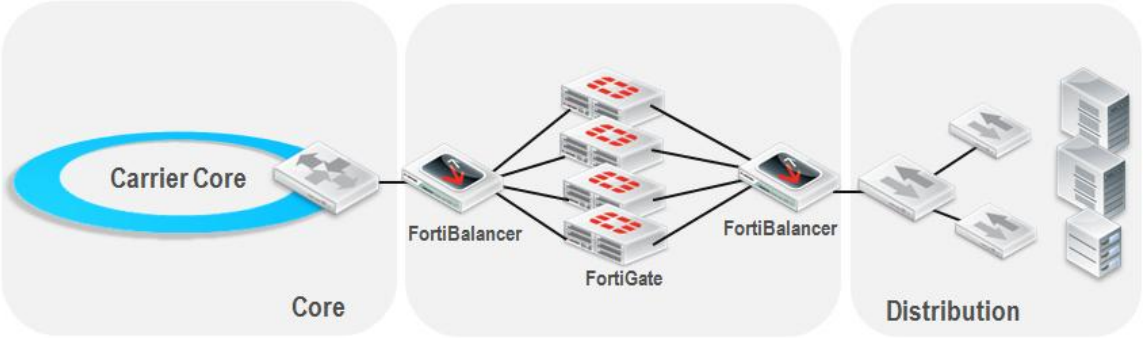

Simply put, the load balancers are responsible for balancing the traffic load across the FortiGates but while maintaining session symmetry. The simple rule is that traffic for a given session needs to be served by the same FortiGate in both directions. For certain protocols a single session may involve multiple separate data and control flows. For appropriate firewall and UTM processing the load balancers must ensure that all flows of a given session are forwarded to the same FortiGate unit. One solution is to configure the load balancers to ensure traffic coming from the same client IP address is predictably forwarded to the same FortiGate unit, a feature known as sticky load balancing.

External load balancers like the FortiBalancer 2000 provide a number of different algorithms that can be used to load balance traffic. Common examples are load balancing based on round robin, weight, and usage. Load balancers decide how to balance traffic by inspecting the packet headers, and in some cases, part of the payload. Depending on the environment and configuration, the load balancer may look at the source and destination MAC addresses, source and destination IP addresses, IP protocol, source and destination UDP and TCP ports. This information is also useful for recognizing flows that belong to the same session.

Since load balancing decisions can be made based on IP address information, special care should be taken when implementing NAT on the FortiGate appliances. If the load balancer algorithm is configured to use both source and destination IP addresses, NAT on the FortiGate appliances will change those addresses likely breaking traffic symmetry. As the addresses change, there is no guarantee that the load balancers at both sides of the FortiGate

units will select the same appliance to serve traffic in both directions for a given session. Fortunately, there are several solutions to this challenge. In most network environments NAT can be configured on the exterior load balancers instead of the FortiGates. This is done by defining one or more Virtual IP (VIP) addresses that represent the group of load balanced appliances. Another solution is to change the load balance algorithm to base its decisions on the IP addresses not affected by NAT. In most data centers only the server IP addresses need to be translated. If that is the case, the load balancing algorithm could be configured to only use the client source IP address. In the scenario depicted in Figure 6 the exterior load balancer would be configured to make its decisions based on the packet source IP address, while the inner load balancer would be configured to use the destination IP address.

Load balancers may also offer several options for handling client IP addresses. A FortiBalancer configured in Reverse Proxy Mode proxies all client connections using its own interface IP address, hiding the client IP addresses. Per contrary, a FortiBalancer configured in Transparent Mode proxies all client connections preserving their original IP addresses. Transparent mode is more desirable for our scenario as the preservation of the client IP addresses facilitates the enforcement of granular IP address based policies on the FortiGate appliances.

Load balancers also provide redundancy. There are many different mechanisms used to check and monitor the health of the units in the load balancing group, but one most common mechanisms is the use of configurable probes with ICMP packets. When a unit fails to respond to probes after a number of attempts, it is assumed to be out of service and it is automatically removed from the group to ensure new sessions are only served by the remaining healthy units. The sessions handled by the failed unit are also redistributed across the remaining units. Some load balancers may recalculate all indexes after a failure, redirecting even those sessions handled by healthy units, and as a result disrupting the user connections. It is desirable that in a failure situation, only new sessions and those served by the failed unit get redistributed. Already established sessions should continue to be served by the same units as long as they don't fail.

FortiGate appliances support TCP session synchronization which ensures session tables are synchronized across FortiGate units. This feature assures that no loss of data occurs during a failover. The TCP session tables are continuously replicated across the FortiGate units. If one of the FortiGate units fails, traffic continues to flow as long as the load balancer redirects the traffic to the units that are still operating. Due to their connectionless nature, UDP and ICMP sessions do not need to be synchronized to failover naturally, therefore only TCP session tables are replicated. Unlike FGCP, TCP synchronization does not include configuration synchronization and does not provide session failover. Session failover is assumed to be done by an external device such as a load balancer. In addition, sessions subject to UTM processing or NAT are not synchronized as they are bound to the FortiGate unit, and units in the cluster may differ in their system and IP address configuration.

For more information on the FortiBalancer please see:
http://www.fortinet.com/products/fortibalancer/

IEEE 802.3ad Link Aggregation

The IEEE 802.3ad standard defines a method of aggregating multiple Ethernet links into a single point-to-point logical channel, delivering increased throughput and providing link redundancy. This type of link aggregation technology can be used to balance the traffic load across a group of FortiGate modules or appliances, delivering linear scalability and link redundancy.

IEEE 802.3ad is an attractive design option for those providers looking for a low latency and high throughput solution, and who do not have very stringent requirements for failure detection. This solution requires the deployment of a switch or router configured with link aggregation. As illustrated in Figure 7 –"FortiGate balancing

with 802.3ad Link Aggregation at the Data Center", the solution consists in connecting a group of FortiGate appliances to the same Link Aggregation Group (LAG) at the distribution switches. Simply put, the switch uses a traffic balancing algorithm to predictably distribute the traffic load across the physical links in the LAG, each one connecting to a separate FortiGate appliance. As with any load balancing solution, firewall and UTM processing requires that flows of a given session are forwarded to the same FortiGate unit. This solution provides some of the highest levels of throughput with minimum latency, and it is mainly limited by the switch capacity.

IEEE 802.3ad requires that all links in the LAG are of the same type and operate at the same speed, and only full duplex is supported. Links can be configured as either Layer 2 or Layer 3 ports.

Figure 7 – FortiGate balancing with 802.3ad Link Aggregation at the Data Center

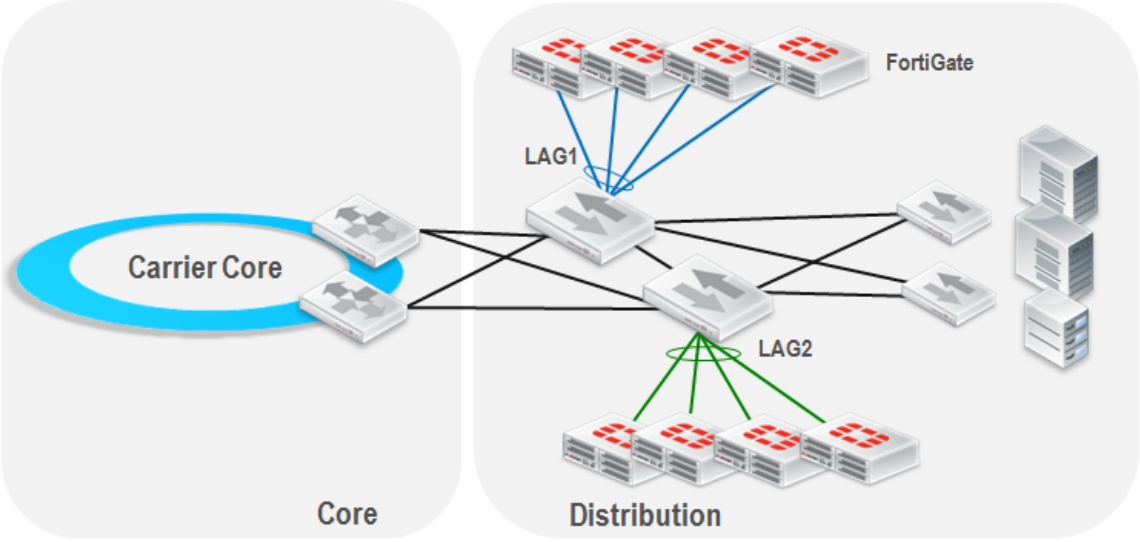

While IEEE 802.3ad is a standard, switch implementations may vary, so it is important to familiarize with the actual functionality supported and as well as the default settings. For instance, most switch implementations allow up to eight physical links per LAG, while others may support fewer or more links. Switches from different vendors also often differ on the default algorithm used for load balancing.

Unlike typical IEEE 802.3ad deployments where both ends are configured, likely dynamically using the Link Aggregate Control Protocol (LACP); in our data center design link aggregation only needs to be configured at the switch side. In fact, the FortiGate units are not aware their links connect to a LAG. For this reason, configuration on the switch is likely to be manual, and not by using LACP.

