

FL MGuard TECHNICAL FAQs

In-depth FAQs for the FL mGuard Security Device

AUTOMATION

Technical Note

2738B

© PHOENIX CONTACT 2011-02-15

Overview

This document provides an in-depth look at the capabilities of the FL mGuard products and how they can be used. The information is presented in a Frequently Asked Questions (FAQ) format.

Questions

1. What does a stateful inspection firewall do?	2
2. What does Stealth mode mean?	2
3. Is the FL mGuard default IP address 1.1.1.1?	2
4. Is the FL mGuard an implicit accept or implicit deny firewall?	2
5. What is routing?	2
6. What is Layer 3 switching?	2
7. What is the difference between static and dynamic routing?	3
8. What is Network Address Translation (NAT)?	3
9. What is a VPN?	3
10. Why is a dedicated piece of hardware needed for a VPN? Aren't laptops equipped with software for this?	3
11. What is a certificate and why is it necessary?	4
12. What VPN functions does the FL mGuard support?	4
13. When can PSKs in a VPN configuration be used?	4
14. Can a tunnel still be set up if the FL mGuard IP address is dynamic?	4
15. Can the FL mGuard connect via VPN tunnel to non-Phoenix Contact software/hardware on the other end?	4
16. What is Network Address Translation Traversal?	4
17. Why are the VPN lifetimes different on the initiator and receiver?	5
18. Is the FL mGuard PCI-compatible with the Windows operating system?	5
19. What are the differences between the RS and PCI variants?	5
20. Does the FL mGuard work with EtherNet/IP?	5
21. Does the FL mGuard work with Spanning Tree?	5
22. Can the FL mGuard work with a Phoenix Contact wireless device?	5
23. How does the FL mGuard work when more ports are needed?	5
24. Is the FL mGuard firmware upgradeable?	6
25. What is the recovery procedure and when should it be used?	6
26. What is IKE Ping and for what do I use it?	6
27. What will tech support ask for when troubleshooting a VPN connection?	6
28. What are "profiles" and how are they used?	6
29. Why can't I PING my FL mGuard through the WAN port?	6
30. When I download a configuration profile which contains a VPN connection, are the installed CERTs also saved?	6
31. What can the IPsec VPN/IPsec Status page really tell me about my connection?	6
32. What devices in my network could be blocking my VPN client from seeing the remote, waiting FL mGuard?	7

Answers

1. What does a stateful inspection firewall do?

A stateful inspection firewall keeps track of the state of network connections, such as TCP streams or UDP communication, as they are traveling through it. The algorithm can distinguish legitimate packets for different types of connections. For example, a TCP packet that has the FIN flag set will not be accepted if a TCP packet with the SYN set has not been seen in that stream. Only packets matching a known connection state will be allowed by the firewall; others will be dropped or rejected.

Setting explicit rules for inbound communication is time consuming when using a non-stateful firewall or simple access-control list. Sometimes, this results in allowing any traffic inside the network, rendering the firewall virtually useless.

With the stateful firewall, the intelligent connection-tracking algorithm works on its own and allows users to only define rules for permitted unsolicited traffic (for example, a PLC that initiates a connection).

2. What does Stealth mode mean?

If the FL mGuard is operated in Stealth mode, it is not necessary to reconfigure the clients connected to the internal interface of the FL mGuard. Just interconnect the FL mGuard between the clients that need to be protected and the outside network. The IP addresses of the clients do not change. All processes, which are listening on TCP and UDP ports, are hidden to the outside network and won't be detected by a port scanner. The FL mGuard is completely transparent and will protect devices regardless of their network configuration or operating system.

3. Is the FL mGuard default IP address 1.1.1.1?

No, the FL mGuard in its default mode of Stealth-AutoDetect does not have an actual IP address and will not initiate traffic. However, its Web server does respond to 1.1.1.1. This explains why a default gateway must be configured and reachable to communicate with the FL mGuard while in Stealth mode. Upon sending a request to <https://1.1.1.1> in a browser, the PC will actually send it to the default gateway. To send to the default gateway, the PC must have a gateway defined and must have a MAC address associated with it. The FL mGuard intercepts any request addressed to 1.1.1.1 on the LAN side and treats it as its own.

