



Secure Application Delivery

Introduction The last several years have seen the usual sea of change in IT — and by extension IT security — systems. Moore's Law, in addition to describing advancements in computing power, has led to more complex algorithms for encryption. The...

White Paper
by Don MacVittie



Introduction

The last several years have seen the usual sea of change in IT—and by extension IT security — systems. Moore’s Law, in addition to describing advancements in computing power, has led to more complex algorithms for encryption. The ability to process more has, for example, driven SSL encryption key sizes from 256 bits to 4 Kb in less than a decade. Applications can be better protected today than they have been historically, assuming the requisite computing power can be dedicated to this protection.

Unfortunately, the number of attacks and attack vectors has been increasing at the same time. With new exploits cropping up on a regular basis, the need for a solid SSL encryption environment to protect applications is greater than ever. Those advancements in computing power and security are required for any organization with sensitive data exchanged on the Internet — and should be required for any organization with an Internet presence, simply because unprotected web pages can be a gateway to protected data.

Now that organizations face distributed denial-of-service (DDoS), sidejacking, SSL man-in-the-middle, and renegotiation attacks, IT security teams need to place protections that cover their organizations’ entire web presence with the adaptability to survive the most insidious attacks. Because we need to protect everything, even new advancements in encryption algorithms — specifically elliptical curve cryptography (ECC)—will not be enough for the industry to ignore the related performance ramifications.

And all of this security is still plagued by the greatest hurdle to SSL adoption in the network space — visibility. Organizations have many valid reasons for looking into the data being transferred, from data leak prevention to advanced layer 4 to layer 7 routing. So the security solution at hand must be capable of providing visibility, preferably 100 percent visibility, into the data moving both into and out of the network.

F5 has the answers to these complex problems, from adaptability to changes between security algorithms to visibility, key management, and integration. While attacks are continuing to evolve, governments and industries are requiring more security, and employees need greater network and data access in spite of growing risk, the strategic point of control that acts as a network gateway could be solving those problems with security services. The F5® BIG-IP® platform offers an array of solutions to drive security home without increasing network latency and while taking considerably less time to manage than separate products would.



Always-On SSL

The need for SSL encryption continues to grow. Access to personally identifiable information (PII) must be protected, as must access to restricted information. Increasingly, websites and remote upload locations must be protected, too. Indeed, many organizations have, or are considering, near 100 percent encryption for connections. Organizations are moving to secure all environments, be they test, QA, or production, to protect against all attacks. Organizations are also starting to realize that all data, applications, and session information needs to be protected if a truly secure environment is the goal.

Historically, only applications — or more accurately, sensitive items of data like credit card numbers and other PII—were SSL protected. The remainder of a website, containing information, session cookies, or static pages, was not. But those unprotected pages became a liability, offering attackers a route into the systems that hosted protected pages.

Organizations must take care to protect their customers' and other sensitive data, and legally they must take "due care," providing sufficient protections to ensure that an attack is not obviously going to get through. So today, many organizations are moving to an always-on SSL posture, and some are moving to SSL for all connections, internal or external.

But encryption comes with costs. RSA-style encryption is expensive in terms of CPU cycles and key maintenance, and manipulating the contents of encrypted streams becomes nearly impossible. Although that is the primary reason encryption exists, the organization implementing encryption has a valid need to modify data for a variety of reasons, from content filtering to load balancing.

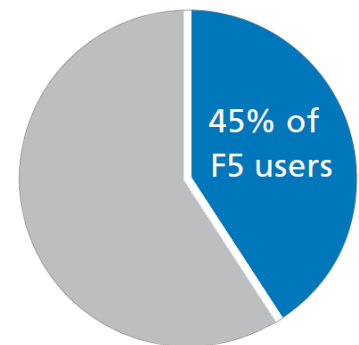
So what is required includes:

- Encryption that still allows the encrypting organization to manipulate streams and/or data.
- Encryption that does not consume a large amount of resources, particularly CPU time on servers.
- Encryption that can utilize a network key store.
- Encryption that can differentiate between internal and external traffic.
- Cipher diversity, the ability to support a wide selection of encryption algorithms.

A security gateway is needed that can provide SSL services with algorithm choices, visibility, integration, and management. A tool that handles encryption of inbound and outbound traffic can make use of a hardware security module (HSM) for key protection, offloads computational resources and the associated operational cost from servers, and either terminates or acknowledges both incoming and outgoing SSL.

