BIG-IP ASM Operations Guide

A Web Application Firewall that Guards Your Critical Apps

With F5© BIG-IP Aplplication Security Manager™ (ASM), organizations gain the flexibility they need to deploy Web Application Firewall services close to apps to protect them wherever they reside—within a virtual software-defined data center, managed cloud service enviornment, public cloud, or traditional data center.

Contents

Figures

Tables

About This Guide

The goal of this guide is to help F5® customers keep their BIG-IP® system healthy, optimized, and performing as designed. It was written by F5 engineers who assist customers with solving complex problems every day. Some of these engineers were customers before joining F5, and their unique perspective and hands-on experience serves the guides F5 customers have requested.

This guide describes common information technology procedures, as well as those which are exclusive to BIG-IP systems. There may be procedures particular to your industry or business that are not identified. While F5 recommends the procedures outlined in this guide, they are intended to supplement your existing operations requirements and industry standards. F5 suggests that you read and consider the information provided to find the procedures to suit your implementation, change-management process, and business-operations requirements. Doing so can result in higher productivity and fewer unscheduled interruptions.

Refer to Feedback and notifications for information on how to help improve future versions of the guide.

Before using this guide

To get the most out of this guide, first complete the following steps, as appropriate to your implementation:

- Install your F5 platform according to its requirements and recommendations. Search the AskF5™ (support.f5.com) for “platform guide” to find the appropriate guide.
- Follow the general environmental guidelines in the hardware platform guide to make sure of proper placement, airflow, and cooling.
- Set recommended operating thresholds for your industry, accounting for predictable changes in load. For assistance contact F5 Professional Services (f5.com/support/professional-services).
- Familiarize yourself with F5 technology concepts and reviewed and applied appropriate recommendations from ***F5 BIG-IP TMOS: Operations Guide***.

Note For information about how to locate F5 product manuals, refer to AskF5 article: K12453464: Finding product documentation on AskF5.

Limits of this guide

This guide does not focus on installation, setup, or configuration of your BIG-IP system or modules. There is a wealth of documentation covering these areas in AskF5 (support.f5.com) The F5 self-help community, DevCentral™ (devcentral.f5.com), is also a good place to find answers about initial deployment and configuration.

The following figure shows where the F5 operations guides can best be applied in the product life cycle.

Figure 0.1: F5 documentation coverage

Glossary

A glossary is not included in this guide. Instead, the Glossary and Terms page (f5.com/glossary) offers an up-to-date and complete listing and explanation of common industry and F5-specific terms.

Customization

Customization may benefit your implementation. You can get help with customization from a subject matter expert, such as a professional services consultant, from F5 Professional Services (f5.com/support/professional-services).

Issue escalation

Refer to Optimizing the Support Experience for issue escalation information.

If you have an F5 websupport contract, you can open a support case by clicking **Contact support** on AskF5 (support.f5.com)

Feedback and notifications

F5 frequently updates the operations guides and new guides may be released as needed. If you would like to be notified when new or updated content is available, or if you have feedback, corrections, or suggestions to improve this guide, email opsguide@f5.com. F5 internal users can file a request using Service-Now.

Configuration utility

The BIG-IP Configuration utility is the name of the graphic user interface (GUI) of the BIG-IP system and its modules. It is a browser-based application you can use to install, configure, and monitor your BIG-IP system.

For more information about the Configuration utility, refer to **Introducing BIG-IP Systems** in ***BIG-IP Systems: Getting Started Guide***.

Note For information about how to locate F5 product manuals, refer to AskF5 article: K12453464: Finding product documentation on AskF5.

Command-line syntax

We show command line input and output in courier font. The corresponding prompt is not included. For example, the following command shows the configuration of the specified pool name:

```
tmsh show /ltm pool my_pool
```

The following table explains additional special conventions used in command-line syntax:

Table 0.1 Command-line syntax

Character	Description
<>	Identifies a user-defined variable parameter. For example, if the command has **<your name>**, type in your name but do not include the brackets.
[]	Indicates that syntax inside the brackets is optional.
...	Indicates that you can type a series of items.

TMOS Shell syntax

The BIG-IP system includes a utility known as the TMOS® Shell (**tmsh**) that you can use to configure and manage the system at the command line. Using **tmsh**, you can configure system features and set up network elements. You can also configure the BIG-IP system to manage local and global traffic passing through the system and view statistics and system performance data.

You can run **tmsh** and issue commands in the following ways:

- You can issue a single **tmsh** command at the BIG-IP system command line using the following syntax:

```
tmsh [command] [module . . . module] [component] (options)
```

- You can open **tmsh** by typing tmsh at the BIG-IP system command line:

```
(tmsh)#
```

At the **tmsh** prompt, you can issue the same command syntax, leaving off **tmsh** at the beginning.

Note You can use the command line utilities directly on the BIG-IP system console, or you can run commands using a remote shell, such as the SSH client or a Telnet client. For more information about command line utilities, refer to the ***Traffic Management Shell (tmsh) Reference Guide***.

Finding other documents

For information about how to locate F5 product manuals, refer to AskF5 article: K12453464: Finding product documentation on AskF5.

Conventions Unique to the BIG-IP ASM Guide

Common terms and concepts

This guide also assumes that you have some familiarity with various Layer 7 (L7) Hypertext Transfer Protocol (HTTP) concepts, such as Uniform Resource Identifier (URI)/Uniform Resource Locator (URL), method, header, cookie, status code, request, response, and parameters.

BIG-IP ASM terminology

Table 0.2 BIG-IP ASM Terminology

Term	Definition
Add All Entities	**Add All Entities** creates an explicit object in a security policy for each entity, such as file type, parameter, or URL. Supports different attributes for each entity. Replaces **Tightening** in previous versions.
Alarm	If selected, the BIG-IP ASM system records requests that trigger the violation in the Charts screen, the system log (**/var/log/asm**), and possibly in local or remote logs (depending on the settings of the logging profile).
Attack signature	Textual patterns which can be applied to HTTP requests and/or responses by BIG-IP ASM to determine if traffic is malicious. For example, the string “<script>” inside an HTTP request triggers an attack signature violation.
Attack signature set	A collection of attack signatures designed for a specific purpose (such as bot detection or Apache).
Block	To prevent a request from reaching a protected web application.
Block	If selected (and enforcement mode is set to **Blocking**), the BIG-IP ASM system blocks requests that trigger the violation.
Blocking response page	A blocking response page can be displayed to a client when a request from that client has been blocked. Also called blocking page and response page.
Bot	A software application, tool, or agent that runs automated tasks over the Internet. Can be either malicious or benign.
Comprehensive policy type	The Comprehensive policy type provides the most granular definitions, includes all security features, and is suited for advanced users or customers with extreme security needs. This policy type includes all elements in the Enhanced policy type, and adds URLs and meta characters, parameters (meta characters and URLs), and dynamic parameters (using statistics). This security policy typically takes longer to deploy.
Content profile	Configurable file which you can add to a security policy to allow BIG-IP ASM to understand and parse specific content formats, such as JSON or XML.

Term	Definition
Denial of Service (DoS)/ Distributed Denial of Service (DDoS)	Denial-of-service (DoS)/distributed-denial-of-service (DDoS) is an attack intended to disrupt a network by consuming excessive bandwidth or cause a machine to stop servicing requests by overloading it. Legitimate users are consequently denied service. In BIG-IP ASM, DoS refers to layer7 DoS attacks (L7 DoS). A DDoS is a DoS attack executed by multiple attackers. In BIG-IP ASM, DDoS is often generically called DoS.
Enforce	To remove an entity or attack signature from staging.
Enforced	Indicates that the relevant object has been removed from staging. If a relevant violation occurs, the request can be blocked (if the overall policy is in Blocking mode.) See Enforcement mode.
Enforcement mode	Security policies can be in one of two enforcement modes: Transparent mode In Transparent mode, **Blocking** is disabled for the security policy, and you cannot set the violations to **Block** on the Blocking screen. Traffic is not blocked even if a violation is triggered. You can use this mode and staging when you first put a security policy into effect to make sure that no false positives occur that would stop legitimate traffic. Blocking mode In Blocking mode, **Blocking** is enabled for the security policy, and you can enable or disable the **Block** flag for individual violations. Traffic is blocked when a violation occurs if the following conditions are met: you configure the system to block that type of violation, the staging period is over, you removed all entities (explicit and wildcard) whose staging period is over from staging, and deleted wildcard entities with tightening (whose tightening period is over) from the security policy. You can use this mode when you are ready to enforce the security policy. You can change the enforcement mode for a security policy on the Policy Properties screen or the Policy Blocking Settings screen.
Enforcement readiness period	For each security policy, you can configure the number of days used as the enforcement readiness period, also called staging. Security policy entities and attack signatures remain in staging for this period of time before the system suggests that you enforce them. Staging allows you to test security policy entities and attack signatures for false positives without enforcing them. The default value of 7 days works for most situations so you typically do not need to change it.
Enhanced policy type	The Enhanced policy type provides additional granularity and security features suited for customers with higher (and, typically, specific) security needs. This policy type includes all elements in the Fundamental policy type, and also includes parameters and lengths (global level), cookies, and methods.
Entities	The elements of a security policy, such as **HTTP** methods, as well as file types, URLs, and/or parameters, which have attributes such as **byte length**. Also refers to elements of a security policy for which enforcement can be turned on or off, such as an attack signature.

Term	Definition
Explicit Entities Learning	The Explicit Entities Learning scheme provides you with three tools for instructing BIG-IP ASM what to learn and how to learn it. Three learning modes apply to the four main entity types that have attributes: URLs, Parameters, File Types, and Redirection Domains. The modes are Add All Entities, Never (Wildcard Only) and Selective.
False positive	An instance when BIG-IP ASM treats a legitimate request as a violation.
File types	Examples of file types are .php, .asp, .gif, and .txt. They are the extensions for many objects that make up a web application. File Types are one type of entity a BIG-IP ASM policy contains.
Fundamental policy type	The Fundamental policy type provides granularity sufficient for most organizations creating a generalized security policy that is easy to maintain. This policy type includes HTTP protocol compliance, evasion techniques, file types and lengths, attack signatures, and the request length exceeds predefined buffer size violation. This is the default setting.
Ignore/Ignore Suggestion	The action of instructing BIG-IP ASM to no longer create learning suggestions for a specific violation or entity.
Illegal request	A request which violates a security policy.
Learn	The action of allowing violations or entities such as File Types, URLs, and Parameters to the security policy. Entities which have been learned can be either enforced or staged.
Learn	If selected, the BIG-IP ASM system generates learning suggestions for requests that trigger the violation.
Learning	The iterative process BIG-IP ASM uses to adapt a policy in order to ultimately prevent false positive violations and accurately profile legitimate interactions with the application. Also called Policy Building.
Learning score	A value from 1 to 100 that determines the confidence level of a suggestion, before it is accepted into the policy. After a suggestion reaches a score of 100%, it can be safely added to the policy.
Learning mode	Learning mode is a setting indicating how learning suggestions are processed by the BIG-IP ASM system. The learning process makes a security policy more accurate by verifying how it complies with traffic requests. If the learning process finds discrepancies between the security policy and the traffic requests, it translates the discrepancies into a learning suggestion for modifying the security policy. **Automatic**: Learning suggestions are accepted when the Learning Score reaches 100%. **Manual**: You must manually accept learning suggestions. **Disabled**: Automatic policy is not running and no suggestions are being created.

Term	Definition
Learning suggestion	The BIG-IP ASM system generates learning suggestions for requests that cause violations and do not pass the security policy checks. You can examine the requests that cause learning suggestions, and then use the suggestions to refine the security policy. In some cases, learning suggestions may contain recommendations to relax the security policy. When dealing with learning suggestions, make sure to relax the policy only where false positives occurred, and not in cases where a real attack caused a violation.
Legal request	A request which has not violated the security policy, or a request which involves an entity that is in staging.
Local Traffic policies	Local Traffic policies allow you to direct HTTP traffic based on rules, such as the presence of a specific URI, and to take action on that traffic based on the rule, such as sending a request to a specific security policy. Also called Central Policy Management (CPM).
Loosening	The process of adapting a security policy to allow specific entities such as File Types, URLs, and Parameters. The term also applies to attack signatures, which can be either manually or automatically disabled—effectively removing the signature from triggering any violations.
Modular blocking	Modular blocking is the process of gradually changing a security policy's mode from Transparent to full Blocking mode.
Never (Wildcard Only)	This learning mode is based on a wildcard representation of file types, parameters, cookies, URLs, and redirection domains. The wildcard, indicated by an asterisk (*****) is never removed from the policy by BIG-IP ASM if the policy is configured in this mode.
Parameters	Parameters consist of "name=value" pairs, such as OrderID=10. The parameters appear in the query string and/or POST data of an HTTP request. Consequently, they are of particular interest to BIG-IP ASM because they represent inputs to the web application.
Policy Builder	The BIG-IP ASM automatic policy building tool. Also called Real Traffic Policy Builder® (BIG-IP 11.6) and Automatic Policy Builder (BIG-IP 12.0).
Policy Diff	The Policy Diff tool allows you to easily view differences between two security policies, and to merge configured elements (such as URLs and Login Pages) from one security policy to another.
Policy Merge	BIG-IP ASM policy merge option combines two security policies. In the merge process, the system compares, and then merges, specific features from one security policy to another.

Term	Definition
Rapid Deployment policy	The Rapid Deployment security policy provides security features that minimize the number of false positive alarms and reduce the complexity and length of the deployment period. By default, the Rapid Deployment security policy includes the following security checks: Performs HTTP compliance checks Stops information leakage Prevents illegal **HTTP** methods from being used in a request Checks response codes Enforces cookie RFC (Request for Comment) compliance Applies attack signatures to requests and responses
Selective Learning	In Selective Learning mode, the BIG-IP ASM system only creates suggestions for adding new policy entities if they have different attributes from the wildcard (for example, a violation created because a parameter had a value length of 70, when the wildcard was set with an allowed value length of 25).
Staging	Staging means that the system applies policy attack signatures to web application traffic, but does not apply the blocking policy action to requests that trigger those attack signatures. The default staging period is seven days. Whenever you add or change signatures in assigned sets, those are also put into staging. You also have the option of putting updated signatures in staging.
Tightening	The process of enforcing entity lists such as File Types, Parameters, and URLs (removing the wildcard), and enforcing those attack signatures which were not triggered during the enforcement readiness period.
TPS/RPS	Transactions per second (TPS)/requests per second (RPS). In BIG-IP ASM, these terms are used interchangeably.
Tuning	Making automatic or manual changes to an existing security policy to reduce false positives and increase the policy's security level.
URI/URL	The Uniform Resource Identifier (URI) specifies the name of a web object in a request. A Uniform Resource Locator (URL) specifies the location of an object on the Internet. For example, in the web address, http://www.siterequest.com/index.html, **index.html** is the URI, and the URL is http://www.siterequest.com/index.html. In BIG-IP ASM, the terms URI and URL are used interchangeably.
Violation	Violations occur when some aspect of a request or response does not comply with the security policy. You can configure the blocking settings for any violation in a security policy. When a violation occurs, the system can **Learn**, **Alarm**, or **Block** a request (blocking is only available when the enforcement mode is set to **Blocking**). Violation names displayed in the Violation List are the names used as reference in the iRule **ASM::custom_violation** command and the **ASM::violation name** command. The violation name is also used in API, F5® iControl® and TMAPI (Target Manager API) code. You can also create user-defined violations if you need them on your system.

HTTP request components

Request syntax:

```
https://www.testsite.com/folder/page.html?parameter1=value1&parameter2=value2
```

Table 0.3 HTTP request components

Components	Request syntax
Protocol	https
Host	www.testsite.com
Path	/folder/page.html (in the BIG-IP ASM system referred to as the URL)
Query string	parameter1=value1&parameter2=value2
Parameters	parameter1, parameter2
Parameter Values	value1, value2

Table 0.4 Additional HTTP request components important to the BIG-IP ASM system

Component	Request syntax
POST	/login.php HTTP/1.1
Referer	http://www.testsite.com/index.php
User-Agent	Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 6.1; WOW64; MS-RTC LM 8)
Content-Type	application/x-www-form-urlencoded
Accept-Encoding	gzip, deflate
Host	www.testsite.com
Connection	Keep-Alive
Content-Length	48
Cookie	SESSION=9899474fd85b473ca496d16b363ec3b8
	parameter1=value1&parameter2=value2
Method	POST
Header	Referrer, Accept-Language, User-Agent, Cookie
Cookie	SESSION
POST Query String/ Request Body	parameter1=value1&parameter2=value2

Introduction

F5® BIG-IP Application Security Manager™ (ASM) enables organizations to protect against the Open Web Application Security (OWASP) Top 10 threats, application vulnerabilities, and zero-day attacks. Leading layer-7 (L7) distributed denial-of-service (DDoS) defenses, detection and mitigation techniques, virtual patching, and granular attack visibility work to thwart sophisticated threats before they reach your servers.

The BIG-IP ASM system allows compliance with key regulatory standards such as Health Insurance Portability and Accountability Act (HIPAA) and Payment Card Industry Data Security Standard (PCI DSS).

With the BIG-IP ASM system, organizations gain the flexibility they need to deploy web application firewall (WAF) services for protecting applications wherever they reside—within a virtual software-defined data center (SDDC), managed cloud service environment, public cloud, or traditional data center.

Web application firewall protection

The BIG-IP ASM system is a WAF. In contrast to a network-layer firewall, which restricts access based on source and destination IP addresses, IP protocols, and TCP/User Datagram Protocol UDP port numbers, a WAF manages traffic based on L7 properties, such as HTTP headers, URIs, parameters, and other web application elements.

WAFs also detect anomalous behaviors from clients attempting to access applications, mitigating the potential impact of automated agents, bots, scanners, and brute-force attacks.

Web applications typically use a two- or three-tier model, in which the web server acts as the presentation layer for the application. The web server receives input from the user and displays outputs, while sending transactions to either an application tier or directly to a back-end database.

In these models, some actions performed on the web application could potentially reach the back-end database, meaning a malicious client can gain unauthorized access to data even if the database servers themselves are not directly accessible to the users.

Benefits of WAF protection

WAFs prevent such access to protect not only the web server, but the application infrastructure as a whole. WAFs are widely used in both Internet-facing applications and internal applications. While most organizations use WAFs in combination with other measures, such as secure coding practices, vulnerability assessments, and software patching and updating, even the most secure application can benefit from WAFs because WAFs do the following:

- Disallow incoming traffic from finding and targeting known vulnerabilities in web applications.
- Enable vulnerabilities to be fixed more quickly and easily, as revealed in web applications by scanners and penetration tests.
- Detect and mitigate malicious bot access to applications.
- Prevent malicious clients from abusing web applications by exploiting flaws in business logic or the application infrastructure.
- Provide effective mitigation much faster than patching application code.

- Reduce the likelihood of an unknown or undiscovered vulnerability being exploited.
- Provide an additional point of control to minimize the risk of developer or administrative error.
- Provide visibility into what types of attacks and scans are targeting the application.
- Prevent malicious and unwanted traffic from consuming server resources.
- Provide a forensic record and event correlation for triaging after a suspected security incident.
- Provide behavioral mitigations, such as web scraping and bot detection, that are very difficult to implement in each application.
- Provide a consistent security posture across all of an organization's applications.

Note F5 recommends that you fix known vulnerabilities in applications when possible.

A BIG-IP ASM security policy consists of multiple parts and layers, all serving the purpose of securing a web application. Some elements of a BIG-IP ASM policy protect your application from specific attacks, while other elements protect against more broad attacks.

Application security policy life cycle

Before you deploy an application security policy, it helps to have an understanding of exactly what you are trying to protect and why. Defining your security problem before you start makes it easier to develop and enforce a security policy. Some use cases require more extensive policy development and tuning, while with other use cases, a simple security policy suffices. For more information about defining your approach to application security, refer to AskF5 article: K07359270: Succeeding with application security.

The application security policy life cycle has three phases: policy deployment, policy tuning, and policy maintenance.

Phase 1: Create and deploy policy

Create a new policy using the template and policy-building mode most appropriate for your web application. Decide whether to automate the policy building process. You can change most decisions made during this phase in later phases.

For information about creating and deploying a policy, refer to ***BIG-IP Application Security Manager: Implementations*** and ***BIG-IP Application Security Manager: Getting Started*** for your system version.

Note For information about how to locate F5 product manuals, refer to the Ask F5 article: K12453464: Finding product documentation on AskF5.

Phase 2: Tune policy

False positive violations are identified and policy settings are adjusted to allow legitimate traffic to pass through to the protected application. This is necessary as some legitimate traffic may not pass set policy rules and may resemble an attack. Over time, the security policy eventually becomes more strict, meaning that policy components that have not triggered any violations are enforced, and acceptance occurs for elements specific to your application, such as file types, parameters, and HTTP methods.

Phase 3: Maintain policy

The security policy adapts to application changes, new security requirements, attack signature updates, and activities, prompted by the review of logs and reports on traffic inspected by the BIG-IP ASM system. During this phase, keep the policy up-to-date and accurate for the application it protects.

A good, finely-tuned security policy deployed according to the security requirements of your application should incur minimal operational costs. A very strict and application-specific security policy can potentially take more time and effort to maintain, especially in light of application changes. A generic, lightweight policy requires very little maintenance, even when applied to multiple or different applications.

Policy tuning details

Of the three phases, policy tuning requires the most attention and skill from an administrator. Administrators can choose to manually tune a security policy or have the BIG-IP ASM system tune the policy automatically.

