BIG-IP® Application Security Manager™: Implementations

Version 11.2.1

IT agility. Your way.

Table of Contents

Legal Notices....7
Acknowledgments....9

Chapter 1:
Automatically Synchronizing Application Security Configurations....13
Overview: Automatically synchronizing ASM systems....14
About device management and synchronizing application security configurations....14
Considerations for application security synchronization....15
Performing basic network configuration for synchronization....15
Specifying an IP address for config sync....15
Establishing device trust....16
Creating a Sync-Failover device group....17
Syncing the BIG-IP configuration to the device group....17
Specifying IP addresses for failover....18
Creating a Sync-Only device group....18
Enabling ASM synchronization on a Sync-Only device group....19
Synchronizing an ASM-enabled device group....19
Implementation result....20

Chapter 2: Manually Synchronizing Application Security Configurations....21
Overview: Manually synchronizing ASM systems....22
About device management and synchronizing application security configurations....22
Considerations for application security synchronization....23
Performing basic network configuration for synchronization....23
Specifying an IP address for config sync....23
Establishing device trust....24
Creating a Sync-Failover device group....25
Syncing the BIG-IP configuration to the device group....25
Specifying IP addresses for failover....26
Enabling ASM synchronization on a device group....26
Synchronizing an ASM-enabled device group....27
Implementation result....27

Chapter 3:
Synchronizing Application Security Configurations Across LANs....29
Overview: Synchronizing ASM systems across LANs....30
About device management and synchronizing application security configurations....31

Considerations for application security synchronization....31
Performing basic network configuration for synchronization....31
Specifying an IP address for config sync....32
Establishing device trust....32
Creating a Sync-Failover device group....33
Syncing the BIG-IP configuration to the device group....33
Specifying IP addresses for failover....34
Creating a Sync-Only device group....34
Enabling ASM synchronization on a Sync-Only device group....35
Synchronizing an ASM-enabled device group....35
Implementation result....36

Chapter 4: Setting Up IP Address Intelligence Blocking....37
Overview: Setting up IP address intelligence blocking....38
Enabling IP address intelligence....38
Setting up IP address intelligence blocking....39
Reviewing IP address intelligence statistics....40
Creating an iRule to log IP address intelligence information....40
Creating an iRule to reject requests with questionable IP addresses....41
IP address intelligence categories....42

Chapter 5: Managing IP Address Exceptions....43
Overview: Managing IP address exceptions....44
Creating IP address exceptions....44
Deleting IP address exceptions....45
Updating IP address exceptions....45

Chapter 6: Enforcing Application Use at Specific Geolocations....47
Overview: Enforcing application use in certain geolocations....48
Enforcing application use in certain geolocations....48
Setting up geolocation enforcement from a request....49

Chapter 7: Configuring Application Security Session Tracking....51
Overview: Tracking application security sessions using login pages....52
Creating login pages....52
Enforcing login pages....53
Setting up session tracking....54
Monitoring user and session information....55

Chapter 8: Tracking Application Security Sessions with APM....57
Overview: Tracking application security sessions using APM....58

Prerequisites for setting up session tracking with APM....58
Creating a VLAN....58
Creating a self IP address for a VLAN....59
Creating a local traffic pool for application security....59
Creating an HTTP class....60
Creating a virtual server to manage HTTPS traffic....60
Creating a security policy automatically....61
Creating an access profile....63
Configuring an access policy....65
Adding the access profile to the virtual server....65
Setting up ASM session tracking with APM....66
Monitoring user and session information....67

Chapter 9:
Automatically Creating Security Policies for AJAX Applications....69
Application security for applications that use AJAX....70
Overview: Creating a security policy for applications that use AJAX....70
Creating a security policy automatically....70
Reviewing security policy status....72
Implementation result....73

Chapter 10: Adding JSON Support to an Existing Security Policy....75
Overview: Adding JSON support to existing security policies....76
Creating a JSON profile....76
Associating a JSON profile with a URL....77
Associating a JSON profile with a parameter....78
Implementation result....78

Chapter 11:
Adding AJAX Blocking Response Behavior to a Security Policy....79
Overview: Adding AJAX blocking and login response behavior....80
Configuring the blocking response for AJAX applications....80

Table of Contents

Legal Notices

Publication Date

This document was published on August 31, 2012.

Publication Number

MAN-0358-03

Copyright

Trademarks

Patents

This product may be protected by U.S. Patent 6,311,278. This list is believed to be current as of August 31, 2012.

Export Regulation Notice

This product may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this product from the United States.

RF Interference Warning

This is a Class A product. In a domestic environment this product may cause radio interference, in which case the user may be required to take adequate measures.

FCC Compliance

This equipment has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This unit generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user, at his own expense, will be required to take whatever measures may be required to correct the interference.

Any modifications to this device, unless expressly approved by the manufacturer, can void the user's authority to operate this equipment under part 15 of the FCC rules.

Canadian Regulatory Compliance

This Class A digital apparatus complies with Canadian ICES-003.

Standards Compliance

This product conforms to the IEC, European Union, ANSI/UL and Canadian CSA standards applicable to Information Technology products at the time of manufacture.

Acknowledgments

This product includes software developed by Bill Paul.

This product includes software developed by Jonathan Stone.

This product includes software developed by Manuel Bouyer.

This product includes software developed by Paul Richards.

This product includes software developed by the NetBSD Foundation, Inc. and its contributors.

This product includes software developed by the Politecnico di Torino, and its contributors.

This product includes software developed by the Swedish Institute of Computer Science and its contributors.

This product includes software developed by the University of California, Berkeley and its contributors.

This product includes software developed by the Computer Systems Engineering Group at the Lawrence Berkeley Laboratory.

This product includes software developed by Christopher G. Demetriou for the NetBSD Project.

This product includes software developed by Adam Glass.

This product includes software developed by Christian E. Hopps.

This product includes software developed by Dean Huxley.

This product includes software developed by John Kohl.

This product includes software developed by Paul Kranenburg.

This product includes software developed by Terrence R. Lambert.

This product includes software developed by Philip A. Nelson.

This product includes software developed by Herb Peyerl.

This product includes software developed by Jochen Pohl for the NetBSD Project.

This product includes software developed by Chris Provenzano.

This product includes software developed by Theo de Raadt.

This product includes software developed by David Muir Sharnoff.

This product includes software developed by SigmaSoft, Th. Lockert.

This product includes software developed for the NetBSD Project by Jason R. Thorpe.

This product includes software developed by Jason R. Thorpe for And Communications, http://www.and.com.

This product includes software developed for the NetBSD Project by Frank Van der Linden.

This product includes software developed for the NetBSD Project by John M. Vinopal.

This product includes software developed by Christos Zoulas.

This product includes software developed by the University of Vermont and State Agricultural College and Garrett A. Wollman.

This product includes software developed by Balazs Scheidler (bazsi@balabit.hu), which is protected under the GNU Public License.

This product includes software developed by Niels Mueller (nisse@lysator.liu.se), which is protected under the GNU Public License.

In the following statement, *This software* refers to the Mitsumi CD-ROM driver: This software was developed by Holger Veit and Brian Moore for use with 386BSD and similar operating systems. *Similar operating systems* includes mainly non-profit oriented systems for research and education, including but not restricted to NetBSD, FreeBSD, Mach (by CMU).

This product includes software developed by the Apache Group for use in the Apache HTTP server project (http://www.apache.org/).

This product includes software licensed from Richard H. Porter under the GNU Library General Public License (© 1998, Red Hat Software), www.gnu.org/copyleft/lgpl.html.

This product includes the standard version of Perl software licensed under the Perl Artistic License (© 1997, 1998 Tom Christiansen and Nathan Torkington). All rights reserved. You may find the most current standard version of Perl at http://www.perl.com.

This product includes software developed by Jared Minch.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (http://www.openssl.org/).

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com).

This product contains software based on oprofile, which is protected under the GNU Public License.

This product includes RRDtool software developed by Tobi Oetiker (http://www.rrdtool.com/index.html) and licensed under the GNU General Public License.

This product contains software licensed from Dr. Brian Gladman under the GNU General Public License (GPL).

This product includes software developed by the Apache Software Foundation (http://www.apache.org/).

This product includes Hypersonic SQL.

This product contains software developed by the Regents of the University of California, Sun Microsystems, Inc., Scriptics Corporation, and others.

This product includes software developed by the Internet Software Consortium.

This product includes software developed by Nominum, Inc. (http://www.nominum.com).

This product contains software developed by Broadcom Corporation, which is protected under the GNU Public License.

This product contains software developed by MaxMind LLC, and is protected under the GNU Lesser General Public License, as published by the Free Software Foundation.

This product includes the Zend Engine, freely available at http://www.zend.com.

This product contains software developed by NuSphere Corporation, which is protected under the GNU Lesser General Public License.

This product contains software developed by Erik Arvidsson and Emil A Eklund.

This product contains software developed by Aditus Consulting.

This product contains software developed by Dynarch.com, which is protected under the GNU Lesser General Public License, version 2.1 or above.

This product contains software developed by InfoSoft Global (P) Limited.

This product includes software written by Steffen Beyer and licensed under the Perl Artistic License and the GPL.

This product includes software written by Makamaka Hannyaharamitu ©2007-2008.

Acknowledgments

Chapter 1

Automatically Synchronizing Application Security Configurations

Topics:

- *Overview: Automatically synchronizing ASM systems*
- *Implementation result*

Overview: Automatically synchronizing ASM systems

This implementation describes how to set up multiple BIG-IP® systems running Application Security Manager™ (ASM) so that they automatically synchronize their security policies and ASM™ configurations. In addition, the ASM devices can fail over to one another if any of the devices goes offline. For synchronizing local traffic configuration data, you can manually synchronize that data as needed.

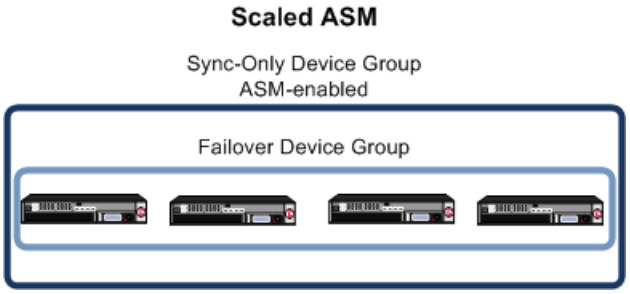

Figure 1: Automatically synchronizing ASM configuration data

In this case, multiple BIG-IP systems are all processing similar traffic for one or more web applications behind a router (or load balancer). All systems are running BIG-IP ASM™ and are in the local trust domain. You organize the systems into two device groups: one Sync-Failover device group for all systems (not ASM-enabled) and one Sync-Only device group with ASM-enabled for all of the systems. The ASM configurations and web applications are automatically duplicated on all of the systems. You can manually synchronize the BIG-IP configuration of the systems in the Sync-Failover device group.

Task summary

Performing basic network configuration for synchronization
Specifying an IP address for config sync
Establishing device trust
Creating a Sync-Failover device group
Syncing the BIG-IP configuration to the device group
Specifying IP addresses for failover
Creating a Sync-Only device group
Enabling ASM synchronization on a Sync-Only device group
Synchronizing an ASM-enabled device group

About device management and synchronizing application security configurations

You can use device management to set up several BIG-IP® systems running Application Security Manager™ (ASM) so that the systems synchronize their security policies and configurations, and fail over to one another if a system goes offline for any reason. By using application security synchronization, you can set up application security and create security policies on one system, and can propagate them to other systems in an application security device group. In BIG-IP ASM™, a *device group* is two or more BIG-IP devices using the same configuration and providing consistent security policy enforcement.

You can set up application security synchronization, for example, behind an Application Delivery Controller where multiple BIG-IP systems running Application Security Manager are deployed as members of a pool. The options and security policies on all of the systems stay in sync regardless of where you update them.