F5 Addresses Security Risks

45% of surveyed companies addressed security risks by deploying F5 solutions.



Source: Survey of 116 F5; BIG-IP users; TVID: D62-134-08A

WHITE PAPER

Secure Application Delivery

This ideal security gateway offers encryption to clients without burdening the server CPU with encryption, and it also offers encryption to servers if such is desired. It should not require separate management of a whole selection of keys, and, while acting as a gateway, it should make use of other services that may require unencrypted data between the offloading of incoming connections and the creation of SSL connections to the back end. In such a scenario, the savings in software on the back end could add up quickly, since not as many servers would be required to serve the same number of connections.

A range of options exists for implementing such an architecture, depending upon the needs of a given organization.

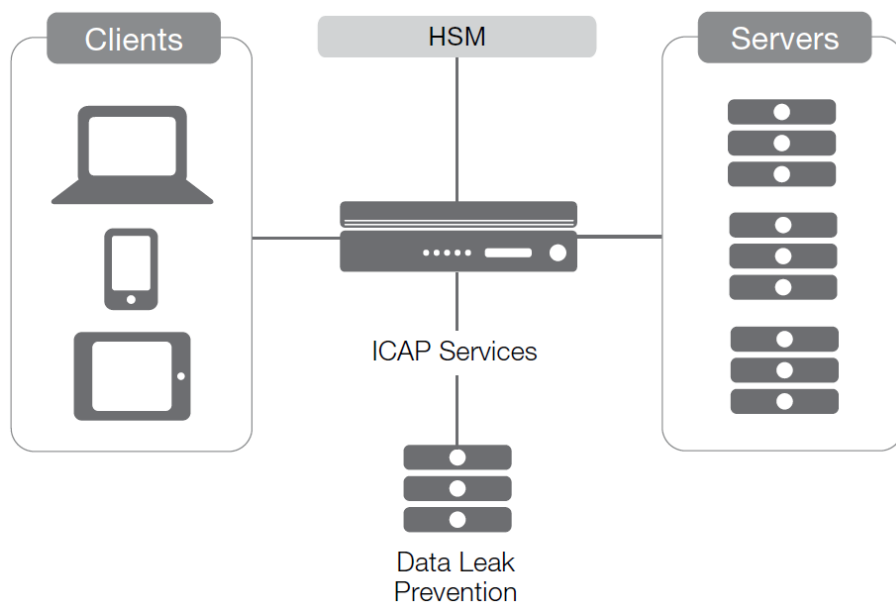


Figure 1: A BIG-IP device provides a platform for advanced SSL services.



Three Styles

The strength of such a solution comes of being placed at a strategic point of control. By sitting at the juncture of the network and the Internet, the ideal device is capable of terminating incoming SSL connections, performing work on the payload of the stream, and then, if necessary, re-encrypting utilizing the internal network's keys. The same is possible with traffic going the other way, although the security measures required for outbound traffic are different from those required for inbound traffic. (Consider the inbound scenario where users need to be authenticated and payloads checked for malicious content, while the outbound stream needs to be checked for sensitive data leaving the data center.) Such a strategic point of control allows for three different styles of encryption manipulation, all valid, depending upon the needs of the organization.

Style one: connection offloading

With termination of the external connection, the gateway device is the endpoint for SSL. SSL is maintained between the device and clients, while the connection to the application server is open and unencrypted. By terminating SSL coming from outside the network, this architecture enables security outside and performance inside. The lack of server-side encryption speeds response times, but results in an insecure environment once an attacker does manage to get inside.