Tasks included in tuning include adjusting policy blocking settings, populating entry lists, and enforcing entities and attack signatures.

Adjusting policy blocking settings

Depending on the policy's initial settings, the BIG-IP ASM system may need to disable specific violations if they block legitimate requests. Conversely, as more traffic passes the policy without triggering an alarm, the BIG-IP ASM system enforces/enables violation mitigation rules.

At the end of the tuning process, the policy should contain all the relevant violations. This is true whether you started with a fully enforced list and tuned it down ("loosened"), or you started with a blank list and enforced violations over time ("tightened").

Modular or gradual blocking

One of the most successful approaches to tuning a BIG-IP ASM policy involves modular blocking, also known as gradual blocking. Gradual blocking allows some violations to be blocked while others are being tuned. This means you can enable policy components with little-to-no tuning, before the tuning process begins.

Populating entity lists

In an HTTP request, all entities are L7. A BIG-IP ASM policy can contain a whitelist of these entities and disallow entities not on that list, it can contain a blacklist of entities to disallow, including wildcards or matching generic patterns.

A BIG-IP ASM policy usually starts with catch-all wildcards (*) for the entity lists (file types, parameters, URIs, cookies, headers, and entities). Depending on the deployment settings, the BIG-IP ASM system suggests specific entities to populate the lists.

After the lists are populated, wildcards can be removed and the list enforced. Not all entity types are relevant to all security and application requirements, so F5 recommends tuning to include only those that are important to you. For more information about using entities in your security policies, refer to AskF5 article: K74535942: Building web application security policies with entities.

Enforcing entities and attack signatures

Most policy entities have attributes in them, such as byte lengths, values, characters, and others. You can configure a security policy to enforce those attributes.

The tuning phase for each entity is marked with a Staging flag, meaning the entity's attributes are only checked, not enforced.

After all attributes are manually or automatically adjusted to fit legitimate application use, the Staging flag can be manually or automatically removed.

Attack signatures behave the same way. All signatures in a policy start in staging. This means that traffic is not blocked if it matches that signature.

After you disable the signatures causing false positives, you can turn off staging and enforce the rest of the signatures.

Additional security checks

The BIG-IP ASM system can deploy additional security checks in order to protect the application. This is true even for malicious traffic that does not match disallowed policy elements or attack signatures. These checks include session tracking, web scraping, brute force protection, and denial-of-service (DoS) protection.

Session tracking includes the ability to track users and sessions according to their activities in the application, blocking them completely, or logging their requests, whether legitimate or malicious.

Web scraping detects non-human behaviors on the website, flagging, and blocking bots with malicious intents, even if all they do is access legitimate resources on the website.

DoS protection guards the application from heavy traffic patterns and rates that are meant to bring down the application or database, even when each request appears to be in order and does not trigger any of the violations in the policy.

Tips and guidelines

Due to the number of available features and capabilities of the BIG-IP ASM system, administrators may feel overwhelmed. F5 recommends that you keep the following tips and guidelines in mind:

Do not allow the path to implementation to become blocked by a desire to instantly build a perfectly secure and tuned environment. Allow for a learning curve and build your security policy to support the needs of your application and organization.

- Do not feel like you must use a feature simply because it exists.
- It is better to see bad traffic than to not see it.
- When zero day hits, it is better to be in Blocking mode with a current policy than to have to build a new policy from scratch.
- Sometimes providing basic protections for many applications is just as important as providing detailed

protection for one.

- The BIG-IP ASM Policy Builder creates an effective security policy and can save you a lot of time.
- The BIG-IP ASM system is designed to learn while in production. If you do not have a robust QA environment, your application users may supply the best source of legitimate traffic from which the BIG-IP system can learn.
- The BIG-IP ASM system has multiple, layered protections for each attack vector. Do not over-invest time or resources on particular mitigations.
- Start with a policy that loosens security restrictions to allow all legitimate behavior and disallow malicious requests.
- Tighten security restrictions over time to incrementally improve protections.

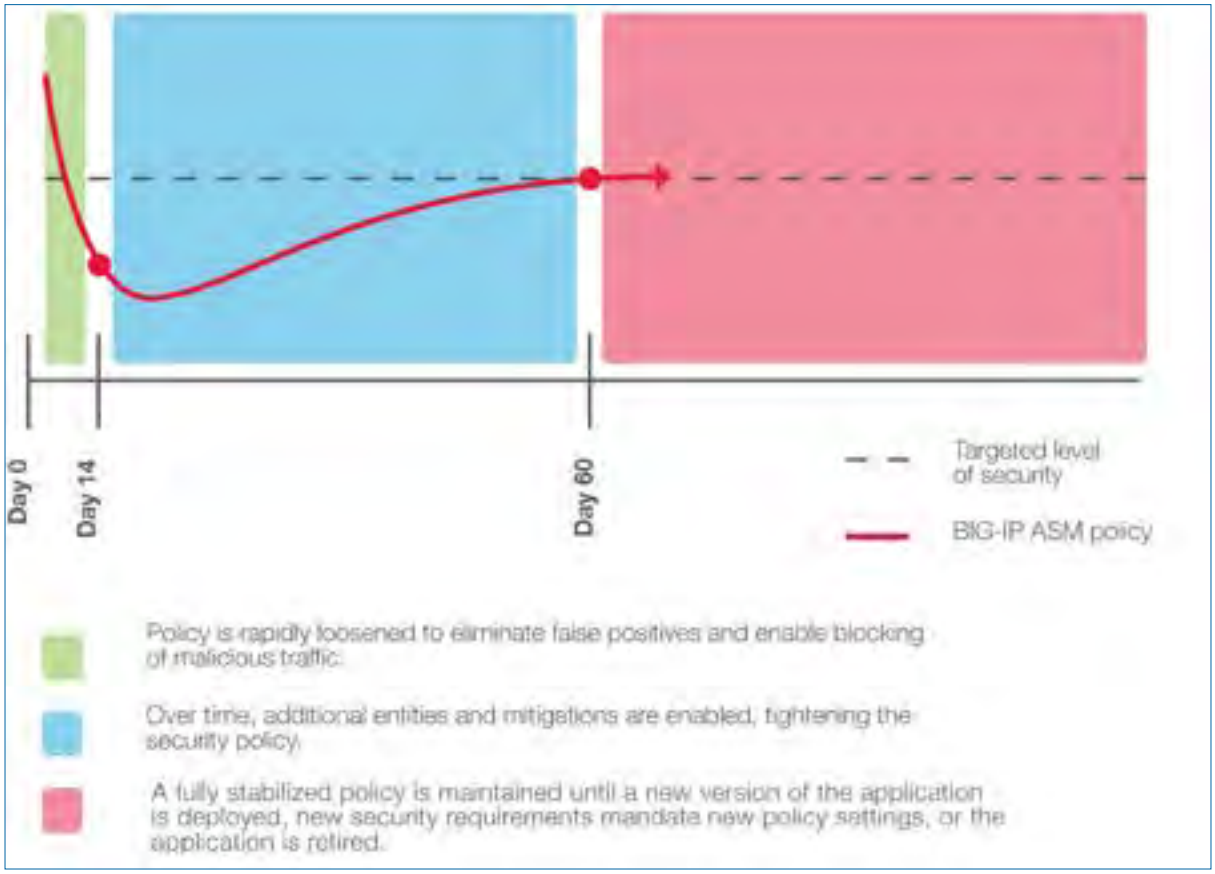

Figure 1.1: Policy adjustment over time

Note BIG-IP ASM security policies exists in either Blocking mode or Transparent mode. There is no "learning mode" to deploy. The BIG-IP ASM system learns the elements of your application as part of an ongoing process.

BIG-IP ASM Event Logging

When appropriately configured and integrated with a security-event management process, the BIG-IP ASM system captures and allows visibility and insights into forensic data. You can use the BIG-IP ASM pre-configured logging options or customize them.

When logging to a remote destination, refer to product documentation to determine whether a custom format is required.

> **Note** F5 technology partner ArcSight sends logs in Common Event Format (CEF), which is a standard for the Security Information and Event Management (SIEM) industry. Many logging and reporting products can properly consume messages in this format.

For more information and guidance, refer to AskF5 article: K9435: Overview of the Storage Format option for a remote logging profile.

Logging display formats

When the BIG-IP ASM logging format and log server are appropriately configured, the system logs every violation identified in a specific request as a single log message.

Violation logs run at the command line

In the following example violation log message at the command line, the violations for a specific request and support ID number appear in bold text for emphasis.

```
Jan 10 09:15:37 bigip.f5demo.com ASM:unit _ hostname="bigip.f5demo.com",management _
ip _ address="10.128.1.245",http _ class _ name
```

```
="/demo/dvwa.vlab.2",web _ application _ name="/demo/dvwa.vlab.2",policy _ name="/demo/
dvwa.vla b.2",policy _ apply _ date="20160110 09:14:34",violations="Illegal meta
character in value,Attack signature detected,Illegal parameter value
length",support _ id="623951107638431468",request _ status="blocked",response _
code="0",ip _ cli ent="10.128.10.1",route _
domain="0",method="GET",protocol="HTTP",query _ string="id=%27OR+1% 3D1 . . .
```

Support ID

The support ID number identifies the request and is associated with the local request logs, manual learning suggestions, and the remotely logged message.

Violations

The violations found in the previous example include the following:

- Illegal meta character in value.
- Attack signature detected.
- Illegal parameter value length.

Violation log in the Configuration utility

In the BIG-IP ASM, under the **Request Details** tab, the **Violations** area shows violations triggered by the request and identifies violation types. The **Learn** column indicates whether the current security policy needs to learn this violation in order to pass it.

The log message also includes the full query string and POST data, which may be used to determine whether a formatted request is legitimate or malicious.

> **Note** Query string data is constrained by the maximum entry limit.

Logging both valid and violation requests may significantly increase the size of your logging files.

By default, the local log storage is finite with a maximum capacity of 3 million records stored across all BIG-IP ASM security policies and 2 GB in database table size.

Log entries are rotated out on a strict age basis. If you log multiple applications locally, it is possible for one application to generate more than its share of messages, filling the log and pushing out entries for other applications before they can be investigated.

The local log provides detailed information for all learning suggestions. If the log rate is too high, it is possible for the BIG-IP ASM system to lose detailed information for learning suggestions that are current for the policy.

Remote logging

Local logging may impact disk performance, especially on systems with slow disk subsystems.

F5 recommends externally logging all requests, or at the very least all legal request, to a security information and event management system for better data retention, searching, event correlation, and scalability.

For logging on remote storage, TCP and UDP are used for log delivery. UDP has lower overhead but can only log up to 1KB of data per log entry, while TCP delivers greater reliability and allows up to 64KB of data per log entry. This greater capacity for data makes TCP a better choice as it allows storing complete log entries.

F5 recommends using session tracking to log only violations rather than all requests. Session tracking allows you to log all requests which cross specific violation thresholds for only specific sessions. This reduces load and data retention requirements by capturing full requests for only suspicious activities without logging the full requests for all sessions.

When using the Remote Storage type **Remote** in BIG-IP 11.6, or **Comma Separated Values** in BIG-IP 12.0, consider the order in which the **Selected Items** are defined in the **Storage Format**. You can also select an option for logging to F5® BIG-IQ® Centralized Management.

By default, the full request is sent in the middle of the log entry. This may cause items such as violations to be lost if the log data limit is met before they appear. You can set the **Maximum Request Size** and **Maximum Query String Size** to fix this problem. See the following table for recommended settings.

Note: BIG-IP ASM 12.1 and later supports multiple remote logging profiles.

Table 2.1 Log format configuration options

BIG-IP ASM 12.0 and later	Limit Request Size	Limit Query String Size	Specify Storage Format
Comma-Separated Values	Yes	Yes	Yes
Key-Value Pairs	Yes	Yes	No
Common Event Format (ArcSight)	Yes	No	No
F5® BIG-IQ® Centralized Management	Yes	Yes	No
BIG-IP ASM 11.6	**Limit Request Size**	**Limit Query String Size**	**Specify Storage Format**
Remote (CSV)	Yes	Yes	Yes
Reporting Server (Key-Value Pairs)	Yes	Yes	No
ArcSight (CEF)	Yes	No	No

Logging bots

The BIG IP 13.0 system includes a new log option for reporting on bots. After the option is enabled, results display on the ASM Reporting page.

To log and view reports on bot violations:

1. Navigate to **Security** > **Event Logs** > **Logging Profiles** and select the **Report Bot Defense** check box.
2. After the option is enabled, assign the log to the virtual server.
3. To view the report, navigate to **Security** > **Event Logs** > **Bot Defense** > **Requests** to locate the report page.

Policy Tuning and Enhancement

The Policy Builder is the automated tool with which you create a security policy. You can run the Policy Builder to build a new security policy, or to update an existing security policy. Policy Builder combines manual and automatic tuning of BIG-IP ASM security policies. It can run in automatic or manual mode, or it can be disabled.

BIG-IP ASM Policy Builder updates

BIG-IP 14.0

Updates to layered security policy management in BIG-IP 14.0 include the following enhancements:

- Enforcing parent policy assignment to a child policy
- Various usability improvements to make it easier to control and understand inheritance attributes

BIG-IP 13.0

Updates to Policy Builder in BIG-IP 13.0 include the following enhancements:

- Compact mode is an entity learning mode designed to effectively manage high traffic loads and increase policy security. Compact mode reduces the amount of learning suggestions, enabling a policy to converge more quickly, and automatically adds disallowed file types.
- Server Technologies is an option that customizes policies to an application. This option enables Policy Builder to identify the *back-end* technologies used by an application and add the relevant signatures to the policy.
- Client Reputation is a technique that improves learning suggestions by using behavioral analysis to assign a reputation score to a source IP or device ID. Policy Builder ignores sources classified as malicious and speeds learning on sources classified as benign.

BIG-IP 12.0

There are several updates to Policy Builder in BIG-IP 12.0, including the following:

- Staging, enforcement, and learning suggestions can be configured manually or by the BIG-IP ASM system.
- Security checks **Learn, Alarm**, and **Block** are now system-wide settings integrated with Policy Builder.
- An improved learning suggestions mechanism handles requests, with or without violations, for manual and automated policy building.

Traffic Learning

Learning suggestions can result in policy changes, such as adding entities, disabling an attack signature, or making the policy more robust based on characteristics of the traffic.

BIG-IP 13.0

In BIG IP 13.0 and later, the Enforcement Readiness Summary is a section of the Learning Suggestion page, making it easier for you to enforce an entity after the enforcement readiness period is over. The Learning Suggestion page also includes new graphs that display Policy Builder converging speed, and more options for filtering and sorting suggestions.

BIG-IP 12.0

In BIG-IP 12.0 - 12.x, the configuration for the Enforcement mode, Learning mode, and learning speed is on the Learning and Blocking Settings page: **Security > Application Security : Policy Building : Learning and Blocking Settings**.

BIG-IP 11.6

In BIG-IP 11.6 - 11.x, learning suggestions are displayed on the Traffic Learning page at **Security > Application Security: Policy Building: Traffic Learning**.

Table 3.1 Examples of Policy Builder Learning Suggestions

Event	Suggestion
An attack signature violation due to a false positive	Disable attack signature.
Traffic entities not in the policy, such as parameters (and their characteristics)	Add entity to policy.

When using the Manual mode in BIG-IP 11.6, you can modify a policy without having any correlated or weighted suggestions from the BIG-IP ASM system.

In BIG-IP 12.0 and later, Policy Builder runs in the background to make it easier for you to make policy modification decisions when in Manual mode.

The BIG-IP ASM system provides suggestions based on the number of requests and how often a signature is triggered. The system also displays a request violation rating based on statistical calculations from different sessions and client IP addresses, with a learning score on a scale from 0 to 100 percent.

If you are using Policy Builder, after a suggestion reaches 100 percent, the BIG-IP ASM system automatically accepts it. If you are using the Manual Policy Builder, the suggestion stays on the Traffic Learning page until you accept it.

When you accept a suggestion, it is added to the policy. You also can delete suggestions. If you delete a suggestion, it is removed and can be learned again. If you ignore the suggestion, it is hidden.

You can accept a suggestion at any time in either Policy Builder or in Manual mode.

Note The default learning suggestion filter uses a 5-100 percent learn score to filter out suggestions with a single occurrence or those not likely to be a false positive.

Configuration guidelines

This section covers guidelines for common configuration tasks involved in setting up a new policy.

Determining and configuring sensitive parameters

Sensitive parameters include private data such as credit card numbers and other personally identifiable information (PII). F5 recommends that these parameter values be masked from logs. In some cases, masking is mandatory for regulatory compliance.

The easiest method to determine whether sensitive information exists is to request a list of the sensitive parameters for your application from the application owner, such as password-based accounts, sensitive session data, and other parameters.

Another method to determine whether sensitive information exists is to use parameter learning to create suggestions in the entity learning section. Parameter learning accumulates parameters present on the site, which you can then examine and determine whether each is sensitive. Parameters can be targeted by name and value.

For more information, refer to **Masking Credit Card Numbers in Logs** in ***BIG-IP ASM: Implementations*** for your system version.

Note For information about how to locate F5 product manuals, refer to the Ask F5 article: K12453464: Finding product documentation on AskF5.

Increased sensitive data masking

Starting with BIG-IP ASM 14.0, you can enable the option to mask data for more policy entities, and the data in some policy entities is masked by default. The name for the data masking option also changed from **Sensitive Data** to **Mask in logs**.

Sensitive data values disclosing personal details about users and credit cards can be masked in:

- Parameters (in BIG-IP ASM 14.0 and later versions, this includes using positional parameters to mask URL segments with personal identification)
- Sensitive parameters
- HTTP headers (BIG-IP ASM 14.0 and later versions)

  Note The **Mask in logs** option is enabled by default for the **Authorization** header.
- Cookie values (BIG-IP ASM 14.0 and later versions)
- JSON profiles
- XML profiles

For example, in BIG-IP ASM 14.0, the **Authorization** header is masked by default in the following request log:

```
GET / HTTP/1.1
Host: blah
Connection: Keep-alive
Authorization: ******
Cookie: BIGIPAuthCookie
```

However, in BIG-IP ASM 13.0, the **Authorization** header is not masked by default in the following request log:

```
GET / HTTP/1.1
Host: blah
Connection: Keep-alive
Authorization: 123xyz
Cookie: BIGIPAuthCookie
```

These data values are masked in the BIG-IP ASM, local, and remote request logs. When at least one of the entities you marked for masking is present in a violation, the BIG-IP ASM system does not show the violation details snippet, even when the violation does not occur within the masked entity.

Sensitive data values are not masked in:

- The response log header and body. In this case, Data Guard applies.
- Antivirus over Internet Content Adaptation Protocol (ICAP).
- Database security.
- Policy Builder.

Creating policy templates

F5 recommends that you create a baseline policy for your environment, which includes the basic security requirements that are embedded in the policy components. You can make these policies into a template in the BIG-IP ASM system configuration and re-use them as a baseline for any future policies that you create for the environment.

Disabling case sensitivity in policies

By default, the BIG-IP ASM system is case-sensitive; however, to reduce the chance of false positives, F5 recommends that you create policies with case sensitivity disabled.

While most web applications do not treat URIs with case sensitivity, some do. In such applications, **example.html** and **Example.html** may see different web pages, each with their own associated content, parameters and security access controls. In such cases, enabling case sensitivity is required.

Disabling Illegal meta character in value

Unless your application requires a strict security posture, you can disable the **Illegal meta character in value** violation. Enforcing against illegal meta characters involves a tuning period during which alarms are tuned down or changed to **Allowed** for a specific parameter or for the policy in general.

Some meta characters are considered risky (such as "<" and " ' "), but most attacks taking advantage of these commonly abused meta characters have matching attack signatures.

Configuring vulnerability scanners

Vulnerability assessment services identify, classify, and report potential security holes or weaknesses in website code. The BIG-IP ASM system integrates with vulnerability assessment services such as IBM Security AppScan Enterprise, Hailstorm, QualysGuard, HP WebInspect, WhiteHat Sentinel, and Quotium Seeker. It also integrates with other vulnerability assessment tools through the use of a generic scanner.

To use vulnerability assessment scanners, you need to identify and configure the scanners' IP addresses in the IP address exception section.

The following table includes configuration options for vulnerability scanners:

Table 3.2 Vulnerability scanner configuration options

Policy Builder trusted IP	Disable
Ignore in Anomaly Detection	Enable.
Ignore in Learning Suggestions	Enable.
Never block this IP Address	Disable. Enable when scanner scans and reports on back-end security vulnerabilities. Illegal Method. Illegal HTTP Response. Request length exceeds defined buffer size. Modified BIG-IP ASM cookie. Access from disallowed User/Session/IP. Access from disallowed geolocation. Attack signature detected. Specific signature detected. Non-human based clients signature set (optional.)
Never log traffic from this IP Address	Disable if evidence of violation must be shown (for security auditors, for example). Enable to reduce spam violations in event log and make it easier to find real-time attacks in production environment. Attack signature detected: Remove attack signatures out of staging, resolve issues with attack signatures triggered by false positives. Illegal Redirection Attempts.

Policy Builder trusted IP	Disable
Ignore IP Address Intelligence	Enable.

For more information, refer to **Using Vulnerability Assessment Tools for a Security Policy** in ***BIG-IP Application Security Manager: Getting Started*** for your system version.

Customizing default response pages

The BIG-IP ASM system uses a default blocking response page to notify users when a request has been blocked by a security policy. It also uses a default page to display when login violations occur.

The default blocking response contains one variable, **<% TS.request.ID() %>,** which BIG-IP dynamically replaces with a support ID reference number for the specific request or response that triggered the blocking event.