When changed to Router mode, the FL mGuard is reachable by the internal IP assigned to it. For the FL mGuard RS-B, which does not support Stealth mode, the default IP address is 192.168.1.1. This can be changed on the Network – Interfaces screen.

4. Is the FL mGuard an implicit accept or implicit deny firewall?

The firewall is an implicit deny firewall, meaning if the packet does not match any of the configured rules (or no rules are configured at all), the packet is dropped.

On the FL mGuard, as with most firewalls, the firewall rule processing stops on the first match, regardless of the action.

5. What is routing?

Routing is the process of moving packets through a series of networks from source to destination. Routing is usually performed by a dedicated device called a router. Routing is a key feature of the Internet because it enables messages to pass from the sending device to routers along the way and eventually reach the target device. Each intermediary router performs a routing algorithm and passes along the message to the next "hop." Part of this process involves analyzing a routing table to determine the best path based on things such as bandwidth and link reliability.

Routing is often confused with bridging or switching, which performs packet forwarding but only on a local network by using MAC addresses. Switches are Layer 2 devices (OSI model) and can only forward data within the same subnet. They have no understanding of the logical Layer 3 IP addresses.

Routers connect two or more subnets, which do not necessarily map one-to-one to the physical interfaces of the router, and forward packets by IP address, allowing the movement of information over much larger networks with millions of devices.

Any FL mGuard product can serve as a router, connecting multiple networks together. The FL mGuard RS-B is a variant without VPN or firewall capabilities; if routing is the only thing a customer needs, the FL mGuard RS-B is a fast and cost-effective solution.

6. What is Layer 3 switching?

Routing is distinguished from bridging by operating at the Network Layer (Layer 3) of the OSI model. The term "switches" typically refers to products that do Layer 2 bridging at wire speed. For example, the bridging functions (forwarding of packets by MAC address) performed by the SFN, LMS, MMS or MCS switches occur on the Data Link Layer (Layer 2).

The term "Layer 3 switch" often is used interchangeably with router, and, functionally, they do the same thing. The difference is that a router typically performs routing functions within the CPU and using firmware, while Layer 3 switches perform these functions with dedicated ASICs or other ICs.

7. What is the difference between static and dynamic routing?

Routing can be accomplished by manually entering the information necessary for packets to reach any part of the inter-network into each router. This is called static routing. Static routing works well for small and mid-sized networks, but does not scale well to larger ones. When using static routing, the routing tables on each router must be updated each time the network topology changes, such as when a network link fails. In many networks, routing is managed automatically through the use of dynamic routing. In dynamic routing, routing protocols create and maintain the routing tables automatically by exchanging routing information with neighbors. Dynamic routing responds much more quickly to network changes (and network failures) than static routing.

8. What is Network Address Translation (NAT)?

NAT (Network Address Translation) is the translation of an Internet Protocol address (IP address) used within one network to a different IP address known within another network. One network is designated the inside network and the other is the outside. Typically, a company maps its local inside network addresses to one or more global outside IP addresses and unmaps the global IP addresses on incoming packets back into local IP addresses. This helps ensure security since each outgoing or incoming request must go through a translation process that also offers the opportunity to qualify or authenticate the request or match it to a previous request.

NAT also conserves the number of global IP addresses that a company needs and lets the company use a single IP address when communicating with the world. NAT routing can also be used to join together identical machines with identical IP addresses into a single, functioning facility network. NAT can be further broken down to 1:1 NAT, where one IP address is mapped directly to another IP address, and IP Masquerading, where all “inside” addresses are translated to the FL mGuard WAN IP address.

9. What is a VPN?

A VPN is a Virtual Private Network. It is a method to set up a secure connection between networks or end devices, regardless of where they are physically located. It is a way to create a private network over an otherwise public network, such as the Internet. Within this private network, devices can communicate as if they were directly connected to one another. In some cases, a VPN is created between two devices already located on the same network simply to provide additional security and privacy.