Figure 7 shows a redundant scenario with two groups of FortiGate appliances, each connected to the LAG configured on the corresponding switch. In this environment, the two groups work independently from each other, and are only shown here to illustrate a secondary path. It is important to note that each FortiGate appliance connects using a single physical link, and that the links of all the appliances converge into a single LAG. Traffic within those links is separated by the use of VLANs, for example one VLAN for internal and another VLAN for the external network of the firewall. A single LAG is used to ensure both directions of traffic are directed to the same FortiGate appliance in the group. Certain switches calculate different hashes per LAG, so traffic symmetry cannot be guaranteed if using more than one LAG. For the same reason this solution is likely to work better on "on a stick" deployment than a "pass-through".

To distribute traffic switches use algorithms that inspect the packet headers. Decisions are based on source and destination MAC addresses, and/or source and destination IP addresses. Some switches may have the ability to use TCP and UDP ports in the calculation of the distribution hashes. In our scenario decisions should be based on source and destination IP addresses only. This is to ensure all traffic flows for a given session are consistently directed to the same FortiGate appliance.

Link aggregation operates above the MAC layer in the Data Link OSI layer, therefore the FortiGate appliances may operate in either NAT/Routed (Layer 3) or Transparent (Layer 2) mode. However, special care should be taken when implementing NAT on the FortiGate appliances. Load balancing decisions are made based on IP address information, consequently any address translation will likely break traffic symmetry. As IP addresses change, there is no guarantee the distribution algorithm will select the same physical link in both directions for a given session. Possible solutions are to implement NAT at a different layer, or to modify the algorithm to base its decisions on the IP addresses not affected by NAT. Another alternative is to configure the FortiGate units in transparent mode when NAT is required.

When deciding between NAT/Route and Transparent, it is important to understand that while both modes have essentially the same feature set, due to its nature some features are not available in transparent mode. DHCP, dynamic routing, SSL VPN, virtual IP are some of the features not available in transparent mode.

With respect to high availability, a solution based on link aggregation is limited to link level redundancy. When a physical link fails it is automatically removed from the LAG and traffic gets redirected over the remaining links. For most switch implementations a link failure will trigger a re-calculation of the forwarding hashes over the remaining links. As a result, existing connections are likely to be forwarded over new links, forcing session resets. The solution does neither provide any mechanism to track the state and health of the units in the group. Should a FortiGate unit fail without bringing its physical interfaces down; unable to detect the failure, the switch will continue forwarding packets to the failed appliance, effectively blackholing traffic. Another limitation is that a link aggregation based solution does not include configuration synchronization and does not provide session failover.

To improve reliability, the FortiGate appliances can be configured with TCP session synchronization to ensure no data is lost in the event of a unit failover. With TCP session synchronization enabled, the TCP session tables are continuously replicated across the FortiGate units. If one of the FortiGate units fails, traffic continues to flow as long as the switch redirects packets to the units that are still operating. As its name indicates, TCP session synchronization replicates TCP session tables. Due to their connectionless nature, UDP and ICMP sessions do not need to be synchronized to failover naturally. In addition, sessions subject to UTM processing or NAT are not synchronized as they are bound to the FortiGate unit, and units in the cluster may differ in their system and IP address configuration.

For more information on the configuration of Link Aggregation, TCP session synchronization and other FortiGate features, please see the FortiOS handbook:
http://docs.fortinet.com/fgt.html

Application Security Services

Email, web servers, and databases are common services customers trust to their hosting and cloud providers. Providers may enhance their service offerings by designing and delivering managed application security services specially architected to protect the customer applications and data. This chapter provides describes the Fortinet products available and the design recommendations a provider can follow to successfully deliver messaging security with hosted antispam and antivirus, web and application firewall, and database security.

The following Fortinet products are used:

- FortiGate
- FortiMail
- FortiWeb
- FortiDB

Secure Messaging

Messaging security is provided with FortiMail, a purpose-built platform that delivers high-performance mail routing, antispam filtering, antivirus and antispyware protection. FortiMail uses a multi-layer security technology to deliver bi-directional filtering. Inbound filtering blocks spam and malware before it affects the customer services and users, while outbound inspection is used to prevent outbound spam and malware from causing other antispam gateways to blacklist the customer users. In addition, FortiMail offers S/MIME, TLS and Identity-Base Encryption (IBE) email encryption.

Depending on the security requirements and the level of integration desired, providers may choose from three deployment options:

- **Gateway Mode:** Provides inbound and outbound proxy mail transfer agent (MTA) services for email gateways. Requires a DNS MX record change to redirect email to the FortiMail device for antispam and antivirus scanning. The FortiMail unit receives messages, scans for viruses and spam, then relays email to its destination email server for delivery.
- **Transparent Mode:** The FortiMail device is inserted in the path of email without requiring any changes to the existing network configuration. Each network interface of the FortiMail device includes a proxy that receives and relays email. Each proxy can intercept SMTP sessions even though the destination IP is not the FortiMail device. FortiMail scans for viruses and spam, and then transmits email to destination email server for delivery.
- **Server Mode:** The FortiMail device acts as a stand-alone messaging server and delivers full-featured SMTP mail server functionality, with support for secure POP3, IMAP and WebMail access. It receives messages, scans for viruses and spam, and then delivers to its users' mailboxes. External MTAs connect to the FortiMail server, as in server mode, the FortiMail device functions as a protected server itself.

FortiMail is available in appliance form factor with the FortiMail-2000B and FortiMail-3000C appliances, the FortiMail-5001A module for the FG-5000 series chassis, and the FortiMail-VM virtual appliance.

Figure 8 – Secure Messaging with FortiMail

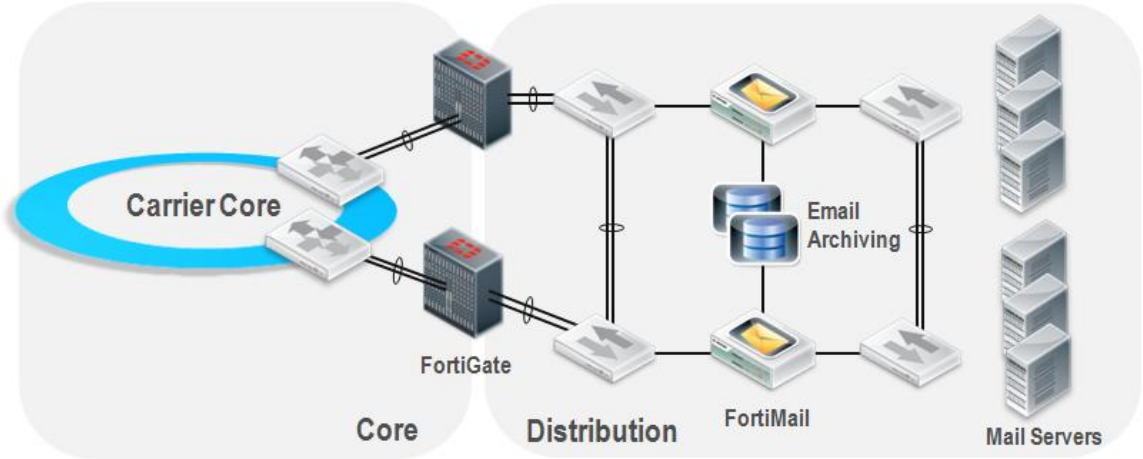

Figure 8 – "Secure Messaging with FortiMail" illustrates the deployment of FortiMail in transparent mode. While the three deployment options may be considered, transparent mode integrates into the carrier network without requiring any changes to the network and without the need of DNS MX record changes. The FortiMail is placed in front of the existing email servers in a similar manner as to a bridge deployment, where no IP address changes are required.

The proposed design delivers multi-layer protection. A first level is provided by a set of redundant FortiGate appliances or chassis with the primary responsibility of providing firewall, denial of service (DoS) protection, and IPS services. The FortiGate units can be implemented using any of the methods described in the previous section, "High Performance Firewall, VPN and IPS". A pair of redundant FortiMail appliances (or modules) is deployed to deliver enhanced antispam and antimalware.

While the FortiGate has messaging security capabilities, the FortiMail platform supports over sixteen different spam detection methods not available on the FortiGate platform, including Bayesian filters, SHASH checksum and spam image analysis. It also supports full Mail Transfer Agent (MTA) features and can perform user-based antispam rules.

As shown in Figure 8, the FortiMail supports a high availability configuration that offers full synchronization of configuration and mail data between two FortiMail systems to ensure maximum availability of email services.

For more information on FortiMail please see:
http://www.fortinet.com/products/fortimail/

Secure Web and Database services

Hosting and cloud providers may combine the FortiWeb and FortiDB platforms to protect the customer web-applications and databases from attacks and data loss.

The FortiWeb is a purpose-built web application firewall that provides layered protection for web-based applications. FortiWeb protects the customer web-applications from application-based attacks like SQL injection and cross-site scripting. In addition, the FortiWeb delivers application acceleration, file compression, SSL hardware offloading, denial of service (DoS) protection, and advanced load balancing. FortiWeb is available in appliance form factor, including the FortiWeb-3000C, and as a virtual appliance, FortiWeb-VM.