You can easily create a custom blocking response page for either general violations or login-specific violations. You can style it to match your organization or application style or branding. F5 recommends that you customize the page so at a minimum it includes contact information and steps that users can take to address the violation.

F5 recommends that you host required assets for the blocking response page outside the BIG-IP environment serving it. For example, you can place assets on Content Delivery Network(CDN)/"sorry server".

Because of potential Cross Origin Resource Sharing (CORS), assets stored on a CDN/"sorry server" may become unavailable to the response page. If this happens, try placing assets like CSS stylesheets inline to the HTML markup as in the following code example:

```
<!doctype html>
    <html>
          <head>
                <meta charset="utf8">
                <meta httpequiv="XUACompatible" content="IE=edge,chrome=1">
                <title>Request Blocked | Company Acme</title>
                <meta name="viewport" content="width=devicewidth">
                <style type="text/css">
          *, *:after, *:before {
          webkitboxsizing: borderbox;
          mozboxsizing: borderbox;
          msboxsizing: borderbox; boxsizing: borderbox; }
    html {
          background: #fff;
          font: bold 16px/24px "Helvetica Neue Lt Std Light", "Tablet Gothic",
```

```
Helvetica, Arial, sansserif;
    color: #000;
    }
    body { width: 100%; margin: 100px auto;}
    .logo { padding: 5px; width: 100%; display: block; textalign: center; margin-
left: 100px; }
    .logo img { padding: 8px; maxheight: 64px; }
    section .message { background: #fafafa; color: #666; width: 410px; margin: 0
auto; padding: 8px 15px 32px 15px; border: 1px solid #e2e2e2;}
```

Note F5 recommends that you serve the blocking response page assets from a local "sorry server" and have the BIG-IP ASM system redirect violating clients with the proper support ID.

F5 recommends that you do not embed images in the BIG-IP ASM response page.

For more information, refer to:

- **Configuring What Happens if a Response is Blocked** in ***BIG-IP ASM: Implementations*** for your system version.
- AskF5 article: K7825: Redirecting a blocking response support ID to an external error page.

Disallowing malicious file types

If your application contains a dynamic list of file types, you can create a list of disallowed file types to block using one illegal file type violation.

For most applications, F5 recommends that you block the following file types: .bat, .bck, .bin, .cfg, .cmd, .com, .config, .dat, .dll, .eml, .exe, .exe1, .exe_renamed, .hta, .htr, .htw, .ida, .idc, .idq, .ini, .log, .nws, .old, .pol, .printer, .save, .shtm, .shtml, .stm, .sys, .wmz. In BIG-IP 13.0, when you use Compact mode for entity learning, these files types are automatically added to your security policy.

To disallow file types, you can enter them individually by navigating to **Security > Application Security > File Types > Disallowed File Types – Create**, or you can import a list of them in the XML policy configuration:

For more information, refer to **Adding File Types to a Security Policy** in ***BIG-IP ASM: Implementations*** for your system version.

Manual learning and policy building

Learning suggestions provide a concise list of application components observed by the BIG-IP ASM system that are not found in the security policy.

The BIG-IP ASM system generates learning suggestions for violation types with corresponding **Learn** check boxes selected on the Blocking Settings page (in BIG-IP 11.6) or the Learning and Blocking Settings page (in BIG-IP 12.0 and later).

When the system receives a request triggering a violation, it updates the Manual Traffic Learning page (in BIG-IP 11.6) or the Traffic Learning page (in BIG-IP 12.0 and later) with learning suggestions based on information from the violating request.

You can review learning suggestions to determine whether the listed requests triggered a legitimate security policy violation or whether you need to update the security policy to allow such requests.

Learning suggestions enable you to create policy entities based on actual traffic. They provide the number of times each violation has occurred and a list of the requests that contain violations. From this information, you can make policy decisions.

Accepting or clearing learning suggestions

As part of tuning your policy, you need to decide whether to accept a learning suggestion, add it to the policy, or clear it and remove it from learning suggestions. Consider the following questions:

- Does the violating entity make sense for the application? For example: Does the application make use of Web Open Font Format (WOFF) file types?
- Can the application owner provide a summary of which entities should be allowed into the application?
- How many learning suggestions exist for the given entity? Numerous occurrences may indicate a false positive?

General guidelines for manual learning suggestions

There is a limit to the number of learning suggestion records that can be stored. To avoid exceeding retention limits and losing suggestion data, F5 recommends that you remove flags from violations after you have captured the necessary information. Further, F5 recommends that you enable learning only for policies that you are actively modifying or configuring, or for applications you are actively tuning.

Learning suggestion records expire at a rate determined by the number of violations set in Blocking Settings and the number of policies using traffic learning. Since learning suggestion requests are not indefinitely retained, F5 recommends that you regularly review suggestions on any policy that the system is actively learning, or use the Policy Builder.

To see the impact of new or updated signatures, enable learning suggestions on all or select violations when updating attack signatures. Doing so can help ensure proper policy configuration and reduce potential false positives. Learning suggestions do not sync or move between BIG-IP ASM systems and therefore can only be found on the active unit that learned the traffic. When a failover occurs, you must fail back to the original system to access learning suggestions. Also, learning suggestions are cleared when a system is upgraded. F5 recommends that you accept or clear all learning suggestions before you upgrade your system.

For more information, refer to **Refining Security Policies with Learning** in ***BIG-IP ASM: Implementations*** for your system version.

Tuning violations

Policy violations are displayed on the Policy Building: Traffic Learning page. An important part of tuning your policy is to fix violations and manage issues associated with false positives. For example, legitimate application traffic may be disrupted if your policy is blocking violation types that create false positives. Even if blocking for a violation is off, false positives can trigger alarms, and an excess of improperly logged violations can create noise that may distract from real issues.

In many cases, violations show both malicious traffic and legitimate traffic. If the system is blocking legitimate traffic (a false positive), you can address these violations by accepting the related suggestion—also known as loosening the policy—even though malicious requests were also detected by the same violation.

This may seem counterintuitive, since the policy, by design, successfully detected and blocked malicious traffic, but it is better to loosen the policy than to create false positives or excessive, improperly logged violations.

Tuning based on deployment environment

Use the following policy tuning guidelines for any violation, based on the deployment environment:

Non-production environment (QA or Test): Only valid traffic requests should exist in a non-production environment; however, policy violations may still occur with a new policy or as a result of an application update.

If vulnerability scanning and penetration testing are performed, make sure to exclude that traffic from **Learning** to distinguish it from legitimate traffic. Address all generated violation types, and accept all resulting learning suggestions for the security policy.

Production environment: After a mature policy is deployed to production, most violations likely represent malicious attacks; however, some false positives may occur, depending on how much tuning you did before deployment. Some violation types are more prone to false positives than others.

Other than volumetric attacks, most real attacks do not cause high volumes of traffic, so a large number of incidents associated with a particular violation type increases the likelihood that false positives exist.

Tuning based on violation type and deployment environment

The following section lists the most common policy violation types and how to treat each, based on whether they occur in a production or non-production environment.

Illegal File Type

Triggered by: Request accesses a file type that is not found in the security policy.

Non-production environment: Accept. This violation typically occurs when new application resources are added or new technology is adopted in the application.

Production environment: Unless these violations occur in conjunction with a new application update, violations of this type in a production environment are generally malicious and you should not accept as a learning suggestion or loosen the policy.

HTTP protocol compliance failed

Triggered by: Request does not comply with a number of HTTP protocol compliance checks.

Non-production environment: Accept immediately.

Production environment: Do not accept. Regard as an attack. If data shows high number of incidents, review several of the incidents for malicious traffic. If samples are malicious, ignore the learning suggestion and do not disable the violation. If the incident count is low, these are very likely malicious. However, application updates or new types of clients may rarely trigger this type of violation. If false positives are evident, disable the improper firing protocol compliance check.

Illegal Method

Triggered by: Request uses the HTTP method not allowed by the security policy.

Non-production environment: Accept immediately.

Production environment: Regard new methods triggered as attacks. If you observe a high number of incidents, review the data for false positives. These false positives are typically caused by a new type of client accessing the application. To permit a new client type, update the policy to allow the new method.

Evasion Technique Detected

Triggered by: the BIG-IP ASM system fails to normalize requests. Normalization is the process of decoding requests that are encoded. Several sub-violations for this violation exist.

Non-production environment: Accept, including all sub-violations.

Production environment: False positives are not uncommon for this violation type in the production environment. If they occur after an application update and the incident count is high, sample the incident data to confirm that it is legitimate traffic. If the traffic is legitimate, accept the learning suggestion.

> **Note** Evasion Technique checks are global for the policy. If you accept the learning suggestion for any subviolation, all Evasion Technique checks are disabled. The risk of disabling this check is mitigated by other security checks, which occur after normalization process.

For more information on this violation type, refer to AskF5 article: K7929: Working with Evasion technique detected violations.

Multiple Decoding Passes sub-violation

Triggered by: Any request requiring at least the configured number of legal decoding passes to achieve normalization. A sub-violation of the Evasion Technique Detected violation, it controls the BIG-IP ASM system's normalization process, in general, as requests are processed. The default value for this setting is 2 decoding passes. This setting impacts operation of the normalization process for the security policy, even if the **Multiple Decoding** sub-violation is not enabled.

If this sub-violation is disabled, requests not exceeding the configured number of decoding passes can still trigger a violation of other security checks. If requests exceed the configured number of decoding passes, a double-encoded request matching a signature still causes a signature violation after decoding. In the same circumstance, if the sub-violation is enabled, such requests generate a violation for both the signature and the sub-violation.

Non-production environment: Accept.

Production environment: False positives may rarely occur due to application updates. Review incidents for false positives. If false positives exist, accept learning suggestion, which disables the sub-violation. Since normalization (including at least two decoding passes) of requests occurs independently of this sub-violation, the policy likely remains effective in mitigating attacks, using multiple encoding passes in an attempt to evade detection.

Attack signature tuning

Attack signatures are rules or patterns that identify attacks or classes of attacks on a web application and its components. When the BIG-IP ASM system receives a request, the system applies the attack signatures associated with the security policy to the request. If, in the request (or response), a pattern matches one or more of the attack signatures, the system generates the **Attack signature detected** violation. If the enforcement mode is Blocking, the system also blocks the request and issues the blocking response page to the client making the request.

Signature overlapping

There's an improvement to the attack signature update process in BIG-IP ASM 13.0: You have the two options, below, for updating signatures in the Learning and Blocking Settings for attack signatures.

1. The default option prior to BIG-IP 13.0: When the BIG IP ASM system updates signatures that are enforced as part of the signature updating process, they are removed from Enforcement mode and put into staging.
2. The default option in BIG-IP 13.0: When the BIG IP ASM system updates signatures that are enforced as part of the signature updating process, they remain in Enforcement mode. This change maintains policy security regardless of the update.

False positives

Indicators of potential false positives include the following:

- High occurrence of a specific attack signature—particularly a new one—triggering high volume of violations. The signature is likely detecting a portion of application code as an attack.
- Newly implemented signature registering immediate violations. When adding new signatures to a policy, either through signature updates or through application of additional current signatures, F5 recommends that you turn on **Learning** for the signatures. If new signatures register immediate violations, review the requests to confirm legitimate blocking or to see if the signature is causing a false positive.

Finding and fixing false positives

Find the request that triggered the violation. F5 recommends that you turn on **Learning** for signatures to help catch the request that triggered the violation.

Review the request for malicious content. If possible, consult the application owner, who may be able to help determine valid traffic from malicious traffic.

Correct the false positive. Your action depends on the nature of the false positive and your ability to fix it. There are common methods for correcting a false positive while maintaining good security posture:

- Disable the attack signature (either globally or on at the parameter level).
- Create a user-defined attack signature (may involve a high level of complexity).
- Open an F5 support case to request attack signature tuning.
  For more information, refer to Policy Tuning and Enhancement.

If the violation is not a false positive, nothing more needs to be done.

Transitioning enforcement mode from Transparent to Blocking

F5 recommends that you transition a security policy's enforcement mode from Transparent to Blocking to put security settings into effect after you have reduced the chances of false positives occurring.

Customizing default response pages

The BIG-IP ASM system uses a default blocking response page to notify users when a request has been blocked by an Advanced policy guideline.

Application language

Every web application has language encoding that determines the character set browsers use to both display the page and send any submitted form data. The BIG-IP ASM system supports multiple language encodings.

You can configure the language for new web applications (or those you want to reconfigure). F5 recommends that you use the default setting, **Auto detect**, when it is available. The BIG-IP AMS system determines the acceptable character set for the application.

> **Note** After you set the web application language, you cannot change it unless you reconfigure the web application completely, losing all settings.

For information about reconfiguring web applications, refer to **Returning a web application to a new, unconfigured state** in ***Configuration Guide for BIG-IP Application Security Manager*** for your system version.

Application language guidelines

- Make sure Policy Builder detects correct character encoding language to avoid false positives due to character encoding.
- If manually setting the application language, make sure you know the character set the application uses for HTML form submission.

For more information, refer to the Ask F5 article: K6335: Overview of encoding language settings for BIG-IP ASM.

The BIG-IP ASM system checks to see if the sequence of bytes is valid within the defined encoding. Sequences not found trigger the **Failed to convert character** violation.

However, the BIG-IP ASM system attempts to parse the data, even when the encoding is wrong. For example, consider a web application configured with ISO-8859-8 that has a website that sends pages in UTF-8 encoding and requests generated from pages that contain UTF-8 encoded data. If the application cannot parse the data as ISO-8859-8, but can validate it as ISO-8859-8 characters in UTF8, the request does not cause a **Failed to convert** violation.

JSON/XML considerations

By default, the BIG-IP ASM system expects HTML form data. If an application contains either XML or JSON protocols, the security policy accounts for this using content profiles.

F5 recommends that you associate an XML/JSON content profile for generic policies that create wildcard entities.

BIG-IP 13.0 includes a default JSON profile that Policy Builder automatically assigns to the JSON payload in the application URL. The JSON payload provides advance structure checks and attack signature detection in key-value pairs.

XML/JSON for content type requests (URL)

For content-type-based XML/JSON, a generic content profile using header-based content validation must be associated with a URL wildcard entity.

For more information, refer to **Adding JSON Support to an Existing Security Policy** in ***BIG-IP ASM: Implementations*** for your system version.

XML/JSON for HTML form data parameter value

For XML/JSON protection at the parameter value level, identify the specific parameter that contains the protocol, and apply a generic content profile; this enables the BIG-IP ASM system to process the traffic appropriately.

Large file uploads

The BIG-IP ASM system has a default limitation of 10 MB for incoming HTTP requests. You can increase this limit up to 30 MB.

When an incoming request exceeds the defined buffer size, the **Request length exceeds the defined buffer size** violation setting determines whether a request is blocked or bypassed.

Set this violation to **Blocking** if the protected environment is not expecting large file uploads. Requests bypass the security policy after inspecting up to 10KB of the request (containing the requests' headers, URIs, and initial request payloads).

For more information, refer to the AskF5 articles: K7935: Error Message: Request length exceeds defined buffer size and K17573: The BIG-IP ASM system does not need to be disabled to allow large file uploads.

Working with external caching server

X-Forwarded-For headers

The traffic that passes through content delivery networks (CDNs) is usually network address translation (NAT). CDNs use the initiating client IP address's HTTP request and add it as a header to that request and send it on to its destination.

You can set the BIG-IP ASM system to trust the IP address contained within the HTTP header instead of the NAT IP address from the CDN. Trusting the X-Forwarded-For (XFF) header (header name may vary depending on the CDN) is generally preferred, as you can see the actual client IP address in the log files.

For a DoS profile, **Insert X-Forwarded-For** should also be enabled. When it's set in the HTTP profile used on the virtual server being protected, it uses a default page to display.

Mandatory HTTP header violation searching for CDN based headers

Some HTTP headers should always appear in traffic that was forwarded by CDN servers. These headers can be added to a policy and checked to verify that traffic came from the CDN.

Session tracking and login pages

You can track user sessions using login pages that are configured in the BIG-IP ASM system, or you can have the policy retrieve the user names from the F5® BIG-IP® Access Policy Manager® (APM) system.

The advantage of using session tracking is that you can identify the user and the unique session for each request, which provides additional data points that you can use in combination with the source IP address to investigate suspicious traffic.

Whether you create them manually or automatically, login pages define the URLs, parameters, and validation criteria required for users to log in to the application. User and session information is included in the system logs so you can track a particular session or user. The system can log activity, or can block a user or session if either triggers more violations than the configured threshold allows.

You can use login pages in the following ways:

- Track a specific user instead of using the authenticated session. Tracking by IP address is impractical for a number of reasons.
- Block suspicious sessions after specific limits are triggered. You can configure limits as a number of violations per unit of time, per user, or per session. This might be an attacker with automated scanner-tools that are generating a high rate of alerts per time frame. If the limits are exceeded, the attacker is blocked for a specific amount of time.

The default blocking response contains one variable, **<% TS.request.ID() %>**, which BIG-IP dynamically replaces with a support ID reference number for the specific request or response that triggered the blocking event.

You can easily create a custom blocking response page for either general violations or login-specific violations. You can style it to match your organization or application style or branding. F5 recommends that you customize the page so at a minimum it includes contact information and steps that users can take to address the violation.

F5 recommends that you host required assets for the blocking response page outside the BIG-IP environment serving it. For example, you can place assets on CDN/"sorry server."

Because of potential Cross Origin Resource Sharing (CORS), assets stored on a CDN/"sorry server" may become unavailable to the response page. If this happens, try placing assets, like CSS stylesheets, inline to the HTML markup as in the following code example:

```
<!doctype html>
    <html>
          <head>
                <meta charset="utf8">
                <meta httpequiv="XUACompatible" content="IE=edge,chrome=1">
                <title>Request Blocked | Company Acme</title>
                <meta name="viewport" content="width=devicewidth">
                <style type="text/css">
          *, *:after, *:before {
          webkitboxsizing: borderbox;
          mozboxsizing: borderbox;
          msboxsizing: borderbox;boxsizing: borderbox;}
    html {
          background: #fff;
          font: bold 16px/24px "Helvetica Neue Lt Std Light", "Tablet Gothic",
Helvetica, Arial, sansserif;
    color: #000;
    }
    body { width: 100%; margin: 100px auto;}
    .logo { padding: 5px; width: 100%; display: block; textalign: center; margin-
left: 100px; }
    .logo img { padding: 8px; maxheight: 64px; }
    section .message { background: #fafafa; color: #666; width: 410px; margin: 0
auto; padding: 8px 15px 32px 15px; border: 1px solid #e2e2e2;}
```

Note F5 recommends that you serve the blocking response page assets from a local “sorry server” and have the BIG-IP ASM system redirect violating clients with the proper support ID.

Monitoring and configuration

User tracking events are monitored through the standard reporting UI pages within BIG-IP ASM security.

For more information, refer to **Configuring Application Security Session Tracking** and **Tracking User Sessions using login pages** in ***BIG-IP ASM: Implementations*** for your system version.

Comparing and merging policies

You can use **Policy Diff** in the BIG-IP ASM system to compare two security policies, view the differences between them, and copy the settings from one policy to the other.

The most common uses for **Policy Diff** are:

- Auditing two policies to determine the differences between them.
- Comparing and merging a production policy with a policy running in a staging environment.

Comparing two policies

After selecting the two policies, you can compare them. You can use the **Policy Diff** functionality to monitor and audit changes, or to move changes from the policies you're comparing to a new copy of the policy, which contains both modifications.

For more information, refer to **Maintaining Security Policies** in ***BIG-IP ASM: Implementations*** for your system version.

Layered policy tuning

When you are interpreting learning suggestions or modifying inheritance settings in parent or security (child) policies, keep in mind the following considerations:

- Parent policies aren’t assigned to a virtual server, and the learning suggestions they provide apply only to the attached security (child) policy, not the parent policy.
- Learning suggestions for policy sections that have mandatory inheritance are disabled on the security (child) policy’s Learning Suggestions page, and display on the parent policy’s page only after they reach a score of 50 percent.
- You can only switch to Blocking enforcement mode from the security (child) policy.

Modifying cookie names

The BIG-IP ASM system uses BIG-IP ASM cookie names to enforce security policies. When clients return these cookies, the BIG-IP ASM system validates them to ensure they are not modified.

Starting in BIG-IP ASM 14.0, you can modify BIG-IP ASM policy and Device ID cookie names. You can mask BIG-IP ASM cookies—to help prevent WAF fingerprinting, for example—or allow users to rename ASM cookies to fit in-house cookie conventions. For more information about modifying cookie names, refer to AskF5 article: K54501322: Modifying ASM cookie names.

Web threat campaigns

Starting in BIG-IP ASM 14.0, you can enhance your security policy by implementing the Web Threat Campaign feature. A threat campaign is an attack associated with a specific malicious actor, attack vector, technique or intent. F5 discovers and investigates these attacks.

The Web Threat Campaign feature requires a separate license. Licensed machines receive dynamic updates, in the form of a binary image file. You can decide if a threat campaign is added in staging after an update. The default is no staging. Threat campaigns are inherited from parent policies to child policies

When you upgrade, the system preserves existing threat campaigns. If a Web Threat Campaign license expires, the system inactivates all campaigns on a device.

Threat campaigns include iRules support that returns:

- A list of the names of the threat campaigns found in the transaction.
- A list of the names of the staged threat campaigns found in the transaction.

Regulatory Compliance

This chapter contains guidance to enhance BIG-IP ASM security policies and improve compliance with various regulatory regimes.