When you set up ASM™ synchronization, in addition to security policies, other settings such as custom attack signatures, logging profiles, SMTP configuration, anti-virus protection, system variables, and policy templates, are synchronized with all devices in the ASM-enabled device group.

Considerations for application security synchronization

When using device management with Application Security Manager™ (ASM™), you need to be aware of the following considerations that apply specifically to application security synchronization.

- A BIG-IP® system with Application Security Manager can be a member of only one ASM-enabled device group.
- All BIG-IP systems in a device group must be running the same version (including hot fix updates) of Application Security Manager (version 11.0 or later).
- The BIG-IP systems in the ASM-enabled device group synchronize application security configuration data and security policies, providing consistent enforcement on all the devices.
- Real Traffic Policy Builder® can run on only one system per security policy. For example, you can set up automatic security policy building on one system that is a member of an ASM-enabled device group, the policy is built on that system and then automatically updated on all of the systems in the device group.

Performing basic network configuration for synchronization

You need to perform basic networking configuration for each of the BIG-IP® systems whose Application Security Manager™ (ASM) configurations you want to synchronize.

1. Install the same BIG-IP system version (including any hot fixes) on each device.
2. Provision LTM® and ASM™ on each device (**System** > **Resource Provisioning**).
3. On each device, create one or more VLANs, depending on your networking configuration (**Network** > **VLANs**).
4. On each device, create a self IP (**Network** > **Self IPs**).

   When creating the self IP, set **Traffic Group** to **traffic-group-local-only (non-floating)**.
5. On each device, create a default gateway, if needed (**Network** > **Routes**).
6. On each device, configure DNS (**System** > **Configuration** > **Device** > **DNS**) and NTP (**System** > **Configuration** > **Device** > **NTP**) so they are set to the same time.
7. Verify connectivity between the devices (self IP address to self IP address). For example, use this command to ensure communications: `ping -I vlan_interface device_self_IP`
8. If your company requires special device certificates, install them on each device (**System** > **Device Certificates** and click **Import**).

The basic networking setup is complete for the BIG-IP ASM systems for which you want to share security policies and configurations.

Specifying an IP address for config sync

Before configuring the config sync address, verify that all devices in the device group are running the same version of BIG-IP® system software.

You perform this task to specify the IP address on the local device that other devices in the device group will use to synchronize their configuration objects to the local device.

Note: You must perform this task locally on each device in the device group.

1. Confirm that you are logged in to the actual device you want to configure.
2. On the Main tab, click **Device Management** > **Devices**.
   This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose ConfigSync.
5. For the **Local Address** setting, retain the displayed IP address or select another address from the list.
   F5 Networks recommends that you use the default value, which is the self IP address for VLAN `internal`. This address must be a non-floating self IP address and not a management IP address.
6. Click **Update**.

After performing this task, the other devices in the device group can sync their configurations to the local device.

Establishing device trust

Before you begin this task, verify that:

- Each BIG-IP® device that is to be part of the local trust domain has a device certificate installed on it.
- The local device is designated as a certificate signing authority.

You perform this task to establish trust among devices on one or more network segments. Devices that trust each other constitute the *local trust domain*. A device must be a member of the local trust domain prior to joining a device group.

By default, the BIG-IP software includes a local trust domain with one member, which is the local device. You can choose any one of the BIG-IP devices slated for a device group and log into that device to add other devices to the local trust domain. For example, devices `A`, `B`, and `C` each initially shows only itself as a member of the local trust domain. To configure the local trust domain to include all three devices, you can simply log into device `A` and add devices `B` and `C` to the local trust domain. Note that there is no need to repeat this process on devices `B` and `C`.

1. On the Main tab, click **Device Management** > **Device Trust**, and then either **Peer List** or **Subordinate List**.
2. Click **Add**.
3. Type an IP address, administrator user name, and administrator password for the remote BIG-IP® device.
   This IP address can be either a management IP address or a self IP address.
4. Click **Retrieve Device Information**.
5. Verify that the certificate of the remote device is correct.
6. Verify that the name of the remote device is correct.
7. Verify that the management IP address and name of the remote device are correct.
8. Click **Finished**.

The device you added is now a member of the local trust domain.

Repeat this task for each device that you want to add to the local trust domain.

Creating a Sync-Failover device group

This task establishes failover capability between two or more BIG-IP devices. If the active device in a Sync-Failover device group becomes unavailable, the configuration objects fail over to another member of the device group and traffic processing is unaffected. You perform this task on any one of the authority devices within the local trust domain.

Repeat this task for each Sync-Failover device group that you want to create for your network configuration.

1. On the Main tab, click **Device Management** > **Device Groups**.
2. On the Device Groups list screen, click **Create**.
   The New Device Group screen opens.
3. Type a name for the device group, select the device group type **Sync-Failover**, and type a description for the device group.
4. In the Configuration area of the screen, select a host name from the **Available** list for each BIG-IP device that you want to include in the device group, including the local device. Use the Move button to move the host name to the **Selected** list.

   The **Available** list shows any devices that are members of the device's local trust domain but not currently members of a Sync-Failover device group. A device can be a member of one Sync-Failover group only.
5. For the **Network Failover** setting:
   - Select the **Enabled** check box if you want device group members to handle failover communications by way of network connectivity.
   - Clear the **Enabled** check box if you want device group members to handle failover communications by way of serial cable (hard-wired) connectivity.

   Serial failover is not available for device groups with more than two members.
6. Click **Finished**.

You now have a Sync-Failover type of device group containing BIG-IP devices as members.

Syncing the BIG-IP configuration to the device group

Before you sync the configuration, verify that the devices targeted for config sync are members of a device group and that device trust has been established.

This task synchronizes the BIG-IP® configuration data from the local device to the devices in the device group. This synchronization ensures that devices in the device group operate properly. When synchronizing self IP addresses, the BIG-IP system synchronizes floating self IP addresses only.

Important: *You perform this task on either of the two devices, but not both.*

1. On the Main tab, click **Device Management** > **Overview**.
2. In the Device Groups area of the screen, in the Name column, select the name of the relevant device group.
   The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
3. In the Devices area of the screen, in the Sync Status column, select the device that shows a sync status of `Changes Pending`.
4. In the Sync Options area of the screen, select **Sync Device to Group**.

5. Click **Sync**.
   The BIG-IP system syncs the configuration data of the selected device in the Device area of the screen to the other members of the device group.

Except for non-floating self IP addresses, the entire set of BIG-IP configuration data is replicated on each device in the device group.

Specifying IP addresses for failover

This task specifies the local IP addresses that you want other devices in the device group to use for failover communications with the local device. You must perform this task locally on each device in the device group.

Note: *The failover addresses that you specify must belong to route domain* `0`.

1. Confirm that you are logged in to the actual device you want to configure.
2. On the Main tab, click **Device Management** > **Devices**.
   This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose Failover.
5. For the Failover Unicast Configuration settings, retain the displayed IP addresses.

   You can also click **Add** to specify additional IP addresses that the system can use for failover communications. F5 Networks recommends that you use the self IP address assigned to the `HA` VLAN.
6. If the BIG-IP® system is running on a VIPRION® platform, then for the **Use Failover Multicast Address** setting, select the **Enabled** check box.
7. If you enable **Use Failover Multicast Address**, either accept the default **Address** and **Port** values, or specify values appropriate for the device.

   If you revise the default **Address** and **Port** values, but then decide to revert to the default values, click **Reset Defaults**.
8. Click **Update**.

After you perform this task, other devices in the device group can send failover messages to the local device using the specified IP addresses.

Creating a Sync-Only device group

You perform this task to create a Sync-Only type of device group. When you create a Sync-Only device group, the BIG-IP system can then automatically synchronize certain types of data such as security policies to the other devices in the group, even when some of those devices reside in another network. You can perform this task on any BIG-IP® device within the local trust domain.

1. On the Main tab, click **Device Management** > **Device Groups**.
2. On the Device Groups list screen, click **Create**.
   The New Device Group screen opens.
3. Type a name for the device group, select the device group type **Sync-Only**, and type a description for the device group.
4. For the **Members** setting, select an IP address and host name from the **Available** list for each BIG-IP device that you want to include in the device group. Use the Move button to move the host name to the **Includes** list.

The list shows any devices that are members of the device's local trust domain.

5. For the **Automatic Sync** setting, select the **Enabled** check box.
6. Click **Finished**.

You now have a Sync-Only type of device group containing BIG-IP devices as members.

Enabling ASM synchronization on a Sync-Only device group

You need to have set up the BIG-IP® systems you want to synchronize in a device trust and a device group. Application Security Manager™ (ASM) must be provisioned on all the systems in the device group.

You must enable ASM™ synchronization on the members of the device group before the BIG-IP system can synchronize security policies and configurations to those members. You do this task on any one system, and the change is later synchronized to all members of the group.

1. On the Main tab, click **Application Security** > **Synchronization**.
   The system displays a list of device groups of which this device is a member.
2. For **Device Group**, select the Sync-Only device group you created.
3. Click **Save**.

After you perform this task, the BIG-IP ASM security policies and configuration data can be synchronized successfully to all devices in the device group.

Synchronizing an ASM-enabled device group

You need to have set up the BIG-IP® Application Security Manager™ (ASM) systems you want to synchronize in a Sync-Failover device group that is ASM™-enabled.

You can manually synchronize security policies and configuration of systems in an ASM-enabled device group.

1. On one system in the ASM-enabled failover device group, create an application security class, then use the Deployment wizard to create a security policy.
   Because the two systems are not in sync, you see a **Changes Pending** status message on the screen.
2. Click the **Changes Pending** message.

   Tip: *You can also click* ***Device Management*** > ***Overview***.

   The Overview screen opens.
3. In the Device Groups area of the screen, in the Name column, select the name of the relevant device group.
   The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
4. In the Devices area of the screen, in the Sync Status column, select the device that shows a sync status of `Changes Pending`.
5. In the Sync Options area of the screen, select **Sync Device to Group**.
6. Click **Sync**.
   The BIG-IP system syncs the configuration data of the selected device in the Device area of the screen to the other members of the device group.
7. Verify that the devices are synchronized.

For example, log in to another device in the device group and verify that the security policy you created also resides on that system. Click **Application Security** > **Security Policies** and see if the policy is listed.

Except for static self IP addresses, the entire set of BIG-IP configuration data including ASM™ security policies and configuration is replicated on one or more devices in the ASM-enabled device group. If the active device is not available, the standby device becomes active and handles traffic.

You can create new security policies or update existing ones on any of the devices in the group, or update the ASM configuration options. You can manually synchronize changes you make on one device with the other devices in the ASM-enabled device group.

Implementation result

You have set up multiple BIG-IP® systems running Application Security Manager™ (ASM) so that they automatically synchronize their ASM security policies and ASM configuration data. In addition, with this implementation, you can manually synchronize the local traffic configuration, as needed.

You can create new security policies or update existing ones on any of the devices in the group, or update the ASM™ configuration options. Any ASM changes you make on one device are automatically synchronized with the other devices in the ASM-enabled Sync-Only device group.

If Attack Signatures **Update Mode** is scheduled for automatic update, the attack signature update settings are synchronized. Each device in the device group updates itself independently according to the configured schedule. If you manually upload attack signatures or click **Update Signatures** to update from the server, the update is propagated to all of the devices in the device group.

Chapter 2

Manually Synchronizing Application Security Configurations

Topics:

- *Overview: Manually synchronizing ASM systems*
- *Implementation result*

Overview: Manually synchronizing ASM systems

This implementation describes how to set up two BIG-IP® systems running Application Security Manager™ (ASM) so that you can synchronize their security policies and configurations. With this implementation, the BIG-IP systems can fail over to one another, and you can manually sync all of the BIG-IP configuration data, including ASM policy data.