FortiWeb appliances can be deployed in four different modes:

- **Inline Transparent:** The FortiWeb appliance is deployed transparently as a layer 2 bridge without requiring a network level redesign. This mode delivers full enforcement.
- **True Transparent Proxy:** layer 2 deployment with no need for network level redesign. The traffic is internally terminated to provide more functionality that pure inspection.
- **Reverse proxy:** Provides additional capabilities such as URL rewrite and advanced routing capabilities. Supports HTTPS offloading and enhanced load balancing.
- **Offline Sniffing:** Monitors environments with zero network footprint and latency. Non-intrusive deployment ideal for initial product evaluations.

Figure 9 – "Secure Web and Database services with FortiWeb and FortiDB" illustrates the deployment of a redundant pair of FortiWeb appliances in inline transparent mode. The diagram presents a multi-layer protection design, where a pair of redundant FortiGate appliances or chassis provide the first level of protection, delivering firewall and IPS services. The FortiWeb appliances provide specialized protection for the customer web servers

against application based threats. The database servers are protected with a second layer of FortiGate appliances or chassis and a pair of FortiDB appliances that provide specialized protection.

The FortiDB product family delivers database and application security, with centralized policy configuration and enforcement, audit policy compliance, and vulnerability management. Providers may leverage the FortiDB platform to protect customer confidential data such as credit card numbers, and help maintain compliance with internal policies and industry and government regulations such as PCI-DSS, SOX, GLBA and HIPAA with automated reporting. With its vulnerability management capabilities, the FortiDB discovers databases, scans them for vulnerabilities and provides remediation advice. The database activity monitoring (DAM) functionality identifies suspicions database activities by privileged users or application users, and alerts on suspicious actions.

Figure 9 – Secure Web and Database services with FortiWeb and FortiDB

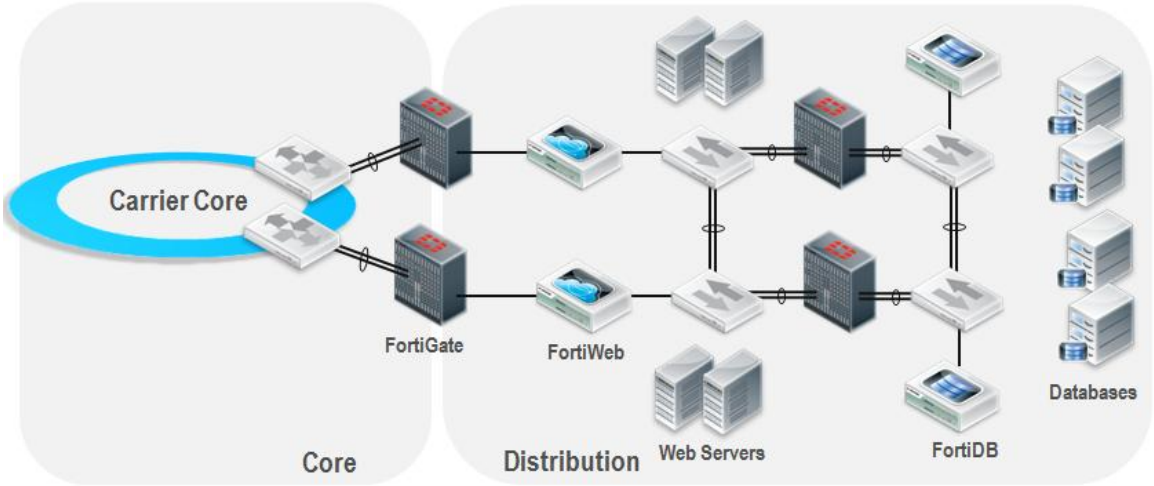

FortiDB is available in appliance form factor, including the FortiDB-2000B appliances. FortiDB is also available in a software version that delivers the same capabilities as the appliances.

FortiDB has three deployment options:

- **Network sniffer:** The FortiDB is implemented transparently without requiring any changes to the network design, and without introducing any latency. This deployment option has no impact on the server.
- **Native Audit:** Selective Audit, only 3-4% performance impact, does not require agents, and captures 100% of events.
- **Network Agents:** Agents send information back to FortiDB appliances. Represents 2-3 % performance impact on the server (not the DB).

While the three deployment options may be considered, the native audit method is the one that provides higher accuracy and completeness, and with a low performance impact. The native audit method is able to capture 100% of local and remote transactions, a common requirement by auditors.

For more information on FortiWeb please see:
http://www.fortinet.com/products/fortiweb/

For more information on FortiDB please see:
http://www.fortinet.com/products/fortidb/

Virtualization

A number of technical and business aspects make virtualization a compelling technology, thus its growing popularity. To hosting and cloud providers, virtualization facilitates the consolidation of services and data centers, significantly reducing cabling and hardware costs, rack space, and decreasing power and cooling requirements. Virtualized services may be implemented and provisioned fast and flexibly, allowing providers to respond to customer demands much quicker. To consumers, public and private cloud services make the deployment of new application and services much faster, with greater flexibility, and at a lower cost. Virtualization also enables customers to easily scale up and down as needed (known as elasticity), while facilitating workload mobility within the data center or across different geographical locations.

Whether the provider delivers a traditional hosting service, an IaaS, PaaS or SaaS cloud service, virtualization is likely to be present in the data center. In fact, these days most hosted services are provided in hybrid environments that combine physical appliances and virtualized services, hence it is crucial to understand how to best integrate both physical and virtualized infrastructures.

Today, many Fortinet products are provided in both physical and virtual machine (VM) form factors to be combined together in the same data center infrastructure. The virtual appliances provide the same functionality as their hardware counterparts, but while providing visibility and security controls between zones within the same virtualized infrastructure. Currently, Fortinet VM products are packaged in Open Virtualization Format (OVF) to be installed on systems running VMware vSphere Hypervisor (ESX/ESXi).

Figure 10 – Virtualized Security Services

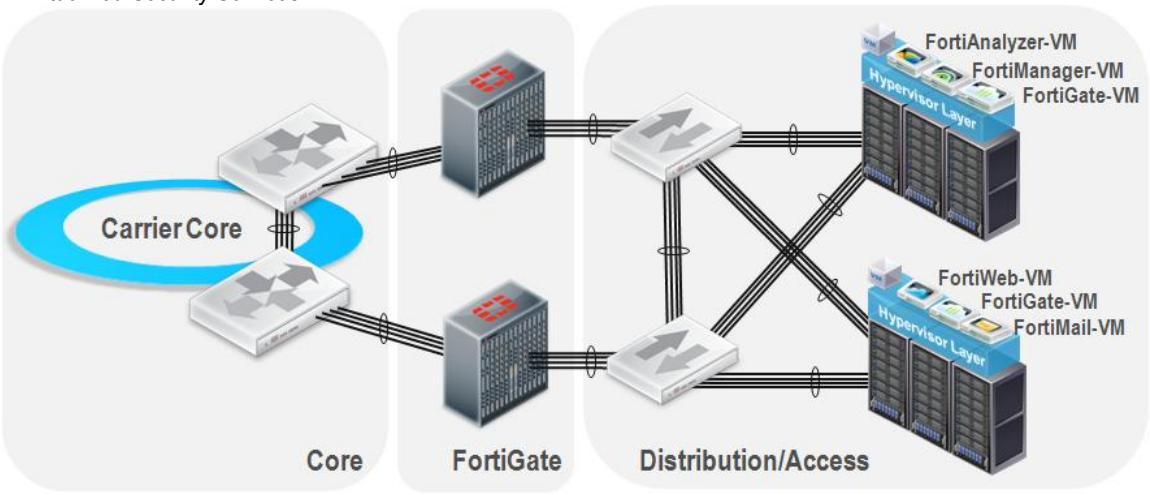

When choosing between physical and virtual appliances, the first thing to understand is that both are complementary technologies, not competing. On one hand, hardware-based appliances deliver predictable performance while performance is harder to estimate on a virtualized environment running hundreds, if not thousands, of different workloads. On the other hand, one of the challenges of virtual environments is that traditional security appliances do not have visibility into the communication flows between virtual machines in the same physical host. Virtual appliances provide visibility and control into the virtualization layer, securing inter-zone and inter-VM traffic, otherwise a blind spot in the virtualized infrastructure. Virtual appliances also bring added deployment flexibility and greater mobility; its elasticity enables customers to easily grow services and horsepower on demand; additionally, customers can easily transfer workloads from one host to another or for between data centers.

Figure 10- "Virtualized Security Services" illustrates the data center infrastructure of a MSSP delivering security services based on a combination of physical and virtual appliances. Physical appliances are the recommended solution for aggregation and edge sections in the network that demand the highest levels of performance. Thus, the design shown uses a pair of redundant FG-5000 chassis between the core and distribution layers to deliver dependable performance for baseline protection. FortiGate-VM, FortiWeb-VM, FortiMail-VM are deployed at the virtualization layer to provide specialized and customizable protection to meet the unique requirements of each customer. FortiAnalyzer-VM and FortiManager-VM provide dedicated management and monitoring for each customer.

The diagram shows a pair of FG-5000 chassis, but other FortiGate platforms could be used as well, as the FortiGate-3950B. These hardware units may be implemented following any of the methods described earlier in the "High Performance Firewall, VPN and IPS" chapter. In this design, the FortiGate are configured with general DoS mitigation, firewall and IPS rules providing a baseline protection to all customers served in the data center or POD, while the virtual appliances deliver specialized and customizable protection per customer.