Review this chapter to determine the type of policy you want to use as your basis for compliance, then create an appropriate security policy or create a unique policy based on your specific security needs.

Regulatory compliance of WAFs

You may be responsible for achieving compliance with one or more regulatory regimes. The difficulty is that regulatory regimes are rarely specifically defined web application firewalls (WAFs).

Many regulatory compliance mandates provide few specific implementation requirements and leave areas open to interpretation. F5 recommends using the BIG-IP ASM system to secure your web application environment with consideration for the security posture of the organization, the sensitivity and threat surface of the applications, and the application security resources available.

Whether a particular implementation satisfies the regulations is often subject to interpretation. F5 recommends consulting organization security policies, auditors, and regulatory bodies to determine whether applications are compliant with a particular regime.

This chapter provides guidelines to ensure that the BIG-IP ASM system is configured and an appropriate policy is created for implicated applications.

These guidelines are intended to:

- Make sure that WAF-specific requirements are met.
- Make sure that the introduction and configuration of the BIG-IP ASM system does not accidentally expose the organization to non-compliance.
- Make sure that the BIG-IP ASM system is used to achieve commonly requested regulatory requirements that are not specific to WAFs, where the BIG-IP ASM system can improve compliance or function as a compensating control for security purposes.

These guidelines are limited to the configuration of the BIG-IP ASM system behavior and security policy.

Regulations that govern the overall device posture and administrative controls are not within the current scope of this document. Examples of what is not covered include overall device hardening, multi-factor requirements for administrative access, and device patching.

Important These are guidelines only. Implementation of these recommendations does not guarantee compliance.

Open Web Application Security Project Top 10

Open Web Application Security Project (OWASP) is a comprehensive resource for secure web application coding guidance, application testing procedures, and vulnerability categorization and mitigation strategies. OWASP periodically updates and publishes a "Top 10" list of vulnerabilities.

Note: This link takes you to a resource outside of AskF5. The third party could remove the content without our knowledge.

The Top 10 list includes vulnerabilities against which organizations should be particularly diligent to protect their applications. The BIG-IP ASM system can detect and report exploit attempts seeking to take advantage of many of these vulnerabilities.

The following table lists BIG-IP ASM system components and capabilities that you may find useful when creating a security policy to mitigate the listed vulnerabilities. Where applicable, the policy type that automatically includes the mechanism is listed in parentheses.

Note: This table includes the Enhanced policy type, which is no longer available in Policy Builder in BIG-IP 13.0.

Table 4.1 OWASP compliance

	Vulnerability	BIG-IP ASM Controls
A1	Injection flaws	Attack signatures Meta character restrictions Parameter value length restrictions
A2	Broken authentication and session management	Brute force protection Credentials stuffing protection Login enforcement Session tracking HTTP cookie tampering protection Session hijacking protection
A3	Sensitive data exposure	Data Guard Attack signatures: predictable resource location and information leakage attacks
A4	XML External Entities (XXE)	Attack signatures: other application attacks—XXE XML content profile: disallow Document Type Definition (DTD) Subset of API protection

	Vulnerability	BIG-IP ASM Controls
A5	Broken access control	File types Allowed and disallowed URLs Login enforcement Session tracking Attack signatures: directory traversal attacks
A6	Security misconfiguration	Attack signatures DAST integration Allowed methods HTML5 Cross-Domain Request Enforcement
A7	Cross-site scripting (XSS)	Attack signatures: XSS attacks Parameter meta characters **HTTPOnly** cookie attribute enforcement Parameter type definitions (such as integer)
A8	Insecure deserialization	Attack signatures: server-side code injection attacks
A9	Using components with known vulnerabilities	Attack signatures DAST integration
A10	Insufficient logging and monitoring	Request and response logging Attack alarm and blocking logging On-device logging and external logging to security information and event management (SIEM) system Event correlation

Payment Card Industry Data Security Standard

A commonly requested compliance assistance for the BIG-IP ASM system is associated with the Payment Card Industry Data Security Standard (PCI DSS). WAFs are specifically referenced, in section 6.6 of the PCI DSS 1.1 (Sept., 2006), as one control mechanism that an organization can implement to verify whether Internet-facing web applications are placing cardholder information at risk.

The current version of the PCI DSS lists WAFs as mechanisms which organizations can use to satisfy section requirements.

Note: This link takes you to a resource outside of AskF5. The third party could remove the content without our knowledge.

You should become familiar with the entire DSS; specifically, section 6.5 describes the application vulnerabilities that developers should take particular care to guard their applications against.

The following table lists requirements and BIG-IP ASM controls that you may find useful when pursuing PCI DSS compliance. Where applicable, the policy type that automatically includes the mechanism is listed in parentheses. This table is similar to the OWASP Top 10 list. Also see OWASP resources and related materials when pursuing PCI DSS compliance.

Table 4.2 PCI DSS compliance

PCI DSS Requirements	Description	BIG-IP ASM Controls
6.5.1	Injection flaws	Attack signatures Meta character restrictions (Enhanced or Comprehensive) Parameter length restrictions
6.5.2	Buffer overflows	Attack signatures Length restrictions (Fundamental or Comprehensive)
6.5.3	Insecure cryptographic storage	Not applicable
6.5.4*	Insecure communications	Implement an HTTP to HTTPS redirect iRule to require encryption of all traffic and use rewrite profile if needed to rewrite explicit HTTP references. Implement server-side re-encryption via the Server SSL profile. Configure the BIG-IP ASM system to enforce URL flows to prohibit HTTP requests (Comprehensive).
6.5.5	Improper error handling	Allowed Response Codes Attack signatures (Response checking required)
6.5.6	"High Risk" vulnerabilities (as defined by the organization under 6.1)	BIG-IP ASM controls may be applicable depending on the specific vulnerability.
6.5.7	Cross-site Scripting	Attack Signatures Parameter meta characters (Comprehensive) Parameter length restrictions
6.5.8	Improper Access Control	Session Tracking Attack signatures (Forceful browsing) File types (Fundamental) URL (Enhanced) URL flows (Comprehensive)
6.5.9	Cross-site Request Forgery	CSRF Protection (Comprehensive)

* Specific requirements regarding SSL/TLS are new as of PCI DSS 3.1. Specifically, the PCI DSS prohibits the use of SSL and TLS version 1.0. For more information, refer to **SSL Traffic Management** in ***BIG-IP System: SSL Administration*** for your system version.

Note For information about how to locate F5 product manuals, refer to the Ask F5 article: K12453464: Finding product documentation on AskF5.

Additional regulatory compliance considerations

The following table lists additional PCI DDS considerations and F5 recommendations.

Table 4.3 Additional Payment Card Industry DDS requirements

Consideration	Recommendation
Logging. By default, the BIG-IP ASM system may log cardholder data. This may contribute to data leakage.	Define all parameters that may include cardholder data as **Sensitive** to prevent the values from being logged. For more information, refer to **Adding Entities to a Security Policy** in ***BIG-IP ASM: Implementations*** for your system version.
RFC compliance. The PCI DSS does not mandate that applications must comply with Request for Comment (RFC). However, non-compliant applications are vulnerable to more threats.	Configure RFC compliance to minimize vulnerabilities.
Cardholder data leakage. The PCI DSS requires that organizations ensure that cardholder data, such as social security and credit card numbers, not be exposed in web applications.	Use Data Guard. For more information, refer to **Protecting Sensitive Data with Data Guard** in ***BIG-IP ASM: Implementations*** for your system version.
Evasion techniques. The BIG-IP ASM system normalizes inputs and blocks requests using common evasion techniques. Make sure that evasion techniques are enforced. Also make sure that, when a technique is allowed to enable application functionality, it does not expose the application to additional risk.	Complete an evasion techniques policy audit. For more information, refer to **Configuring Security Policy Blocking** in ***BIG-IP ASM: Implementations*** for your system version.

PCI report

The BIG-IP ASM system includes a PCI report that you can run to help assess whether a given security policy is providing adequate controls for an application. Compliance is defined as properly securing the application. The PCI report is a useful data point, but must be augmented by testing the application with web application vulnerability scanners, manual testing of suspected sections of an application, and consultation with your organization's PCI auditors.

Health Information Portability and Accountability Act

Among other reasons, the Health Information Portability and Accountability Act (HIPAA) was mandated to control the ways in which security and privacy of protected health information (PHI) must be ensured by regulated organizations. HIPAA is a broad and non-specific regulatory regime, and security policy authorities and security administrators are given a wide margin to determine how to implement the requirements. While WAFs are not explicitly referenced, web applications are important vectors for possible exposure of PHI and it is important to

properly protect them. Additionally, BIG-IP ASM administrator must comply with specific requirements.

HIPAA is primarily focused on:

- Preventing accidental PHI exposure
- Ensuring that PHI is only accessed by authorized individuals
- Auditing and accounting all access to PHI and the systems that manage PHI
- Ensuring that PHI is properly encrypted at all times, both during transmission and at rest
- Disclosing any breach that compromises PHI in an accurate, complete, and timely manner

The following table lists the requirements and BIG-IP ASM controls that you may find useful when pursuing HIPAA compliance. Where applicable, the **Policy Learning type level** that automatically includes the mechanism is listed in parentheses.

Table 4.4 HIPPA compliance

Requirement	BIG-IP ASM Controls
PHI Encryption. Network communications that may include PHI must be encrypted at all times.	Implement an HTTP to HTTPS iRule to require encryption of all traffic. Implement server side re-encryption via serverssl profile. Configure the BIG-IP ASM system to enforce URL Flows to prohibit HTTP requests (**Complete**). Ensure strong encryption.
Data Leakage.	Data Guard may be customized to block or mask requests including PHI if the administrator identifies leakage in a web application. PHI and PHI keywords are organization-specific. You are encouraged to implement Data Loss Prevention technologies to best protect PHI.
Logging. By default, BIG=-IP ASM may incorrectly log PHI data which could lead to data leakage.	Define all parameters that may include cardholder data as Sensitive to prevent the values from being logged.
Auditing. HIPAA is specifically concerned with “who accessed what, when, and how.” This includes comprehensive logging of all policy changes and configuration changes to security devices in the environment.	BIG-IP Audit Logging BIG-IP ASM Change Logging Session Tracking

National Institute of Standards and Technology

National Institute of Standards and Technology (NIST) does not require WAFs to specifically address the requirements stated in the following documents, but you can use the BIG-IP ASM system to mitigate some of the concerns and requirements in them:

- NIST Information Technology Laboratory
- Guidelines on Securing Public Web Servers

Another part of NIST regulations requires that your BIG-IP system is securely hardened. For more information, refer to the NIST Special Publication 800-53.

Note: These links take you to resources outside of AskF5. The third party could remove the content without our knowledge.

The following table lists the requirements and BIG-IP ASM controls that you may find useful when pursuing NIST compliance. Where applicable, the **Policy Learning type level** that automatically includes the mechanism is listed in parentheses.

Table 4.5 NIST compliance

Consideration	BIG-IP ASM Controls
Protocol Compliance	On the Policy Settings page, enable **HTTP protocol compliance checks** to allow the BIG-IP ASM system to enforce proper usage of the HTTP protocol, limiting automated and other non-browser clients from accessing the web application.
Evasion Techniques	On the Policy Settings page, enable **Evasion Techniques** to allow the BIG-IP ASM system to find attempts to circumvent scanning and matching mechanisms, enhancing the BIG-IP ASM system capabilities in enforcing security on the application.
Login Enforcement and Session Tracking	Configure the login page(s) for the application. Enable **Session Tracking** and **Login Enforcement**. This allows the BIG-IP ASM system to control which parts of the application are accessible only after a user successfully logs in.
Brute Force Protection	After the login pages are configured, add **Brute Force Protection** to your policy to make sure that authenticated users are not using brute-forced credentials.
Attack Signatures	**Attack Signatures** are enabled by default. The policy finds and enforces controls for attempts to access predictable resources, such as active content libraries and administrative folders.
File Types, URLs, and Parameters	Enable Policy Builder to learn and populate a security policy, inspect traffic (requests and responses), and build the application tree. This enables the BIG-IP ASM system to enforce the policy and block access to URLs not specifically added and allowed during the policy-building process. You can choose to populate your policy with these application elements either automatically or manually.

Common Deployment Topologies

The BIG-IP ASM system supports a variety of deployment topologies to secure applications, while it properly accommodates unique network requirements, protected applications, and operational requirements.

This chapter provides an overview of the BIG-IP ASM system platforms and several common topology options, including considerations for each. Specific and detailed configuration instructions are outside the scope of this document, but may be found in articles and manuals on AskF5 and by searching DevCentral.

Note A DevCentral login is required to access content.

Platforms and licenses

BIG-IP ASM is available on physical and virtual platforms, including the following:

- BIG-IP hardware appliances
- BIG-IP Virtual Edition
- BIG-IP VIPRION

Your platform choice determines the availability of topologies and scaling options. There are several factors to consider when you are deciding on the appropriate platform, include the following:

- SSL/TLS encryption requirements
- Rate of environment growth
- Number of applications
- Policy-building methods
- Deployment environment (traditional data center, cloud, or private or public hybrid cloud)

BIG-IP ASM can be licensed as follows:

- As a standalone BIG-IP ASM-only platform.
- As an add-on license to F5® BIG-IP® Local Traffic Manager™ (LTM).

Hardware appliances

The BIG-IP ASM system is commonly deployed with Device Service Clustering, as a redundant system. However, other deployment configurations are available.

For more information, refer to ***BIG-IP Device Service Clustering: Administration*** for your system version.

Note For information about how to locate F5 product manuals, refer to the Ask F5 article: K12453464: Finding product documentation on AskF5.

Hardware sizing

The BIG-IP ASM system performs best on hardware devices that are properly sized for the protected applications, and that have the following:

- Cryptographic hardware for offloading SSL/TLS computation
- Hardware compression for offloading to which customers can add blades, and on which BIG-IP software runs compression computations
- Solid-state drives for superior local I/O
- Enough available RAM, particularly when using traffic learning on multiple applications
- Enough CPU resources

Important RAM and CPU sizing recommendations for web application firewalls vary widely by number and type of application, security policy settings, and network conditions. Consult your F5 account team for assistance in identifying the properly sized appliance for your security requirements.

Upgrading hardware

Within BIG-IP hardware appliance series, each appliance is available in two models, which differ in CPU, SSL capacity, and other characteristics:

- Base model (n0n0)
- Scale-Up model (n2n0)

Using a software license, you can upgrade most BIG-IP hardware appliances from a base appliance to a scale-up model in that series. For most appliance series, you can purchase the base model and later increase performance by purchasing the scale-up model. For example, if you purchased a 5050, using a software license key you can upgrade that model to a 5250, which nearly doubles the capacity of the platform.

F5 FIPS models

Some F5 hardware appliances are available in a Federal Information Processing Standard (FIPS) model for organizations storing their SSL/TLS private keys in a certified hardware security module (HSM). F5 FIPS models offer FIPS 140-2 level 2 compliance and also an interface to network HSMs from Thales and Safenet.

For more information, refer to ***BIG-IP Platform: FIPS Administration*** for your system version.

vCMP support

Some hardware appliances support F5® Virtual Clustered Multiprocessing™ (vCMP) virtualization technology to allow multiple virtual BIG-IP instances on a single platform. A vCMP guest instance functions identically to a hardware appliance for the topologies described, as long as adequate resources are assigned to the guest instance.

For more information, refer to ***vCMP for Appliance Models: Administration*** for your system version.

BIG-IP ASM Virtual Edition

BIG-IP ASM is available as a standalone, virtual BIG-IP ASM-only appliance.

The BIG-IP ASM system can be deployed in cloud environments or as a virtual appliance on a variety of hypervisors. (To find supported hypervisor types and versions, refer to the virtual edition documentation for your version.)

Reasons to choose a virtual appliance include the following:

- To deploy the BIG-IP ASM system in a cloud environment.
- To reduce acquisition cost.
- You do not expect high volumes of SSL/TLS traffic in your environment.
- Your environment requires rapid deployment, high levels of instantiation, or mobile deployment.

Capacity concerns

Most applications protected by the BIG-IP ASM system are encrypted with SSL/TLS. F5 recommends that you carefully consider capacity.

A virtual edition is licensed to perform up to 500 SSL TPS for 2048-bit keys per vCPU allocated to the virtual appliance. Actual TPS varies depending on the type and speed of the CPUs.

Depending on your license, you can allocate up to eight vCPUs. Applications that require high volumes of SSL/TLS traffic may require either cryptographic offload (hybrid topology for example) or scaling the virtual edition appliances using a pool topology or clustering.

Additionally, each virtual edition appliance is available in a variety of throughput levels, from 25Mb/sec to10Gb/sec, depending on hypervisor. You can easily upgrade these throughput levels. When allocating disk space, CPU, and RAM, you can increase the resources, as available, and increase the capacity and performance of each BIG-IP ASM virtual appliance.

For more information, refer to the AskF5 articles: K14810: Overview of BIG-IP VE license and throughput limits and BIG-IP Virtual Edition Datasheet.

VIPRION

The F5® VIPRION® platform is a scalable chassis to which customers can add blades, and on which BIG-IP software runs. The platform presents multiple blades as a single logical device, to scale rapidly and without disruption through the installation of additional blades. Since flow on any blade can be processed on the computing resources for that blade or any other blade, the design provides robust processing power for a BIG-IP ASM deployment.

The VIPRION platform is a frequent choice to protect high performance and high capacity applications, or to protect a large number of applications.

vCMP support

With an optional, additional license, VIPRION supports vCMP virtualization technology to allow multiple BIG-IP instances on a single platform. In examples provided in this chapter, a vCMP guest instance can function identically to a hardware appliance for the topologies described, as long as adequate resources are assigned to the guest instance.

Sizing

The BIG-IP ASM system performs best on VIPRION systems offering the following:

- Cryptographic hardware for offloading SSL/TLS computation
- Hardware compression for offloading HTTP compression computations.
- Solid-state drives for superior local I/O
- Enough available RAM, particularly when using traffic learning on multiple applications
- Enough CPU resources

Important RAM and CPU sizing recommendations for web application firewalls vary widely by number and type of application, security policy settings, and network conditions. Consult your F5 account team for assistance in identifying the properly sized VIPRION for your security requirements.

Common topology options

The three types of topology are single-tiered, multi-tiered, and hybrid.

- **Single-tiered** is the most common topology. It involves deploying the BIG-IP ASM system in a simple active-standby device pair, with protected application traffic directed to the active unit.
- **Multi-tiered** involves deploying multiple BIG-IP ASM appliances in parallel with traffic steered to BIG-IP ASM devices through BIG-IP LTM appliances.
- **Hybrid** involves deploying BIG-IP ASM VE in a distributed topology, combined with VIPRION or hardware appliances, performing traffic distribution and hardware-based cryptographic and compression offload.

Important If in order to meet specific business goals you need to deploy an advanced topology for an atypical or complex application delivery environment, F5 strongly encourages you to consult your F5 account team. You can also get architectural and implementation assistance from F5 Professional Services or an F5 Guardian Professional Services Partner.

Modes

You can deploy the BIG-IP ASM system in either **routed mode**—with or without secure network address translation (SNAT)—or in a **one-armed mode** (with SNAT).

You can configure both modes at the same time on a single BIG-IP system, on an app-by-app basis.

You can use both modes with single-tiered or multi-tiered topologies.

Note Requests and responses must go through the BIG-IP system. This means that if you do not use network address translation (NAT) on the source IP address of the client, the default gateway of the server needs to be the BIG-IP system. If you use SNAT for all traffic from the client to an IP address of the BIG-IP system, all responses are sent back to the IP. To keep track of the original client IP address, you can enable the X-Forwarded-For feature of the HTTP profile. This adds the client IP address to the HTTP header that was sent to the web server.

Routed mode

In route mode, the BIG-IP ASM system is in the routing path of the web servers, and all traffic to the server flows through the system.

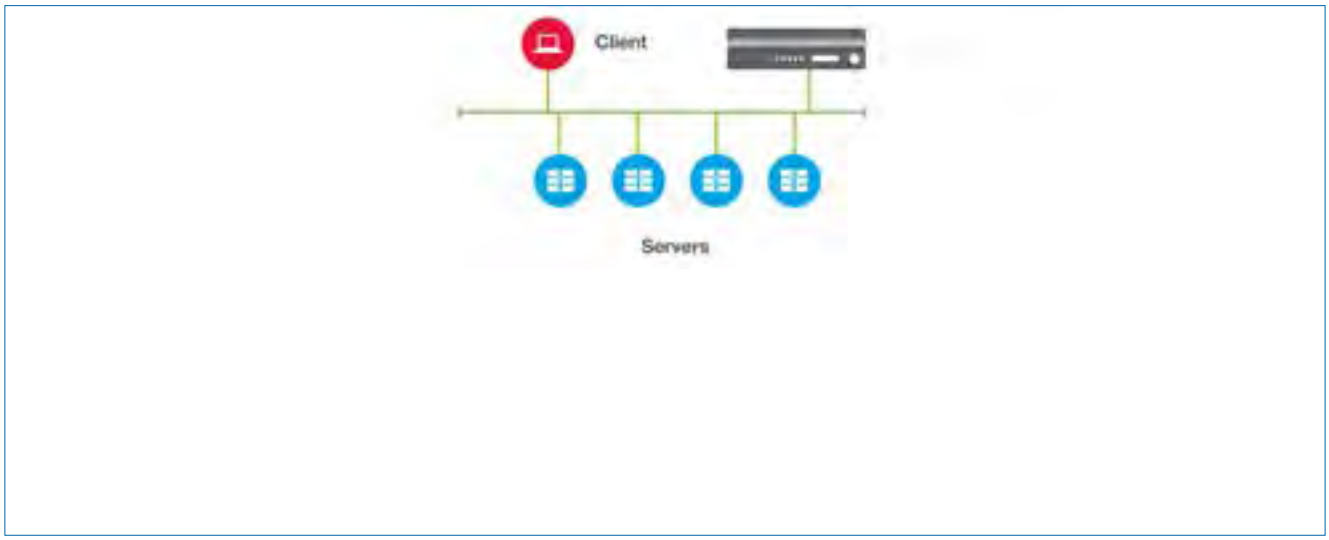

Figure 7.1: BIG-IP ASM routed mode deployment

- Servers detect the actual client IP address in the IP header for security and logging purposes. Since all communication traverses the BIG-IP system, there are no alternate, unprotected routes to the protected applications.
- The BIG-IP system must be configured to allow administrative and non-application traffic.
- Response traffic must be routed through the BIG-IP system. You usually accomplish this by setting the default gateway of the web server to the floating self-IP of the BIG-IP system.

