



Dynamic Perimeter Security with IP Intelligence

Dynamic perimeter security and Internet host reputation evaluation have emerged as a primary security focus for businesses. In answer to this, F5 introduces the IP Intelligence service that delivers a database of over 1 million malicious Internet addresses. With IP Intelligence at the security perimeter, organizations gain near real-time protection against phishing, attackers, and scanners.

White Paper
by David Holmes



Introduction

Defending the integrity of a modern enterprise network perimeter is a difficult task. Security architects must contend with threats such as customer asset losses due to phishing, data breaches from hacking penetration, and advanced persistent threats via malware distribution.

Phishing

Phishing has become one of the top threats for financial organizations. An APWG Internet Policy Committee report on phishing observed that after a decrease in 2010, phishing attacks were on the rise again in 2011.¹ Today, financial organizations are particularly sensitive to the financial losses and damage to public perception associated with phishing, and they struggle to find solutions that effectively combat it.

Data breaches

Data breaches continue to be high-profile news stories as organizations are penetrated by attackers from anonymous proxies around the globe. An eye-opening report by researchers from Verizon and the U.S. Secret Service show that in 2010, 50 percent of data breaches utilized some form of hacking, and 49 percent incorporated some form of malware.²

Advanced persistent threat

Modern malware can infect nearby hosts, attack external targets, generate spam, and participate in advanced persistent threat (APT) activities.

Threat Vector	Result
Malware	Breaches, data loss
Phishing	Customer asset loss
Scanners	Network reconnaissance
Botnets	Advanced persistent threat

Figure 1: Today's threats can cause a host of problems that can result in disastrous losses.



WHITE PAPER

Dynamic Perimeter Security with IP Intelligence

With these threats in the mix, defending a network perimeter is a formidable undertaking. After deploying firewalls, scanning for vulnerabilities, mitigating network DDoS attacks, and integrating web application firewalls (WAFs), there's still a significant problem: no matter how secure the enterprise network, hosts on the outside Internet may not be safe. But which hosts? How can outbound connections be authorized? How can inbound requests be intelligently evaluated?

The answer is by reputation. By intelligently evaluating the reputation of Internet hosts, organizations can prevent malicious hosts from disrupting business functions, stealing data, or probing resources. They must consider the reputation of every outbound Internet destination address. Some of those addresses map to malicious hosts such as phishing proxies or botnet servers. Incoming connections can be from active, malicious, or suspect addresses such as anonymous exit nodes or scanners.

To help IT organizations tackle this project, F5 is offering a new set of context-oriented services, including IP Intelligence, that will be key to making dynamic decisions about traffic management.

Welcome to IP Intelligence

The F5 product portfolio has long been known for providing intelligence and agility to the network. New to its BIG-IP system is BIG-IP Global Delivery Intelligence, which enhances delivery decisions based on context. The BIG-IP Global Delivery Intelligence with IP Intelligence service secures an enterprise's perimeter against malicious Internet hosts. BIG-IP Global Delivery Intelligence also provides location-based data to services such as those provided by BIG-IP Global Traffic Manager (GTM).

WHITE PAPER

Dynamic Perimeter Security with IP Intelligence

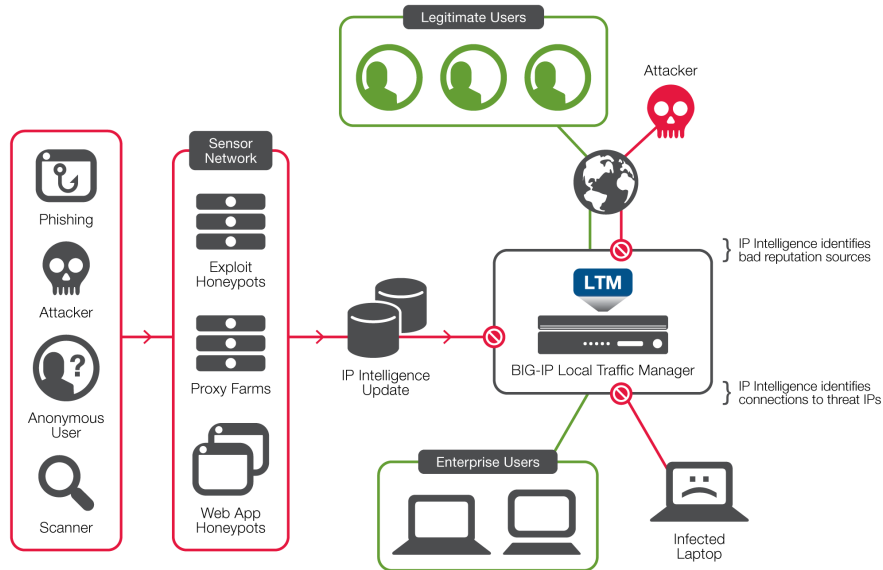


Figure 2: IP Intelligence gathers reputation data for use by F5 solutions.