The virtual appliances integrate to the network using the same technologies available to physical appliances including virtual domains (VDOMs), administrative domains (ADOMs), and VLANs. These are discussed in more detailed in the following section.

Multi-tenancy

MSSP services are built on infrastructures shared across multiple customers or tenants. The success of such multi-tenant environments depends on the provider ability to ensure that tenant resources and data remain properly secured, protected and available at any time.

A secure multi-tenant environment requires the appropriate isolation and separation of customer services and data to ensure no data is compromised, lost or leaked. The necessary mechanisms should be in place to guarantee the continuous availability and accessibility of tenant applications and data. In addition, it is desirable that each tenant is given adequate visibility and control into the customer assets with policy management, monitoring and reporting capabilities.

The Fortinet products described in this document provide a number of features that can be leveraged to successfully secure a multi-tenant environment:

- VPN: SSL or IPSec VPN technologies can be implemented to provide secure access to the tenant various applications, services and data hosted in the provider data center. VPN access can be offered to both mobile clients and selected customer locations, the later with site to site tunnels.
- Virtual Domains (VDOMs): A feature that allows the logical partition of a FortiGate unit into two or more virtual units functioning as independent devices. Each VDOM maintains its own separate virtual interfaces, routing tables, VPN services and UTM protection profiles. Per tenant VDOMs can be implemented to achieve isolation between customers, and to allow them to deploy their own protection profiles. This feature is fundamental to the MSSP as it facilitates the delivery of different service plans addressing the unique needs of each tenant.
- VLANs: VLANs allow the separation of tenant and provider traffic over the same Ethernet LAN. Per tenant VLANs can be defined to separate the customer traffic from other tenants and the provider.
- Virtual appliances: Virtual appliances deliver the same functionality as their physical counterparts but in a virtual machine format. The provider may allow tenants to run their own set of virtual appliances to address their specific security and management requirements. Tenants may choose the security services to be enabled, while maintaining control of their own configurations and protection profiles.

Figure 11 “Multi-tenancy at the Data Center” illustrates how the different technologies are used in the data center design.

The FortiGate appliances or modules deployed at the distribution layer may be configured to provide two distinct levels of protection, baseline protection for all data center users, and per tenant customizable VDOM protection. Baseline protection refers to generic DoS mitigation, firewall and IPS rules that provide a common protection profile for all customers served in the data center or POD. The FortiGate appliances may also be configured to provide customizable edge protection to each tenant via the definition of dedicated VDOMs. Each VDOM functions as a complete separate FortiGate unit with its own configuration. The tenant decides what security services to enable and how security policies should look like. In the example illustrated in Figure 11, Tenant 1 has a dedicated VDOM configured with customized firewall policies, and a VPN service to allow the secure access to the tenant resources and data hosted in provider premises. Tenant 2 has a separate VDOM with firewall and IPS services configured.

Figure 11 – Multi-tenancy at the Data Center

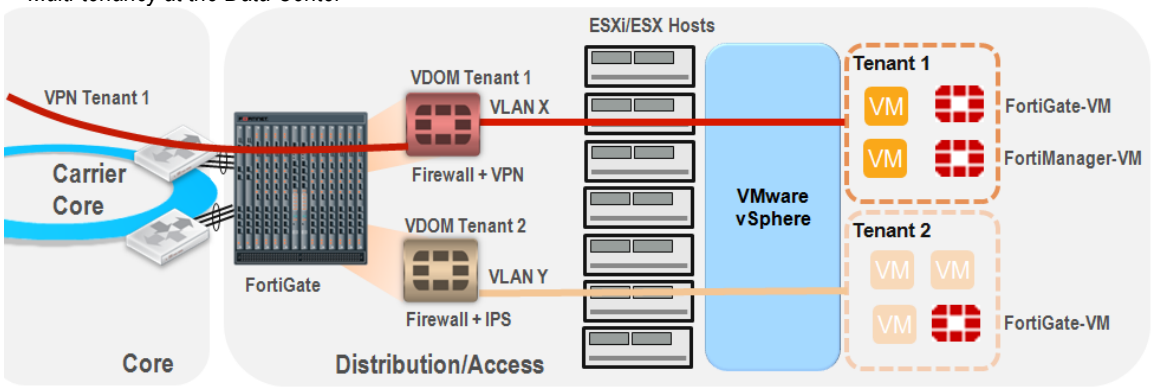

When using VDOMs, FortiGate hardware resources such as memory, disk storage, and CPU are shared across multiple VDOMs as needed. It is important to limit the amount of resources taken by each VDOM to avoid resource starvation and to ensure compliance with Service Level Agreements (SLAs). The FortiGate platforms allow the definition of per VDOM maximum and minimum resource limits, a feature called per-VDOM resource settings. The maximum level is the highest amount of that resource that a VDOM can use if it is available. The minimum level defines a guaranteed level of resources available to the VDOM at any time. The FortiGate can also be configured with global resource limits that control the maximum resources used by the entire FortiGate unit. Both resource settings can be used together for better control and flexibility.

Figure 11 includes a virtual environment using VMware vSphere. In this scenario each tenant runs and controls a set of virtual appliances, which may include one or multiple instances of FortiGate-VM, FortiMail-VM, FortiWeb-VM, FortiAnalyzer-VM, and FortiManager-VM. The use of virtual appliances gives the tenant visibility and control into the virtualized infrastructure. It accelerates the deployment of services and facilitates the mobility of services for higher availability and reliability. FortiGate-VMs can be used to control traffic between different virtual security zones, and to provide granular protection to the customer applications and services running under a virtualized architecture. FortiMail-VM and FortiWeb-VM provide specialized protection for email and web applications. FortiAnalyzer-VM and FortiManager-VM allows the tenant to maintain visibility and control over the management, monitoring and reporting functions.

Traffic belonging to different tenants can be separated by the use of dedicated VLANs. Likewise, it is a good practice for the MSSP to ensure management and control plane traffic is separated from the user data traffic with

the use of dedicated VLANs. The tenant VLANs act as the nexus between virtual appliances and the tenant VDOMs.

With respect to device management, logging and reporting, the MSSP provider may either use share FortiAnalyzer and FortiManager appliances, or provide per tenant FortiAnalyzer-VM and FortiManager-VM instances as described earlier.

FortiManager-3000C, FortiManager-5001A, FortiAnalyzer-2000B and FortiAnalyzer-4000B are the preferred platforms for MSSPs as they handle the largest number of devices. When using shared physical appliances like these, the provider may segregate the management, logging and reporting functions by defining one or multiple Administrative domains (ADOMs) per tenant. This feature allows the MSSP administrator to create separate administrative groups of devices based on geography, business unit or tenant. Each VDOM is managed and monitor by its own administrators. In this way, a tenant ADOM can include a group of VDOMs and virtual appliances that belong to that customer. When tenant administrators log in, they see only the devices and VDOMs configured for that tenant. The MSSP administrator retains visibility and control on all administrative domains and the devices within those domains.

Figure 12 – Administrative Domains

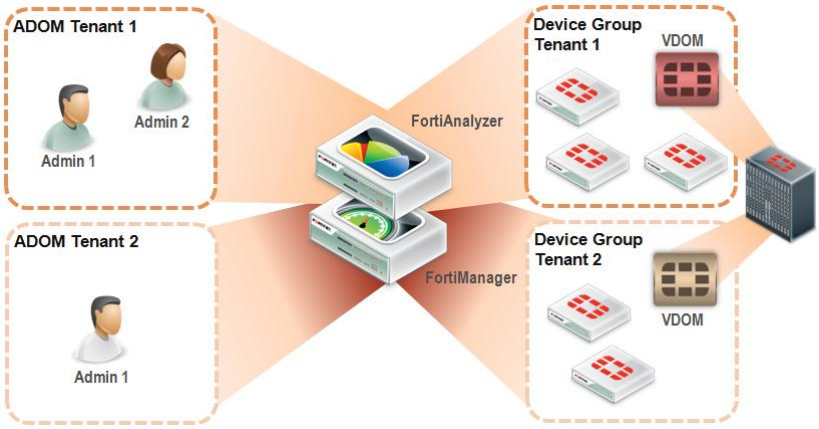

Figure 12 “Administrative Domains” illustrates the concept of administrative domains in a multi-tenant environment. The diagram shows two tenant ADOMs, each with its own administrators. Each tenant ADOM is associated with the appliances and VDOMs assigned to that tenant.

For more information on the configuration of VLANs, VDOMs and other FortiGate features, please see the FortiOS handbook:
http://docs.fortinet.com/fgt.html

Carrier Edge Solutions

Carrier edge services are high performance security services that the MSSP may offer to subscribers to ensure secure and clean access. To that end, the MSSP may build a portfolio of service plans by combining technologies such as firewall, web filtering, intrusion prevention, DoS mitigation, and VPN. The combination of these technologies is what ensures secure and clean access to subscribers. The firewall, IPS and DoS protection services protect subscribers against unauthorized access, intrusions, external threats, and service disrupting attacks. The web filtering component helps avoid loss of productivity, exposure to web-based threats, and loss of confidential information. Additionally, VPN technologies like IPSec may be used to implement secure tunnels between the CPEs and the provider edge, ensuring data confidentiality and providing protection against data

interception attacks. These services may be offered as added-value features included in the subscription, or as advanced services at a premium.