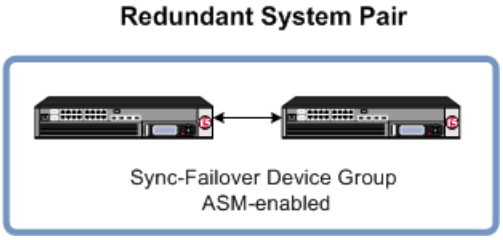

Figure 2: Manually synchronizing ASM configuration data

The two BIG-IP systems are set up for redundancy: one active and the other standby. Both systems are in the local trust domain and in the same Sync-Failover device group. If one system is unavailable, the other system begins to process application traffic. You can manually synchronize the systems. The ASM™ configurations and security policies are duplicated on both systems.

You can use this implementation as the basis for more complex configurations. For example, if you have multiple redundant pairs each supporting a different web application, you can use this implementation to set up each pair. You could create a Sync-Failover device group for each pair and then synchronize the data within each pair only. In this configuration, you all devices reside in the local trust domain.

Task summary

Performing basic network configuration for synchronization
Specifying an IP address for config sync
Establishing device trust
Creating a Sync-Failover device group
Syncing the BIG-IP configuration to the device group
Specifying IP addresses for failover
Enabling ASM synchronization on a device group
Synchronizing an ASM-enabled device group

About device management and synchronizing application security configurations

You can use device management to set up several BIG-IP® systems running Application Security Manager™ (ASM) so that the systems synchronize their security policies and configurations, and fail over to one another if a system goes offline for any reason. By using application security synchronization, you can set up application security and create security policies on one system, and can propagate them to other systems in an application security device group. In BIG-IP ASM™, a *device group* is two or more BIG-IP devices using the same configuration and providing consistent security policy enforcement.

You can set up application security synchronization, for example, behind an Application Delivery Controller where multiple BIG-IP systems running Application Security Manager are deployed as members of a pool. The options and security policies on all of the systems stay in sync regardless of where you update them.

When you set up ASM™ synchronization, in addition to security policies, other settings such as custom attack signatures, logging profiles, SMTP configuration, anti-virus protection, system variables, and policy templates, are synchronized with all devices in the ASM-enabled device group.

Considerations for application security synchronization

When using device management with Application Security Manager™ (ASM™), you need to be aware of the following considerations that apply specifically to application security synchronization.

- A BIG-IP® system with Application Security Manager can be a member of only one ASM-enabled device group.
- All BIG-IP systems in a device group must be running the same version (including hot fix updates) of Application Security Manager (version 11.0 or later).
- The BIG-IP systems in the ASM-enabled device group synchronize application security configuration data and security policies, providing consistent enforcement on all the devices.
- Real Traffic Policy Builder® can run on only one system per security policy. For example, you can set up automatic security policy building on one system that is a member of an ASM-enabled device group, the policy is built on that system and then automatically updated on all of the systems in the device group.

Performing basic network configuration for synchronization

You need to perform basic networking configuration for each of the BIG-IP® systems whose Application Security Manager™ (ASM) configurations you want to synchronize.

1. Install the same BIG-IP system version (including any hot fixes) on each device.
2. Provision LTM® and ASM™ on each device (**System** > **Resource Provisioning**).
3. On each device, create one or more VLANs, depending on your networking configuration (**Network** > **VLANs**).
4. On each device, create a self IP (**Network** > **Self IPs**).

   When creating the self IP, set **Traffic Group** to **traffic-group-local-only (non-floating)**.
5. On each device, create a default gateway, if needed (**Network** > **Routes**).
6. On each device, configure DNS (**System** > **Configuration** > **Device** > **DNS**) and NTP (**System** > **Configuration** > **Device** > **NTP**) so they are set to the same time.
7. Verify connectivity between the devices (self IP address to self IP address). For example, use this command to ensure communications: `ping -I vlan_interface device_self_IP`
8. If your company requires special device certificates, install them on each device (**System** > **Device Certificates** and click **Import**).

The basic networking setup is complete for the BIG-IP ASM systems for which you want to share security policies and configurations.

Specifying an IP address for config sync

Before configuring the config sync address, verify that all devices in the device group are running the same version of BIG-IP® system software.

You perform this task to specify the IP address on the local device that other devices in the device group will use to synchronize their configuration objects to the local device.

Note: You must perform this task locally on each device in the device group.

1. Confirm that you are logged in to the actual device you want to configure.
2. On the Main tab, click **Device Management** > **Devices**.
   This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose ConfigSync.
5. For the **Local Address** setting, retain the displayed IP address or select another address from the list.
   F5 Networks recommends that you use the default value, which is the self IP address for VLAN `internal`. This address must be a non-floating self IP address and not a management IP address.
6. Click **Update**.

After performing this task, the other devices in the device group can sync their configurations to the local device.

Establishing device trust

Before you begin this task, verify that:

- Each BIG-IP® device that is to be part of the local trust domain has a device certificate installed on it.
- The local device is designated as a certificate signing authority.

You perform this task to establish trust among devices on one or more network segments. Devices that trust each other constitute the *local trust domain*. A device must be a member of the local trust domain prior to joining a device group.

By default, the BIG-IP software includes a local trust domain with one member, which is the local device. You can choose any one of the BIG-IP devices slated for a device group and log into that device to add other devices to the local trust domain. For example, devices `A`, `B`, and `C` each initially shows only itself as a member of the local trust domain. To configure the local trust domain to include all three devices, you can simply log into device `A` and add devices `B` and `C` to the local trust domain. Note that there is no need to repeat this process on devices `B` and `C`.

1. On the Main tab, click **Device Management** > **Device Trust**, and then either **Peer List** or **Subordinate List**.
2. Click **Add**.
3. Type an IP address, administrator user name, and administrator password for the remote BIG-IP® device.
   This IP address can be either a management IP address or a self IP address.
4. Click **Retrieve Device Information**.
5. Verify that the certificate of the remote device is correct.
6. Verify that the name of the remote device is correct.
7. Verify that the management IP address and name of the remote device are correct.
8. Click **Finished**.

The device you added is now a member of the local trust domain.

Repeat this task for each device that you want to add to the local trust domain.

Creating a Sync-Failover device group

This task establishes failover capability between two or more BIG-IP devices. If the active device in a Sync-Failover device group becomes unavailable, the configuration objects fail over to another member of the device group and traffic processing is unaffected. You perform this task on any one of the authority devices within the local trust domain.

Repeat this task for each Sync-Failover device group that you want to create for your network configuration.

1. On the Main tab, click **Device Management** > **Device Groups**.
2. On the Device Groups list screen, click **Create**.
   The New Device Group screen opens.
3. Type a name for the device group, select the device group type **Sync-Failover**, and type a description for the device group.
4. In the Configuration area of the screen, select a host name from the **Available** list for each BIG-IP device that you want to include in the device group, including the local device. Use the Move button to move the host name to the **Selected** list.

   The **Available** list shows any devices that are members of the device's local trust domain but not currently members of a Sync-Failover device group. A device can be a member of one Sync-Failover group only.
5. For the **Network Failover** setting:
   - Select the **Enabled** check box if you want device group members to handle failover communications by way of network connectivity.
   - Clear the **Enabled** check box if you want device group members to handle failover communications by way of serial cable (hard-wired) connectivity.

   Serial failover is not available for device groups with more than two members.
6. Click **Finished**.

You now have a Sync-Failover type of device group containing BIG-IP devices as members.

Syncing the BIG-IP configuration to the device group

Before you sync the configuration, verify that the devices targeted for config sync are members of a device group and that device trust has been established.

This task synchronizes the BIG-IP® configuration data from the local device to the devices in the device group. This synchronization ensures that devices in the device group operate properly. When synchronizing self IP addresses, the BIG-IP system synchronizes floating self IP addresses only.

Important: *You perform this task on either of the two devices, but not both.*

1. On the Main tab, click **Device Management** > **Overview**.
2. In the Device Groups area of the screen, in the Name column, select the name of the relevant device group.
   The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
3. In the Devices area of the screen, in the Sync Status column, select the device that shows a sync status of `Changes Pending`.
4. In the Sync Options area of the screen, select **Sync Device to Group**.

5. Click **Sync**.
   The BIG-IP system syncs the configuration data of the selected device in the Device area of the screen to the other members of the device group.

Except for non-floating self IP addresses, the entire set of BIG-IP configuration data is replicated on each device in the device group.

Specifying IP addresses for failover

This task specifies the local IP addresses that you want other devices in the device group to use for failover communications with the local device. You must perform this task locally on each device in the device group.

Note: *The failover addresses that you specify must belong to route domain* `0`.

1. Confirm that you are logged in to the actual device you want to configure.
2. On the Main tab, click **Device Management** > **Devices**.
   This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose Failover.
5. For the Failover Unicast Configuration settings, retain the displayed IP addresses.
   You can also click **Add** to specify additional IP addresses that the system can use for failover communications. F5 Networks recommends that you use the self IP address assigned to the `HA` VLAN.
6. If the BIG-IP® system is running on a VIPRION® platform, then for the **Use Failover Multicast Address** setting, select the **Enabled** check box.
7. If you enable **Use Failover Multicast Address**, either accept the default **Address** and **Port** values, or specify values appropriate for the device.
   If you revise the default **Address** and **Port** values, but then decide to revert to the default values, click **Reset Defaults**.
8. Click **Update**.

After you perform this task, other devices in the device group can send failover messages to the local device using the specified IP addresses.

Enabling ASM synchronization on a device group

You need to have already set up the BIG-IP®systems you want to synchronize in a device group. Application Security Manager™ (ASM) must be provisioned on all the systems in the device group.

You perform this task to enable ASM™ synchronization on a device group, before you can synchronize security policies and configurations to device group members. You perform this task on any one device in the device group, and the change is later synchronized to all members of the group.

1. On the Main tab, click **Application Security** > **Synchronization**.
   The system displays a list of device groups of which this device is a member.
2. For **Device Group**, select the device group whose members you want to synchronize.
3. Click **Save**.

After you perform this task, the BIG-IP ASM security policies and configuration data can be synchronized successfully to all devices in the device group.

Synchronizing an ASM-enabled device group

You need to have set up the BIG-IP® Application Security Manager™ (ASM) systems you want to synchronize in a Sync-Failover device group that is ASM™-enabled.

You can manually synchronize security policies and configuration of systems in an ASM-enabled device group.

1. On one system in the ASM-enabled failover device group, create an application security class, then use the Deployment wizard to create a security policy.
   Because the two systems are not in sync, you see a **Changes Pending** status message on the screen.
2. Click the **Changes Pending** message.

   Tip: *You can also click* ***Device Management*** > ***Overview***.

   The Overview screen opens.
3. In the Device Groups area of the screen, in the Name column, select the name of the relevant device group.
   The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
4. In the Devices area of the screen, in the Sync Status column, select the device that shows a sync status of `Changes Pending`.
5. In the Sync Options area of the screen, select **Sync Device to Group**.
6. Click **Sync**.
   The BIG-IP system syncs the configuration data of the selected device in the Device area of the screen to the other members of the device group.
7. Verify that the devices are synchronized.

   For example, log in to another device in the device group and verify that the security policy you created also resides on that system. Click **Application Security** > **Security Policies** and see if the policy is listed.

Except for static self IP addresses, the entire set of BIG-IP configuration data including ASM™ security policies and configuration is replicated on one or more devices in the ASM-enabled device group. If the active device is not available, the standby device becomes active and handles traffic.

You can create new security policies or update existing ones on any of the devices in the group, or update the ASM configuration options. You can manually synchronize changes you make on one device with the other devices in the ASM-enabled device group.

Implementation result

You have now set up two BIG-IP® systems running Application Security Manager™ (ASM) so that you can synchronize their security policies and configurations. With this implementation, you manually synchronize the ASM and BIG-IP configurations.

The two BIG-IP systems are in the same Sync-Failover device group. If one system becomes unavailable, the other system begins processing application traffic.