One-armed mode

In one-armed mode, only application traffic flows through the BIG-IP ASM system, and the server-side connection uses a SNAT. The BIG-IP ASM appliance is logically in line with the web application traffic flow, but not physically in line with all traffic to and from the web servers.

Note Requests and responses must go through the BIG-IP system. This means that if you do not use NAT on the source IP address of the client, the default gateway of the server needs to be the BIG-IP system. If you do use SNAT for all traffic from the client to an IP address of the BIG-IP system, all responses are sent

back to the IP. To keep track of the original client IP address, you can enable the X-Forwarded-For feature of the HTTP profile. This adds the client IP address to the HTTP header that was sent to the web server.

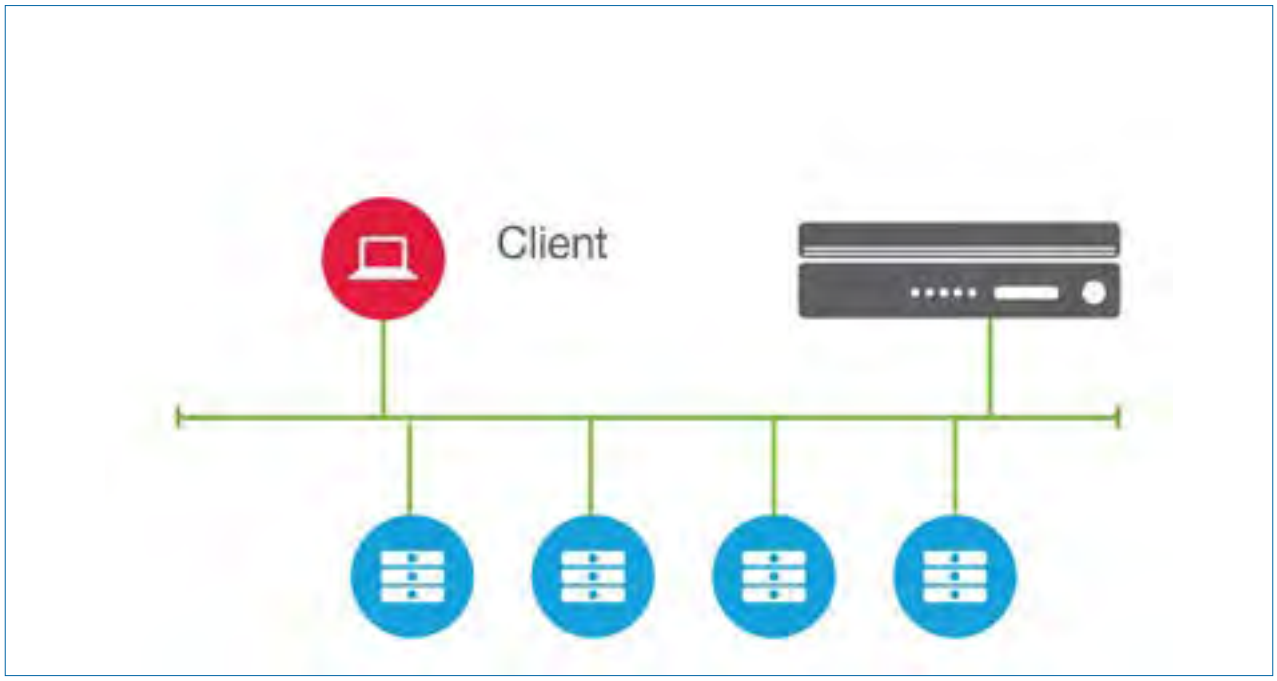

Figure 7.2: BIG-IP ASM single-arm deployment

- No changes to routing are required on the servers.
- Only application traffic is sent through the BIG-IP system, which reduces traffic traversing the device.
- Servers detect the IP address of the BIG-IP system in the TCP/IP header, which may complicate logging.

- There is more than one path to the protected application. You need additional security controls such as firewalls to ensure that malicious users do not access the application.

Tip If you expect a large number of concurrent connections, use a SNAT Pool instead of SNAT Automap to prevent port exhaustion on the BIG-IP system. For more information, refer to the AskF5 article: K7820: Overview of SNAT features.

Single-tiered topology

Combining BIG-IP LTM and BIG-IP ASM on a single BIG-IP platform has several advantages from an architectural perspective, such as:

- Simplified traffic flow
- Centralized administration on one device
- Encryption/decryption of traffic (SSL) on one device

Depending on project requirements, you may not need BIG-IP LTM features. In this case, you can deploy the BIG-IP ASM system only. (This is unlikely; most need BIG-IP LTM.)

Most deployments implement a redundant system configuration to ensure device high-availability.

Multi-tiered topology

The BIG-IP LTM load balances to multiple BIG-IP systems. Traffic terminates on the BIG-IP LTM where the BIG-IP ASM devices are configured in a pool.

- You can horizontally scale the BIG-IP ASM devices.
- You do not have to change licensing or the virtual server profile configuration on the load balancer.
- You can add or remove a BIG-IP ASM system from the pool without changing the network infrastructure.
- The BIG-IP ASM and BIG-IP LTM software versions do not need to match.
- You can use the BIG-IP ASM Virtual Edition (VE) to alleviate WAF processing, which is CPU-intensive.
- You can deploy BIG-IP ASM systems in an N+1 redundant system configuration.
- You have to manage more devices.

The BIG-IP LTM uses a virtual server with a pool assigned. The pool members are virtual servers of the back-end BIG-IP ASM systems.

- The system uses cookie persistence (Cookie Insert method) on the front-end BIG-IP LTM system. This means SSL needs to be terminated on the front-end BIG-IP LTM.
- The system uses only L7 monitors on the front-end because L4 or Internet Control Message Protocol (ICMP) monitors check only the state of the back-end BIG-IP ASM virtual servers and not the real application.

- Use sync-only device-groups on BIG-IP ASM stand-alone systems to ensure that the BIG-IP ASM policy is synchronized among all back-end BIG-IP ASM members.
- Configured thresholds are individual, per BIG-IP ASM pool member.

BIG-IP ASM systems are behind BIG-IP LTM, so statistics on the BIG-IP ASM system are individual, per stand-alone instance. The same is true for BIG-IP ASM event logs; therefore, F5 recommends that you use a central logging system if you need a centralized view of all event log data.

- On BIG-IP ASM systems, **Auto Last Hop** should be disabled on a per virtual server basis. Otherwise **Mac Masquerade** needs to be configured on BIG-IP LTM to ensure seamless failover.

Hybrid Topology with Virtual Edition

Other than some hardware offload considerations, such as SSL/TLS, BIG-IP ASM scaling is based on the available CPU and RAM, as well as the I/O capacity of the platform. Therefore, you may deploy the BIG-IP ASM system on a pool of VE instances, with a hardware appliance or VIPRION directing the traffic.

In this configuration, you can size the appropriate hardware appliance for L7 traffic management while scaling up capacity requirements for the BIG-IP ASM system by adding resources to the virtual appliances, increasing the number of virtual appliances, or by doing both. Doing so allows a flexible and efficient right-sizing of capacity.

The hybrid topology allows organizations with a robust and rapidly expanding virtual environment to use commodity computing resources and to rapidly deploy new BIG-IP ASM instances that do not need to be shipped or physically installed after purchase. VE licensing pools and other volume and on-demand purchasing and licensing options allow you to simply deploy another instance from a pool of pre-purchased licenses on demand.

- Rapidly deploy new BIG-IP ASM instances on demand.
- Add and remove commodity computing resources on demand.
- Use commodity computing resources for computing-constrained tasks while using hardware acceleration for application delivery tasks.
- Provide hardware crypto offload.
- Use BIG-IP LTM to direct traffic directly to the pools.
- Use hardware DDoS L3/L4 capabilities on physical devices, but scale L7 DDoS capabilities on commodity computing.
- Requires more complex deployment than a simple redundant system.
- Requires management of more than one system.
- Key management requires more operational overhead when re-encryption is required.
- If virtual BIG-IP ASM systems are not also licensed for BIG-IP LTM, traffic must traverse BIG-IP LTM twice (or to another tier of BIG-IP LTM) for distribution to the web server pools.

- SSL offload on hardware reduces VE load, but sends traffic unencrypted to the VE and to the servers. This does not apply to all applications in all environments. Using F5® OneConnect™ greatly reduces TCP overhead, improves traffic distribution, and reduces impact of the topology's added built-in network latency.
- SSL offload with re-encryption performs much better on VE when you use Elliptic Curve Cryptography (ECC) ciphers and certificates exclusively for the back-end encryption, all the way through to the web servers.
- If you are offloading or bridging SSL at an external BIG-IP LTM, F5 recommends that you enable an HTTP compression profile to perform hardware compression on the hardware at the outer tier. This improves the overall performance of BIG-IP ASM VEs.
- If VE appliances do not load balance to the web servers, traffic between BIG-IP LTM and BIG-IP ASM VE requires you to use SNAT and/or **Auto Last Hop**. If the web servers need to detect the original client IP address in the TCP/IP header, consider a different option.
- You can "re-SNAT" to the original client IP from a header. This requires you to enable **X-Forwarded-For** on the external BIG-IP LTM virtual server, enable **Trust XFF** in BIG-IP ASM policies, and set up an iRule on the load balancing BIG-IP LTM virtual server to use SNAT to return traffic to the original client IP address.
- With additional licensing, VE can perform crypto offload to a hardware appliance, a VIPRION system, or a third-party network HSM platform.

Note The crypto offload feature is outside the scope of this document. For more information, refer to **Overview: Implementing external cryptographic server offload** in ***BIG-IP Local Traffic Manager: Implementations*** for your system version.

Common Management Tasks

This chapter contains guidelines for common tasks involved in managing your BIG-IP ASM system.

You can deploy the BIG-IP ASM system with the following configuration types:

- **Stand-alone**

  High availability (HA) is not available.

- **Redundant system (Sync-Failover device group)**

  Multiple BIG-IP devices share the same BIG-IP version and configuration.

- **Redundant system (Sync-Only device group)**

  Multiple BIG-IP devices share a configuration, particularly specific application security policies.

Guidelines

Redundant systems guidelines

- A BIG-IP system can be a member of only one device group.
- All BIG-IP systems in a device group must use the same BIG-IP ASM version, including hotfix updates.
- BIG-IP systems in the device group synchronize application security configuration data and security policies, providing consistent enforcement on all the devices. However, event logs, reporting, and learning suggestions are not synchronized.
- Policy Builder can run on only one system per security policy. When you set up Policy Builder on one member of a device group, the policy is built on that BIG-IP system and then automatically updated on the other systems in the device group.
- The BIG-IP ASM system considers one VIPRION platform (with multiple blades) as one device. You need add only the master blade to the device trust and group.
- You cannot use connection mirroring on virtual servers that have a BIG-IP ASM policy attached.
- You must have consistent system times across all units using network time protocol (NTP).
- For BIG-IP systems in a group, you should enable ports 22 and 443 for communication among self IPs, and ports 1026 UDP and TCP for syncing between devices. Make sure to configure the **Port Lockdown** setting and use a self IP address rather than the management IP address for configuration synchronization (ConfigSync).
- We also need ports 1026 UDP and TCP for syncing between devices. For more information, refer to AskF5 article: K13250: Overview of port lockdown behavior (10.x - 11.x).
- When troubleshooting ConfigSync issues on a redundant system, review the **/var/log/ltm** file.

For more information, refer to **Synchronizing Application Security Configurations Across LANS** in ***BIG-IP ASM: Implementations*** for your system version.

Note For information about how to locate F5 product manuals, refer to the Ask F5 article: K12453464: Finding product documentation on AskF5.

Attack signatures guidelines

BIG-IP ASM attack signatures are an evolving set of protections that must be kept up-to-date to provide the best available protection against new and emerging threats and to ensure minimal false positives.

F5 publishes signature update files on a regular basis. You must establish a process to keep attack signatures updated and manage the update process within the policies.

The easiest way to update attack signatures is using automatic delivery mode through scheduled or manual updates. However, for air-gapped systems with no access to the Internet (either direct or via a proxy), use the manual delivery mode. Manual delivery mode allows you to download the update file manually from F5 Downloads and then upload that file using the BIG-IP ASM Configuration utility.

For more information, refer to **Working with Attack Signatures** or **Updating signatures manually** in ***BIG-IP ASM: Implementations*** for your system version, and ***BIG-IP ASM: Custom Signature Reference*** for your system version. Also refer to AskF5 article: K8217: Managing BIG-IP ASM attack signatures.

Signature staging guidelines

The BIG-IP ASM system places all new attack signatures into staging.

BIG-IP 13.0

In BIG-IP ASM 13.0, when the system updates existing signatures that are in Enforced mode, the existing version of the signature remains in Enforced mode while the updated version is placed in staging, where it remains until it has passed QA. This change ensures policy security throughout the entire signature updating process.

Versions prior to 13.0

In versions of the BIG-IP ASM system that precede 13.0, when the system updates existing signatures that are in Enforced mode as part of the updating process, they are removed from Enforced mode and put into staging.

If you update signatures to add protection or mitigation for a specific, newly identified vulnerability, after updating, you decide whether the new signature should be transitioned immediately into Enforced mode or whether it should remain in staging and be allowed to proceed as defined in staging settings.

Important If a signature is triggered on a request while in staging, the request isn't blocked.

Deleting inactive, allowed entities

Starting with BIG-IP ASM 13.1, you can reduce security policy inflation and simplify policy maintenance by deleting inactive entities. Policy Builder detects entities that your policy has not observed in traffic for more than 90 days

and displays them as learning suggestions on the following page: **Security** > **Application Security** > **Policy Building** > **Traffic Learning**. You can then accept or delete the **Delete inactive entity** suggestion**.**

You can also change the number of days Policy Builder uses to determine an entity is inactive from the default setting of 90 days on the following page: **Security** > **Application Security** > **Policy Building** > **Learning and Blocking settings** > **Policy Building Process (advanced settings)** > **Options**.

Policy Builder only monitors allowed entities for inactivity. Further, if you retain the default, pure wildcard (*), the system does not monitor allowed entities that Policy Builder adds using the wildcard.

The system generates learning suggestions to delete the following types of allowed, inactive entities:

- File type
- HTTP URL
- WebSocket URL
- Parameter
- Cookie
- Redirection domain
- Hostnames

For more information about learning settings, refer to **Reviewing learning settings** in ***BIG-IP Application Security Manager: Implementations*** for your system version.

Updating geolocation

The BIG-IP ASM system can protect against attacks based on the geolocation of the source IP during a request. Geolocation data is used to control the accessibility of requests in the **Geolocation Enforcement** page and to detect suspicious L7 denial-of-service (DoS) attacks from specific geolocations. Geolocation data changes regularly, so F5 releases a monthly download based on collected data.

You must manually update geolocation information using the **TMOS Shell (tmsh)** utility on your BIG-IP system.

For more information, refer to AskF5 article: K11176: Downloading and installing updates to the IP geolocation database.

Remove only the **/shared/tmp/GeoIP.old**/ path if the installation of all four **RPM** packages is successful, and you are able to verify proper functionality by querying an IP address using the **geoip_lookup** tool.

Important F5 strongly recommends performing the **md5sum** validation whenever you download F5 data over the Internet, particularly data to use on your BIG-IP devices. Without this validation check, compromised releases or corrupt archives can corrupt your configuration or otherwise cause adverse effects to your BIG-IP devices.

Maintaining the IP intelligence database

F5® IP Intelligence Services is a subscription-based database which allows BIG-IP systems to recognize IP addresses, including known botnets, anonymous proxies, and malicious actors.

In the BIG-IP ASM system, the IP intelligence database allows you to:

- Block a specific category of IP addresses.
- Log and report on the IP intelligence categories.
- Provide additional details for BIG-IP ASM scoring functionality to help better understand if a request is likely an attack (a score of 4-5) or a false positive (a score of 1-2).

Subscribing to and configure updates

To use IP Intelligence Services, you purchase an add-on license (1-year or 3-year options are available). Some functionality is available without the license through the Configuration utility. However, without the license, the database cannot be downloaded and automatic categorization is not available. You can purchase a license through the F5 sales team,

The IP intelligence database is updated very frequently, so the default minimum download interval is five minutes.

> **Note** The IP intelligence database is very large and can take a long time to download for the first time, depending on your connection. Later updates are incremental and thus quicker to download and apply.

You can update the IP intelligence database using a direct IP connection or through an intermediate proxy.

> **Important** DNS lookups on the BIG-IP system must first be configured for the database to successfully download.

For more information, refer to **Enabling IP Address Intelligence** or **Setting Up Address Intelligence Blocking** in ***BIG-IP ASM: Implementations*** for your system version. Also refer to the AskF5 article: K13875: Managing IP reputations and the IP Address Intelligence database.

Verifying updates

After you've downloaded and installed updates to the IP intelligence database, you should verify the installation by checking the last time an update was received.

To display date and time of updates in the Configuration utility

1. Navigate to **Security > Application Security > IP Address > IP Address Intelligence**.
2. Click **Intelligence last updated**.

To display date and time of updates using tmsh at the command line

- Type the following command:

```
tmsh show sys iprep-status
```

For more information, refer to AskF5 articles: K13776: Determining the IP intelligence subscription expiration date and K13653: The IP Intelligence Service database cannot be updated.

Checking IP address reputation

You can check the reputation of a specific IP address using the IP intelligence database.

To check reputation of a specific IP address at the command line

- Use the following command syntax:

```
iprep _ lookup <IP address>
```

Output displays similar to the following example:

```
iprep _ lookup 1.1.1.1  opening database in /var/IpRep/F5IpRep.dat size of IP
reputation database = 4163298  iprep threats list for ip = 1.1.1.1 is:  bit 4
– Scanners  bit 5 – Denial of Service
```

Adding an IP address to the IP Address Whitelist

You can eliminate false positives by adding specific IP addresses to the **IP Address Whitelist**.

To add an address to the IP Address Whitelist in the Configuration utility

1. Navigate to **Security > Application Security > IP Addresses > IP Address Intelligence**.
2. In the **IP Address Whitelist** box, type the IP address.
3. Click **Save**.

For more information, refer to **Setting Up IP Address Intelligence Blocking** in ***BIG-IP ASM: Implementations*** for your system version.

Logs

IP intelligence functionality on the BIG-IP system is performed by the **iprepd** process, which logs to the **/var/log/iprepd/iprepd.log** file.

Backing up your BIG-IP configuration information

F5 strongly recommends performing regularly scheduled backups of your BIG-IP system. At a minimum, backups must be performed before and after any major change to the system, such as a configuration change or upgrade.

You can use the Configuration utility to trigger backups, or you can use a central management platform such as F5® BIG-IQ® Centralized Management or F5® Enterprise Manager™. Only the configuration information is backed up.

Event logs, reporting, and learning suggestions are not backed up and cannot be saved from the local system. For this reason, these items are also not synchronized between devices or blades in a cluster.

Tip To store event logs across upgrades and configuration restore operations, F5 recommends using remote logging and sending the log data to an external **syslog** server or Security Information and Event Management (SIEM) for storage. For more information, refer to Common Management Tasks.

Restoring from a backup

Before restoring a BIG-IP system from a backup, you should:

- Make sure you have a valid BIG-IP system license for that system.
- Provision BIG-IP ASM on the system.

For more information, refer to AskF5 article: K13945: The BIG-IP ASM MySQL database is not installed completely if the BIG-IP ASM is not provisioned when the UCS is loaded.

To successfully install a user configuration set (UCS) archive file on a BIG-IP system, perform one of the following actions:

- Restore the UCS archive to the same system from which it was saved.
- Have the license associated with the serial number of a new system. To do so, contact F5 Technical Support.

Note F5 Technical Support associates a license file with a new serial number only on an as-needed basis, in the event of a Return Materials Authorization (RMA).

- Relicense the BIG-IP system after restoring the UCS archive.
- Save the license file prior to restoring the configuration from another system, and then copy the license file back.
- Install the UCS archive by using tmsh **no-license** option.

To install UCS by using tmsh no-license option

- Type the following command syntax:

```
tmsh load sys ucs [ucs file name] no-license
```

For more information, refer to AskF5 articles: K13132: Backing up and restoring BIG-IP configuration files (11.x - 13.x) and K12880: Configuring a replacement BIG-IP device after a Return Materials Authorization.

Troubleshooting BIG-IP ASM

Using platform logs

When troubleshooting your BIG-IP system, you should know the locations of the different log files and understand their contents.

The following table describes the contents of each platform log related to the BIG-IP ASM system.