Technically speaking, the carrier edge services are implemented at the consumer aggregation edge of the carrier Point of Presence (POP). These are complex network environments designed to serve large numbers of consumers and to aggregate high volumes of traffic, hence they tend to have very stringent performance and high availability requirements. To start with, due to the high volumes of connections and traffic, these environments are highly demanding in terms of performance metrics such as throughput, number of concurrent connections, connections per second, packets per second, and latency. In addition, these environments are expected to be highly redundant to avoid service disruptions and to ensure continuous compliance with service level agreements. Other advanced features often required include advanced routing support like BGP, and traffic engineering capabilities such as traffic optimization and quality of service (QoS).

Figure 13 – Carrier Edge Security Services

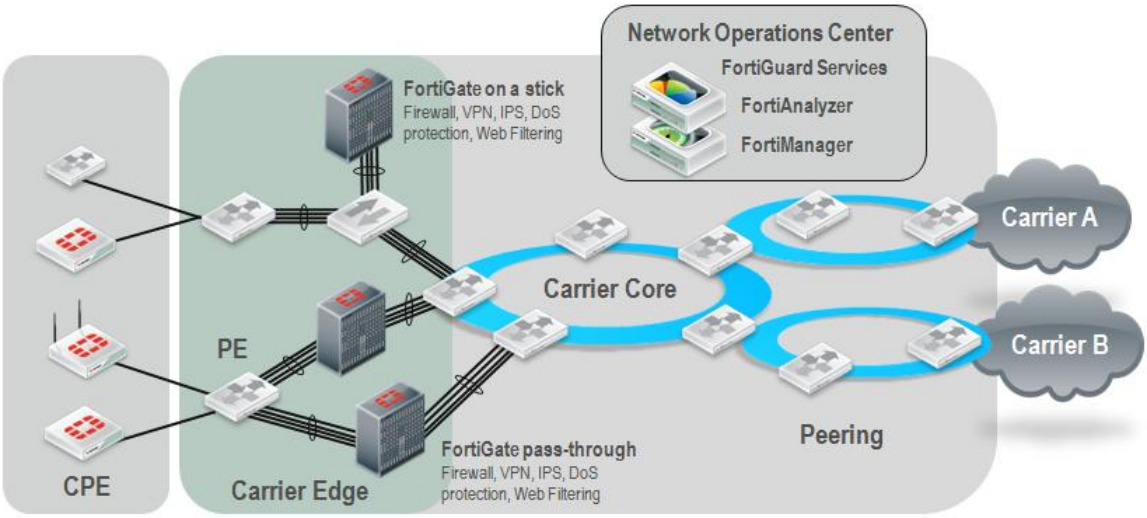

Figure 13 “Carrier Edge Security Services” illustrates the delivery of carrier edge services using the high-end FortiGate-5000 product family. While other platforms may be considered, the FortiGate-5000 is the recommended solution for this type of highly demanding environments. This solution is designed to be highly reliable, and to deliver the highest possible levels of throughput, number of concurrent connections, connections per second, and packets per second.

As Figure 13 shows, the FortiGate units are deployed at the carrier edge in the providers POP, between the PE routers and the carrier core routers. The carrier edge is an ideal location for this type of services because it allows serving large number of subscribers with a common shared infrastructure.

In this scenario, the MSSP may leverage FortiGate’s security services for firewall, IPS, DoS mitigation, web filtering and site-to-site VPN. Depending on the service offering model adopted, the MSSP may offer services that are preconfigured for generic users or customizable per subscriber. In the first option the MSSP applies a generic protection profile to all subscribers of a given service. For example, a single generic DoS protection profile may be configured to address the most common denial of service attacks, and then be applied to all users that have subscribed to the service. The second option are customizable services, where the subscriber can decide what services to enable and how the configure them. For instance, the subscriber of web filtering service is given control on how to treat traffic for given URL categories. Services may be customized using different methods, including by defining per-tenant protection profiles or by assigning pet tenant VDOMs.

There are different common approaches that can be followed to deploy the FortiGate units at the carrier edge. The design depicted in Figure 13 shows two common designs, “on a stick” and “pass-through”. Both designs may implement multiple redundant units for high availability.

The “on a stick” approach, also known as “one arm”, consists in connecting the FortiGate unit to a switch between the PE router and the carrier core routers using multiple high speed links. These physical links, often 10Gb Ethernet, may be bundled together in Link Aggregation Groups (LAGs), providing greater scalability and link redundancy. This approach provides the best flexibility as services can be very easily inserted without the need to re-connect any of the physical links. It is just a matter of configuring VLANs and LAGs to manage how the traffic will flow in and out of the FortiGate unit. One important consideration in the “on a stick” design is that packets will likely traverse the switch multiple times as they flow between the subscriber and carrier networks. Consequently this design may introduce more latency and may be limited by the forwarding capacity of the switch, i.e. packets per second.

The “pass-through” design consists in placing the FortiGate unit in the path between the PE router and carrier core routers. Multiple high speed links, i.e. 10Gb Ethernet, can be used to connect the FortiGate appliance. Links can be bundled together in LAGs for better scalability and link redundancy. This design delivers better latency as there is one less hop in the path. However, the insertion of the FortiGate unit will more likely require adjusting the physical links between the PE router and the carrier core.

When configuring LAGs in platforms that offer hardware acceleration, it is important to understand what ports to use to preserve traffic acceleration. In certain platforms, traffic over a group of ports served by different network processors (NPs) may be handled in software, instead of being accelerated in hardware. The general best practice is to bundle ports served by the same NP. Please consult the hardware documentation of the platform used for more information.

This document presents several recommended deployment options for both “on a stick” and “pass-through” designs at the carrier edge. The proposed designs deliver high availability by combining path redundancy, device failover, link backup, and other mechanisms available on FortiGate platforms. Multiple device failover options are discussed including clustering, stateless failover, stateful active-passive and active-active failover. These failover implementations allow the provider to achieve N+1 redundancy by adding one or more FortiGate appliances or modules as backup, operating in either stand-by or active mode.

As illustrated in Figure 13, carrier edge services are centrally managed and monitor by the MSPP Network Operations Center (NOC) or Security Operations Center (SOC) as well. The NOC/SOC is responsible for provisioning services, monitoring of UTM devices, and the proactive alerting on conditions that may require special attention, such as link failures and security incidents. FortiGuard services, FortiManager and FortiAnalyzer are fundamental components that enhance the NOC/SOC operation.

The FortiGuard Subscription Services keep the UTM devices up to date with the latest antivirus, antispam, IPS, and Web Filtering definitions, protecting the MSSP subscribers against the latest content and network level threats. The FortiManager platform integrates seamlessly with the FortiAnalyzer products delivering a single point of command, control, analysis, and reporting. The FortiManager platform provides centralized policy-based provisioning, configuration, and update management for thousands of FortiGate, and FortiMail appliances, as well as FortiClient security agents. The FortiAnalyzer product family provides network-wide visibility by aggregating log and event data from large numbers of FortiGate and FortiMail security appliances. The “Centralized Management and Reporting” chapter provides more details on these important operational aspects.

There are different ways the FortiGate platforms can be implemented at the carrier edge. The deployment options below are the ones that deliver the highest levels of performance:

- ELBC (FortiGate-5000 only)
- FortiSwitch-5203B (FortiGate-5000 only)
- 802.3ad Link Aggregation

Enhanced Load Balance Cluster (ELBC)

ELBC is a load balancing feature available on the FortiGate-5140 Chassis equipped with a FortiSwitch-5003A or FortiSwitch-5003B and up to twelve (12) FortiGate-5001A or FortiGate-5001B modules acting as worker blades. ELBC balances the traffic across the FortiGate-5001A/B worker blades in the cluster, providing the highest levels of scalability for firewall, IPS, DoS protection, and web filtering.

ELBC is the solution that delivers the highest levels of throughput, concurrent connections, connections per second, and packets per second. As an example, a FortiGate-5140 Chassis equipped with FS-5003B and twelve FG-5001B blades may deliver 160Gbps of stateful firewall, 120 million concurrent connections, and 1.4 million of connections per second. ELBC is ideal for those providers looking for high performance security services at the carrier edge. It should be noted that some functions including VPN, multicast, and certain complex NAT scenarios may not be supported depending on how ELBC and the FortiGate blades are configured.

With ELBC, all packets enter and leave the system throughout the FortiSwitch-5003A/B front panel ports, and not the ports of the FortiGate-5001A/B blades. As packets enter the system, the FortiSwitch blade applies a load balancing algorithm to determine to which FortiGate blade to forward the packet. The load balancing algorithm calculates a hash key value based on the source and destination addresses of the packet. Each FortiGate blade servers a hash key value assigned to it. The load balancing algorithm ensures that traffic for the same source and destination address pair is served by the same FortiGate blade in both directions. ELBC provides high availability by monitoring the health of the blades, and by detecting failures and automatically redistributing traffic to the remaining worker blades.

In a FortiGate-5140 Chassis configured with ELBC, the FG-5001A/B worker blades can operate in Transparent (Layer 2) or NAT/Route (Layer 3) mode, and a single chassis can support mixed modes across VDOMs. ELBC can also be deployed “on a stick” or “pass-through”.