Chapter 3

Synchronizing Application Security Configurations Across LANs

Topics:

- *Overview: Synchronizing ASM systems across LANs*
- *Implementation result*

Overview: Synchronizing ASM systems across LANs

This implementation describes how to set up multiple BIG-IP® systems running Application Security Manager™ (ASM) so that you can synchronize their security policies and configurations for disaster recovery. You can use this implementation to synchronize BIG-IP ASM™ security policies and configurations on systems that reside in different network segments or LANs, such as those in separate offices or data centers. Note that traffic must be routable between the network segments. If a disaster occurs at one of the offices and both devices are disabled, the latest security policies are still available on the systems in the other location.

This implementation also configures failover between systems in a redundant pair on a particular network segment. If one of the devices in a pair goes offline for any reason, the other device in the pair begins processing the application traffic.

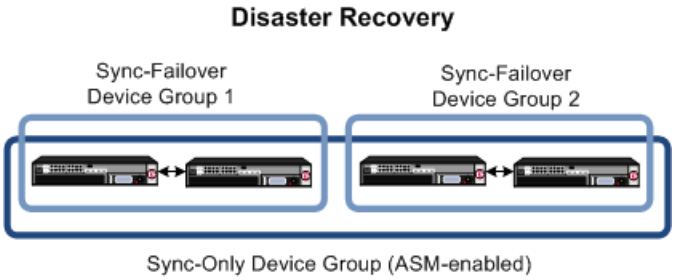

Figure 3: Automatically synchronizing ASM configuration data across LANs

In the figure, two sets of BIG-IP systems are set up for redundancy: one active and the other standby. Each pair is in a different network segment (LAN), and there can be additional pairs, as needed. Each LAN has one pair of devices, where both have the same default routing, but routing is not the same for the devices in the other LAN.

All of the systems are running ASM and are in the trust domain. Three device groups are set up: one Sync-Failover device group for each pair (not ASM-enabled), and one Sync-Only device group with ASM enabled using automatic synchronization for all of the systems. The systems automatically duplicate the ASM configurations and security policies on all of the systems. You can manually synchronize the BIG-IP configurations of each pair of systems when needed.

Task summary

Performing basic network configuration for synchronization
Specifying an IP address for config sync
Establishing device trust
Creating a Sync-Failover device group
Syncing the BIG-IP configuration to the device group
Specifying IP addresses for failover
Creating a Sync-Only device group
Enabling ASM synchronization on a Sync-Only device group
Synchronizing an ASM-enabled device group

About device management and synchronizing application security configurations

You can use device management to set up several BIG-IP® systems running Application Security Manager™ (ASM) so that the systems synchronize their security policies and configurations, and fail over to one another if a system goes offline for any reason. By using application security synchronization, you can set up application security and create security policies on one system, and can propagate them to other systems in an application security device group. In BIG-IP ASM™, a *device group* is two or more BIG-IP devices using the same configuration and providing consistent security policy enforcement.

You can set up application security synchronization, for example, behind an Application Delivery Controller where multiple BIG-IP systems running Application Security Manager are deployed as members of a pool. The options and security policies on all of the systems stay in sync regardless of where you update them.

When you set up ASM™ synchronization, in addition to security policies, other settings such as custom attack signatures, logging profiles, SMTP configuration, anti-virus protection, system variables, and policy templates, are synchronized with all devices in the ASM-enabled device group.

Considerations for application security synchronization

When using device management with Application Security Manager™ (ASM™), you need to be aware of the following considerations that apply specifically to application security synchronization.

- A BIG-IP® system with Application Security Manager can be a member of only one ASM-enabled device group.
- All BIG-IP systems in a device group must be running the same version (including hot fix updates) of Application Security Manager (version 11.0 or later).
- The BIG-IP systems in the ASM-enabled device group synchronize application security configuration data and security policies, providing consistent enforcement on all the devices.
- Real Traffic Policy Builder® can run on only one system per security policy. For example, you can set up automatic security policy building on one system that is a member of an ASM-enabled device group, the policy is built on that system and then automatically updated on all of the systems in the device group.

Performing basic network configuration for synchronization

You need to perform basic networking configuration for each of the BIG-IP® systems whose Application Security Manager™ (ASM) configurations you want to synchronize.

1. Install the same BIG-IP system version (including any hot fixes) on each device.
2. Provision LTM® and ASM™ on each device (**System** > **Resource Provisioning**).
3. On each device, create one or more VLANs, depending on your networking configuration (**Network** > **VLANs**).
4. On each device, create a self IP (**Network** > **Self IPs**).

   When creating the self IP, set **Traffic Group** to **traffic-group-local-only (non-floating)**.
5. On each device, create a default gateway, if needed (**Network** > **Routes**).
6. On each device, configure DNS (**System** > **Configuration** > **Device** > **DNS**) and NTP (**System** > **Configuration** > **Device** > **NTP**) so they are set to the same time.
7. Verify connectivity between the devices (self IP address to self IP address). For example, use this command to ensure communications: `ping -I vlan_interface device_self_IP`

8. If your company requires special device certificates, install them on each device (**System** > **Device Certificates** and click **Import**).

The basic networking setup is complete for the BIG-IP ASM systems for which you want to share security policies and configurations.

Specifying an IP address for config sync

Before configuring the config sync address, verify that all devices in the device group are running the same version of BIG-IP® system software.

You perform this task to specify the IP address on the local device that other devices in the device group will use to synchronize their configuration objects to the local device.

Note: *You must perform this task locally on each device in the device group.*

1. Confirm that you are logged in to the actual device you want to configure.
2. On the Main tab, click **Device Management** > **Devices**.
   This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose ConfigSync.
5. For the **Local Address** setting, retain the displayed IP address or select another address from the list.
   F5 Networks recommends that you use the default value, which is the self IP address for VLAN `internal`. This address must be a non-floating self IP address and not a management IP address.
6. Click **Update**.

After performing this task, the other devices in the device group can sync their configurations to the local device.

Establishing device trust

Before you begin this task, verify that:

- Each BIG-IP® device that is to be part of the local trust domain has a device certificate installed on it.
- The local device is designated as a certificate signing authority.

You perform this task to establish trust among devices on one or more network segments. Devices that trust each other constitute the *local trust domain*. A device must be a member of the local trust domain prior to joining a device group.

By default, the BIG-IP software includes a local trust domain with one member, which is the local device. You can choose any one of the BIG-IP devices slated for a device group and log into that device to add other devices to the local trust domain. For example, devices `A`, `B`, and `C` each initially shows only itself as a member of the local trust domain. To configure the local trust domain to include all three devices, you can simply log into device `A` and add devices `B` and `C` to the local trust domain. Note that there is no need to repeat this process on devices `B` and `C`.

1. On the Main tab, click **Device Management** > **Device Trust**, and then either **Peer List** or **Subordinate List**.
2. Click **Add**.
3. Type an IP address, administrator user name, and administrator password for the remote BIG-IP® device.
   This IP address can be either a management IP address or a self IP address.

4. Click **Retrieve Device Information**.
5. Verify that the certificate of the remote device is correct.
6. Verify that the name of the remote device is correct.
7. Verify that the management IP address and name of the remote device are correct.
8. Click **Finished**.

The device you added is now a member of the local trust domain.

Repeat this task for each device that you want to add to the local trust domain.

Creating a Sync-Failover device group

This task establishes failover capability between two or more BIG-IP devices. If the active device in a Sync-Failover device group becomes unavailable, the configuration objects fail over to another member of the device group and traffic processing is unaffected. You perform this task on any one of the authority devices within the local trust domain.

Repeat this task for each Sync-Failover device group that you want to create for your network configuration.

1. On the Main tab, click **Device Management** > **Device Groups**.
2. On the Device Groups list screen, click **Create**.
   The New Device Group screen opens.
3. Type a name for the device group, select the device group type **Sync-Failover**, and type a description for the device group.
4. In the Configuration area of the screen, select a host name from the **Available** list for each BIG-IP device that you want to include in the device group, including the local device. Use the Move button to move the host name to the **Selected** list.

   The **Available** list shows any devices that are members of the device's local trust domain but not currently members of a Sync-Failover device group. A device can be a member of one Sync-Failover group only.
5. For the **Network Failover** setting:
   - Select the **Enabled** check box if you want device group members to handle failover communications by way of network connectivity.
   - Clear the **Enabled** check box if you want device group members to handle failover communications by way of serial cable (hard-wired) connectivity.

   Serial failover is not available for device groups with more than two members.
6. Click **Finished**.

You now have a Sync-Failover type of device group containing BIG-IP devices as members.

Syncing the BIG-IP configuration to the device group

Before you sync the configuration, verify that the devices targeted for config sync are members of a device group and that device trust has been established.

This task synchronizes the BIG-IP® configuration data from the local device to the devices in the device group. This synchronization ensures that devices in the device group operate properly. When synchronizing self IP addresses, the BIG-IP system synchronizes floating self IP addresses only.

Important: *You perform this task on either of the two devices, but not both.*

1. On the Main tab, click **Device Management** > **Overview**.
2. In the Device Groups area of the screen, in the Name column, select the name of the relevant device group.
   The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
3. In the Devices area of the screen, in the Sync Status column, select the device that shows a sync status of `Changes Pending`.
4. In the Sync Options area of the screen, select **Sync Device to Group**.
5. Click **Sync**.
   The BIG-IP system syncs the configuration data of the selected device in the Device area of the screen to the other members of the device group.

Except for non-floating self IP addresses, the entire set of BIG-IP configuration data is replicated on each device in the device group.

Specifying IP addresses for failover

This task specifies the local IP addresses that you want other devices in the device group to use for failover communications with the local device. You must perform this task locally on each device in the device group.

Note: *The failover addresses that you specify must belong to route domain* `0`.

1. Confirm that you are logged in to the actual device you want to configure.
2. On the Main tab, click **Device Management** > **Devices**.
   This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose Failover.
5. For the Failover Unicast Configuration settings, retain the displayed IP addresses.

   You can also click **Add** to specify additional IP addresses that the system can use for failover communications. F5 Networks recommends that you use the self IP address assigned to the `HA` VLAN.
6. If the BIG-IP® system is running on a VIPRION® platform, then for the **Use Failover Multicast Address** setting, select the **Enabled** check box.
7. If you enable **Use Failover Multicast Address**, either accept the default **Address** and **Port** values, or specify values appropriate for the device.

   If you revise the default **Address** and **Port** values, but then decide to revert to the default values, click **Reset Defaults**.
8. Click **Update**.

After you perform this task, other devices in the device group can send failover messages to the local device using the specified IP addresses.

Creating a Sync-Only device group

You perform this task to create a Sync-Only type of device group. When you create a Sync-Only device group, the BIG-IP system can then automatically synchronize certain types of data such as security policies to the other devices in the group, even when some of those devices reside in another network. You can perform this task on any BIG-IP® device within the local trust domain.

1. On the Main tab, click **Device Management** > **Device Groups**.
2. On the Device Groups list screen, click **Create**.
   The New Device Group screen opens.
3. Type a name for the device group, select the device group type **Sync-Only**, and type a description for the device group.
4. For the **Members** setting, select an IP address and host name from the **Available** list for each BIG-IP device that you want to include in the device group. Use the Move button to move the host name to the **Includes** list.

   The list shows any devices that are members of the device's local trust domain.
5. For the **Automatic Sync** setting, select the **Enabled** check box.
6. Click **Finished**.

You now have a Sync-Only type of device group containing BIG-IP devices as members.

Enabling ASM synchronization on a Sync-Only device group

You need to have set up the BIG-IP® systems you want to synchronize in a device trust and a device group. Application Security Manager™ (ASM) must be provisioned on all the systems in the device group.

You must enable ASM™ synchronization on the members of the device group before the BIG-IP system can synchronize security policies and configurations to those members. You do this task on any one system, and the change is later synchronized to all members of the group.

