TMOS Secure Development and Implementation

Overview TMOS—the foundation and architecture for F5's application delivery controllers running on the BIG-IP platform—brings a wealth of security to existing application delivery networks. Design and operational features, such as a full TCP and application proxy, optimized IP stacks, and virtual network segmentation, are just a small subset of the security features available on all TMOS-based appliances. But before the secure implementation of TMOS can be discussed, the question is often asked "How is TMOS itself secured?"

Challenge When creating any security-enabled network device, security of the device itself must be questioned. A gate provides no security to a house if the gap between the bars is large enough to drive an SUV through. Many highly effective exploits have been found against the very software and hardware that is designed to protect against those exploits. Attacks against anti-virus software are among the most common. If you can attack the guards, then you don't need to worry about being stealth.

With TMOS, F5 is very conscious of building and maintaining a secure and robust application delivery platform, and has implemented many different checks and counter-checks ensuring that TMOS provides the secure networking environment it was designed for. From providing security to the customer's application delivery network, through mandatory and routine checks against the stack source code to provide its own internal security, application delivery security starts with a secure application delivery controller.

Solution **Hardware and Software: Secure Symmetry on the Outside**

BIG-IP and TMOS are designed from the ground up so that both the hardware and software work together to provide the highest level of security. The BIG-IP application delivery networking product's hardware and software architecture is as follows:

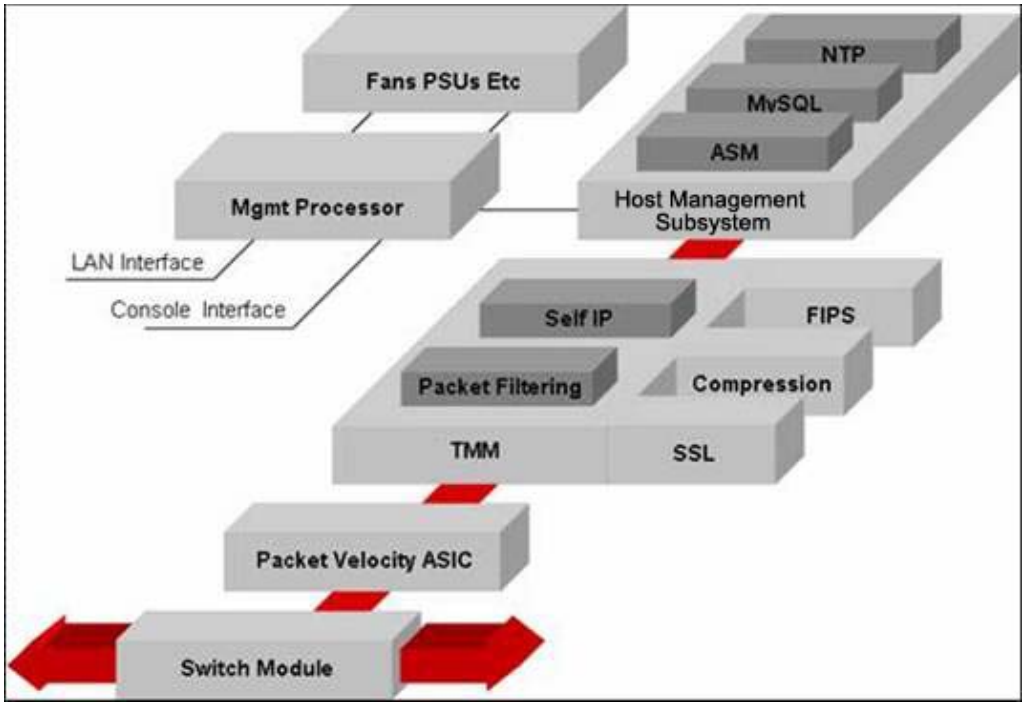

The Switch Module, where all application delivery traffic enters and exits, connects to the PVA (Packet Velocity ASIC), F5's custom-engineered L4 load balancing ASIC switching fabric. Traffic that can be handled within the PVA never goes any further; at this step, all packet and connection management occurs at the hardware level by the PVA through the Switch Module. Traffic enters through the switch into the PVA, where the appropriate logic and transformations are applied before the traffic is sent back out through the Switch Module. Generically speaking, this is typically referred to as the fastL4 profile.

For traffic and which is not handled by the PVA, it is simply passed through the PVA onto the next layer, which is F5's primary traffic management processing system called TMM (Traffic Management Microkernel). TMM handles all of BIG-IP's local traffic functionality such as intelligent load balancing, compression, SSL, iRules, packet filters, etc. (with the exception of L4-only load balancing, which can be handled in either the PVA or TMM). The TMM can manage traffic using several optional hardware acceleration modules such as SSL, FIPS, and Compression and has entirely dedicated hardware. TMM is also responsible for delivering traffic to the Host Management subsystem as necessary for products such as BIG-IP Global Traffic Manager (GTM).

Traffic Management Subsystems

In addition to TMM, BIG-IP incorporates multiple subsystems for various traffic management tasks, each with their respective areas of expertise. Two of these subsystems are the Host Management subsystem and the SCCP subsystem.