Table 7.1 Platform log contents

Log	Description
/var/log/ltm	Contains general BIG-IP LTM log entries, such as availability of pool members, high availability, and **config-sys** entries.
/var/log/asm	Contains critical messages from processes that the BIG-IP ASM system uses (MYSQL, policy building engine, Enforcer). Each process has its own log file under **/var/log/ts** (see log, following).
/var/log/audit	Contains audit trails for BIG-IP user activities such as failed login attempts and certain security events involving HTTPS, secure shell (SSH), or other configuration changes.
/var/log/ts/...	Contains information about various BIG-IP ASM system daemons and policy building.
/var/log/pktfiler	Contains results from implementation of packet filters and packet filter rules.
/var/log/dosl7	Contains BIG-IP ASM system DDoS event information such as attack started/stopped and bot signature update status.
/var/log/iprepd/iprepd.log	Contains entries for the IP Intelligence Services updates.
/var/log/icrd	Contains entries for F5® iControl® API calls, their results, and failures.
/var/log/restjavad.0.log	REST API Java framework used by the BIG-IP system and BIG-IQ Centralized Management to communicate writing logs. Log entries include communication attempts and failures, such as SSH handshake failures, certification, and authentication failures, and time skews.
/var/log/tmm[0...]	Contains traffic events, such as unexpected resets, DDoS drops, and other events. Typically very large. File should be analyzed only with F5 support assistance.

You can configure the logging level for most platform logs at **System\Logs\Configuration\Options**.

Checking BIG-IP ASM system health

There are several ways to check if your BIG-IP ASM system is up and running.

Checking if BIG-IP daemons are running

- Check logs for error messages and illegal requests.

- Check for TS-cookie in header.
- Upload a QKView file to BIG-IP iHealth.

Checking if BIG-IP daemons are running

To check if BIG-IP daemons are running, using tmsh

1. Type the following commands:

```
tmsh show sys service asm
tmsh show sys service mysql
```

2. Make sure that all daemons are running.

For detailed information on core BIG-IP ASM services, and the impact to the BIG-IP ASM system operation if the service is not running, refer to the AskF5 article: K14020: BIG-IP ASM daemons (11.x - 14.x).

Checking logs for error messages and illegal requests

1. Check for error messages in the following logs:

```
/var/log/ltm
/var/log/asm
/var/log/dosl7/dosl7d.log
/var/log/mysql.out
```

2. Check the request log for illegal requests.

For more information, refer to Using platform logs.

Checking reporting tools for violations and attacks

The BIG-IP ASM system provides several reporting tools to assist you in troubleshooting violations and attacks, including:

- Application security overview
- DoS attacks report
- Brute force attacks report
- Web scraping statistics
- Session tracking status

Event correlation report

Added to the BIG-IP ASM system in BIG-IP 13.1, the event correlation report is an additional reporting tool designed to facilitate incident response.

When the system applies one or more correlation heuristics to illegal requests, it triggers an incident. An incident is comprised of requests that appear to be part of an attack. The system identifies each request as either a false positive violation or malicious activity. Incidents originate from a single source: either a device ID, or, if not available, a source IP.

The event correlation report is a list of these incidents. For details about how to view incidents, refer to the online help: **Event Logs** > **Event Correlation**.

For more information about reporting tools, refer to **ASM Reporting Tools** in ***BIG-IP Application Security Manager: Implementations*** for your system version.

Note For information about how to locate F5 product manuals, refer to the Ask F5 article: K12453464: Finding product documentation on AskF5.

Checking for TS-cookie in header

If the security policy is configured on a virtual server, you see the TS-cookie in place in the HTTP response from the BIG-IP ASM system. In addition, the server-header of the response is removed. Both happen in **Blocking** and **Transparent** mode.

```
HTTP/1.1 200 OK
Date: Tue, 17 Nov 2015 13:01:53 GMT
LastModified: Mon, 21 Apr 2014 07:13:20 GMT ETag: "970c1784f7883bf94780"
AcceptRanges: bytes ContentLength: 376
KeepAlive: timeout=300, max=500 Connection: KeepAlive
ContentType: text/html; charset=UTF8 SetCookie:
S01425764=01ae97ff1e19f24a9b73757578683232de31fec9e054b6f04b167d6a7d88f4f7333871255d;
Path=/

<html>
...
...
</html>
```

For more information, refer to the AskF5 article: K6850: Overview of ASM cookies.

Uploading a QKView file to iHealth

BIG-IP iHealth® is freely available to customers who run BIG-IP 10.x and later, or Enterprise Manager 2.x and later. iHealth enables you to verify operation of your BIG-IP system and ensures that your hardware and software function at peak efficiency by providing information covering your hardware, software, licensing, configuration, best practices, and known issues.

iHealth is a hosted application that parses a QKView file. The QKView provides a running snapshot of your BIG-IP system, with up-to-the-minute configuration and diagnostic information. You can download a QKView from your BIG-IP system and then upload the file to the iHealth system.

For more information, refer to **BIG-IP iHealth** in ***F5 BIG-IP TMOS: Operations Guide***.

Bypassing BIG-IP ASM

If the BIG-IP ASM system service has insufficient resources or is down, you can allow web application traffic to bypass it.

Starting in BIG-IP ASM 10.2.3, you can configure the following system variables using the following Configuration utility commands:

- **bypass_upon_load**
- **bypass_upon_asm_down**

For more information, refer to AskF5 article: K15093: The BIG-IP ASM system bypass_upon_load and bypass_upon_asm_down variables should not be enabled.

If you enable the **bypass_upon_load** variable (set the value to **1**), web application traffic bypasses the BIG-IP ASM system when there are insufficient resources for BIG-IP ASM service.

If you enable the **bypass_upon_asm_down** variable (set the value to **1**), web application traffic bypasses the BIG-IP ASM system when any of the following conditions occur:

- BIG-IP ASM service is stopped.
- BIG-IP ASM service is restarted; traffic bypasses the BIG-IP ASM system from the time the BIG-IP ASM service is down until the service resumes processing.
- BIG-IP ASM service performs a core dump; traffic bypasses the BIG-IP ASM system until the BIG-IP ASM service resumes processing.

If you enable either or both of the system variables, all web application traffic that is protected by BIG-IP ASM security policies is no longer directed through the BIG-IP ASM system for security checks when BIG-IP ASM service experiences the previously-described conditions; traffic is forwarded directly to the origin web servers. As a result, the web application may be at risk of security threats and false positives.

Caution F5 recommends that you enable these system variables with careful consideration to the security impact on your application environment.

Handling unexpected HTTP responses

In rare cases, you may need to review an entire HTTP transaction to determine the reason the expected response is not returned to the client. The following flowchart shows the steps you can follow to decide whether the BIG-IP ASM system is processing an HTTP response correctly:

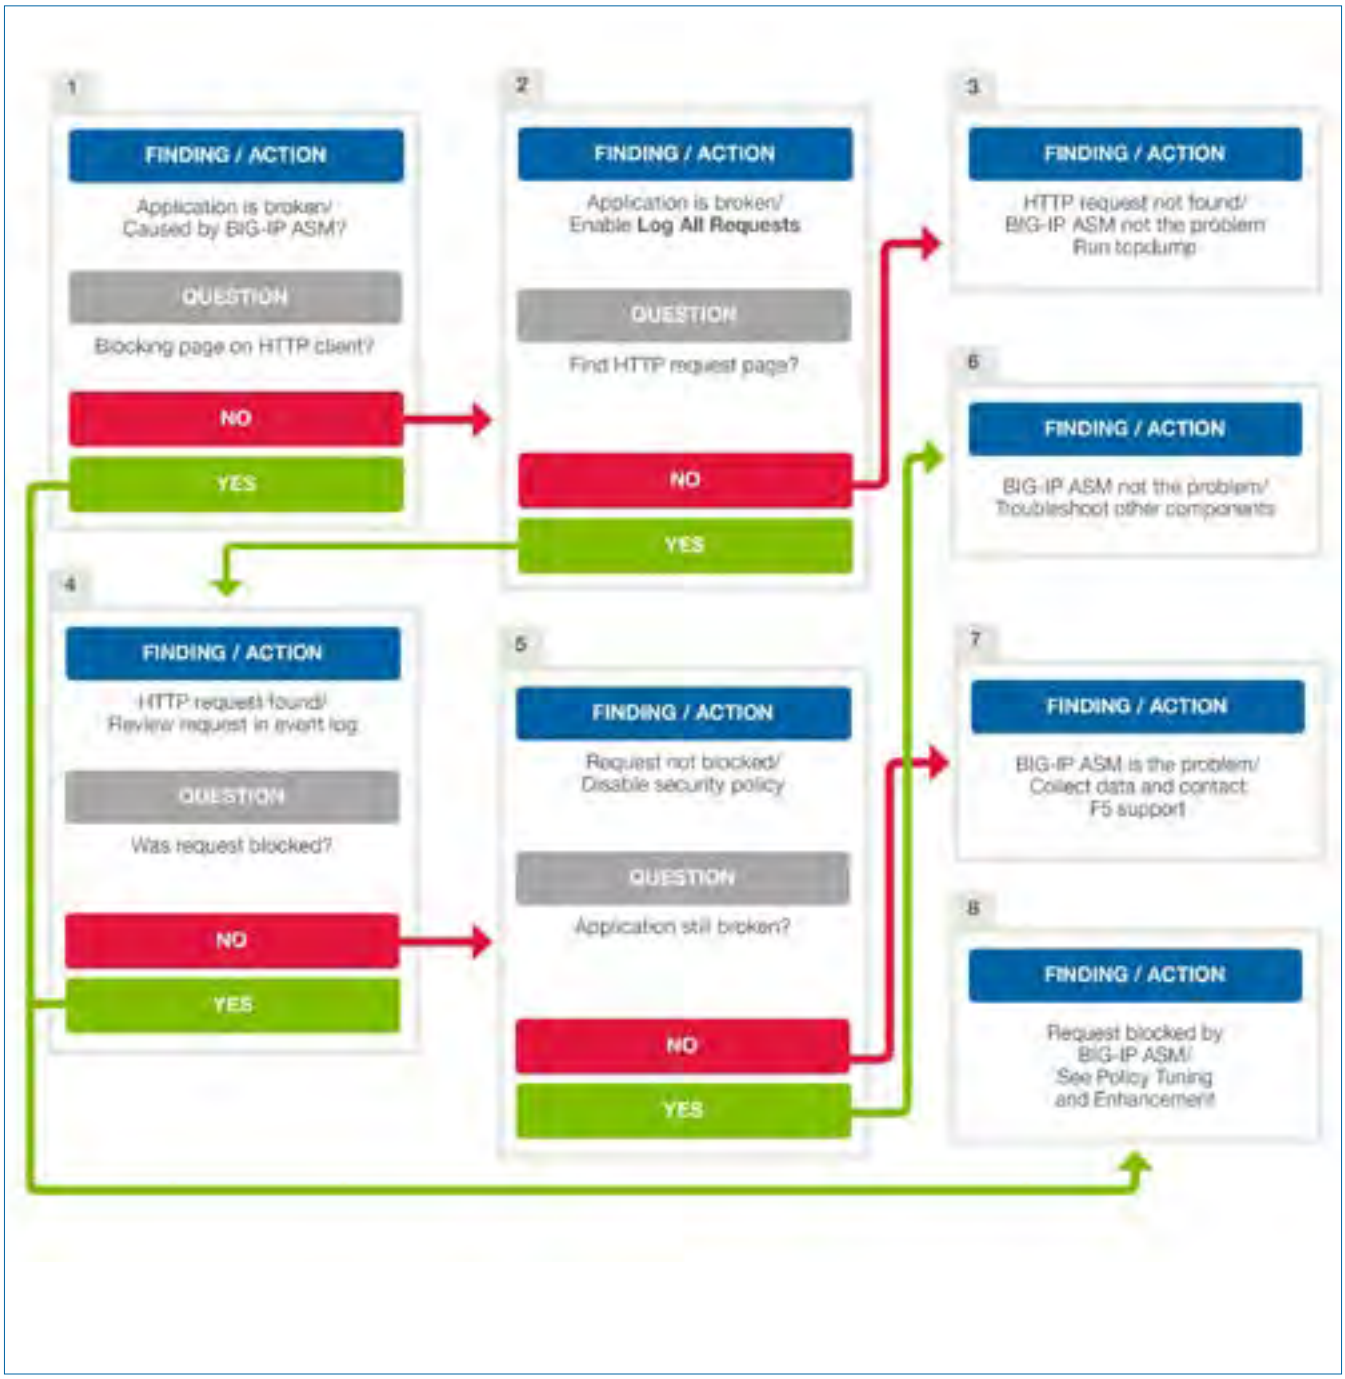

Figure 7.1: Troubleshooting flow chart for unexpected HTTP responses

1	
Finding	Application is broken. Page is blank. Response is blocked. Connection is reset.
Action	Find out if the BIG-IP ASM system is causing the problem.
Question	Does the request trigger the BIG-IP ASM blocking response page on the HTTP client?
YES	Request is blocked by BIG-IP ASM security policy. Go to step 8.
NO	Go to step 2.
2	
Finding	Application is broken. Page is blank. Response is blocked. Connection is reset. BIG-IP ASM blocking response page is not displayed on HTTP client.
Action	Enable **Log All Requests** in the event log on the virtual server running the BIG-IP ASM system (refer to BIG-IP ASM Event Logging). Try to match the HTTP request in question with the event log. (Refer to Capturing HTTP request/response.)
Question	Can you find the HTTP request page?
NO	HTTP request did not arrive at the BIG-IP ASM system. Go to step 3.
YES	Go to step 4.
3	
Finding	HTTP request does not arrive at the BIG-IP ASM system.
Action	The BIG-IP ASM system is not the problem. To continue investigating, run **tcpdump** on the virtual server running BIG-IP ASM to see if the HTTP request reaches the BIG-IP system. Use the following syntax at the command line: `tcpdump –I 0.0:nnn –s 0 –w /var/tmp/asm_client.cap host <virtual server IP address> and port <virtual server port>`

4	
Finding	Matched HTTP request appears in the BIG-IP ASM event log.
Action	Review the logged request by selecting the HTTP request in the event log.
Question	Does the request log show that the request was blocked?
NO	Go to step 5.
YES	Request blocked by BIG-IP ASM security policy. Go to step 8.
5	
Finding	HTTP request appears in BIG-IP ASM request log. HTTP request not blocked by the BIG-IP ASM system.
Actions	Do one of the following: Remove the BIG-IP ASM security policy from the virtual server. Disable the BIG-IP ASM system for a specific IP address URL using an iRule or local traffic policy. For iRule examples, refer to Wiki: iRules API on DevCentral. **Note** A DevCentral login is required to access this content.
Question	Does the issue continue after the removal of BIG-IP ASM security policy?
NO	Go to step 7.
YES	The BIG-IP ASM system is not the problem. Go to step 6.
6	
Finding	The BIG-IP ASM system is not the problem.
Action	Continue troubleshooting to find out if BIG-IP LTM is the problem.
7	
Finding	The BIG-IP ASM system is the problem. Request is not blocked by the BIG-IP ASM system in a security policy.
Action	Run **tcpdump** on client side and server side connections. Use the following command syntax: `tcpdump –I 0.0:nnn –s –w /var/tmp/asm _ clientserver.cap \(host <virtual server IP address> and port <virtual server port>\) or \(host <pool member IP address> and port <pool member port>` Contact F5 support. Refer to F5 Support resources.
8	
Finding	Request is blocked by the BIG-IP ASM system. Possible false-positive violation.
Action	Tune your BIG-IP ASM security policy. Refer to Policy Tuning and Enhancement.

Monitoring performance

Performance can be monitored in several ways. You can use performance graphs, iHealth, or the **top** command.

Performance graphs

The BIG-IP system provides performance graphs that show historical running performance data for up to one month. BIG-IP 12.0 and later also provides real-time dashboard style performance analytics, including CPU utilization, memory utilization, and bypass information**.**

Regularly check these graph to monitor the overall health and capacity of the BIG-IP system. They can show early signs of potential issues, allowing you to act on problems before they occur.

iHealth

iHealth provides an alternative for viewing Round Robin Database (RRD) based performance graphs, and it is easy to compare performance data between different BIG-IP ASM systems, and snapshots of different time periods. iHealth provides heuristic diagnostics based on log messages and statistical data to help a customer identify performance issues.

For more information, refer to **BIG-IP iHealth** in ***F5 BIG-IP TMOS: Operations Guide***.

top command

Table of processes (**top**) is a Linux/Unix command-line interface (CLI) tool for process reporting, and is used for troubleshooting CPU and memory performance issues on the BIG-IP ASM system.

You can start **top** in **Interactive** mode and watch the performance data change.

Monitoring CPU usage

F5 recommends that you regularly check the CPU usage of the BIG-IP ASM system.

Normal CPU usage at the peak of any CPU/thread on the system should not exceed 75 percent under normal traffic volume.

However, some administrative processes, such as those related to learning, may cause higher loads on the last CPU core. This is not generally a problem on most platforms that use process scheduling and process separation across CPU cores.

Note For BIG-IP 11.5 and later, on systems with a CPU using Hyper-Threading Technology, Traffic Management Microkernel (TMM) runs on half of the cores/threads on the system. For more information, refer to the AskF5 article: K5003: Data plane and control plane tasks use separate logical cores when the BIG-IP system CPU uses Hyper-Threading Technology.

If CPU usage is too high, you need to identify the process that is causing the problem. Use the performance graph to view CPU usage per core/thread and see whether CPU over-use occurs on multiple cores/threads or just one.

Some important BIG-IP ASM-related processes such as TMM or the BIG-IP ASM traffic processing daemon (**bd)** are multi-threaded and tend to spread their threads across all even-numbered CPU cores. If you find excessive CPU usage on all even-numbered cores, these multi-threaded processes are likely the problem.

The top command provides dynamic CPU usage by each process; using this command, you can display additional threads and cores information.

To display additional threads and cores using the top command

1. Remove the current **/root/.toprc** file.
2. Run the **top** command in **Interactive** mode.
3. To get the command to show threads, type **H**.
4. To get the command to show individual cores, type **1** .
5. Type **f**
6. Type **j**.
7. To add **Last used cpu** as a column, press **Enter**.

   The system displays the following output:

```
top 02:18:26 up 2 days, 18:36, 1 user, load average: 0.13, 0.23, 0.19
Tasks: 659 total, 1 running, 656 sleeping, 0 stopped, 2 zombie
Cpu0 : 9.1%us, 2.0%sy, 4.0%ni, 84.8%id, 0.1%wa, 0.0%hi, 0.0%si, 0.0%st
Cpu1 : 8.9%us, 2.5%sy, 3.7%ni, 84.8%id, 0.1%wa, 0.0%hi, 0.1%si, 0.0%st
Cpu2 : 7.5%us, 1.7%sy, 3.4%ni, 87.4%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st
Cpu3 : 5.9%us, 1.1%sy, 1.4%ni, 90.1%id, 1.1%wa, 0.4%hi, 0.0%si, 0.0%st
Cpu4 : 15.4%us, 3.8%sy, 0.7%ni, 80.2%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st
Cpu5 : 10.7%us, 2.7%sy, 0.8%ni, 85.8%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st
Cpu6 : 7.7%us, 1.8%sy, 0.8%ni, 89.7%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st
Cpu7 : 26.7%us, 5.2%sy, 0.1%ni, 67.0%id, 0.0%wa, 0.3%hi, 0.6%si, 0.0%st
Mem: 16528548k total, 14496992k used, 2031556k free, 569104k buffers
Swap: 1048572k total, 0k used, 1048572k free, 2747692k cached
  PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ RUSER P DATA COMMAND
20273 root 20 0 2038m 66m 26m S 8.7 0.4 20:31.20 root 1 338m avrd
17088 root RT 0 8915m 127m 101m S 8.7 0.8 330:27.58 root 0 206m tmm.0
17348 root RT 0 8915m 127m 101m S 8.7 0.8 309:40.64 root 1 206m tmm.1
```

```
17350 root RT 0 8915m 127m 101m S 8.7 0.8 312:17.03 root 2 206m tmm.2
17347 root RT 0 8912m 125m 99m S 8.7 0.8 310:49.88 root 6 206m tmm.6
17349 root RT 0 8912m 125m 99m S 8.7 0.8 312:14.12 root 7 206m tmm.7
17351 root RT 0 8915m 127m 101m S 7.0 0.8 287:02.05 root 3 206m tmm.3
17089 root RT 0 8912m 125m 99m S 7.0 0.8 314:44.24 root 4 206m tmm.4
17346 root RT 0 8912m 125m 99m S 7.0 0.8 308:54.54 root 5 206m tmm.5
7417 root 20 0 26776 20m 12m S 5.2 0.1 123:34.95 root 4 4140 csyncd
23050 root 20 0 3060 1384 792 R 5.2 0.0 0:00.06 root 6 828 top
5907 root 20 0 159m 118m 32m S 3.5 0.7 91:54.95 root 5 84m mcpd
  17 root 20 0 0 0 0 S 1.7 0.0 11:10.78 root 3 0
ksoftirqd/3
5044 root 20 0 39332 21m 13m S 1.7 0.1 0:05.33 root 1 4268 eventd
5652 root 20 0 28480 27m 19m S 1.7 0.2 6:34.82 root 5 4184 clusterd
9214 root 20 0 2220 884 744 S 1.7 0.0 50:28.49 root 0 248 LCDd
  1 root 20 0 2904 1368 1164 S 0.0 0.0 0:00.91 root 3 260 init
  2 root 20 0 0 0 0 S 0.0 0.0 0:00.02 root 6 0 kthreadd
  3 root RT 0 0 0 0 S 0.0 0.0 0:00.05 root 0 0 migration/0
```

To see utilization data for each virtual server CPU

1. Navigate to **Statistics > Module Statistics: Local Traffic > Virtual servers**.
2. See the data in the **CPU Utilization Avg.** and **ASM CPU Utilization Avg.** columns.