1. On the Main tab, click **Application Security** > **Synchronization**.
   The system displays a list of device groups of which this device is a member.
2. For **Device Group**, select the Sync-Only device group you created.
3. Click **Save**.

After you perform this task, the BIG-IP ASM security policies and configuration data can be synchronized successfully to all devices in the device group.

Synchronizing an ASM-enabled device group

You need to have set up the BIG-IP® Application Security Manager™ (ASM) systems you want to synchronize in a Sync-Failover device group that is ASM™-enabled.

You can manually synchronize security policies and configuration of systems in an ASM-enabled device group.

1. On one system in the ASM-enabled failover device group, create an application security class, then use the Deployment wizard to create a security policy.
   Because the two systems are not in sync, you see a **Changes Pending** status message on the screen.
2. Click the **Changes Pending** message.

   Tip: *You can also click* ***Device Management*** > ***Overview***.

   The Overview screen opens.
3. In the Device Groups area of the screen, in the Name column, select the name of the relevant device group.
   The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.

4. In the Devices area of the screen, in the Sync Status column, select the device that shows a sync status of `Changes Pending`.
5. In the Sync Options area of the screen, select **Sync Device to Group**.
6. Click **Sync**.
   The BIG-IP system syncs the configuration data of the selected device in the Device area of the screen to the other members of the device group.
7. Verify that the devices are synchronized.

   For example, log in to another device in the device group and verify that the security policy you created also resides on that system. Click **Application Security** > **Security Policies** and see if the policy is listed.

Except for static self IP addresses, the entire set of BIG-IP configuration data including ASM™ security policies and configuration is replicated on one or more devices in the ASM-enabled device group. If the active device is not available, the standby device becomes active and handles traffic.

You can create new security policies or update existing ones on any of the devices in the group, or update the ASM configuration options. You can manually synchronize changes you make on one device with the other devices in the ASM-enabled device group.

Implementation result

You have set up disaster recovery for multiple BIG-IP® systems running Application Security Manager™ (ASM). Each office or data center has an active system and a standby that takes over if the active system should fail. You must manually synchronize the BIG-IP configuration from one system to the other if you change the configuration.

You can create new security policies or update existing ones on any of the devices in the group, or update the ASM™ configuration options (**Application Security**>**Options**). Any changes you make on one device are automatically synchronized with the other devices in the ASM-enabled Sync-Only device group.

If Attack Signatures **Update Mode** is scheduled for automatic update, the attack signature update settings are synchronized. Each device in the device group updates itself independently according to the configured schedule. If you manually upload attack signatures or click **Upload Signatures** to update from the server, the update is propagated to all of the devices in the device group.

Chapter 4

Setting Up IP Address Intelligence Blocking

Topics:

- *Overview: Setting up IP address intelligence blocking*
- *IP address intelligence categories*

Overview: Setting up IP address intelligence blocking

In Application Security Manager™, you can use IP address intelligence blocking in a security policy to block requests from IP addresses that have questionable reputations. IP addresses from which attacks or spam have originated are included in an IP intelligence database, along with the category describing the problem. The BIG-IP® system must connect to the IP intelligence database before you can use IP address intelligence blocking.

You can configure a security policy to log (alarm) or block requests from IP addresses of questionable reputation, and to perform different actions depending on the categories of problems. For example, you can block requests from IP addresses associated with Windows exploits and log requests from scanners.

You can create a whitelist of IP addresses that might be in the database, and allow them to access the web application regardless of their IP reputation. This is a way to ensure that traffic from known sources is not blocked because of IP address intelligence data.

You can also use iRules® to instruct the system how to use IP address intelligence information.

Task Summary

These are tasks for setting up IP address intelligence blocking in a security policy.

Enabling IP address intelligence
Setting up IP address intelligence blocking
Reviewing IP address intelligence statistics
Creating an iRule to log IP address intelligence information
Creating an iRule to reject requests with questionable IP addresses

Enabling IP address intelligence

The requirements for using IP address intelligence are:

- The system must have an IP Intelligence license.
- The system must have an Internet connection either directly or through a proxy server.
- The system must have DNS configured (go to **System** > **Configuration** > **Device** > **DNS**).

Important: *IP address intelligence is enabled by default. You only need to enable it if it was previously disabled.*

To enable IP address intelligence on the BIG-IP® system, you enable auto-update to connect the system to the IP intelligence database.

1. Log in to the command line for the BIG-IP® system.
2. To determine whether IP intelligence is enabled, type the following command: `tmsh list sys db iprep.autoupdate`
   If the value of the `iprep.autoupdate` variable is `disable`, IP intelligence is not enabled. If it is `enable`, your task is complete.
3. At the prompt, type `tmsh modify sys db iprep.autoupdate value enable`
   The system downloads the IP intelligence database and stores it in the binary file, `/var/IpRep/F5IpRep.dat`. It is updated every 5 minutes.

4. If the BIG-IP system is behind a firewall, make sure that the BIG-IP system has external access to `vector.brightcloud.com` using port `443`. That is the IP Intelligence server from which the system gets IP Intelligence information.
5. (Optional) If the BIG-IP system connects to the Internet using a forward proxy server, set these system database variables.
   a) Type `tmsh modify sys db proxy.host value` *`hostname`* to specify the hostname of the proxy server.
   b) Type `tmsh modify sys db proxy.port value` *`port_number`* to specify the port number of the proxy server.
   c) Type `tmsh modify sys db proxy.username value` *`hostname`* to specify the user name to log in to the proxy server.
   d) Type `tmsh modify sys db proxy.password value` *`password`* to specify the password to log in to the proxy server.

The IP address intelligence feature remains enabled unless you disable it with the command `tmsh modify sys db iprep.autoupdate value disable`.

You can create iRules® to instruct the system how to handle traffic from IP addresses with questionable reputations, or use Application Security Manager™ to configure IP address intelligence blocking.

Setting up IP address intelligence blocking

Before you can set up IP address intelligence blocking, your system must have IP address intelligence enabled.

You can configure a security policy to log and block requests from source IP addresses that, according to an IP intelligence database, have a bad reputation and could cause a potential attack.

1. On the Main tab, click **Application Security** > **IP Addresses** > **IP Address Intelligence**.
   The IP Address Intelligence screen opens.
2. In the **Current edited policy** list near the top of the screen, verify that the edited security policy is the one you want to work on.
3. For the **IP Address Intelligence** setting, select the **Enabled** check box.
4. For the **IP Address Whitelist** setting, specify any IP addresses you want to allow, even if they are found in the IP intelligence database.
   a) Type the **IP Address** and **Subnet Mask** of the address to consider safe.
   b) Click **Add**.

   The addresses that you typed are added to the list.
5. In the IP Address Intelligence Categories area, select **Alarm** or **Block**, or both, for the categories of IP addresses you are interested in.

   - Select **Alarm** to cause the system to log the IP address intelligence data (IP address intelligence category and status) on the Requests screen whenever a request is from a source IP address in that category.
   - Select **Block** to stop requests sent from a source IP address that matches that category

   

   Note: *If these settings are not available, click* ***Policy*** > ***Blocking*** *and for the violation* `Access from malicious IP address`*, select the* ***Alarm*** *and* ***Block*** *settings.*

6. Click **Save**.

The system matches source IP addresses to those in the IP address intelligence database. When a match is found, the violation `Access from malicious IP address` occurs. The system determines what category of reputation the IP address has, then logs or blocks the IP address according to how the IP Address Intelligence categories are set.

Reviewing IP address intelligence statistics

Before you can view IP address intelligence statistics, your system must have IP address intelligence enabled.

After you set up IP intelligence blocking on the Application Security Manager™, you can review statistics concerning how many requests were received from IP addresses with questionable reputations. You can also view the requests from those IP addresses.

1. On the Main tab, click **Application Security** > **Reporting** > **Charts**.
   The Charts screen opens, where you can view graphical reports.
2. In the Charts area, next to **View by**, click **IP Address Intelligence**.
   The chart shows details about IP addresses that were used to send the illegal requests, grouped according to their reputation in the IP intelligence database.
3. Hover over the pie chart or look at the Details table it to see the categories of IP addresses with questionable reputations.
4. Under Chart Path on the left, click **View Requests** to see the requests from IP addresses in the IP intelligence database.
   The Requests list opens.
5. Click any request to view details about the request.
   The screen expands to show more information about the request. IP address intelligence information is shown in the **Source IP Address** field in the request details. The details include the category of the malicious IP address and information about when the IP intelligence database was last updated.
6. If you have set up remote logging, you can also review IP intelligence data on the remote logger.

By reviewing the IP address intelligence data, you can examine requests from potentially malicious IP addresses.

Based on the statistics and IP address intelligence categories that the IP addresses fall into, you can adjust what happens (alarm or block) when the system receives requests from IP addresses in different categories.

Creating an iRule to log IP address intelligence information

Before you can create an iRule to log IP address intelligence information, your system must have IP address intelligence enabled.

You use iRules® to log IP address intelligence categories to the file `/var/log/ltm`. This is an example of the type of iRule you can write.

1. On the Main tab, click **Local Traffic** > **iRules**.
   The iRule List screen opens, displaying any existing iRules.
2. Click **Create**.
   The New iRule screen opens.
3. In the **Name** field, type a name between 1 and 31 characters, such as `my_iRule`.
4. In the **Definition** field, type the iRule using Tool Command Language (Tcl) syntax.

For example, to log all IP addresses and any associated IP address intelligence categories, type the following iRule:

```
when CLIENT_ACCEPTED {
             log local0. "IP Address Intelligence for IP address
[IP::client_addr]:
             [IP::reputation [IP::client_addr]]"
           }
```

5. Click **Finished**.
   The new iRule appears in the list of iRules on the system.

When traffic is received from an IP address with a questionable reputation and that is included in the IP intelligence database, the system prints the IP address intelligence information in the `/var/log/ltm` log.

For complete and detailed information about iRules syntax, see the F5 Networks DevCentral web site, `http://devcentral.f5.com`.

Creating an iRule to reject requests with questionable IP addresses

Before you can create an iRule to reject requests based on an IP address reputation, your system must have IP address intelligence enabled.

You can use iRules® to reject requests from IP addresses that have questionable reputations and are listed in the IP intelligence database. This is an example of the type of iRule you can write.

1. On the Main tab, click **Local Traffic** > **iRules**.
   The iRule List screen opens, displaying any existing iRules.
2. Click **Create**.
   The New iRule screen opens.
3. In the **Name** field, type a name between 1 and 31 characters, such as `my_iRule`.
4. In the **Definition** field, type the iRule using Tool Command Language (Tcl) syntax.

   For example, to reject requests from IP addresses listed in the IP intelligence database because they could be Windows Exploits or Web Attacks, type the following iRule:

```
when HTTP_REQUEST {
    set ip_reputation_categories [IP::reputation [IP::client_addr]]
    set is_reject 0
    if {($ip_reputation_categories contains "Windows Exploits")} {
       set is_reject 1
    }
    if {($ip_reputation_categories contains "Web Attacks")} {
       set is_reject 1
    }
    if {($is_reject)} {
        log local0. "Attempted access from malicious IP address
[IP::client_addr]
        ($ip_reputation_categories), request was rejected"
        HTTP::respond 200 content
        "<HTML><HEAD><TITLE>Rejected Request</TITLE>
        </HEAD><BODY>The request was rejected. <BR>
        Attempted access from malicious IP address</BODY></HTML>"
    }
}
```

5. Click **Finished**.
   The new iRule appears in the list of iRules on the system.

When traffic is received from an IP address with a questionable reputation that is included in the IP intelligence database, the system prints the IP address intelligence information in the `/var/log/ltm` log.

For complete and detailed information about iRules syntax, see the F5 Networks DevCentral web site, `http://devcentral.f5.com`.

IP address intelligence categories

Along with the IP address, the IP intelligence database stores the category that explains the reason that the IP address is considered untrustworthy.