IP Intelligence—which is built into F5 ADCs positioned at the perimeter of the network—maintains information about over 1 million malicious addresses on the Internet, and can block connections to and from those addresses. This database of addresses is refreshed every five minutes from the cloud to minimize the threat window and keep an organization's data—and its reputation—safe.

Behind the IP Intelligence technology is a global threat sensor network composed of:

- Semi-open proxy farms
- Exploit honeypots
- Naive user simulation
- Web application honeypots
- Third-party sources

These sensor components capture and contribute incident data to a threat analysis network, which produces the IP Intelligence host database, with a classification for each entry.

Phishing

Phishing remains the top security concern for consumer financial organizations. User credentials to online bank accounts are still a primary phishing target, but the latest trend is for attackers to capture credentials that lead not only to e-commerce sites, but to any site where there is potential financial gain. For example, registered accounts with which users enter credit card numbers to purchase retail items are especially attractive targets. So far, banks have been covering these losses instead of passing them on the consumer; but the costs of phishing are growing every year, and organizations must seek better ways to mitigate this ongoing threat.

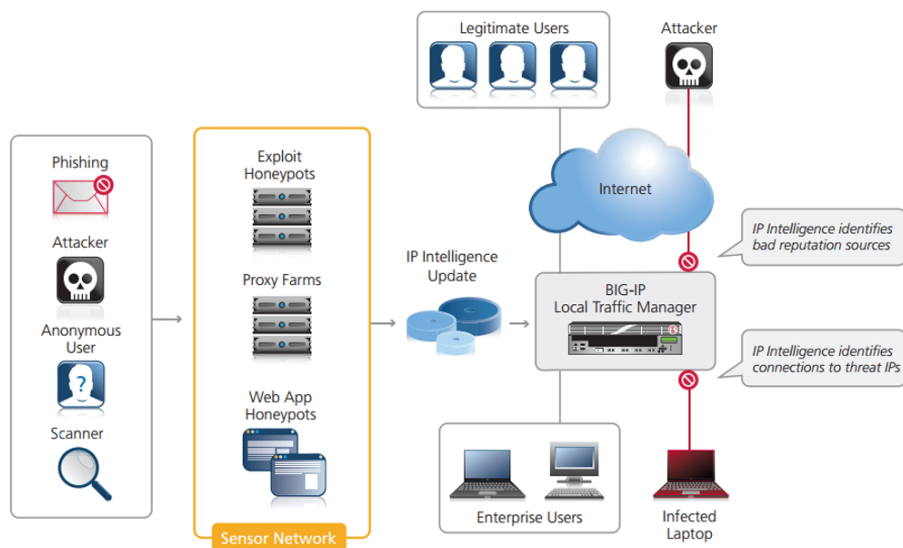


Figure 3: IP Intelligence separates legitimate users from phishing proxies.

Phishing is to some degree a social engineering technique, in that it fools the victim into thinking they are connecting to their bank. Because most sites with financial assets use SSL certificates to secure the connection, the phishing attacker must trick the victim into connecting to a phishing proxy site for which they have a certificate that looks like the user's target site, but whose URL doesn't really match. This prevents the browser from displaying an alert. As long as the victim is not paying close attention to the browser bar, the attack will succeed at the phishing proxy.

IP Intelligence tracks known phishing proxies, allowing BIG-IP users to prohibit malicious requests from phishing sites, such as man-in-the-middle attacks, or to respond with an alert.

Anonymous Proxies

Anonymous proxy networks, like The Onion Routing (TOR) project, mask network traffic source information using network graphs and multiple levels of encryption. Traffic that passes through the anonymous network will arrive at its destination, but without the original source address. The payload has also been mixed with other traffic and bounced around enough nodes to add delay and repudiation on the behalf of the sender. Traffic coming from these exit nodes may have originated anywhere and is, of course, very difficult to trace.