Devices set up VPN communication by first authenticating their partner, generally with certificates or a Pre-Shared Key (a PSK functions as a password). After the authentication takes place, the VPN sets up encryption policies and decides what method will be used to keep

the data secure. They also negotiate other things, like how often to refresh security information. Lastly, they exchange encryption keys.

Now, traffic will come from a node on one side of the VPN tunnel, be encrypted, and then sent to the other end of the tunnel. The tunnel can span any physical distance and go over numerous Internet routers to get to the other side. Once the other side receives the data, it decrypts it and sends it along to the node for which it is intended.

VPNs replace slow and costly “point-to-point” dial-up connections by using the Internet as a way to get from point A to B. They are great at securely (and quickly) getting data from remote sites, customer locations, branch offices, etc.

10. Why is a dedicated piece of hardware needed for a VPN? Aren't laptops equipped with software for this?

While it's true that there are software options for VPN connectivity out there, the FL mGuard has several advantages in comparison:

Dedicated hardware means better performance – even the best software solutions only average about 35 Mbps, and it can be much lower depending on the encryption and compression being used. The FL mGuard offers up to 70 Mbps.

Software running on a PC siphons resources for other PC functions – Using an FL mGuard to handle the VPN means that valuable PC CPU and memory aren't siphoned away from critical tasks. Software VPNs can use over 90% of CPU time during heavy traffic, resulting in potential issues for the control network.

Software running on a PC can crash – Errors, such as Dr. Watson crashes or “program not responding” messages result in application crashes. If the PC's software VPN crashes, then connectivity is lost. The FL mGuard, which is dedicated to secure connectivity, handles the VPN connectivity to ensure the link stays up.

FL mGuard can handle many simultaneous tunnels – Software VPN applications can typically only handle a single, outgoing connection. A single FL mGuard can be licensed to handle up to 250 concurrent tunnels, and it can both initiate and receive the connections.

FL mGuard can serve non-Windows devices – Software solutions are overwhelmingly made for Windows® operating systems only, but sometimes data needs to be tunneled to or from an I/O or PLC. The FL mGuard is compatible with these end devices and will carry the data to them.

The FL mGuard is faster and more reliable than software solutions. It won't slow down the performance of a PC, it can handle more connections, and it works with non-Windows devices.

11. What is a certificate and why is it necessary?

Certificates function as electronic passports – they are issued by trusted authorities, verify identities, and can cause issues if they are misplaced. A certificate is a small, 1 or 2 kB file that contains a very long, encrypted signature created and signed by a CA (Certificate Authority). These certificates are exchanged during the authentication stage of a VPN connection and can also be used in lieu of a standard login password to authenticate to some secure devices. The benefit is that they are much more secure than a typical password, which is often easily guessed or never changed from the default.

The FL mGuard supports the use of x.509 certificates for VPN authentication and for logging in to the browser-based management system.

12. What VPN functions does the FL mGuard support?

The FL mGuard supports all major VPN protocols and, because of this, is interoperable with most existing VPN infrastructures. The FL mGuard supports the IPsec protocol with the following configuration options:

Encryption: DES, 3DES, AES-128, AES-192 and AES-256

Authentication: x.509 certificates and PSK (Pre-Shared Keys)

Hashing: SHA-1 and MD5.

Additional VPN features that are common for IT environments and are supported by the FL mGuard IPsec implementation include Perfect Forward Secrecy (PFS), Dead Peer Detection (DPD), TCP Encapsulation, NAT-Traversal (NAT-T), and hub-and-spoke communication.

The FL mGuard also supports a separate set of firewall rules for VPN-based traffic, as opposed to the local firewall. The FL mGuard supports both Transport and Tunnel mode VPN connections and also NATing for both local and remote VPN addresses.

In addition to IPsec, the FL mGuard supports both L2TP and PPTP for use in legacy applications.

13. When can PSKs in a VPN configuration be used?

It is generally preferred to always use x.509 certificates for VPN authentication, because of the higher degree of security. However, PSKs are supported as long as the configuration does not require Aggressive mode, which is not currently supported by the FL mGuard as it is less secure than Main mode. Main mode protects the identity of both sides of the IKE transaction, whereas Aggressive mode does not.