Category	Description
Windows exploits	IP addresses that have exercised various exploits against Windows resources using browsers, programs, downloaded files, scripts, or operating system vulnerabilities.
Web attacks	IP addresses that have launched web attacks of various forms.
Botnets	IP addresses of computers that are infected with malicious software and are controlled as a group, and are now part of a botnet. Hackers can exploit botnets to send spam messages, launch various attacks, or cause target systems to behave in other unpredictable ways.
Scanners	IP addresses that have been observed to perform port scans or network scans, typically to identify vulnerabilities for later exploits.
Denial of Service	IP addresses that have launched Denial of Service (DoS) attacks. These attacks are usually requests for legitimate services, but occur at such a fast rate that targeted systems cannot respond and become bogged down or unable to service legitimate clients.
Infected Sources	IP addresses that issue HTTP requests with a low reputation index score, or are known malware sites.
Phishing	IP addresses that are associated with phishing web sites that masquerade as legitimate web sites.
Proxy	IP addresses that are associated with web proxies that shield the originator's IP address (such as anonymous proxies).

Chapter 5

Managing IP Address Exceptions

Topics:

- *Overview: Managing IP address exceptions*

Overview: Managing IP address exceptions

An *IP address exception* is an IP address that you want the system to treat in a specific way for a security policy. For example, you can specify IP addresses from which the system should always trust traffic, IP addresses for which you do not want the system to generate learning suggestions for the traffic, and IP addresses for which you want to exclude information from the logs. You can use the IP address exception feature to create exceptions for IP addresses of internal tools that your company uses, such as penetration tools, manual or automatic scanners, or web scraping tools. You can add an IP address exception, and instruct the system how to handle traffic coming from that address.

You can view a centralized list of IP address exceptions, and you can add new IP address exceptions to the list. The list of IP address exceptions shows exceptions that you add directly to the list, or those which you add from other locations, as shown by the following examples:

- When creating a security policy, you can specify IP addresses that you want the Policy Builder to always trust.
- When creating a security policy that is integrated with a vulnerability assessment tool, you can configure the scanner IP address as an IP address exception.
- When setting up anomaly detection (such as for DoS, brute force, and web scraping protections), you can specify IP addresses that the system should consider legitimate (called *whitelists*).
- When setting up IP address intelligence, you can add IP addresses that the system should allow even if the IP address is in the IP intelligence database.

The IP Address Exceptions list shows in one location all of the IP exceptions configured for this security policy. You can view or modify IP exceptions both from the centralized IP exception list and from the specific feature screens.

This implementation describes how to create, delete, and update the list of IP address exceptions.

Creating IP address exceptions

For each security policy, you can create a list of IP address exceptions, and indicate how you want the system to handle the traffic from these IP addresses.

1. On the Main tab, click **Application Security** > **IP Addresses** > **IP Address Exceptions**.
   The IP Address Exceptions screen opens, and displays a centralized list of configured IP address exceptions.
2. Click **Create**.
   The New IP Address Exception screen opens.
3. In the **IP Address** field, type the IP address that you want the system to trust.

   

   Note: *To add a route domain, type %n after the IP address where n is the route domain identification number.*

4. In the **Netmask** field, type the netmask of the IP address exception.
   If you omit the netmask value, the system uses a default value of `255.255.255.255`.
5. To consider traffic from this IP address as being safe, for the **Policy Builder trusted IP** setting, select **Enabled**.
   The system adds this IP address to the **Trusted IP Addresses** setting on the Automatic Configuration screen for the Policy Builder.

6. To ignore this IP address when performing DoS, brute force, and web scraping detection, for the **Ignore in Anomaly Detection** setting, select **Enabled**.
   The system adds this IP address to the **IP Address Whitelist** setting on the anomaly detection screens for DoS attacks, brute force, and web scraping.
7. If you do not want the system to generate learning suggestions for traffic sent from this IP address, for the **Ignore in Learning Suggestions** setting, select **Enabled**.

   Note: *Application Security Manager does not generate learning suggestions for requests that result in the web server returning HTTP responses with 400 or 404 status codes unless the security policy is configured to learn and block traffic (here both the* ***Ignore in Learning Suggestions*** *check box and the* ***Never block this IP Address*** *check box need to be disabled).*

8. To never block traffic from this IP address, for the **Never block this IP Address** setting, select **Enabled**.
   If the check box is cleared, a system in blocking mode blocks requests sent from this IP address according to the violation settings on the Policy Blocking Settings screen.
9. To prevent the system from logging requests (either legal or illegal) from this IP address, for the **Never log requests from this IP** setting, select **Enabled**.
10. To consider traffic from this IP address to be legitimate even if it is found in the IP Intelligence database, for the **Ignore IP Intelligence** setting, select **Enabled**.
    The system adds this IP address to the **IP Address Whitelist** setting on the IP Address Intelligence screen.
11. Click **Create**.
    The IP Address Exceptions screen opens and shows all of the exceptions configured for the security policy including the one you created.

You can view and manage all of your IP address exceptions from the centralized IP Address Exceptions screen.

Deleting IP address exceptions

If you no longer want an IP address on the exceptions list, you can delete the IP address exceptions.

1. On the Main tab, click **Application Security** > **IP Addresses** > **IP Address Exceptions**.
   The IP Address Exceptions screen opens, and displays a centralized list of configured IP address exceptions.
2. Select the IP address exception you want to delete and click **Delete**.
   The IP address exception is deleted from the list.
3. You can also delete IP address exceptions from the anomaly detection whitelists, the IP address intelligence whitelist, and the policy building configuration. On any of these screens, select the IP address, and click **Delete**.
   The system removes the IP address from the whitelist on the screen. However, the IP address remains on the IP Address Exceptions screen with the related setting changed. For example, if you deleted the IP address from an anomaly detection whitelist, the Anomaly Detection column for that IP address in the exceptions list changes from Ignore IP to say Include IP.
4. In the editing context area, click **Apply Policy** to immediately put the changes into effect.

Updating IP address exceptions

You can update IP address exceptions from the centralized list of IP address exceptions.

1. On the Main tab, click **Application Security** > **IP Addresses** > **IP Address Exceptions**.
   The IP Address Exceptions screen opens, and displays a centralized list of configured IP address exceptions.
2. Click the IP address of the IP address exception you want to modify.
   The IP Address Exception Properties screen opens.
3. Change the settings as needed.
4. Click **Update**.
5. In the editing context area, click **Apply Policy** to immediately put the changes into effect.

Chapter 6

Enforcing Application Use at Specific Geolocations

Topics:

- *Overview: Enforcing application use in certain geolocations*
- *Enforcing application use in certain geolocations*
- *Setting up geolocation enforcement from a request*

Overview: Enforcing application use in certain geolocations

Geolocation software can identify the geographic location of a client or web application user. *Geolocation* refers either to the process of assessing the location, or to the actual assessed location.

For applications protected by Application Security Manager™, you can use geolocation enforcement to restrict or allow application use in specific countries. You adjust the lists of which countries or locations are allowed or disallowed in a security policy. If an application user tries to access the web application from a location that is not allowed, the `Access from disallowed GeoLocation` violation occurs. By default, all locations are allowed, and the violation learn, alarm, and block flags are enabled.

Requests from certain locations, such as RFC-1918 addresses or unassigned global addresses, do not include a valid country code. The geolocation is shown as **N/A** in both the request, and the list of geolocations. You have the option to disallow N/A requests whose country of origination is unknown.

Enforcing application use in certain geolocations

Before you can set up geolocation enforcement, you need to create a security policy. If the BIG-IP®system is deployed behind a proxy, you might need to set the **Trust XFF Header** option in the security policy properties. Then the system identifies the location using the address from the XFF header instead of the source IP address.

You can set up a security policy to allow or disallow access to the web application by users in specific countries, areas, or from anonymous proxies.

1. On the Main tab, click **Application Security** > **Policy** > **Geolocation Enforcement**.
2. In the **Current edited policy** list near the top of the screen, verify that the edited security policy is the one you want to work on.
3. In the **Geolocation List** setting, use the move buttons to adjust the lists of allowed and disallowed geolocations. To restrict traffic from anonymous proxies, move **Anonymous Proxy** to the disallowed geolocations list.

   If no geolocations are assigned, the list displays the word **None**. The screen shows the value **N/A** in the list of geolocations for cases where a user is in a location that cannot be identified, for example, if using RFC-1918 addresses or unassigned global addresses.

   

   Tip: *You can approach geolocation enforcement by specifying either which locations you want to disallow or which locations you want to allow.*

4. Click **Save**.
5. In the editing context area, click **Apply Policy** to immediately put the changes into effect.

Now, if a user in a disallowed location attempts to access the web application, the security policy (if in blocking mode) blocks the user and displays the violation `Access from disallowed Geolocation`.

Setting up geolocation enforcement from a request

You can restrict application use in certain geolocations by using the Requests list. This is an easy way to restrict users in a certain country from accessing the web application. By examining illegal request details, you can disallow the locations from which frequent problems are originating.

1. On the Main tab, expand **Application Security** and click **Reporting**.
   The Requests screen opens and shows all illegal requests that have occurred for this security policy.
2. In the Request List, click anywhere on a request.
   The screen displays details about the request including any violations associated with the request, and other details, such as the geolocation.
3. In the Request Details area, next to **Geolocation**, the country is displayed, and if the country is not on the disallowed geolocation list, you see **Disallow this Geolocation**.
   The system asks you to verify that you want to disallow this geolocation. When you verify that you do, the system adds the country to the geolocation disallowed list.
4. Apply the change to the security policy: on the Main tab, click **Policy**, and then click **Apply Policy**.
5. On the menu bar, click **Geolocation Enforcement**.
   The Geolocation Enforcement screen opens, and you can see that the country was added to the disallowed geolocations list.

Now, if a user in a disallowed location attempts to access the web application, the security policy (if in blocking mode) blocks the user and displays the violation `Access from disallowed Geolocation`.

Chapter 7

Configuring Application Security Session Tracking

Topics:

- *Overview: Tracking application security sessions using login pages*

Overview: Tracking application security sessions using login pages

You can track sessions using login pages configured from within Application Security Manager™ (ASM™), or have the policy retrieve the user names from Access Policy Manager®(APM™). This implementation describes how to set up session tracking for a security policy using login pages. The advantage of using session tracking is that you are able to identify the user, session, or IP address that instigated an attack.

When creating login pages for the application, you 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 block a user or session if either generates too many violations.

If you configure session awareness, you can view the user and session information in the application security charts.

Task Summary

Following are tasks for tracking application security sessions using login pages.

Creating login pages
Enforcing login pages
Setting up session tracking
Monitoring user and session information

Creating login pages

In your security policy, you can create a login page to specify a login URL that presents a site that users must pass through to gain access to the web application. The login URL commonly leads to the login page of the web application.

1. On the Main tab, click **Application Security** > **Sessions and Logins**.
2. In the **Current edited policy** list near the top of the screen, verify that the edited security policy is the one you want to work on.
3. Click **Create**.
   The New Login Page screen opens.
4. For the **Login URL** setting, specify a URL that users must pass through to get to the application.
   a) From the list, select the type of URL: **Explicit** or **Wildcard**.
   b) Select either **HTTP** or **HTTPS** based on the type of traffic the web application accepts.
   c) Type an explicit URL or wildcard expression in the box. When you click in the field, the system lists URLs that it has seen, and you can select a URL from the list.
      Type explicit URLs in the format `/login`, and wildcards without the slash, such as `*.php`.
5. From the **Authentication Type** list, select the method the web server uses to authenticate the login URL's credentials with a web user.

Option	Description
HTML Form	The web application uses a form to collect and authenticate user credentials. If using this option, you also need to type the user name and password parameters written in the code of the HTML form.

Option	Description
HTTP Basic Authentication	The user name and password are transmitted in Base64 and stored on the server in plain text.
HTTP Digest Authentication	The web server performs the authentication; user names and passwords are not transmitted over the network, nor are they stored in plain text.
NTLM	Microsoft® LAN Manager authentication (also called Integrated Windows Authentication) does not transmit credentials in plain text, but requires a continuous TCP connection between the server and client.