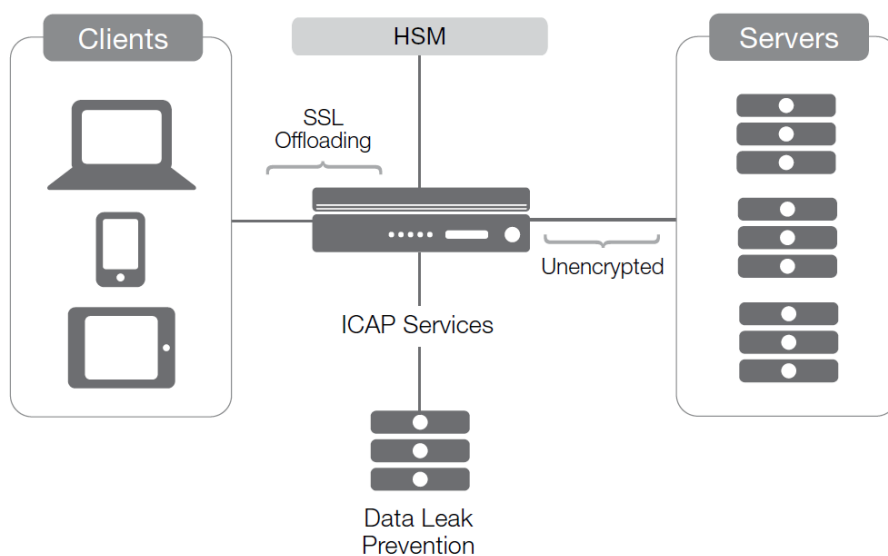


Figure 2: A BIG-IP device terminates inbound SSL without slowing performance farther inside the network.



Style two: SSL visibility

With both inside and outside traffic encrypted, the device at the strategic point of control can manage traffic and adapt and optimize content while the entire network is protected.

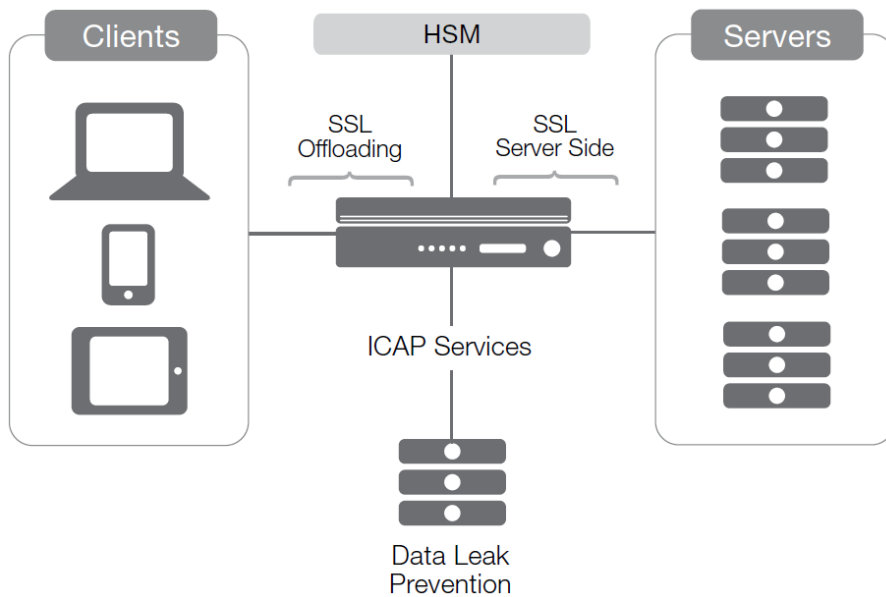


Figure 3: A BIG-IP device can also manage traffic and optimize content inside the network.

Style three: SSL transparent inspection

Increasingly, IT teams are driving toward a solution that maintains a single round of encryption for the entire connection, from client to server, yet still permits intervening data manipulation to facilitate use of a variety of tools, such as those for data leak prevention (DLP), pre-access authentication, and load balancing.