A VPN connection that does not cross a NAT router and that uses fixed names or IP addresses on both ends of the tunnel is eligible to use PSKs for authentication. All others must use x.509 certificates.

14. Can a tunnel still be set up if the FL mGuard IP address is dynamic?

An FL mGuard might get its external IP address through DHCP or, as part of some network changes, IT might request this address be changed from time to time. This does not pose a problem for using VPN tunnels due to the two following features:

DynamicDNS (aka DDNS or DynDNS) – a service whereby the FL mGuard updates its current IP address to a public DNS Server. The other end of the tunnel then looks up and keeps track of the IP Address by its name (e.g., mGuard-75.acmecocom). Any time the FL mGuard has an IP change it immediately notifies the DynDNS server.

The second feature is the ability to configure the receiving end of the tunnel to use “%any” as the accepted name/IP address of the initiator. This means the remote site can be configured to look for a tunnel to get initiated from any IP address. The initiator must still correctly authenticate to establish the tunnel, so there isn’t a security risk in accepting a tunnel using the “%any” variable.

A dynamic IP or NAT’d IP is not a problem. The FL mGuard’s easy-to-implement solutions eliminates these common IP issues.

15. Can the FL mGuard connect via VPN tunnel to non-Phoenix Contact software/hardware on the other end?

Yes, because the FL mGuard does not use proprietary methods to establish a tunnel, it is possible to tunnel to any number of non-Phoenix Contact devices. The FL mGuard supports a wide range of encryption options (AES, DES, 3DES), both certificate-based and PSK authentication, and standard IEEE functions (for example, Dead Peer Detection and Perfect Forward Secrecy). The only unsupported mode is Aggressive mode (as opposed to Normal and Quick modes); this is because Aggressive mode is not as secure.

Phoenix Contact has several interoperability guides available to walk you through setting up a connection with another vendor’s equipment.

16. What is Network Address Translation Traversal?

Network Address Translation Traversal (NAT-T) is designed to solve the problems inherent in using IPsec VPNs with NAT routing. The problem is that NAT must change information in the packet headers, namely source/destination IP addresses, in order to serve its function. Part of IPsec’s duty is to ensure nothing within the Encapsulated Security Packet (ESP) has been modified, including the IP information. This means that a NAT’d packet contains different information, such as IP and checksums, at its destination than it did at its origin, causing the VPN gateway to discard it.

This conflict between NAT and IPsec created the need for NAT-T, defined by the IEEE in RFC 3947 and

3948. NAT-T adds a UDP header that encapsulates the ESP header (it sits between the ESP header and the outer IP header). NAT-T also puts the sending devices' original IP address into a NAT-OA (Original Address) payload. This gives the receiving device access to that information so that the source and destination IP addresses and ports can be checked and the checksum validated by the receiver. This also solves the problem of the embedded source IP address not matching the source address on the packet.

The FL mGuard implementation of IPSec fully supports NAT-T.

17. Why are the VPN lifetimes different on the initiator and receiver?

The FL mGuard supports "rekey fuzz" and "rekey margin" parameters, which allow the actual keying lifetimes to vary from their configured values. The rekey margin value specifies how long before expiration an attempt to rekey should be made, while the rekey fuzz value specifies the maximum percentage by which the rekey margin can be randomly varied. This means that the key lifetimes on either end of the tunnel will likely be different. Having different rekeying times is useful for devices with multiple VPN connections.

18. Is the FL mGuard PCI-compatible with the Windows operating system?

Yes, the FL mGuard PCI can operate either as a Windows device or independently from the operating system. In the first case, the FL mGuard PCI card just appears as a normal network card, where drivers are installed and Windows manages it as it does a normal NIC card.

In the second case, the FL mGuard PCI is unknown and uncontrolled by the operating system. It is seen as an unknown Ethernet controller and managed via the browser-based management system. This OS-independent mode is the factory default; to change to a Windows-management method, simply move a jumper on the PCI card itself.