6. In the Access Validation area, define at least one validation criteria for the login page response. If you define more than one validation criteria, the response must meet all the criteria before the system allows the user to access the application.

   See the online help for definitions of the criteria.
7. Click **Create** to add the login page to the security policy.
   The new login page is added to the login pages list.
8. Add as many login pages as needed for your web application.
9. In the editing context area, click **Apply Policy** to immediately put the changes into effect.

The security policy now has one or more login pages associated with it.

You can now configure how the login pages are enforced, including the authentication URLs, logout URLs, and whether or not the login pages have time limits.

Enforcing login pages

Login enforcement settings prevent forceful browsing by users to restricted parts of the web application by forcing users to pass through one URL (known as the *login URL*) before viewing a different URL (known as the *target URL*). You use the login enforcement settings to specify how the security policy enforces login pages including the expiration time, authenticated URLs, and logout URLs. You can also use authenticated URLs to enforce idle time-outs on applications that are missing this functionality.

1. On the Main tab, click **Application Security** > **Sessions and Logins** > **Login Enforcement**.
   The Login Enforcement screen opens.
2. If you want the login URL to be valid for a limited time, set **Expiration Time** to **Enabled**, and type a value, in seconds.
3. Specify the target URLs that users can access only by way of the login URLs:
   a) For the **Authenticated URLs** setting, type the target URL name in the format `/private.php` (wildcards are allowed).
   b) Click **Add** to add the URL to the list of authenticated URLs.
   c) Add as many authenticated URLs as needed.
4. Optionally, specify the URLs used to log out of the web application:
   a) For the **Logout URLs** setting, type the URL in the format `/logout.html` (explicit URLs only).
   b) Click **Add**.
   c) Add as many logout URLs as needed.
5. Click **Save**.

If you specify authenticated URLs and a user tries to bypass them, the system now issues the `Login URL bypassed` violation. If a user session is idle and exceeds the expiration time, the system now issues the `Login URL expired` violation, and the user can no longer reach the authenticated URLs. For both login

violations, if the enforcement mode is blocking, the system now sends the Login Page Response to the client (see **Application Security** > **Policy** > **Response Pages**).

Setting up session tracking

You can use session tracking to track, enforce, and report on user sessions and IP addresses. To perform tracking, you enable session awareness and indicate how to associate the application user name with the session. You can also determine whether to track violations and perform logging or blocking actions based on the number of violations per user, session, and IP address.

1. On the Main tab, click **Application Security** > **Sessions and Logins** > **Session Tracking**.
   The Session Tracking screen opens.
2. For **Session Awareness**, select the **Enabled** check box.
3. Use the **Application Username** setting to specify the login pages for the application:
   a) From the list, select **Use Login Pages**.
   b) Move the login pages for the application from the Available list to the Selected list.
      If the login page is not listed, click **Add** to create it.
4. For **Track Violations and Perform Actions**, select the **Enabled** check box.
5. In the **Violation Detection Period** field, type the number of seconds that indicates the sliding time period to count violations for violation thresholds. The default is `900` seconds.
6. If you want the system to block all activity for a user, session, or IP address when the number of violations exceeds the threshold, specify one or more of the following settings on the Block All tab.

Note: *For the system to block requests, the security policy Enforcement Mode must be set to blocking (see **Policy** > **Blocking** > **Settings**) and some violations must be set to block.*

Option	Description
Blocked URLs	Specify which URLs to block after the number of violations exceeds the enabled thresholds. To block all URLs, select **Block all URLs**. To block authenticated URLs protected by login pages, select **Block Authenticated URLs**.
Username Threshold	Select **Enable** and specify the number of violations allowed before the system starts to block this user's activity.
Session Threshold	Select **Enable** and specify the number of violations allowed before the system starts to block activity for this HTTP session.
IP Address Threshold	Select **Enable** and specify the number of violations allowed before the system starts to block the activity of this IP address.
Block All Period	Specify how long to block users, sessions, or IP addresses if the number of violations exceeds the threshold. To block the user, session, or IP address indefinitely, click **Infinite**. Otherwise, click **User-defined** and type the number of seconds to block the traffic. The default is `600` seconds.

7. If you want the system to log activity when the number of user, session, or IP address violations exceeds the threshold during the violation detection period, specify one or more of the following settings on the Log All Requests tab.

Option	Description
Username Threshold	Select **Enable** and specify the number of violations allowed before the system starts logging this user's activity for the log all requests period.
Session Threshold	Select **Enable** and specify the number of violations allowed before the system starts logging activity for this HTTP session for the log all requests period.
IP Address Threshold	Select **Enable** and specify the number of violations allowed before the system starts logging the activity of this IP address for the log all requests period.
Log All Requests Period	Specify how long the system should log all requests when any of the enabled thresholds is reached. Type the number of seconds in the field.

8. If you want more tolerant blocking for selected violations, such as those prone to false positives, specify one or more of the following settings on the Delay Blocking tab.

*Note: For the system to block requests, the security policy Enforcement Mode must be set to blocking (see **Policy** > **Blocking** > **Settings**) and the specified violations must be set to block.*

Option	Description
Username Threshold	Select **Enable** and specify the number of violations a user must cause before the system begins blocking this user for the delay blocking period.
Session Threshold	Select **Enable** and specify the number of violations users must cause (during the violation detection period) before the system begins blocking this HTTP session for the delay blocking period.
IP Address Threshold	Select **Enable** and specify the number of violations allowed before the system begins blocking this IP address for the delay blocking period.
Delay Blocking Period	Type the number of seconds that the system should block the user, session, or IP address when any of the enabled thresholds is reached.
Associated Violations	Move the violations for which you want delay blocking from the **Available** list into the **Selected** list. If the selected violations occur, the system does not block traffic until one of the enabled thresholds is reached. At that point, the system blocks traffic causing those violations for the user, session, or IP address, but allows other transactions to pass.

9. Click **Save**.

After you set up session tracking, if any enabled threshold exceeds the number of violations during the detection period, the system starts the configured actions (block all, log all requests, and delay blocking).

Monitoring user and session information

To monitor user and session information, you first need to set up session tracking for the security policy.

You can use the reporting tools in Application Security Manager™ to monitor user and session details, especially when you need to investigate suspicious activity that is occurring with certain users, sessions, or IP addresses.

1. On the Main tab, click **Application Security Reporting**.
   The Requests screen opens and shows all illegal requests that have occurred for this security policy.

2. In the Requests List, click anywhere on a request.
   The screen displays details about the request including any violations associated with the request and other details, such as the source IP address, user name, and session ID.
3. In the General Details area, next to the **Username**, **Source IP Address**, or **Session ID**, click the **Show Session Awareness details** link.
   The screen displays the session awareness action flags that you can set.
4. Update the settings for your selections, as appropriate.

Option	Description
Log All Requests	When set to **Enabled**, the system immediately begins to log activity for the user, session, or IP address and continues for the log activity period (600 seconds by default).
Delay Blocking	When set to **Enabled**, the system is immediately more tolerant of blocking selected violations (configured using **Policy** > **Session Awareness**. The delay lasts for the delay blocking period (600 seconds by default).
Block All	When set to **Enabled**, the system blocks all activity for this user, session, or IP address until further notice.

5. On the menu bar, click **Session Tracking Status**.
   You can see the list of action flags that you previously set. You can also add or release action flags from the Session Awareness screen.
6. To see a graphical view of the violations, from the Charts menu, choose Charts.
   The Charts screen opens where you can view pie charts and bar charts.
7. In the Charts area, next to **View by**, click the viewing criteria for the report you want to see.

   For example, you can view information about illegal requests by user name, session ID, or IP address. Then you can filter the Requests list by the top violator and examine request details for the user, session, or IP address.
8. Examine the charts and review the data you need. Click **Export** to create a PDF of any charts you want to save.

After you set up session tracking, you can monitor the specific requests that cause violations by examining each request and reviewing graphical charts. From the Requests list, you can also set up logging, delay blocking, or block all requests for a specific user, session, or IP address.

Chapter 8

Tracking Application Security Sessions with APM

Topics:

- *Overview: Tracking application security sessions using APM*
- *Prerequisites for setting up session tracking with APM*

Overview: Tracking application security sessions using APM

You can track sessions using login pages configured from within Application Security Manager™ (ASM™), or have the policy retrieve the user names from Access Policy Manager®(APM™). This implementation describes how to set up session tracking for a security policy using APM to verify user credentials. Then, you can set up session awareness from within ASM to identify the user, session, or IP address that instigated an attack.

If you configure session tracking, you can view the user and session information in the application security charts.

Prerequisites for setting up session tracking with APM

In order to set up session tracking from within Application Security Manager™ (ASM™) so that the security policy retrieves the user names from Access Policy Manager ®(APM™), you need to perform basic these system configuration tasks according to the needs of your networking configuration:

- Run the setup utility and create a management IP address.
- License and provision ASM, APM, and Local Traffic Manager™ (LTM™).
- Configure a DNS address (**System** > **Configuration** > **Device** > **DNS**).
- Configure an NTP server (**System** > **Configuration** > **Device** > **NTP**).
- Restart ASM (at the command line, type `tmsh restart /sys service asm`).

If you need more information about basic networking configuration on the BIG-IP® system, refer to the BIG-IP documentation.

Task summary

Use the following tasks to set up application security session tracking with APM authentication integrated.

Creating a VLAN
Creating a self IP address for a VLAN
Creating a local traffic pool for application security
Creating an HTTP class
Creating a virtual server to manage HTTPS traffic
Creating a security policy automatically
Creating an access profile
Configuring an access policy
Adding the access profile to the virtual server
Setting up ASM session tracking with APM
Monitoring user and session information

Creating a VLAN

VLANs represent a collection of hosts that can share network resources, regardless of their physical location on the network.

1. On the Main tab, click **Network** > **VLANs**.
   The VLAN List screen opens.
2. Click **Create**.
   The New VLAN screen opens.
3. In the **Name** field, type a unique name for the VLAN.
4. For the **Interfaces** setting, click an interface number from the **Available** list, and use the Move button to add the selected interface to the **Untagged** list. Repeat this step as necessary.
5. Click **Finished**.
   The screen refreshes, and displays the new VLAN in the list.

Creating a self IP address for a VLAN

Ensure that you have at least one VLAN configured before you create a self IP address.

Self IP addresses enable the BIG-IP® system, and other devices on the network, to route application traffic through the associated VLAN.

1. On the Main tab, click **Network** > **Self IPs**.
   The Self IPs screen opens.
2. Click **Create**.
   The New Self IP screen opens.
3. In the **Name** field, type a unique name for the self IP.
4. In the **IP Address** field, type an IP address.

   This IP address should represent the address space of the VLAN that you specify with the **VLAN/Tunnel** setting.

   The system accepts IP addresses in both the IPv4 and IPv6 formats.
5. In the **Netmask** field, type the network mask for the specified IP address.
6. From the **VLAN/Tunnel** list, select the VLAN to associate with this self IP address. If creating a self IP address for an address space:
   - On the internal network, select the VLAN that is associated with an internal interface or trunk.
   - On the external network, select the VLAN that is associated with an external interface or trunk.
7. Use the default values for all remaining settings.
8. Click **Finished**.
   The screen refreshes, and displays the new self IP address in the list.

The BIG-IP system can now send and receive TCP/IP traffic through the specified VLAN.

Creating a local traffic pool for application security

You can use a local traffic pool with Application Security Manager™ system to forward traffic to the appropriate resources.