Troubleshooting memory usage

Monitor the memory usage to establish a baseline that is specific to your BIG-IP system and your network traffic.

Memory considerations when using CLI tools

TMM has a fixed memory size, which the system allocates when you start the system. When you use the **top** command, the system output appears as a large VIRT (virual size of a process) memory allocation.

For performance reasons, **bd** does not release memory back to the operating system. Typical **bd** process is as follows:

1. You start the system.
2. The **bd** process immediately loads the configuration and allocates initial buffers.
   This process uses a small amount of resident memory.
3. If no memory is available, the **bd** process sends a request to the system to allocate new memory for a fresh buffer.
4. The **bd** process retains this memory to use for a future request.
5. On the first day in production, **bd** memory increases and peaks as traffic throughput increases.
6. On subsequent days, **bd** memory stabilizes, unless peak throughput increases.

If **bd** memory continues to increase linearly, you should investigate.

The following graphs show typical BIG-IP ASM system memory use over time with respect to traffic.

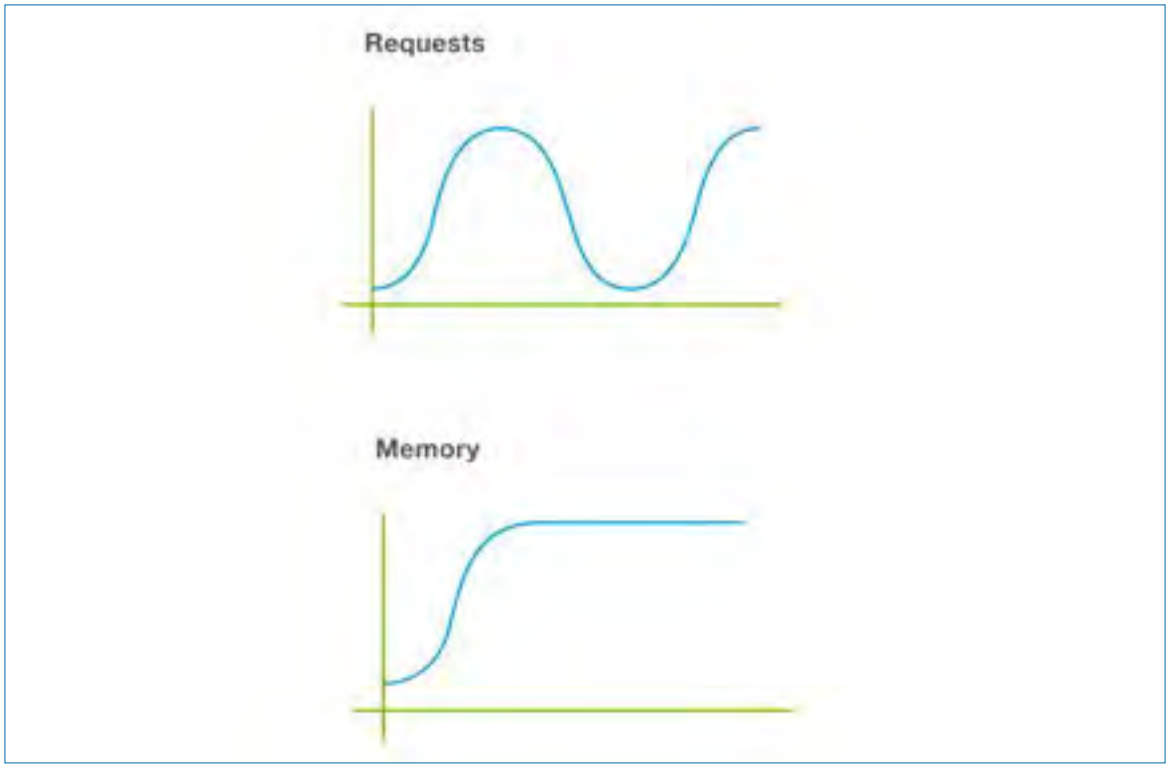

Figure 7.2: BIG-IP ASM memory use over time with respect to traffic

To identify a potential memory leak in a BIG-IP ASM process, use the **top** command to collect continuous snapshots for a particular interval, and monitor a process's memory use.

To use the top command to monitor memory

- Use the following command syntax:

```
top –b –n <number of snapshots> -d <interval of snapshots in seconds> > /
var/tmp/top _ output.txt
```

For example, the following example captures 60 iterations at 10 second intervals and log the output to **/var/tmp/top_output.txt**.

```
top –b –n 60 –d 10 > /var/tmp/top _ output.txt
```

Using these settings, the capture takes 10 minutes (60 x 10 seconds).

Optimizing the Support Experience

F5 technical support commitment

F5® strives to continuously improve its support service and create closer customer relationships. Designed to provide assistance with specific break-fix issues and ongoing maintenance of F5 products, F5 professional support services are consistently high-quality.

This means:

- F5 network support engineers conduct themselves professionally at all times.
- F5 is committed to providing the best customer experience possible.
- F5 treats customers are with respect and give them every consideration possible.
- F5 aims to provide resolutions the first time, every time.
- You can ask for manager escalation for unresolved or “site down” issues.

Some technical support issues arise from configuration errors, either within the BIG-IP® system or with other devices in the network. In other cases, a misunderstanding of BIG-IP capabilities can lead to support questions and issues. Although F5 does everything possible to prevent defects in BIG-IP hardware and software, these issues may still arise periodically. Regardless of the root cause of a problem, the goal is to resolve any issues quickly.

F5 technical support offerings

A variety of technical support offerings are available to provide the right level of support for any organization.

F5 Standard and Premium Support include remote assistance from F5 network support engineers, both online and over the phone.

Premium Plus customers receive priority status at F5, with fast, easy access to a dedicated team of senior-level, F5-certified network support engineers and a Technical Account Manager.

To learn more, refer to F5 Technical Support Offerings or send email to services@f5.com.

Professional services

Take advantage of the full range of F5 Professional Services to help you design, customize, and implement a solution that is right for your IT infrastructure and which supports your business goals.

Professional Services (f5.com/support/professional-services) provides information on a wide range of F5 Professional Services offerings and Professional Services Partners. You can use our online forms to request Consulting Services OnDemand for custom, shorter scope consulting engagements, or iRules® OnDemand to get fast access to iRules scripts tailored to your specific needs.

You can make an online request for specific support services by filling out a request form:

- Consulting request form (https://www.f5.com/services/professional-services/request-f5-professional-services).

GUARDIAN Professional Services Partners

F5 GUARDIAN® Professional Services Partners are authorized as installation providers and are also available to assist you. F5 GUARDIANs are selected because they have the skills and experience required to ensure successful implementations of F5 BIG-IP installations.

Refer to F5 GUARDIAN Professional Service Partners (f5.com/support/professional-services#guardian) for a complete list of partners.

F5 certification

F5 Certified® exams test the skills and knowledge necessary to be successful when working with today's application delivery challenges. Our technically relevant and appropriate exams deliver consistently reproducible results that guarantee excellence in those that achieve certification.

Certification levels

F5 Certified! is the F5 certification program, with a progressive program of four levels (Administrator, Specialist, Expert, and Professional), each of which build on the skills and knowledge demonstrated on previous exams.

C1 – F5 Certified BIG-IP Administrator (F5-CA)

The starting point for all certifications: a certified BIG-IP Administrator has basic network and application knowledge to be successful in application delivery.

C2 – F5 Certified Technology Specialists (F5-CTS)

The Technology Specialist certification assures employers that the candidate is fully qualified to design, implement, and maintain that specific product and its advanced features.

C3 – F5 Certified Solution Expert (F5-CSE)

The Solution Expert focuses on how F5 technologies combine with industry technology to create real-world business solutions.

C4 – F5 Certified Application Delivery Engineer (F5-CADE)

The Application Delivery Engineer certification exam and requirements are still under development.

C5 – F5 Certified Application Delivery Architect (F5-CADA)

The Application Delivery Architect certification exam and requirements are still under development.

Certificate expiration

F5 certifications are valid for two (2) years. Three months before the expiration date, the holder becomes recertification-eligible and can register for the exam necessary to re-certify. Only the last exam in the highest level

certification achieved needs to be retaken.

Certification beta program

F5 uses beta exams in the creation of all our exams and to maintain their relevancy and accuracy after production. Beta exams are open to all and give candidates an opportunity to have an impact on the F5 Certified program. While beta exams are twice as long, they cost less than regular exams and give candidates the chance to leave feedback on the exam. Beta exams are critical to our exam development process and a great way to change the F5 Certified program for the better.

Get involved

There are a several ways to get involved with the F5 certification beta program:

- Beta participation. Interested in taking our beta exams? Contact us at **F5Certification@f5.com** to learn more.
- Exam development. Contact us at **F5Certification@f5.com** if you're interested in helping us create our Certified exams.
- LinkedIn community. Join us on LinkedIn (https://www.linkedin.com/company/f5-networks) for answers to frequently asked questions, community developed resources, and more.

Note: This link takes you to a resource outside of F5, and it is possible that the document may be removed without our knowledge.

Visit F5 Credential Manager System (certification.f5.com) for information or follow the steps to get registered.

Self-help

F5 offers a number of resources to assist in managing and supporting your F5 systems:

- AskF5™ (support.f5.com)
- Downloads (downloads.f5.com) User name and password required.
- Security Updates (interact.f5.com/AskF5-SubscriptionCenter.html)
- BIG-IP iHealth® (f5.com/support/tools/ihealth)
- TechNews (interact.f5.com/AskF5-SubscriptionCenter.html)
- RSS feeds (https://support.f5.com/csp/article/K9957)
- DevCentral (devcentral.f5.com/) User name and password required.
- F5 Training Programs and Education (f5.com/education/training)

AskF5

AskF5 (support.f5.com) is a great resource for thousands of articles and other documents to help you manage your F5 products more effectively. Step-by-step instructions, downloads, and links to additional resources give you the means to solve known issues quickly and without delay, and to address potential issues before they become reality.

Whether you want to search the knowledge base to research an issue, or you need the most recent news on your F5 products, AskF5 is your source for product manuals, operations guides, and release notes, including the following:

- F5 announcements
- Known issues
- Security advisories
- Recommended practices
- Troubleshooting tips
- How-to documents
- Changes in behavior
- Diagnostic and firmware upgrades
- Hotfix information
- Product life cycle information

Downloads

Downloads are available from the F5 website. F5 strongly recommends that you keep your F5 software up-to-date, including hotfixes, security updates, OPSWAT updates, BIG-IP Application Security Manager™ (ASM®) signature files, and geolocation database updates. All software downloads are available from F5 Downloads (https://downloads.f5.com).

Security updates

You can receive timely security updates and BIG-IP ASM attack signature updates from F5. When remote vulnerabilities are discovered, F5 implements, tests, and releases security hotfixes for any vulnerable supported version, and sends an email alert to the F5 Security mailing list. F5 encourages customers with an active support account to subscribe to this list. For more information, refer to AskF5 article: K41942608: Overview of AskF5 security advisory articles.

BIG-IP iHealth

The BIG-IP iHealth® (iHealth.f5.com) diagnostic viewer is among the most important preventative tools to verify the proper operation of your BIG-IP system. It ensures hardware and software are functioning at peak efficiency and helps detect and address issues that may potentially affect F5 systems. BIG-IP iHealth is not integrated within the BIG-IP system. It is hosted by F5 and can be accessed with any web browser.

F5 recommends you generate a BIG-IP iHealth QKView file on the BIG-IP system and upload it to iHealth on a weekly basis in order to benefit from the many regularly occurring diagnostic updates. Uploading QKView files to iHealth also provides F5 technical support with access to your QKView files if you open a support case.

By reviewing the iHealth output, many of the issues commonly experienced by customers can be resolved without the need for opening a support case with F5.

TechNews

Communications Preference Center provides two email publications to help keep administrators up-to-date on various F5 updates and other offerings:

- **TechNews Weekly eNewsletter** Up-to-date information about product and hotfix releases, new and updated articles, and new feature notices.
- **TechNews Notifications** Do you want to get release information, but not a weekly eNewsletter? Sign up to get an HTML notification email any time F5 releases a product or hotfix.
- **Security Alerts** Receive timely security updates and ASM attack signature updates from F5.

AskF5 recent additions and updates

You can subscribe to F5 RSS feeds to stay informed about new documents pertaining to your installed products or products of interest. The Recent additions and updates page on AskF5 provides an overview of all the documents recently added to AskF5.

New and updated articles are published over RSS. You can configure feeds that pertain to specific products, product versions, and/or document sets. You can also aggregate multiple feeds into your RSS reader to display one unified list of all selected documents.

DevCentral

DevCentral™ (devcentral.f5.com) is an online forum of F5 employees and customers that provides technical documentation, discussion forums, blogs, media and more, related to application delivery networking. DevCentral is a resource for education and advice on F5 technologies and is especially helpful for iRules and iApps® developers. Access to DevCentral is free, but registration is required.

As a DevCentral member, you can do the following:

- Ask forum questions.
- Rate and comment on content.
- Contribute to “wikis.”
- Download lab projects.
- Join community interest groups.
- Solve problems and search for information.
- Attend online community events.
- View educational videos.

F5 training programs and education

F5 provides training programs and education, including traditional classroom learning opportunities, live online training, and free, self-paced online courses to help you get the most out of your investment. F5 Education (f5.com/education/training) provides links to course schedules, pricing, and registration details. It also has information about alternative training solutions such as virtual and web-based training for those who cannot attend training in person.

- **In-person courses**: F5 courses are available in multiple training facilities across five continents. Each one combines instructor presentations, classroom discussions, and interactive labs. The hands-on learning environment helps provide a fast track to accomplishing your goals.
- **Virtual instructor-led training**: Remote on-line courses mirror classroom training. Participants watch the remote instructors’ live lecture online, participate in discussions, and perform lab exercises using remote desktop control.
- **Free online training**: You can use the self-paced Getting Started series of free, web-based courses to learn how to deploy F5 solutions to address your most common application delivery problems.

Engage F5 Support

F5 Support is designed to provide support for specific break-fix issues for customers with active support contracts. For more information about F5 scope of support, refer to Support Policies.

F5 Support resources

F5 Support resources are available 24 hours a day, seven days a week, and are distributed around the world in multiple support centers. Live support is provided by our professional network support engineers. Hours of availability may vary depending on the service contract with F5.

Contact numbers

Standard, Premium, and Premium Plus Support customers can open and manage cases by calling one of the contact numbers listed below.

North America

North America: 1-888-882-7535 or (206) 272-6500

Traffix® Support Only: 1-855-849-5673 or (206) 272-5774

Outside North America

Outside North America, Universal Toll-Free: +800 11 ASK 4 F5 or (800 11275 435)

Additional contact numbers by country

Australia: 1800 784 977

China: 010 5923 4123

Egypt: 0800-000-0537

Greece: 00-800-11275435

Hong Kong: 001-800-11275435

India: 000-800-650-1448; 000-800-650-0356 (Bharti Air users)

Indonesia: 001-803-657-904

Israel: 972-37630516

Japan: 81-3-5114-3260 or 0066-33-812670

Malaysia: 1-800-814994

New Zealand: 0800-44-9151

Philippines: 1-800-1-114-2564

Saudi Arabia: 800-844-7835

Singapore: 6411-1800

South Africa: 080-09-88889

South Korea: 002-800-11275435

Taiwan: 00-800-11275435

Thailand: 001-800-12-0666763

United Arab Emirates: 8000-3570-2437

United Kingdom: 44-(0)8707-744-655

Vietnam: 120-11585

Open a support case

F5 provides several resources to help find solutions to problems. Before opening a support case with F5 technical support, check to see if the issue you are encountering is already documented.

The following is a list of resources to consult before opening a support case with F5:

- Deployment guides and white papers provide information about specific deployment configurations.
- AskF5 provides many articles including known issues, how-to guides, security issues, release notes, and general information about products. Many of the issues customers encounter are already documented on this site.
- BIG-IP iHealth enables customers to upload QKView files in order to verify operation of any BIG-IP system.

Gather information to open a support case

If your issue cannot be solved using the resources listed, and you need to open a support case, you must first gather several pieces of important information about your issue. Providing full and accurate information helps speed the path to resolution. The required information for the majority of situations is summarized below:

- The serial number or base registration key of the specific BIG-IP system requiring support. For more information, refer to AskF5 article: K917: Finding the serial number or registration key of your F5 device.
- A full description of the issue. A clear problem statement is the best tool in helping to troubleshoot issues. Your description should include as much of the following information as you can provide.
- Occurrences and changes: The date and times of initial and subsequent recurrences. Did this issue arise at implementation or later? Were there any changes or updates made to the BIG-IP system prior to the issue arising? If so, what were they?
- Symptoms: Ensuring your list of symptoms is as detailed as possible gives more information for support personnel to correlate with.
- Scope of the problem: Note whether the problem is system-wide or limited to a particular configuration feature, service, or element (such as VLAN, interface, application service, virtual server, pool, and so on).
- BIG-IP component: The feature, configuration element, or service being used when the problem occurred (for example: portal access, network access, authentication services, VDI, Exchange).

- Steps to reproduce: The steps to reproduce the problem as accurately and in as much detail as possible. Include expected behavior (what should happen) as well as actual behavior (what does happen).
- Errors: Complete text of any error messages produced.
- Environment: Current usage of the system. (Is this unit in production? If so, is there currently a workaround in place?)
- Browsers: Types and versions, if applicable.
- Changes: System changes made immediately prior to the problem's first occurrence. This may include upgrades, hardware changes, network maintenance, and so on. Have any changes been made to resolve the problem? If so, what were they?
- Issue Severity: A description of the impact the issue is having on your site or case severity
  - Severity 1: Software or hardware conditions on your F5 device are preventing the execution of critical business activities. The device does not power up or is not passing traffic.
  - Severity 2: Software or hardware conditions on your F5 device are preventing or significantly impairing high-level commerce or business activities.
  - Severity 3: Software or hardware conditions on your F5 device are creating degradation of service or functionality in normal business or commerce activities.
  - Severity 4: Questions regarding configurations ("how to"), troubleshooting non-critical issues, or requests for product functionality that are not part of the current product feature set.
- Contact and availability information including alternate contacts authorized to work on the problem with F5 Support. When there are more personnel available to work with F5 Support, the resolution of your issue may be expedited.
- Remote access information, if possible.
- A QKView file obtained while problem symptoms are manifesting. A QKView of the system before the occurrence is also useful. F5 recommends archiving QKView files regularly. For more information, refer to **BIG-IP iHealth** in the ***TMOS Operations Guide***.

Note For information about how to locate F5 product manuals, refer to AskF5 article: K12453464: Finding product documentation on AskF5.

- Product-specific information: Software versions and types of equipment in use.
- Platform and system. Version and provisioned software modules of the affected system.

To locate platform and system information using tmsh at the command line

- Type the following command:

```
tmsh show /sys hardware
```

Output appears similar to the following example:

```
<SNIP some of the output>
Platform
Name  BIG-IP 3900
BIOS Revision  F5 Platform: C106 OBJ-0314-03 BIOS (build: 010) Date: 02/15/12
Base MAC       00:01:d7:be:bf:80
 System Information
Type                  C106
Chassis Serial        f5-jspv-lzxw
Level 200/400 Part    200-0322-02 REV C
Switchboard Serial
Switchboard Part Revision
Host Board Serial
Host Board Part Revision
```

To copy software version and build number information at the command line

1. Type the following command:

```
cat /VERSION
```

   Output appears similar to the following example:

```
Product: BIG-IP
Version: 11.6.0
Build: 0.0.401
Sequence: 11.6.0.0.0.401.0
BaseBuild: 0.0.401
Edition: Final
Date: Mon Aug 11 21:08:03 PDT 2014
Built: 140811210803
Changelist: 1255500
JobID: 386543
```

2. Highlight and copy the output information and include it with your support case.

To copy provisioned module information at the command line

1. Type the following command:

```
tmsh list /sys provision
```

Output appears similar to the following example:

```
sys provision afm { }
sys provision am { }
sys provision apm {
level nominal
}
sys provision asm { }
sys provision avr { }
sys provision fps { }
sys provision gtm { }
sys provision lc { }
sys provision ltm {
level minimum
}
sys provision pem { }
sys provision swg { }
```

2. Highlight and copy the output information and include it with your support case.

Open a support case

If you cannot find the answer to your problem using the resources listed above, you can open a support case online, using F5 Support (f5.com/support).

Before you open a support case, you need to log in to F5. If you do not have an F5 login, you'll need to register for one.

To register for support access

1. Navigate to login.f5.com.
2. Click **Register for an F5 Support Account**.
3. Enter your email address.

4. Enter your contact information. If you have a support contract, click **I have a support contract and need access to MySupport**.
5. Enter your serial number or registration key in the **Serial Number or Registration Key (optional)** field.

   After you've submitted your information, your service contract is reviewed. If your information is accurate you receive an email from MySupport, and you can use this to open your case.

Send information to Support

After you have the information listed in Gather information to open a support case, transfer it to F5 technical support following the steps in Share diagnostic files with F5 technical support. For more information, refer to AskF5 article: K2486: Providing files to F5 Technical Support.