19. What are the differences between the RS and PCI variants?

There are some hardware differences between them, but there are many other similarities. Both have the same cryptographic CPU, flash and memory capacities, 10/100 speed LAN and WAN ports.

The RS has an alarm contact and RJ11 serial port that the PCI card does not. The PCI card is rated to a slightly wider temperature rating. The firmware and security functionality is identical between them, so they both offer high protection levels.

20. Does the FL mGuard work with EtherNet/IP?

When in Stealth mode, the firewall can easily be configured to work with multicast traffic. Out of the box, it will pass multicast data, IGMP queries, and reports from the LAN to the WAN side without any issues.

When the FL mGuard is in router mode or when using a VPN, multicast traffic cannot be forwarded. The devices on the LAN side of the FL mGuard can communicate via multicast between one another, but that traffic can't get routed to the WAN side. This is due to the way in which routing works; the same algorithms that limit broadcast data and make routing so desirable in traffic management prevent multicast traffic from being routed across segments.

EtherNet/IP using point-to-point connections instead of multicast connections work regardless of the mode being used.

21. Does the FL mGuard work with Spanning Tree?

Because it is a Layer 3 device, the FL mGuard doesn't actively participate in bridging (a Layer 2 function). Because of this, the FL mGuard will not send Spanning Tree BPDUs, attempt to be the root bridge, disable one of its ports or automatically attempt other Spanning Tree BPDUs functions.

However, the FL mGuard can be configured to either pass BPDUs through or block them, depending on the physical topology and user demands. This allows the FL mGuard to participate in a ring topology by permitting BPDUs to pass through. It can also allow for multiple root bridges on different subnets by blocking the BPDUs.

The default behavior is to block Spanning Tree BPDUs.

22. Can the FL mGuard work with a Phoenix Contact wireless device?

Yes, the FL mGuard can be used in conjunction with Phoenix Contact wireless devices as well as wired devices. For example, two wireless radios can be attached to an FL mGuard via the WAN port; a VPN tunnel can then be established across one or more of these wireless links. This is useful if the radios themselves do not support any encryption or a key long enough to meet the desired security level.

Because the FL mGuard uses IEEE-compliant protocols, it is an interoperable device suitable for use in many different network topologies, both wired and wireless.

23. How does the FL mGuard work when more ports are needed?

There are several application scenarios where the FL mGuard can protect and service more than one device through the LAN port. In these cases, simply connect a switch, for example, the Lean Managed Switch (LMS) to the FL mGuard LAN port. Then, connect the end devices, or additional switches, to the LMS to extend the FL mGuard firewall protection or VPN connectivity to an entire subnet of devices. The FL mGuard's high-performing hardware can service over a hundred devices easily and without issues.

24. Is the FL mGuard firmware upgradeable?

Yes, it is upgradeable by several means. Firmware can be downloaded from the Internet and installed manually, or the FL mGuard can go to the Internet itself to download and install minor updates or new releases. Additional licenses (for example, to add additional VPN tunnels) can also be installed with a few mouse clicks.

25. What is the recovery procedure and when should it be used?

The recovery procedure is used when the internal (LAN) IP address of the FL mGuard is unknown or when Web and/or SSH access to the FL mGuard is disabled, rendering management access impossible. The recovery procedure consists of depressing the “Rescue” button in a specific pattern. Consult the user manual for the specific pattern of each variant.

The recovery procedure resets the FL mGuard RS-B internal IP address back to 192.168.1.1. For all other variants the recovery procedure puts the device in Stealth mode and makes it accessible through IP 1.1.1.1.

The recovery procedure does not make any changes to VPN connections or configurations, firewall rules or passwords.

26. What is IKE Ping and for what do I use it?

IKE Ping is a tool to verify that the VPN initiation traffic is getting from the initiator to the receiver. It is very helpful in ensuring that any firewall or port forwarding rules are correctly set up on a router in front of the receiving FL mGuard. IKE Ping sends a single, dummy VPN initiation packet to the IP address or name that you specify. When this packet arrives, the receiver will reply to it; let the sender know that the packet arrived; and that there are no routers or firewalls blocking the VPN traffic. To use IKE Ping, go to the Support >> Tools web page, IKE Ping tab, then enter the IP address or name used to access the receiver and click the “Ping” button.