Figure 14 – ELBC at the Carrier Edge

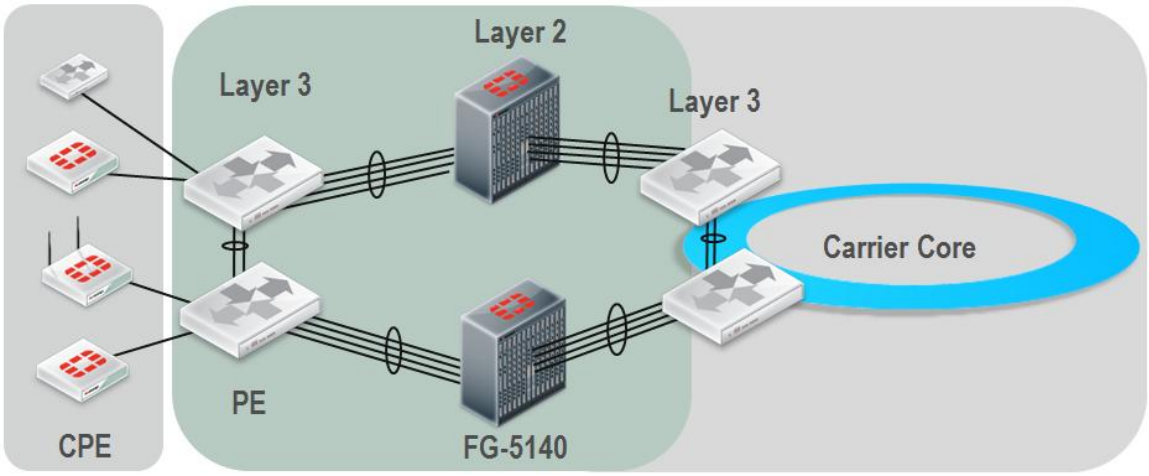

Figure 14 “ELBC at the Carrier Edge” shows one of the most recommended deployments for ELBC, where the FortiGate units in the ELBC cluster are configured in transparent (Layer 2) mode. Configuring the cluster in transparent mode has multiple benefits. To start with, the FortiGate units integrate seamlessly into the network, without requiring a network redesign or change of IP addresses. The carrier may also take full advantage of dynamic routing between the PE and the core routers to deliver enhanced path redundancy, path load balancing, and policy based routing. Additionally, a layer 2 cluster allows the forwarding of non-IP packets, such as those of IS-IS. As an ELBC cluster in NAT/Route mode operates as a layer 3 router, it does not forward non-IP packets. In the event NAT/Route mode is preferred, the FortiGate cluster may be configured with BGP to make it part of the routing infrastructure.

In this design multiple levels of redundancy may be implemented. Figure 14 shows two parallel paths, each with its own FG-5140 chassis. Path selection can be determined by the dynamic routing protocols running between the PE and the core routers. As the FortiGate cluster is configured in transparent mode, it does not participate in the routing decisions. The FG-5140 chassis can be also deployed with dual FortiSwitch-5003A/B blades, providing intra-chassis redundancy. Finally, the actual physical links can be grouped in Link Aggregation Groups (LAGs). The example shown in Figure 14 uses a 4-Port-LAG, where four fabric ports are aggregated for the internal interface, and other four ports for the external interface, giving 40GB of bandwidth to each. The 4-Port-LAG configuration allows for a single internal/external interface pair, or service group. Other configuration options are available such as 2-Port-LAG and basic mode, the former allowing two service groups with 20GB per aggregated interface, and the later delivering four service groups with 10GB of bandwidth per interface.

FortiSwitch-5203B

The FortiSwitch-5203B module provides load balancing capabilities for the FortiGate-5140B and FortiGate-5060 Chassis equipped with one or more FortiGate-5001B modules acting as worker blades in an HA over chassis (HAOC) cluster. The FortiSwitch-5203B runs FortiOS code and is able to deliver firewall and VPN services, while load balancing the UTM functions across the FortiGate-5001B worker blades and the FortiSwitch-5203B itself. In addition to the FortiGate-5140B chassis and FortiGate-5060 chassis, FortiSwitch-5203B boards may be installed in selected versions of the NEBS-compliant FortiGate-5140-R chassis.

The FortiSwitch-5203B is a good option for providers looking for a solution that delivers high performance for firewall, VPN and UTM functions, and that require dynamic routing, multicast, and complex NAT deployments not supported by ELBC. In terms of capacity the FortiSwitch-5203B solution performance is less than ELBC but more than FGCP.

The deployment of the FortiSwitch-5203B requires the configuration of the HA over chassis mode on the board itself and the FortiGate-5001B blades. In this mode, the FortiSwitch-5203B operates as the primary unit, and it is the only responsible for load balancing the UTM sessions across the blades in the cluster. At this time, HA over chassis is supported for blades installed in the same chassis only.

Similarly to ELBC, in an HAOC cluster all packets enter and leave the system throughout the FortiSwitch-5203B front panel ports, and not the ports of the FortiGate-5001B blades. Unlike ELBC, session setup, firewall, stateful inspection, VPN, and session helpers are handled by the FortiSwitch-5203B itself, and are not load balanced across the worker blades. UTM functions like IPS and antivirus are balanced across the various FortiGate-5001B blades and the FortiSwitch-5203B itself. The FortiSwitch-5203B uses a weighted load balancing algorithm to decide where to forward an UTM session. It is possible to adjust weights to influence load balancing, for instance, setting the weight of the primary unit (always the FortiSwitch-5203B) to 0 will ensure all UTM sessions are load balanced to the worker blades, leaving the FortiSwitch-5203B dedicated to session setup and traffic load balancing. High availability is provided by monitoring the health of the blades in the cluster, and by detecting

failures and automatically redistributing traffic to the remaining worker blades. In an HAOC cluster, only a FS-5203B blade can become the primary unit.

As illustrated in Figure 15 – "FS-5203B Deployment at the Carrier Edge", in the FortiSwitch-5203B solution the FG-5001B worker blades can operate in Transparent (Layer 2) or NAT/Route (Layer 3) mode, and a single chassis can support mixed modes across VDOMs. This solution can also be deployed "on a stick" or "pass-through".

The FortiSwitch-5203B solution provides greater deployment flexibility compared with ELBC.
The FG-5001B worker blades can be configured in NAT/Route mode, allowing them to participate in dynamic routing, even with protocols other than BGP. The FG-5001B worker blades can also be deployed in transparent mode, allowing routing protocols to pass between the PE and the core routers.

Figure 15 – FS-5203B Deployment at the Carrier Edge

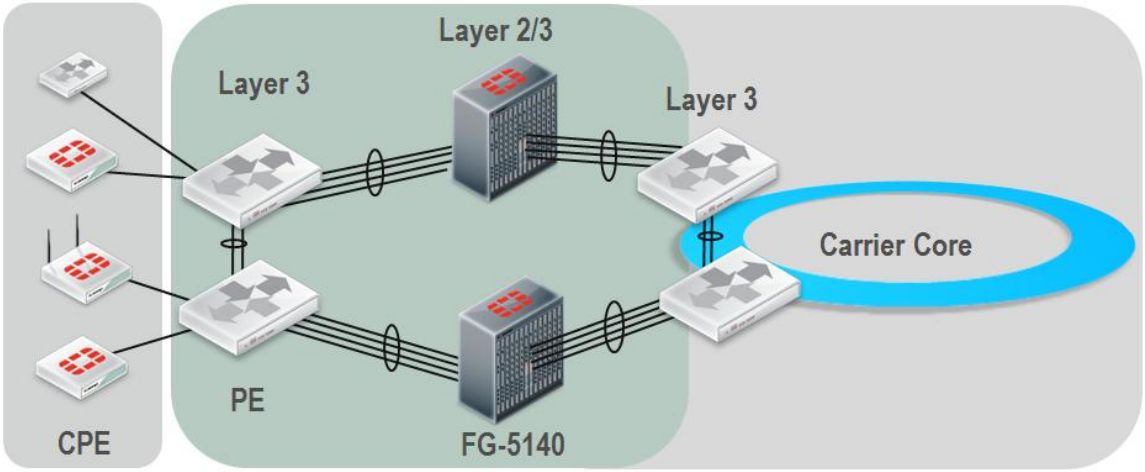

When deciding between NAT/Route and Transparent, it is important to understand that while both modes have essentially the same feature set, due to its nature some features are not available in transparent mode. DHCP, dynamic routing, SSL VPN, virtual IP are some of the features not available in transparent mode.

Multiple levels of redundancy may be implemented as well. Two parallel paths can be implemented as illustrated in Figure 15, each one with its own FG-5000 chassis. In that scenario path selection may be determined by a dynamic routing protocol. Optionally, two FS-5203B blades and multiple FortiGate worker blades may be deployed in the same FG-5000 chassis for intra-chassis redundancy. In this case, the primary FS-5203B blade can also load balance sessions to the backup FS-5203B blade as well as the worker blades. Finally, an additional level of redundancy may be achieved by grouping physical links in Link Aggregation Groups (LAGs).

IEEE 802.3ad Link Aggregation

The IEEE 802.3ad standard defines a method of aggregating multiple Ethernet links into a single point-to-point logical channel, delivering increased throughput and providing link redundancy. This type of link aggregation technology can be used to balance the traffic load across a group of FortiGate modules or appliances, delivering linear scalability and link redundancy.