The Host Management subsystem is responsible for the BIG-IP system's management processes. This includes SSH for secure command-line access, an HTTPS web server for secure Web GUI access, SNMP, NTP, and a variety of other management services. Other functions handled in the Host Management subsystem are health checks, configuration management, and logging. Some software modules, such as GTM, are also hosted on the Host Management system. These modules are tightly integrated with the TMM so as to provide accelerated performance and complementary functionality; however, it is important to note that TMM and the Host Management subsystem traffic management are two completely different systems, each with their own uniquely implemented kernels, user environments, and network protocol stacks.

Connected to the Host Management subsystem is the Management Processor (SCCP – Switch Card Control Processor). This is another dedicated system-within-a-system that provides "lights out" management and other supporting functions for the TMM, such as managing the L2 switch and over-the-network re-installation.

BIG-IP employs two network connection entry points; one is the Switch Module (the Ethernet ports on the front of the unit for application delivery traffic), the other being the LAN Interface which connects to the Management Processor. Each of these entry points can potentially provide access to administrative services on the Host Management subsystem. As discussed above, the TMM is in the path between the Switch Module and the HMS, so TMM is an enforcement point which can limit network and application access to the HMS. For example, TMM can limit which ports are seen as "open" on each administrative IP address (Self IP), regardless of what may actually be open on the HMS. Furthermore, TMM has a full packet filter engine that can be used to limit access to HMS in a very granular way. By default, BIG-IP is configured to allow a few default services such as SSH and HTTPS on the switch module's internal VLAN and management interface, and is configured to allow no access via the external VLAN.

The main entry point to access any part of BIG-IP is the Switch Module, and access through the switch can be controlled by TMM in a highly granular manner. Please note that access through the Management Port is considered trusted, and no access controls can be applied; this is an important consideration when designing BIG-IP into your network. If you have a trusted and secure management-only network, this is where the Management Port would be connected. If you do not have a trusted and secure management network, it's recommend that the management port is not

used; instead, access is granted only through the switch ports managed by TMM and/or the local serial console.

Multiple Protocol Stacks

Across both the hardware ASIC and the software-managed connections of TMM, TMOS utilizes a "multi-stack" approach: multiple protocol stacks for multiple functions. As described earlier, fastL4 is a specific type of protocol stack implemented by the ASIC for extremely fast hardware acceleration and delivery. Likewise, the fastL4 stack is one type of software protocol stack that can also be managed in software, providing accelerated Layer 4 management via TMM. This multi-stack architecture allows BIG-IP, via configuration options set by the administrator, to choose the correct stack(s) for a particular purpose, be it using TCP Express and the full proxy for SSL acceleration and compression, using Fast HTTP for optimized HTTP-only delivery, or hardware/software fastL4 for pure TCP and UDP acceleration.

In addition to providing the right tool for the job, this multi-stack architecture also provides multiple layers of security for all packets and connections managed by TMOS. In most network appliances, and even down to networked operating systems, one protocol stack is used for all connections. Not only does this approach limit traffic optimization options, it also exposes all parts of the appliance to a single-point-of-failure attack: exploit one security flaw in the protocol stack, DoS the entire appliance. By BIG-IP leveraging multiple, independent, and unique protocol stacks (with each stack implementing its own security measures in unique ways), a very robust security system is created, with multiple layers relying on and protecting one another. Exploit one flaw in one stack on a BIG-IP, and the chances of that exploit impacting the overall BIG-IP traffic management system are minimal, because that exploit will have to pass through parts of other stacks which may be immune to the attack.

Popping The Circular Stack

When each of these industry-leading technologies are brought together in one appliance under one extremely sophisticated traffic management system, TMOS, the complete security solution reveals itself. A single connection may take advantage of any of the above technologies. Consider the following example:

Smith Co. has two SMTP email servers behind a BIG-IP LTM virtual server (VIP), with both SMTP servers configured in one pool. The administrator's primary concern with the SMTP VIP is optimization and availability of the mail servers, so she applies the fastL4 profile to the VIP because she's looking for fast and redundant access for incoming connections to their mail network. This VIP also applies some very basic source IP identification for inbound SMTP connection requests, and applies SMTP proxy logic to specific connections that originate from unknown sources, such as providing basic sanity checking to ensure the SMTP connection matches RFC requirements. In other words, the VIP matches each new connection against a list of known "good" IP addresses, and if the connection does not originate from one of those addresses, it sanitizes the SMTP connection. As a final security measure, the administrator has applied a set of packet filters that restrict which internal hosts and VLANs are allowed to access the SMTP VIP from the internal network.

This example seems very ordinary. This is a standard SMTP connection system that many F5 customers use today to provide load balancing, basic delivery logic, and transport security to their existing mail delivery network. But under the hood, there are many pieces of the cogs moving together in very specific and unique ways to guarantee the most secure and available delivery network for Smith Co.'s mail network. Even though the administrator has applied the fastL4 profile, she's also instructed the VIP to perform application-specific tasks, which will invoke the TCP Express full proxy profile. As each one of these seemingly simple configurations change from customer to customer, from network to network, and from implementation to implementation, each

one of these cogs will move together in their own unique ways to provide the same high level of securely available delivery.