27. What will tech support ask for when troubleshooting a VPN connection?

Generally speaking, most configuration, troubleshooting and log information can be obtained using the “Snapshot” feature of the FL mGuard. A Snapshot can be downloaded from the Support >> Advanced web page, Snapshot tab. Getting a Snapshot from both ends of the tunnel, i.e., the initiator and receiver, will provide almost all the required information to verify that the tunnel is configured correctly at both ends. An IKE Ping test may also be requested to verify that the initiator is able to successfully get the VPN traffic through any firewalls/routers to the receiver.

28. What are “profiles” and how are they used?

Profiles are essentially a collection of all of the parameter settings you can make in your FL mGuard, including firewall rules, IP addresses, certificates and

VPN settings. Profiles, which can be saved, uploaded and downloaded on the Management >> Configuration Profiles web page, are useful in a number of areas.

For example, if you have several, similar FL mGuards deployed to many end customers using similar setups, you can build a “master” profile with the IP settings, VPN tunnel configuration, firewall rules, etc., just the way you want. Then, by saving and downloading this master profile, you can use it to quickly configure and commission any future FL mGuards by simply uploading and restoring their configuration.

Saving a profile is also a good way to create a local backup of a configuration. If something should happen to an FL mGuard that causes it to be replaced, this replacement can be quickly shipped to a site and the desired profile loaded for a quick recovery.

29. Why can't I PING my FL mGuard through the WAN port?

By default, the FL mGuard will only respond to pings on the LAN side. This is a common security measure to ensure that the “outside” network is unaware of the FL mGuard or “inside” network. To enable ping response on the WAN port, simply go to the Network Security >> Packet Filter web page, Advanced Tab and change the “ICMP via primary external interface...” from “Drop” to “Allow Ping Requests.”

30. When I download a configuration profile which contains a VPN connection, are the installed CERTs also saved?

Yes, the configuration profile (the ATV file) contains all installed certificates along with all of your configuration settings. Snapshots, on the other hand, contain the configuration settings, including references to certificates, but NOT any actual certificates.

31. What can the IPsec VPN/IPsec Status page really tell me about my connection?

They indicate the status of any enabled VPN connections. The “ISAKMP State” refers to the Authentication of the tunnel

- if you have an “Established” state it means your FL mGuards have successfully authenticated the tunnel.
- if it is blank, either the VPN packets are not getting through or there is a mismatch in Authentication (PSK or Certificates) or the algorithms used to encrypt them.

The “IPSec State” refers to the Networking part of the tunnel. If the “ISAKMP State” is established but the “IPSec State” is not, it means there is a mismatch between the “local” and “remote” networks defined for the two ends of the tunnel. What is “local” for FL mGuard A must be “Remote” for FL mGuard B and vice versa.

32. What devices in my network could be blocking my VPN client from seeing the remote, waiting FL mGuard?

You can verify the end-to-end initiator to receiver using the IKE Ping tool to ensure that VPN traffic is not blocked along the way. If the traffic is getting blocked, you will need to do a little troubleshooting to see where.

First, to see if the traffic is being blocked on the initiator's network, try an IKE Ping to pxcmguard2.dnsalias.net; this name resolves to a Phoenix Contact FL mGuard with IP address 206.192.98.131. If you get a response to your IKE Ping, you can be confident the initiator's network is allowing the VPN traffic (UDP port 500) out without an issue and you can move to troubleshoot on the receiver side. If there is no response, port UDP 500 may be blocked on the initiator network. Coordinate with the IT department or use TCP Encapsulation to resolve the issue.

If the IKE Ping test to the public FL mGuard works but the IKE Ping to your receiver's public IP address fails, then the problem is likely a router/firewall on the receiver side that is not correctly set up to permit or forward VPN traffic to the FL mGuard.