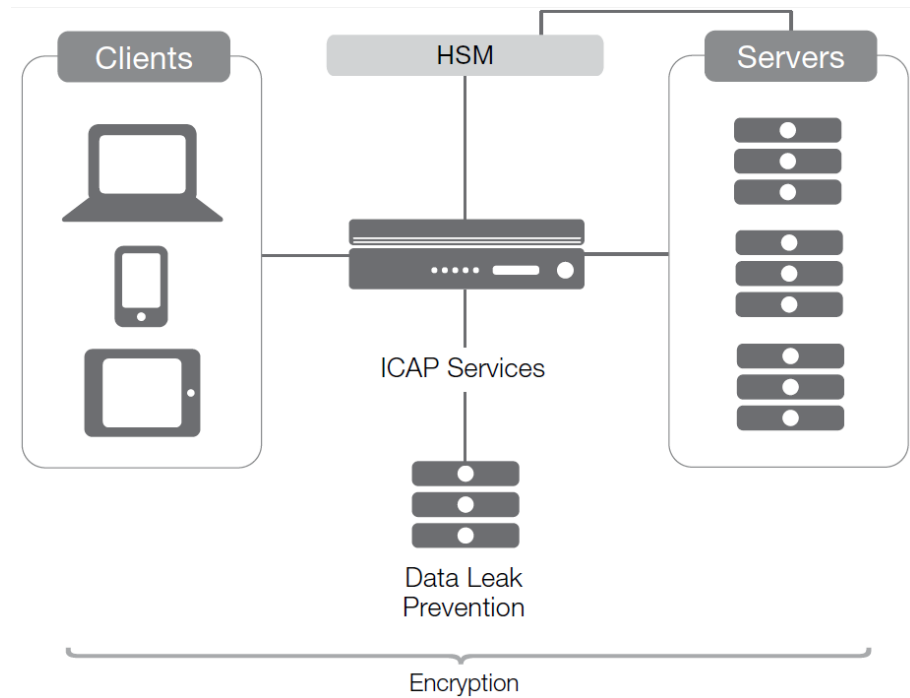


Figure 4: Traffic is encrypted client-to-server, while the BIG-IP device still enables data manipulation and visibility.

Enter SSL visibility. In this scenario, the server and the device at the strategic point of control share an SSL certificate, granting the gateway device access to streams as they pass through it, without requiring complete termination and recreation of SSL connections.

SSL Transformation

The thing is, the weight of encryption processing in this third scenario falls back upon the servers. This would have been a huge negative in the past, but with the advent of ECC encryption, the CPU cycles required to do such encryption are significantly reduced, as is the space required to store keys. That is a function of the ECC algorithm, which is capable of offering the same level of security with much smaller keys so the processing cost on the client is significantly reduced. (While in the past, many have said, “The client isn’t the problem,” today we’re facing an unprecedented number of mobile clients, and CPU usage is one of the things that eats battery life. So now everyone has to care about the client.)



WHITE PAPER

Secure Application Delivery

While all three public key cryptography systems are secure, efficient, and commercially viable, they differ in the kind of mathematical problems on which they are based. Not only does this affect how vulnerable they are to brute force attacks often used by attackers, it can also lead to differences in the size of the keys generated by the algorithms to provide a certain level of security. The National Institute of Standards and Technology (NIST) provides guidelines for minimum key sizes for each, according to the level of security required. And of course, if an organization has cipher biodiversity, it has several different encryption algorithms to choose from. That means that if a viable hack for RSA key lengths up to 2 Kb comes out tomorrow, it will be possible to simply change the algorithm utilized by the servers and be able to sleep at night.

But ECC is more secure with the same key length. Indeed, in most cases ECC can do with a very short key what contemporary algorithms offer with very long ones. This efficiency is a function of the math utilized inside the ECC algorithms (and there are several). The idea that ECC reduces CPU overhead for encryption while simultaneously maintaining or improving security, with smaller key sizes, makes it a powerful new tool that enables secure client-to-server communications.

The F5 platform takes advantage of this important new tool. At this time, BIG-IP devices offer DSA for signing and ECC for encryption. That's on top of the RSA, AES, and 3DES algorithms that F5 has always supported.

With ECC helping to address the CPU expense of encryption, that leaves the other big issue that encryption has traditionally struggled with: visibility. If everything is encrypted, how do devices for data loss prevention (DLP) and external authentication get access to the unencrypted content?

SSL Visibility

Enter the networked HSM tool. While HSM solutions are increasingly common, their use is generally limited to storage of credentials. Yet they can also be an enabling technology, given the right connections and usage. Sharing certificates between a highly flexible device at the edge of the network and servers in the heart of the network offers the ability for the device at the edge to view the unencrypted traffic without terminating the streams. By storing symmetric keys, asymmetric private keys (including digital signatures), and more, the gateway device protects critical information while making it available across the network. That provides the ability to send login information to LDAP devices, incoming streams through virus scanners, and outgoing data through DLP devices, for example.



WHITE PAPER

Secure Application Delivery

Hardware security modules supported by F5 are hardened devices with tamper evidence and NIST 140-2 Level 3 support. By placing them as security gateways at a network's strategic point of control, an organization can achieve not only efficient SSL services but visibility into the encrypted data. The result is a complete system with enhanced security that still has all of the capabilities of a terminating, bidirectional proxy.