Share diagnostic files with F5 technical support

You can provide files to F5 Support using BIG-IP iHealth or Support Files, the F5 file transfer tool. Support Files complies with global data protection standards to safeguard the data you send.

Prerequisites

Two categories of customers provide files to F5 Technical Support:

- Permanent account holders—Customers who have an F5 support account, including a user email and password for the AskF5 website
- Temporary users—Associates of permanent account holders who assist them with uploading or downloading files for an F5 service request

Obtaining a Support Files tool user name and password for a temporary user

If you are a permanent account holder who wants a temporary user to upload or download files for one of your service requests, you must provide them with a user name and password for logging in to the Support Files website.

The user name is the case ID from your service request, and the password is in the activity notes of your service request. You can also request the password directly from F5 Technical Support via email. Temporary passwords cannot be provided over the phone.

Note Temporary user credentials are only active for a specific service request.

To locate a password for a temporary user in the activity notes of your service request

1. Open the **Service Request Details** view.
2. In the **Activities** section, locate the temporary password which should appear in the following format:

```
Case file access password : XXyy=zZzZ123
```

Uploading QKView files to BIG-IP iHealth

BIG-IP iHealth allows you to quickly diagnose the health and proper operation of your BIG-IP system, and provides a convenient location for you to send diagnostic data for case resolution with F5 Technical Support.

If you are running BIG-IP 10.x or later and need to provide a QKView file to F5, the preferred way to do so is to upload the file to the BIG-IP iHealth website. For more information, refer to K12878: Generating diagnostic data using the qkview utility.

Uploading and downloading files using support files

Accessing support files

To access Support Files as a permanent account holder

1. In the upper-right corner of the **AskF5** home page, click **My Support**.

   If you have not already logged in, you are prompted to do so.

2. Under **Service Requests**, click the service request for which you want to upload or download files.

3. On the right side of the **Service Request Details** section, click **Manage Attachments**.

   The F5 Support Files site opens in a new window.

To access Support Files as a temporary user

- Navigate to the Support Files website and log in using the user name and password provided by a permanent account holder.

Uploading files using a web browser

1. On the **Home** folder page, click the folder for the service request you want.
2. Click **Incoming (upload to F5)**.
3. In the upper-right corner of the **Upload to F5** page, click the **Upload files** icon.
4. In the **Upload files** dialog box, click **Browse**.
5. In the **Open** dialog box, navigate to each file you want to upload, point at it with your cursor, select the check box that displays, and when you are finished selecting files, click **Open**.

   Either a success or failure notification displays. When an upload fails, close the notification and try

again.

6. In the **Upload Files** dialog box, click **Done**.

Note You cannot see the files in the folder after uploading them.

Downloading files using a web browser

1. On the **Home** folder page, click the folder for the service request you want.
2. Click **Outgoing (download from F5)**.
3. Select the check box next to the files you want to download.
4. In the upper-right corner of the **Download from F5** page, click the **Download selected items** icon.
5. Depending on which browser you have, you may be prompted to save your files to the location of your choice, or they may simply download to your Downloads folder.

Uploading files using SFTP

For permanent account holders, your user name and password are the same as your AskF5 credentials. For temporary users, your user name is the service request number, and your password must be the correct passphrase for that service request.

Important Support Files does not support the Secure Copy (SCP) protocol.

Note Support Files supports the SFTP protocol, but only a subset of features provided by many SFTP clients. The SFTP server does not support or allow setting file ownership or permissions, updating timestamps, or creating symlinks.

Note The supportfiles.f5.com RSA server key MD5 fingerprint is **MD5: 04:a6:4b:9b:d4:eb:48:97:15:e6:7f:90:64:bf:35:96** and the SHA256 fingerprint is **SHA256:stQCq50hEwDPfRMeRf/Ya9dXcm1KCdx5I5lIOwODgNU**.

1. From the F5 device, SFTP to the **supportfiles.f5.com** site using the following syntax:

```
sftp user@host
```

For example:

```
sftp c123456@supportfiles.f5.com
```

or

```
sftp 1-12345678@supportfiles.f5.com
```

Note On the first attempt to connect, you must accept the host key. You should compare that output with the fingerprints listed in this article.

2. When prompted for the password, enter the email address of the user associated with the case.
3. To upload the requested files, use the following command syntax for each file:

```
put <name _ of _ file> <SR _ number>/INCOMING/
```

For example:

```
put mybigip.conf C123456/INCOMING/
```

4. To exit the **SFTP** utility when all files have been uploaded, type the following command:

```
quit
```

For more information, refer to the manual or man pages for the **SFTP** utility by typing **man sftp** at the command line.

You may also use external SFTP applications to upload files if they are on a workstation or other system.

Downloading files using SFTP

To download a file from Support Files, you must use the exact file name and location (path) provided by your F5 Technical Support representative.

Important Support Files does not support the SCP protocol.

Note Support Files supports the SFTP protocol, but only a subset of features provided by many SFTP clients. The SFTP server does not support or allow setting file ownership or permissions, updating timestamps, or creating symlinks.

1. From the F5 device, SFTP to the **supportfiles.f5.com** site using the following syntax:

```
sftp user@host
```

For example:

```
sftp C123456@supportfiles.f5.com
```

or

```
sftp 1-12345678@supportfiles.f5.com
```

2. When prompted for the password, enter the password.
3. To list the files available for download, type the following command:

```
ls C123456/OUTGOING
```

4. To download the requested files, use the following command syntax for each file:

Note The **<full_path_to_file>** section indicates the full path to the file.

```
get <full _ path _ to _ file>
```

For example:

```
get C123456/OUTGOING/ Hotfix-BIGIP-12.1.0.0.40.1434-ENG.iso
```

5. To exit the SFTP utility when all files have been downloaded, type the following command:

```
quit
```

You may also use external SFTP applications to download files if the files are on a workstation or on another system.

Collecting BIG-IP ASM Data

Collecting BIG-IP ASM data

This section discusses BIG-IP ASM data you may need to gather for your support case, as well as tools to collect that data.

F5 Technical Support is the single point of contact for security vulnerability questions. For more information, refer to AskF5 article: K4602: Overview of the F5 security vulnerability response policy.

Using asmqkview utility

F5 Technical Support requires **asmqkview** output in all BIG-IP ASM related cases.

The **asmqkview** script automatically collects configuration and diagnostic information from BIG-IP ASM systems. The output includes all data collected in QKView files and other data for BIG-IP ASM systems in a single file, which you can then provide to F5 Technical Support to aid in troubleshooting.

To run asmqkview in BIG-IP 11.6

- Use the following command syntax:

```
asmqkview -s0 --add-proxy-log output: /var/tmp/<hostname>.qkview
```

To run asmqkview in BIG-IP 12.0

- Use the following command syntax:

```
qkview -s0 -o asm-request-log output: /var/tmp/<hostname>.qkview
```

For more information, refer to AskF5 article: K6824: Overview of the asmqkview script.

Exporting security policy

You may be asked to provide an exported version of your security policy for F5 Technical Support to review.

You can export a security policy and save it in a file. The exported security policy can be used as backup, or you can import it onto another system.

To export a security policy using the Configuration utility in BIG-IP 13.x

1. Navigate to **Security** > **Application Security** > **Security Policies** > **Policies List**.
2. Select the policy you want to export.
3. Click **Export** and then click the appropriate export format options.

- To save the security policy as an XML file, for **Export policy format,** click **XML Format**. To reduce

the size of the file, for **Compact format,** click **Enabled**.

- To save the security policy as a policy archive file (PLC file), click **Binary Format**.

4. Click **Export Policy**.

To export a security policy using the Configuration utility in BIG-IP 11.6 - 12.x

1. Navigate to **Security > Application Security : Security Policies**.

   The **Active Policies** page opens.

2. In the **Active Security Policies** list, click the security policy that you want to export, then click **Export**.

   Note You can also export security policies from the **Inactive Policies** list using the same method.

   The **Select Export Method** pop-up dialog opens.

3. Click an export method.

- To save the security policy as an XML file, click **Export security policy in XML format**. To reduce the size of the file, select the **Compact format** check box.
- To save the security policy as a policy archive file (PLC file), click **Binary export of the security policy**.
- If the security policy integrates with a vulnerability assessment tool, select the **Include Vulnerability Assessment configuration and data** check box.

4. Click **Export**.

   The system exports the security policy in the format you specified.

For more information, refer to **Importing and Exporting Security Policies** in ***BIG-IP ASM: Implementations*** for your system version.

Note For information about how to locate F5 product manuals, refer to the Ask F5 article: K12453464: Finding product documentation on AskF5.

Capturing HTTP request/response

You may be asked to provide an overview of the HTTP request/response from the client side. This can help F5 Technical Support understand the problem and provide a baseline for reproduction of the issue in the lab.

F5 recommends using one of the following tools to collect the data:

- HttpWatch (save as .hwl)

- Fiddler (save as .saz)
- Chrome (save as cURL)

Important HttpWatch and Fiddler are commercial software and a license is required.

In the following example, a Chrome browser is used to save a single HTTP request into a **cURL** command and the entire transaction into an HAR archive file.

1. Navigate to **View > Developer > Developer tools > Network** tab.
2. Refresh the page.
3. Locate the HTTP request in question.
4. Copy as cURL.
5. Send the content in clipboard to F5.
6. Save as .har file.
7. Navigate to **View > Developer > Developer tools > Network tab.**
8. Refresh the page.
9. Save as .har file with content.
10. Send the output to F5 Technical Support.

Important Information leakage may occur when capturing sensitive application transactions on production traffic.

Using tcpdump utility

You can use **tcpdump** to capture client-side and server-side traffic of HTTP request/response to help F5 reproduce and troubleshoot your issue.

To capture traffic using tcpdump

- Use the following command syntax:

```
tcpdump -i 0.0:nnn -s0 -w /var/tmp/asm _ traffic _ catpure.cap host

<Virtual server IP address> or host <Pool member IP address or Host more
pool member IP address> -vvv
```

Important Information leakage may occur when capturing sensitive application transactions on production traffic.

For more information, refer to the following AskF5 articles:

- K411: Overview of packet tracing with the tcpdump utility

- K6546: Recommended methods and limitations for running tcpdump on a BIG-IP system
- K7227: Considerations when using the tcpdump utility with tagged VLAN traffic
- K3637: Capturing internal TMM information with tcpdump
- K2289: Using advanced tcpdump filters

Saving UCS file

You may be asked to supply a user configuration set (UCS) file so that F5 Technical Support can load your configuration to reproduce your issue in the lab.

To save a UCS using tmsh at the command line

- Type the following command:

```
tmsh save sys ucs asm_support
```

Output is logged to the **/var/local/ucs/asm_support.ucs** file.

For more information, refer to AskF5 article: K4423: Overview of UCS archives.

Important A typical UCF archive contains user accounts, passwords, critical system files, and SSL private keys. However, you can explicitly exclude SSL private keys from a UCS archive during the backup process. If your UCS archive contains SSL private keys, you must store backup UCS archives in an environment that is as secure as where you store your private keys.

Appendix

BIG-IP AAM dynamic caching integration

You can deploy both the BIG-IP ASM system and F5® BIG-IP® Application Acceleration Manager® (AAM) for a single web application.

Caching considerations

BIG-IP ASM system overrides the eligibility for caching entities that it inspects in the following instances:

- Policy Builder is running and the policy is unstable.
- A violation or learning suggestion is generated.
- Frame cookies are either created or removed due to an issue such as dynamic parameter extraction.
- The system processes a request triggering policy tightening suggestions (file types, URLs, or parameters).
- The system processes requests for pages that have extractions configured (including pages with dynamic sessions).
- The system processes requests for pages that are within a Flow Access (including URLs embedded in the cookie and logout URLs).
- Applying a BIG-IP ASM security policy invalidates the BIG-IP AAM cache.

The Policy Builder periodically applies a policy that invalidates the BIG-IP AAM cache. Use BIG-IP AAM only after automatic learning is complete and disabled. When protecting websites where performance improvements are critical, you can also choose a policy type (such as the Rapid Deployment Policy template) that does not require continuous learning.

For more information, refer to AskF5 article: K16565: Configuring a Web Acceleration profile for use with a BIG-IP ASM-enabled virtual server (11.4.0 and later).

Integrated BIG-IP APM session tracking and event logging

BIG-IP ASM system integrates with BIG-IP APM to provide combined session tracking and event logging. This allows you to configure your security policy to capture user names from the BIG-IP APM login process. It also allows you to see application security violations that are associated with a user session.

For more information, refer to **Tracking User Sessions** in ***BIG-IP ASM: Implementations*** for your system version.

Note For information about how to locate F5 product manuals, refer to the Ask F5 article: K12453464: Finding product documentation on AskF5.

Using multiple decoding passes with evasion technique

The **Multiple Decoding: Passes** option for evasion techniques allows you to configure the number of encoding passes that the system should use to decode multiple encoded characters. The decoding passes are performed on URI and parameter input.

You can locate the **Multiple Decoding** option by navigating to **Policy > Blocking > Evasion Techniques**.

Note As part of the normalization process that BIG-IP ASM uses, Multiple Decoding is performed whether or not the **Blocking** properties are enabled. You can enable **Blocking** properties to **Alarm** or **Block** an evasion technique when one is detected.

You can configure the **Multiple Decoding** option to perform two to five passes.

The number of specified decoding passes determines how the system responds. It either flags the results as an evasion technique, triggers an attack signature, or both.

When the **Blocking** properties are not set for **Multiple Decoding**, the ability to identify signatures is reduced to the number of encoded passes.

The following table shows the actions the BIG-IP ASM system takes on each of the decoding passes:

Table A.1 Decoding pass actions

Decoding Pass(es)	Behavior
2 to 3	Triggers an evasion technique but doesn't trigger an attack signature if one exists
4	Triggers an evasion technique and an attack signature if one exists
5	Triggers an attack signature if one exists, but does not trigger an evasion technique

For example, **Multiple Decoding** set to perform three decoding passes converts "a%252fb" to "a/b" after the second pass. On the third decoding pass, the system responds with the appropriate alarm or block action you have configured.

- On first pass, the system looks at the hexadecimal "%252f."

  In ASCII, "25" translates to "**%**," so the first pass decoding result is "%2f."

- On second pass, the system decodes "2f" to "/" in ASCII

  The two decoding passes of a "%252fb" result in "a/b."

- On third pass, the system takes action. The encoding attempt of the characters "a/b" results in action specified by the **Learn, Alarm**, and **Block** settings of the **Evasion Technique Detected** category on the Blocking Policy page.

For example, take the following string:

```
/nameandcolor.asp? aaa=%2525253cSCRIPT%2525253e
```

Using the previous string as an example within a URI, the string translates from hexadecimal to ASCII is as follows:

Table A.2 Hexadecimal to ASCII translation

Hexadecimal	ASCII
25	%
3C	<
3e	>

The following table shows the decoding results for this URI string, based on the number of passes configured:

Table A.3 Decoding passes and results

Passes	Result
2 to 3	First decoding pass results in “25253c” and “25253e.” Second decoding pass results in “253c” and “253e.” Third decoding pass results in “<” and “>.” System triggers an evasion technique and responds with the appropriate action. The number of configured passes is not high enough for the system to check **/nameandcolor.asp?aaa=<SCRIPT>** against the list of known attack signatures, which could possibly trigger an attack signature.
4	Fourth decoding pass results in the system checking the fully decoded results against the known attack signatures. The system triggers both an evasion technique and an attack signature.
5	Fifth decoding pass results in the system triggering on the detected matching attack signature. No evasion technique violation is reported.

Note A higher number of decoding passes impacts system performance. F5 recommends that you set **Blocking** properties when you use the lower settings (two to three passes).

Resource Materials

BIG-IP ASM product documentation

BIG-IP ASM product documentation provides step-by-step instructions for how to create a security policy and add available protections.

- BIG-IP ASM documentation for 11.6.
- BIG-IP ASM documentation for 12.0.
- BIG IP ASM documentation for 13.0
- BIG-IP ASM documentation for 14.0

AskF5 articles

The following articles contain information you may find useful.

Table A.4 Useful AskF5 Articles

For information about	See this article
Opening a support case	K6825: Information required when opening a support case for BIG-IP ASM
Updating BIG-IP ASM attack signatures	K8217: Updating the BIG-IP ASM attack signatures
Understanding BIG-IP ASM cookies	K6850: Overview of BIG-IP ASM cookies
Using local traffic policies	K15085: Overview of the Local Traffic Policies feature (11.4.0 - 12.0.0)
Configuring the language encoding	K6335 Overview of encoding language settings for BIG-IP ASM
Sending SNMP traps to communicate a blocked request and violation	K7738 Configuring the BIG-IP ASM system to send SNMP traps to communicate a blocked request and request violation
Redirecting response page to an external server	K7825: Redirecting a blocking response support ID to an external error page
Working with evasion technique violations	K7929: Working with Evasion technique detected violations
Using the wildcard entity in the BIG-IP ASM system	K8623: Using the wildcard entity in BIG-IP ASM
Understanding the BIG-IP ASM system and caching	K14880: BIG-IP ASM may prevent object caching
Understanding the BIG-IP ASM system cookies	K6850: Overview of BIG-IP ASM cookies
BIG-IP ASM daemons	K14020: BIG-IP ASM daemons (11.x - 14.x)

Legal Notices

Trademarks

AAM, Access Policy Manager, Advanced Client Authentication, Advanced Firewall Manager, Advanced Routing, AFM, APM, Application Acceleration Manager, Application Security Manager, Applications without Constraints, ARX, AskF5, ASM, BIG-IP, BIG-IP EDGE GATEWAY, BIG-IQ, BIG-IP iControl, Cloud Extender, CloudFucious, Cloud Manager, Clustered Multiprocessing, CMP, COHESION, Data Manager, DDoS Frontline, DDoS Hybrid Defender, DDoS SWAT, Defense.net, DevCentral, DevCentral [DESIGN], DNS Express, DSC, DSI, Edge Client, Edge Gateway, EDGE MOBILE, EDGE MOBILITY, EdgePortal, ELEVATE, EM, Enterprise Manager, ENGAGE, F5, F5 Agility, F5 iApps, F5[DESIGN], F5 Certified [DESIGN], F5 iControl, F5 LINK CONTROLLER, F5 Networks, F5SalesXchange [DESIGN], F5Synthesis, f5Synthesis, F5Synthesis[DESIGN], F5 TechXchange [DESIGN], F5 TMOS, Fast Application Proxy, Fast Cache, FCINCO, Global Traffic Manager, GTM, GUARDIAN, Herculon, iApps, IBR, Intelligent Browser Referencing, Intelligent Compression, IPv6 Gateway, iCall, iControl, iHealth, iQuery, iRules, iRules OnDemand, iSeries, iSession, L7 RateShaping, LC, Link Controller, Local Traffic Manager, LTM, LineRate Operating System, LineRate Point, LineRate Precision, LineRate Systems [DESIGN], LROS, LTM, Message Security Manager, MSM, OneConnect, Packet Velocity, PEM, Policy Enforcement Manager, Protocol Security Manager, PSM, Real Traffic Policy Builder, Ready Defense, SalesXchange, ScaleN, Signalling Delivery Controller, Silverline, Silverline Threat Intelligence, SDC, SSL Acceleration, SSL Everywhere, SSL Orchestrator, SDAC (except in Japan), StrongBox, SuperVIP, SYN Check, SYNTHESIS, TCP Express, TDR, TechXchange, TMOS, TotALL, Traffic Management Operating System, Traffix Systems, Traffix Systems (DESIGN), Transparent Data Reduction, UNITY, VAULT, vCMP, VE F5 [DESIGN], Versafe, Versafe [DESIGN], VIPRION, Virtual Clustered Multiprocessing, WAF Express, WebSafe, We Make Apps Go [DESIGN], We Make Apps GO, and ZoneRunner, are trademarks or service marks of F5 Networks, Inc., in the U.S. and other countries, and may not be used without express written consent.

All other product and company names herein may be trademarks of their respective owners.

Patents

This product may be protected by one or more patents. See the F5 Patents page (https://www.f5.com/about/guidelines-policies/patents).

Notice

THE SOFTWARE, SCRIPTING, AND COMMAND EXAMPLES ARE PROVIDED “AS IS,” WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE, SCRIPTING AND COMMAND EXAMPLES, OR THE USE OR OTHER DEALINGS WITH THE SOFTWARE, SCRIPTING, AND COMMAND EXAMPLES.

Publication Date

This document was published in December 2018.

Copyright

Change list

Date	Chapter/Section	Change	Reason
June 2016	Deployment Examples	Add chapter	New content
July 2016	BIG-IP ASM Event Logging	Add note regarding support for multiple remote logging profiles	12.1 update
September 2016	Policy Tuning and Enhancement	Move Table 4.2 to correct location Add List of Tables	Error User request
June 2017	BIG-IP ASM Event Logging Policy Tuning and Enhancement Regulatory Compliance	Add updates for 13.0 release	13.0 release
	Deployment Examples BIG-IP ASM API/Web Services Protection	Remove chapters covering deployment	Content outside scope of operations Content covered in other guides. For more information, refer to Phase 1: Create and deploy policy.
August 2017	Optimizing the Support Experience	Remove reference to F5 Dropbox	Dropbox use discontinued
September 2017	Optimize the Support Experience	Add section on providing files to F5 Technical Support using F5 Support Files	New
January 2018	Regulatory Compliance	Update Table 4.1	Update
April 2018	Common Management Tasks	Update for BIG-IP 13.1	BIG-IP 13.1 release
	Troubleshooting BIG-IP ASM		
December 2018	Policy and Tuning Enhancement	Update for BIG-IP 14.0	BIG-IP 14.0 release