These attributes make anonymous networks popular with those who want to hide their identities. Sometimes the senders are oppressed citizenry evading dictatorial review; but more often they are hackers or other malicious agents.




Figure 4: An attacker can mask his or her source by using anonymous proxy network.

While some legitimate use cases do exist, one fact cannot be denied: the best malicious hackers in the world use anonymous networks to hide their locations. In June 2011, one of the world's most wanted hackers was finally apprehended when he neglected to mask his source address a single time.

It's critical for enterprises to prevent these individuals from connecting to their networks by implementing a policy prohibiting, or at least discouraging, connections from known anonymous networks. Because IP Intelligence knows, in near real time, which Internet addresses represent the exit nodes of the anonymous network relays, it can successfully block connections from them.



Scanners

Scanners are more dangerous than they appear at first glance. Scanners attempt to connect to different hosts and ports at the enterprise. They often connect but fail to authenticate, or they speak the wrong protocol and appear to give up and go away. The real result though, is that they are gathering reconnaissance data for use in later attacks.

Scanners perform reconnaissance for:

- Conventional network attacks
- Low-bandwidth asymmetric application attacks
- Web application security vulnerabilities

Scanners and Application Attacks

Scanners have always been used to lay the foundation for a conventional network attack, but they've evolved to perform reconnaissance for two forms of application attacks. The first is a new class of low-bandwidth asymmetric attacks. Attackers use scanners to recursively query a website and measure the time interval that certain queries take. With this information, the attacker can turn around and simply request the most computationally expensive queries that the site offers, leading to a denial of service at the database level. This attack is extremely difficult to detect because it does not involve a heavy spike in incoming traffic.

The second type of reconnaissance has seen a massive uptick in the last two years. Attackers have been probing web applications for vulnerabilities such as SQL injections and cross-site scripting. They can then catalogue and sell the data they retrieve, or else use it in a subsequent application attack.

In the BIG-IP system, IP Intelligence stops this class of attacks by preemptively preventing known scanning hosts from connecting to the site. If scanners attempt to evade detection by using anonymous proxies, they will still be denied because IP Intelligence knows about those networks as well. If the scanner attempts to evade detection by setting up a new scanning host, that host address will become known through the IP Intelligence threat analysis network, and will subsequently appear in the IP Intelligence database within five minutes.

WHITE PAPER

Dynamic Perimeter Security with IP Intelligence

Botnet Command and Control

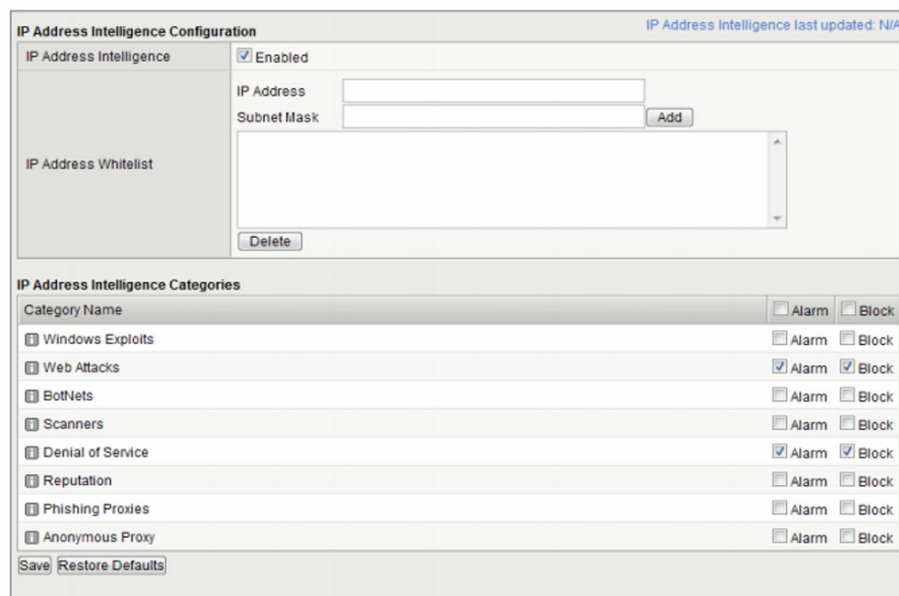
Today's malware can generate enormous amounts of spam, infect nearby hosts, attack external targets via DDoS, and participate in advanced persistent threat (APT) activities. If a device inside the enterprise perimeter is infected with malware and becomes part of a botnet, it may show itself by attempting to connect to a command-and-control (C&C) host on the Internet in order to receive attack orders. The threat analysis network behind IP Intelligence tracks known botnet C&C hosts and includes them in its reputation database.