The possibilities from this strategic point of control are immense. Utilizing the F5® iRules® scripting language, an organization can keep sensitive information inside the data center. For example, a community-maintained iRule shared on F5® DevCentral™ can catch most credit card information and scrub it on the way out of the building.

In another example, IPv6 gateway functionality or even SPDY translation can be placed on the F5 device.

Many of these options are not possible in a normal SSL environment. By having access to the incoming connection, however, the BIG-IP device also gains access to the payload of the data stream, allowing it to act intelligently on that data, not just on the stream. And because the device doesn't have to terminate the connection, performance is not degraded as significantly as it would be if two separate SSL connections had to be maintained.

The BIG-IP device also can direct a request for a secured page to authentication, only returning the page after the user has successfully logged in and the authorization server confirms that the user has rights to the page in question. The key here is that authentication, authorization, and accounting (AAA) are occurring, and on failure, the user can be redirected to a page removed from the core network. That means unauthorized users never get past the BIG-IP device to the internal network, and attackers masquerading as valid users never reach critical systems to attempt exploits. Their connections don't go that far, since they are stopped and authenticated before entering the internal network. In DDoS scenarios, a high-performance device that can redirect attacking connections to a quarantine network will help keep the public-facing network available to actual users, thwarting the intent of the DDoS attack.

FIPS 140-2 Support

Often, key protection is the biggest problem introduced by a large encryption methodology. The need to keep certificates up-to-date, manage which device has which certificate, and determine when certificates can be moved from machine to machine — particularly in a virtualized environment — can limit the number of certificates deployed and even affect architecture decisions.



WHITE PAPER

Secure Application Delivery

For those organizations requiring FIPS 140-2 Level 2 support, many BIG-IP device models can have this support added in. Those organizations wanting to secure information such as keys and certificate pass phrases but that do not require FIPS 140-2 support can do so using built-in Secure Vault technology.

BIG-IP devices also are capable of acting as HSMs or of utilizing one already in the network. Since HSMs provide secure storage of certificates, utilizing one offers a single point of reference for managing key usage and expirations, simplifying certificate management and freeing IT staff to handle other important business functions.

Additional Functionality, One Device

Every BIG-IP platform comes with maximum encryption throughput licensed as part of the cost, but another high-value advantage to this F5 security solution is the ability to add or turn on additional functionality. Built-in DDoS protection keeps websites online and more secure during an attack, and the ability to add a variety of security and performance modules means IT needs to manage only one device, not many, to meet the organization's application delivery needs. Simplified management and faster response times come with a single point of processing — the strategic point of control.

Conclusion

The future of IT includes encryption everywhere to avoid the insecurity introduced by having some encrypted and some non-encrypted web pages. This implies a heavy burden on infrastructure and the need to implement encryption everywhere while still intelligently interacting with traffic for functions such as load balancing and data leak prevention, among a host of others.

F5 offers a solution that can simplify and accelerate a highly secure infrastructure. With a host of capabilities and the ability to look at traffic as it passes without terminating SSL connections, BIG-IP devices also enable calls to ICAP-enabled devices for specialized processing based upon the industry or technology in use. With support for HSM, too, the F5 solution is highly adaptable to the needs of the organization. F5 BIG-IP platforms also are able to fend off DDoS attacks at the required levels, stopping attacks while continuing to route valid users to applications.

In all, the BIG-IP product family delivers granular control, scalability, and flexibility not available with other traffic management or security technologies.

WHITE PAPER

Secure Application Delivery



F5 Networks, Inc.
401 Elliott Avenue West, Seattle, WA 98119
888-882-4447 www.f5.com

Americas
info@f5.com

Asia-Pacific
apacinfo@f5.com

Europe/Middle-East/Africa
emeainfo@f5.com

Japan
f5j-info@f5.com

©2015 F5 Networks, Inc. All rights reserved. F5, F5 Networks, and the F5 logo are trademarks of F5 Networks, Inc. in the U.S. and in certain other countries. Other F5 trademarks are identified at f5.com. Any other products, services, or company names referenced herein may be trademarks of their respective owners with no endorsement or affiliation, express or implied, claimed by F5. WP-2900 0113