IEEE 802.3ad is an attractive design option for those providers looking for a low latency and high throughput solution, and who do not have very stringent requirements for failure detection. This solution requires the deployment of a switch or router configured with link aggregation. As illustrated in Figure 16 –"FortiGate balancing with 802.3ad Link Aggregation at the Carrier Edge", the solution consists in connecting a group of FortiGate blades

to a Link Aggregation Group (LAG) configured on a switch or router. Simply put, the switch or router uses a traffic balancing algorithm to predictably distribute the traffic load across the physical links in the LAG, each one connecting to a separate FortiGate blade. As with any load balancing solution, firewall and UTM processing requires that flows of a given session are forwarded to the same FortiGate blade. This solution provides some of the highest levels of throughput with minimum latency, and it is mainly limited by the capacity of the switch or router configured with link aggregation.

IEEE 802.3ad requires that all links in the LAG are of the same type and operate at the same speed, and only full duplex is supported. Links can be configured as either Layer 2 or Layer 3 ports.

While IEEE 802.3ad is a standard, switch implementations may vary, so it is important to familiarize with the actual functionality supported and as well as the default settings. For instance, most switch implementations allow up to eight physical links per LAG, while others may support fewer or more links. Switches from different vendors also often differ on the default algorithm used for load balancing.

Figure 16 – FortiGate balancing with 802.3ad Link Aggregation at the Carrier Edge

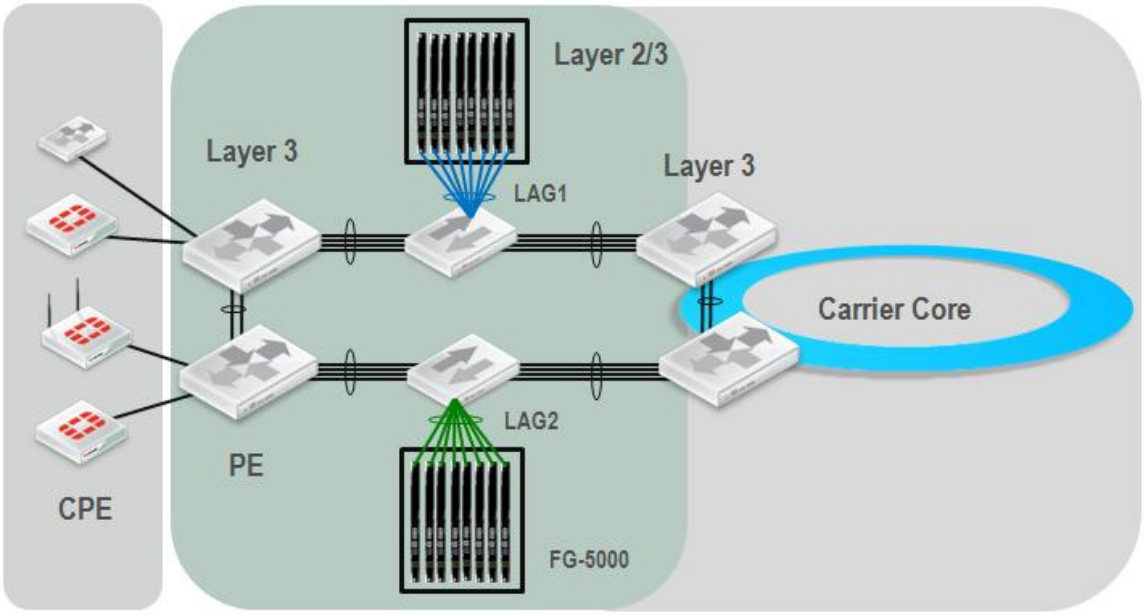

Unlike typical IEEE 802.3ad deployments where both ends are configured, likely dynamically using the Link Aggregate Control Protocol (LACP); in our carrier edge design link aggregation only needs to be configured at the switch side. In fact, the FortiGate units are not aware their links connect to a LAG. For this reason, configuration on the switch is likely to be manual, and not by using LACP.

Figure 16 shows a redundant scenario with two FG-5000 chassis hosting eight FortiGate blades each. Every blade connects to the LAG configured on the switch serving the corresponding chassis. In this environment, the two chassis work independently from each other, and are only shown here to illustrate a secondary path. It is important to note that each FortiGate blade connects using a single physical link, and that the links of all the blades converge into a single LAG. Traffic within those links is separated by the use of VLANs, for example one VLAN for the subscriber's side and another VLAN for the carrier network side. A single LAG is used to ensure both directions of traffic are directed to the same FortiGate blade in the group. Certain switches calculate different hashes per LAG, so traffic symmetry cannot be guaranteed if using more than one LAG. For the same reason this solution is likely to work better on "on a stick" deployment than a "pass-through".

To distribute traffic switches use algorithms that inspect the packet headers. Decisions are based on source and destination MAC addresses, and/or source and destination IP addresses. Some switches may have the ability to use TCP and UDP ports in the calculation of the distribution hashes. In our scenario decisions should be based on source and destination IP addresses only. This is to ensure all traffic flows for a given session are consistently directed to the same FortiGate blade.

Link aggregation operates above the MAC layer in the Data Link OSI layer, thus the FortiGate blade may operate in either NAT/Routed (Layer 3) or Transparent (Layer 2) mode. However, special care should be taken when implementing NAT on the FortiGate blades. Load balancing decisions are made based on IP address information, consequently any address translation will likely break traffic symmetry. As IP addresses change, there is no guarantee the distribution algorithm will select the same physical link in both directions for a given session. Possible solutions are to implement NAT at a different layer, or to modify the algorithm to base its decisions on the IP addresses not affected by NAT. Another alternative is to configure the FortiGate units in transparent mode when NAT is required.

When deciding between NAT/Route and Transparent, it is important to understand that while both modes have essentially the same feature set, due to its nature some features are not available in transparent mode. DHCP, dynamic routing, SSL VPN, virtual IP are some of the features not available in transparent mode.

With respect to high availability, a solution based on link aggregation is limited to link level redundancy. When a physical link fails it is automatically removed from the LAG and traffic gets redirected over the remaining links. For most switch implementations a link failure will trigger a re-calculation of the forwarding hashes over the remaining links. As a result, existing connections are likely to be forwarded over new links, forcing session resets. The solution does neither provide any mechanism to track the state and health of the units in the group. Should a FortiGate blade fail without bringing its physical interfaces down; unable to detect the failure, the switch will continue forwarding packets to the failed appliance, effectively blackholing traffic. Another limitation is that a link aggregation based solution does not include configuration synchronization and does not provide session failover.

To improve reliability, the FortiGate blades can be configured with TCP session synchronization to ensure no data is lost in the event of a unit failover. With TCP session synchronization enabled, the TCP session tables are continuously replicated across the FortiGate blades. If one of the FortiGate blades fails, traffic continues to flow as long as the switch redirects packets to the blades that are still operating. As its name indicates, TCP session synchronization replicates TCP session tables. Due to their connectionless nature, UDP and ICMP sessions do not need to be synchronized to failover naturally. In addition, sessions subject to UTM processing or NAT are not synchronized as they are bound to the FortiGate unit, and units in the cluster may differ in their system and IP address configuration.

For more information on the configuration of Link Aggregation, TCP session synchronization and other FortiGate features, please see the FortiOS handbook:
http://docs.fortinet.com/fgt.html

CPE Solutions

A customer premise equipment (CPE) solution consist in the deployment of standalone UTM appliances at the customer premises that deliver a number of services including routing, firewall, VPN, intrusion protection, and web filtering, to mention a few. These services are designed to protect the subscriber's computing resources and data at the various customer sites, to provide secure connectivity between customer locations, and to protect the subscriber from Internet threats. The CPE devices are provided by the MSSP and may be owned or leased by the subscriber. One of the main characteristics of this service is that the MSSP maintains and monitors the CPE

devices. If the provider allows it, the subscriber may be given certain degree of visibility and control into the CPE devices.

The CPE based service comes with multiple benefits to the subscriber. To start with, the subscriber gets access to a vast range of advanced technologies and resources at a lower cost. As the CPE devices are provided and maintained by the MSSP, the subscriber does not have to spend time and money in selecting hardware, training personnel, and building and maintaining an operational infrastructure. The subscriber also benefits from highly trained personnel and sophisticated infrastructures that the customer could not possibly match on site. In addition, this type of services is often delivered under monthly or yearly subscriptions, simplifying budgeting and helping the subscriber keep expenses under control. This class of services also accelerates the deployment of services, allowing the subscriber to quickly respond to market and regulatory changes.

As illustrated in Figure 17 “CPE Security Services”, the CPE devices may be deployed at several of the customer locations, such as main offices and branches. The CPE equipment is chosen based on a number of factors, including the type of security services required, throughput, and number of users.