An administrator can designate virtual servers for outgoing traffic to watch for specific destination ports and then trap and log connections to known botnet C&C hosts. The log indicates which internal devices are infected so administrators can initiate remediation measures.

There are new botnets every day, and their C&C hosts are constantly changing. This necessitates the use of an intelligent threat analysis network like the one behind IP Intelligence. The automatic updates that occur every five minutes to the IP Intelligence database within the BIG-IP system minimize the threat window.

IP Intelligence Control

With BIG-IP Application Security Manager (ASM), IP Intelligence will block incoming connections whose source addresses match its database. A whitelist can be configured by the administrator to bypass IP Intelligence for the selected address in the list.



IP Address Intelligence Configuration IP Address Intelligence last updated: N/A

IP Address Intelligence ☒ Enabled

IP Address:
 Subnet Mask:

IP Address Whitelist

IP Address Intelligence Categories

Category Name	Alarm	Block
Windows Exploits	<input type="checkbox"/>	<input type="checkbox"/>
Web Attacks	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
BotNets	<input type="checkbox"/>	<input type="checkbox"/>
Scanners	<input type="checkbox"/>	<input type="checkbox"/>
Denial of Service	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Reputation	<input type="checkbox"/>	<input type="checkbox"/>
Phishing Proxies	<input type="checkbox"/>	<input type="checkbox"/>
Anonymous Proxy	<input type="checkbox"/>	<input type="checkbox"/>

Figure 5: IP Intelligence includes a whitelist for fine-grain control.



WHITE PAPER

Dynamic Perimeter Security with IP Intelligence

BIG-IP ASM presents IP Intelligence usage in an easy-to-use GUI. IP Intelligence is also available for use with F5's iRules scripting language, which allows IT organizations to customize their evaluation to their needs.

```
when CLIENT_ACCEPTED {

    # Setup High-Speed Logging

    set hsl [HSL::open -proto UDP -pool syslog_server_pool]

}

when HTTP_REQUEST {

    set ip_reputation_categories [IP::reputation [IP::client_addr]]

    set is_reject 0

    if {($ip_reputation_categories contains "Windows Exploits")} {

        set is_reject 1

    }

    if {($ip_reputation_categories contains "Web Attacks")} {

        set is_reject 1

    }

    if {($is_reject)} {

        HSL::send $hsl "Attempt access from malicious ip address [IP
::client_addr]
($ip_reputation_categories), request was rejected"

        HTTP::respond 200 content "
```



WHITE PAPER

Dynamic Perimeter Security with IP Intelligence

```
The request was rejected.  
Attempt access from malicious ip address  
  
" } }
```

Integration with Content Delivery Networks

IP Intelligence can provide its defensive services even when used with a content delivery network (CDN). IP Intelligence can evaluate the original IP address in the X-Forwarded-For (XFF) header. Other solutions, such as intrusion prevention systems (IPSs) or conventional firewalls, examine the source address of the packets and mistakenly evaluate the reputation of the CDN's proxy address rather than correctly evaluating the reputation of the original client source address.

Conclusion

Attack methodologies change. Threat vectors change. Once-powerful defenses become obsolete. As the pace of these changes increases, it's unrealistic to wait for defense methods to catch up, or for new products to be developed and releases pushed out. Financial enterprises need near real-time host protection, where the bad actors are known to the rest of the world even as they act.

The BIG-IP system and the IP Intelligence service at the security perimeter provide the up-to-date network threat intelligence that is independent of attack type. If a new criminal actor attacks this today, the IP Intelligence database will begin providing protection from the hosts involved within minutes.

IP Intelligence provides protection against phishing, defends against anonymous attackers and scanners, and blocks botnet control. By leveraging IP Intelligence in their BIG-IP deployments, financial organizations can minimize the threat window and protect valuable enterprise assets.

¹ [Global Phishing Survey: Trends and Domain Name Use in 2H2010](#). APWG Internet Policy Committee, April 2011.

² [2011 Data Breach Investigations Report](#). Verizon RISK Team, 2011.

WHITE PAPER

Dynamic Perimeter Security with IP Intelligence



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-SEC-IP-24505 0113