Policing Ourselves: Security on the Inside

As anyone in security will tell you, the security of a system is only as strong as its weakest link. As stated above, a gate with holes is a useless gate. From the inside-out, TMOS was designed and built with security in mind. By focusing on secure implementations for both internal and external components, such as user-level access to the LTM and the devices behind it, and packet filters for external connections as well as for subsystem communication, TMOS treats security with the same level of urgency across the entire platform.

Code Scanning and Cleansing

While there are many pieces to a truly secure system, two of the most important are design and coding. There is near-uniform agreement that a completely secure source code is the holy grail of application and network security. If all applications were secure from the get-go, we probably wouldn't need advanced security devices such as application firewalls. F5, like many software development houses, focuses heavily on writing secure code and has invested heavily in training the internal development staff on how to write it securely. However, in the world of software and network exploits, tiny mistakes can have huge ramifications.

Therefore, to augment secure development initiatives - and to provide an additional layer of security to the TMOS source code - F5 has implemented a sophisticated 3rd party scanning application, one which analyzes nightly source code for a number of critical flaws. At compile time, the code scanning application looks for flaws such as security bugs and defects, "build breaker" bugs, crashing bugs such as memory leaks and corruption, and unpredictable behavior introduced by new code. Beyond critical bug and flaw detection, source code scanning can also find non-fatal flaws such as data integrity issues and performance bottlenecks.

When used as a tool alongside secure design and implementation, secure source code scanning and evaluation can find critical flaws long before they become a catastrophe, and still while the cost to fix those flaws is relatively inexpensive. Secure code development and code evaluation are critical pieces to the TMOS design and development stages which augment each other.

After Source Scanning: Black Box Testing

Beyond source code evaluation, application "black box" testing is also another tool used to secure TMOS prior to general market release. "Black box" testing is a method in which application and platform testing is done mostly in the dark, without any additional knowledge beyond what a standard human attacker would have access to (in contrast to source code scanning). Black box testing and analysis can be inserted anywhere in the software development lifecycle all the way through release. This type of security evaluation adds additional layers of protection by continuing security impact testing beyond what source code scanning can reveal. Cars today have both airbags and seatbelts for that same reason: each provides a different form of protection, both are important, and when used together provide a more secure environment than when used independently.

F5 has implemented both internally developed and external 3rd party tools which focus on locating and eradicating very specific flaws and exploits against very specific application platforms and protocols. This, in conjunction with source code scanning, adds to the robust security evaluation TMOS undergoes during every development release.

Local Application Level Security

Like developmental security, deployed appliance security is a key element of F5's secure platform strategy. TMOS runs across three separate operating system platforms: the Host Management subsystem for the core operating application functions, such as the web-based GUI and SSH; SCCP for the management applications and interface; and TMM, for application delivery traffic management. Each of these unique environments requires their own security systems to keep the applications and data managed by BIG-IP secure. Beyond the TMM security addressed above, both the HMS and SCCP management systems have their own security subsystems installed. For network level security, all systems on BIG-IP rely on the same stack-level protection, but the operations management systems require additional protection.

The two key areas of operating system security are user and process management. These key areas apply to both local and remote users and processes, and likewise, the security mechanisms in place for both BIG-IP management systems provide security for each. BIG-IP employs both local and network application level security, often referred to as Host Intrusion Protection (HIP), which enforces access control and containment restrictions, protecting against user and system exploits which bypass standard system-level access restrictions. Exploits against any application installed on BIG-IP, including kernel-level exploits, are always treated as critical issues; patches are released as soon as possible and out of cycle from major software releases. Between these public exploits and patches, BIG-IP's application security HIP subsystem maintains a secure working environment for both users and applications. All applications run in a restricted environment and are constantly monitored for restricted actions and behavior. BIG-IP administrators are immediately notified of application violations via one of BIG-IP's many notification systems such as iControl.

Conclusion

F5 provides application delivery network security to protect the most valuable application assets. In order to provide organizations with reliable and secure access to corporate applications, F5 must carry the secure application paradigm all the way down to the core elements of BIG-IP. It's not enough to provide security to application transport; the transporting appliance must also provide a secure environment. As stated earlier, a secure application delivery network is only as secure as its weakest link. F5 has taken multiple steps to ensure that the BIG-IP application delivery networking controller is an extremely secure link in the application delivery networking chain.

About F5

F5 Networks is the global leader in Application Delivery Networking. F5 provides solutions that make applications secure, fast and available for everyone, helping organizations get the most out of their investment. By adding intelligence and manageability into the network to offload applications, F5 optimizes applications and allows them to work faster and consume fewer resources. F5's extensible architecture intelligently integrates application optimization, protects the application and the network, and delivers application reliability—all on one universal platform. Over 10,000 organizations and service providers worldwide trust F5 to keep their applications running. The company is headquartered in Seattle, Washington with offices worldwide. For more information, go to www.f5.com.