Figure 17 – CPE Security Services

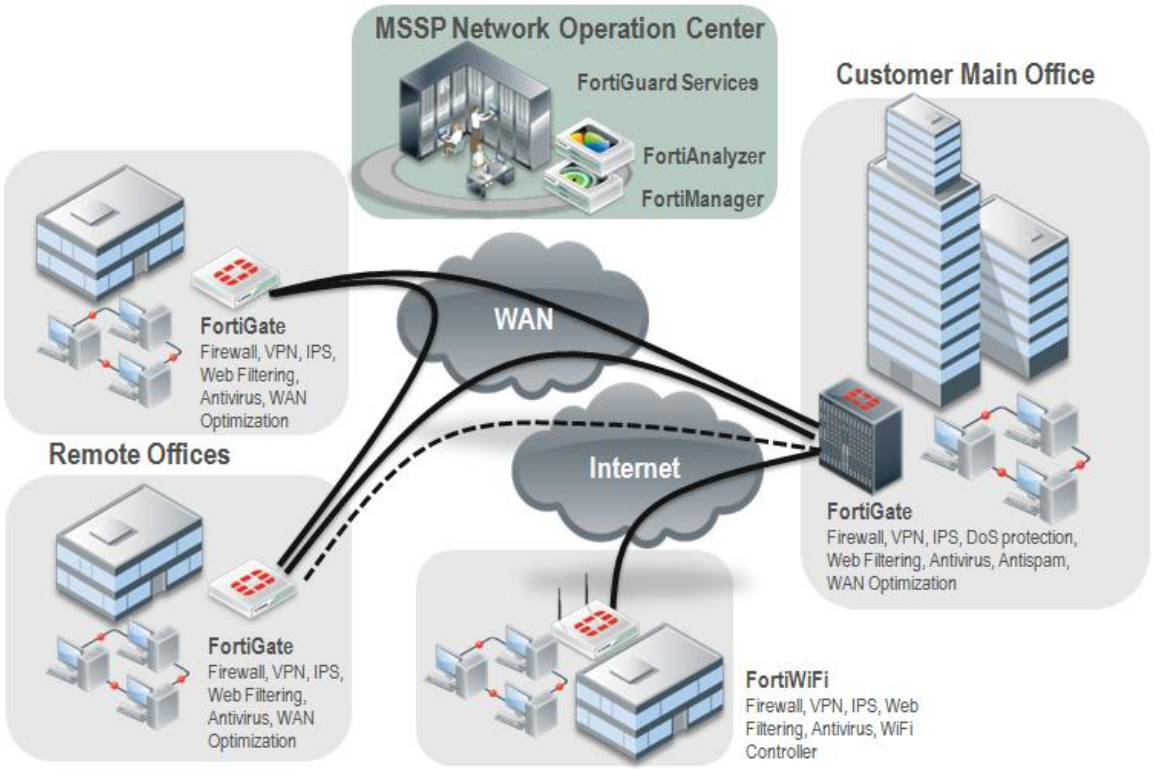

The MSSP may design service plans that offer different levels and types of protection. The service plans may include firewall, IPS, antivirus, and DoS mitigation technologies to protect subscribers against unauthorized access, intrusions, viruses and other external threats. Web filtering may be provided to protect the customer from loss of productivity, exposure to web-based threats, and loss of confidential information. VPN technologies like IPSec and SSL VPN may be leveraged for multiple purposes. IPSec VPN may be used to implement secure tunnels between the several customer sites, ensuring the confidentiality of data over a private WAN or the open Internet. IPSec VPNs may also be implemented to provide a secure backup to WAN connections over the Internet. IPSec VPN and SSL VPNs may be implemented in selected locations to provide secure access to mobile and home-based users. Internet and backup connectivity may be provided by leveraging the 3G/4G modem and

analog V.90 modem support available in various FortiGate and FortiWiFi models. In addition, the MSSP may complement these services with caching, WAN optimization and quality of service (QoS) to boost application performance without requiring expensive bandwidth upgrades.

As mentioned earlier, the MSSP is responsible for monitoring and managing the CPE devices. For most MSSPs those responsibilities lay on the hands of the NOC or SOC personnel. Common tasks performed by the NOC/SOC include the provisioning of services, the monitoring of CPE devices, and the proactive alerting on conditions that may require special attention, such as link failures and security incidents. FortiGuard services, FortiManager and FortiAnalyzer are fundamental components that enhance the NOC/SOC operation.

The FortiGuard Subscription Services keep the CPE devices up to date with the latest antivirus, antispam, IPS, and Web Filtering definitions, protecting the MSSP subscribers against the latest content and network level threats. The FortiManager platform integrates seamlessly with the FortiAnalyzer products delivering a single point of command, control, analysis, and reporting. The FortiManager platform provides centralized policy-based provisioning, configuration, and update management for thousands of FortiGate, FortiWiFi, and FortiMail appliances, as well as FortiClient security agents. The FortiAnalyzer product family provides network-wide visibility by aggregating log and event data from large numbers of FortiGate and FortiMail security appliances. The "Centralized Management and Reporting" chapter provides more details on these important operational aspects.

Subscribers are typically given access to personalized portals that provide them with visibility and certain control into their CPE equipment and associated services. These portals often include information on device configuration, current activity, network performance, logs, events, and customer advisories. The MSSP may gather this information from the FortiManager and FortiAnalyzer platforms, and then feed it into the subscriber portals via an API. Optionally, the MSSP may use administrative domains (ADOMs) to delegate the necessary management, logging and reporting functions to the subscriber. Using role-based administration, the MSSP administrator defines the privileges granted to the subscriber and limits access to the CPE devices in the subscriber domain.

With regards to the selection of CPE platforms, the MSSP may find convenient to standardize service based on a limited number of platform options. As mention previously, the CPE equipment should be chosen based on a number of factors, including but not limited to the type of security services required, performance, and number of users. It is important that all the key service parameters are considered when selecting the service platforms. The following is an example of what type of Fortinet platforms may be chosen for locations of different sizes. Please note every environment is different, so use the example as a reference only.

Location Relative Size	Fortinet Platform
Small office	FortiGate-50B, FortiGate-51B, FortiWiFi-50B, FortiGate-60C, FortiWiFi-60C, FortiGate-80CM, FortiWiFi-80CM, FortiWiFi-81CM, FortiGate-110C, FortiGate-111C
Medium office	FortiGate-310B, FortiGate-620B, FortiGate-3040B, FortiGate-3140B, FortiGate-3950B
Main office, large office	FortiGate-3950B, FortiGate-5000 Chassis

For more information on the FortiGate appliances please see:
http://www.fortinet.com/products/fortigate/

Centralized Management and Reporting

Achieving operational efficiency is a fundamental goal for MSSPs managing complex infrastructures as those required for providing network and CPE based services. In such environments, centralized management, reporting and monitoring becomes highly desirable. Most MSSPs count with Network Operations Center (NOC) or Security Operations Center (SOC) responsible for operation activities. Common tasks performed by the NOC/SOC include the provisioning of services, the monitoring of devices, and the proactive alerting on conditions that may require special attention, such as link failures and security incidents. FortiGuard services, FortiManager and FortiAnalyzer are fundamental components that enhance the NOC/SOC operation. FortiManager-3000C, FortiManager-5001A, FortiAnalyzer-2000B and FortiAnalyzer-4000B are the preferred platforms for MSSPs as they handle the largest number of devices.

The FortiGuard Subscription Services keep the unified threat management (UTM) devices up to date with the latest antivirus, antispam, IPS, and Web Filtering definitions, protecting the MSSP subscribers against the latest content and network level threats. The FortiGuard services are delivered by a world-wide network of servers that form the FortiGuard Distribution Network (FDN). The Fortinet Global Threat Research Team continuously monitors world-wide virus, spyware and vulnerability activities. As new vulnerabilities are found, signatures are created and automatically (or manually) pushed to the subscribed devices. FortiGuard services are fundamental for delivering antispam, web filtering, antivirus, and IPS protection to MSSP subscribers with the most up-to-date information.

The FortiManager platform provides centralized policy-based provisioning, configuration, and update management for thousands of FortiGate, FortiWiFi, and FortiMail appliances, as well as FortiClient security agents. Additionally, FortiManager appliances may locally host security content updates for the managed devices and agents, minimizing web filtering rating request response time and maximizing network protection. The FortiManager platform integrates seamlessly with the FortiAnalyzer products delivering a single point of command, control, analysis, and reporting.

The FortiAnalyzer product family provides network-wide visibility by aggregating log and event data from large numbers of FortiGate and FortiMail security appliances. It provides advanced features such as event correlation, forensic analysis, and vulnerability assessment. In addition, it provides the MSSP with a vast set of customizable reports on network capacity and utilization, policy violations, attack patterns, and compliance with industry standards and governmental regulations.

For more information on the FortiManager please see:
http://www.fortinet.com/products/fortimanager/

More information on the FortiAnalyzer may be found at:
http://www.fortinet.com/products/fortianalyzer/

Related Information

Fortinet Solutions for Managed Service Providers (MSSP)
http://www.fortinet.com/solutions/mssp.html

Fortinet Solutions for Cloud Providers
http://www.fortinet.com/solutions/cloud.html
Fortinet Solutions for Telecommunications Carriers
http://www.fortinet.com/solutions/telecom_carriers.html

FortiOS and FortiGate Technical Documentation
http://docs.fortinet.com/fgt.html

Fortinet Knowledge Base
http://kb.fortinet.com/kb/microsites/microsite.do

FortiGate appliances
http://www.fortinet.com/products/fortigate/

FortiGate Multi-Threat Security Matrix
http://www.fortinet.com/doc/FortinetMatrix.pdf