Note: *You can optionally create a pool as part of creating a security policy using the Deployment wizard.*

1. On the Main tab, click **Local Traffic** > **Pools**.
   The Pool List screen opens.

2. Click **Create**.
   The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. In the Resources area, for the **New Members** setting, add to the pool the application servers that host the web application:
   a) Type an IP address in the **Address** field.
   b) In the **Service Port** field, type a port number (for example, type `80` for the HTTP service), or select a service name from the list.
   c) Click **Add**.
5. Click **Finished**.

The BIG-IP® system configuration now includes a local traffic pool containing the resources that you want to protect using Application Security Manager™.

Creating an HTTP class

HTTP classes, also called application security classes, can specify which incoming HTTP traffic to route to the Application Security Manager™ for security inspection.

Note: *Creating an HTTP class is optional. When you create a security policy using the Deployment wizard, the system automatically creates an HTTP class with application security enabled.*

1. On the Main tab, click **Local Traffic** > **Profiles** > **Protocol** > **HTTP Class**.
2. Click **Create**.
   The New HTTP Class Profile screen opens.
3. In the **Name** field, type a name for the HTTP class.

   Tip: *This name is also the name of the security policy in Application Security Manager.*
4. From the **Application Security** list, select **Enabled**.
5. Retain the default values for the other settings.
6. Click **Finished**.

The system adds the HTTP class profile, and also creates a security policy with the same name as the class in the Application Security Manager.

Creating a virtual server to manage HTTPS traffic

You can create a virtual server to manage HTTPS traffic.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen displays a list of existing virtual servers.
2. Click the **Create** button.
   The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. For the **Destination** setting, select **Host** and in the **Address** field, type the IP address for the virtual server.
5. In the **Service Port** field, type `443` or select **HTTPS** from the list.
6. From the **Configuration** list, select **Advanced**.

7. From the **HTTP Profile** list, select **http**.
8. For the **SSL Profile (Client)** setting, from the **Available** list, select **clientssl**, and using the Move button, move the name to the **Selected** list.
9. (Optional) From the **SSL Profile (Server)** list, select **serverssl**.

Note: *This setting ensures that there is an SSL connection between the HTTP virtual server and the external HTTPS server.*

10. From the **SNAT Pool** list, select **Auto Map**.
11. In the Resources area, for the **HTTP Class Profiles** setting, move the application security class that you created into the **Enabled** list.
12. From the **Default Pool** list, select the pool that is configured for application security.
13. Click **Finished**.

The HTTPS virtual server appears in the Virtual Server List screen.

Creating a security policy automatically

Before you can create a security policy, you must perform the minimal system configuration tasks including defining a VLAN, a self IP address, and other tasks required according to the needs of your networking environment.

Application Security Manager™ can automatically create a security policy that is tailored to secure your web application.

1. On the Main tab, click **Application Security** > **Security Policies**.
   The Active Policies screen opens.
2. Click the **Create** button.
   The Deployment wizard opens to the Select Local Traffic Deployment Scenario screen.
3. For the **Local Traffic Deployment Scenario** setting, specify a virtual server to use for the security policy.
   - Select **Existing Virtual Server** and click **Next** to use an existing virtual server (as long as it does not have an HTTP Class profile associated with it).
   - Select **New Virtual Server** and click **Next** to create a new virtual server and pool with basic configuration settings.

   The virtual server represents the web application you want to protect. The system automatically creates an HTTP Class with the same name as the virtual server.

   The Configure Local Traffic Settings screen opens.
4. Configure the new or existing virtual server, and click **Next**.
   The Select Deployment Scenario screen opens.
5. For **Deployment Scenario**, select **Create a policy automatically** and click **Next**.
   The Configure Security Policy Properties screen opens.
6. From the **Application Language** list, select the language encoding of the application, or select **Auto detect** and let the system detect the language.

Important: *You cannot change this setting after you have created the security policy.*

7. If the application is not case-sensitive, clear the **Security Policy is case sensitive** check box. Otherwise, leave it selected.

Important: *You cannot change this setting after you have created the security policy.*

8. Click **Next**.
   The Configure Attack Signatures screen opens.
9. To configure attack signatures, move the systems used by your web application from the **Available Systems** list into the **Assigned Systems** list.
   The system adds the attack signatures needed to protect the selected systems.
10. For the **Signature Staging** setting, verify that the default option **Enabled** is selected.

Note: *Because the Real Traffic Policy Builder® begins building the security policy in Blocking mode, it is a good idea to keep signature staging enabled to make sure that false positives do not occur.*

New and updated attack signatures remain in staging for 7 days, and are not enforced (according to the learn, alarm, and block flags) during that time.

11. Click **Next**.
    The Configure Automatic Policy Building screen opens.
12. For **Policy Type**, select an option to determine the security features to include in the policy.

Options	Description
Fundamental	Creates a security policy enforcing HTTP request protocol compliance, evasion techniques, allowed file types (including length checks), attack signatures, the violation Request Length Exceeds Defined Buffer Size, and host names.
Enhanced	Creates a security policy with all the elements of the **Fundamental** policy type; also checks for global parameters (including length checks), cookies, and allowed methods to the security policy.
Comprehensive	Creates a security policy with all the elements of the **Enhanced** policy type; also checks for allowed URLs, meta characters on URLs, meta characters on parameters, URL parameters (instead of global parameters), and dynamic parameters.

    A bulleted list on the screen describes which security features are included in each type.

13. For **Rules**, move the slider to set the Policy Builder learning speed.

Option	Description
Fast	Use for a small number of requests from a small number of sessions; for example, useful for web sites with less traffic. However, there is a greater chance of adding false entities to the security policy.
Medium	Use for a medium number of requests, or if you are not sure about the amount of traffic on the application web site. This is the default setting.
Slow	Use for a large number of requests from many sessions; for example, useful for web sites with lots of traffic. This option creates the most accurate security policy, but takes Policy Builder longer to collect the statistics.

    Based on the option you select, the system sets greater or lesser values for the number of different user sessions, different IP addresses, and length of time before it adds and enforces elements in the security policy.

14. For **Trusted IP Addresses**, select which IP addresses to consider safe:

Options	Description
All	Specifies that the policy trusts all IP addresses. For example, if the traffic is in a corporate lab or preproduction environment where all of the traffic is trusted; the policy is created faster.
Address List	Specifies networks to consider safe. Fill in the **IP Address** and **Netmask** fields, then click **Add**. This option is typically used in a production environment where traffic could come from untrusted sources. The IP Address can be either an IPv4 or an IPv6 address.

If you leave the trusted IP address list empty, the system treats all traffic as untrusted. In general, it takes more untrusted traffic, from different IP addresses, over a longer period of time to build a security policy.

15. If you want the security policy to automatically detect JSON and XML protocols, select the **JSON/XML payload detection** check box.
    This option is available only for the **Enhanced** and **Fundamental** policy types.
    If requests contain legitimate XML or JSON data, the Policy Builder creates content profiles in the security policy according to the data it detects.
16. If you want to display a response page when an AJAX request does not meet the security policy, select the **AJAX blocking response behavior** check box.
17. Click **Next**.
    The Security Policy Configuration Summary opens where you can review the settings to be sure they are correct.
18. Click **Finish** to create the security policy.
    The Automatic Policy Building Status screen opens where you can view the current state of the security policy.

The Policy Builder starts and automatically begins building the security policy by examining the traffic to the web application. The system sets the enforcement mode of the security policy to Blocking, but it does not block requests until the Policy Builder processes sufficient traffic, adds elements to the security policy, and enforces the elements.

Tip: *This is a good point at which to test that you can access the application being protected by the security policy.*

Creating an access profile

You create an access profile to provide the access policy configuration for a virtual server that establishes a secured session.

1. On the Main tab, click **Access Policy** > **Access Profiles**.
   The Access Profiles List screen opens.
2. Click **Create**.
   The New Profile screen opens.
3. Type a name for the access profile.
4. To configure timeout and session settings, select the **Custom** check box.
5. In the **Inactivity Timeout** field, type the number of seconds that should pass before the access policy times out. Type `0` to set no timeout.

If there is no activity (defined by the **Session Update Threshold** and **Session Update Window** settings in the Network Access configuration) between the client and server within the specified threshold time, the system closes the current session.

6. In the **Access Policy Timeout** field, type the number of seconds that should pass before the access profile times out because of inactivity.

   Type `0` to set no timeout.

   You must select the associated **Custom** check box before you can configure this setting.

7. In the **Maximum Session Timeout** field, type the maximum number of seconds the session can exist.

   Type `0` to set no timeout.

   You must select the associated **Custom** check box before you can configure this setting.

8. In the **Max Concurrent Users** field, type the maximum number of users that can use this access profile at the same time.

   Type `0` to set no maximum.

   You must select the associated **Custom** check box before you can configure this setting.

9. In the **Max Sessions Per User** field, type the maximum number of concurrent sessions that one user can start.

   Type `0` to set no maximum.

   You must select the associated **Custom** check box before you can configure this setting.

10. In the **Max In Progress Sessions Per Client IP** field, type the maximum number of concurrent sessions that one client IP address can support.

    Type `0` to set no maximum.

    You must select the associated **Custom** check box before you can configure this setting.

11. Select the **Restrict to Single Client IP** check box to restrict the current session to a single IP address. This setting associates the session ID with the IP address.

    You must select the associated **Custom** check box before you can configure this setting.

    With this setting enabled, upon a request to the session, if the IP address has changed, the request is redirected to a logout page, the session ID is deleted, and a log entry is written to indicate that a session hijacking attempt was detected. If such a redirect is not possible, the request is denied and the same events occur.

12. To configure logout URIs, in the Configurations area, type each logout URI in the **URI** field, and then click **Add**.

13. In the **Logout URI Timeout** field, type the delay in seconds before logout occurs for the customized logout URIs defined in the **Logout URI Include** list.

14. In the SSO across Authentication Domains area, use the **Domain Mode** setting to select whether users log in to a single domain or multiple domains.

15. If you selected **Multiple Domains**, then in the **Primary Authentication URI** field, type the primary URI for authentication.

16. If the policy requires a secure cookie, in the **Cookie Options** area select the **Secure** check box to add the **secure** keyword to the session cookie. If you are configuring an LTM access scenario where you are using an HTTPS virtual server to authenticate the user, and then sending the user to an existing HTTP virtual server to use applications, clear this check box.

17. If the access policy requires a persistent cookie, in the **Cookie Options** area select the **Persistent** check box.

    This sets cookies if the session does not have a webtop. When the session is first established, session cookies are not marked as persistent, but when the first response is sent to the client after the access policy completes successfully, the cookies are marked persistent. Persistent cookies are updated for the

expiration timeout every 60 seconds. The timeout is equal to session inactivity timeout. If the session inactivity timeout is overwritten in the access policy, the overwritten value will be used to set the persistent cookie expiration.

18. From the **SSO Configuration** list, select the SSO configuration.
19. In the **Domain Cookie** field, specify a domain cookie, if required.
20. In the Language Settings area, add and remove accepted languages, and set the default language.

    A browser uses the highest priority accepted language. If no browser language matches the accepted languages list, the browser uses the default language.
21. Click **Finished**.

The access profile appears in the Access Profiles List.

To provide functionality with an access profile, you must configure the access policy. The default access policy for a profile denies all traffic and contains no actions. Click **Edit** in the **Access Policy** column to edit the access policy.

Configuring an access policy

You configure an access policy to provide authentication, endpoint checks, and resources for an access profile.

1. On the Main tab, click **Access Policy** > **Access Profiles**.
   The Access Profiles List screen opens.
2. Click the name of the access policy you want to edit.
3. On the menu bar, click **Access Policy**.
4. For the **Visual Policy Editor** setting, click the **Edit access policy for Profile** ***policy_name*** link.
   The visual policy editor opens the access policy in a separate window or tab.
5. Click the **[+]** sign anywhere in your access profile to add your new policy action item.
   An Add Item window opens, listing Predefined Actions that are grouped by General Purpose, Authentication, and so on.
6. From the General Purpose area, select **Logon Page** and click the **Add Item** button.
   The Logon Page Agent popup screen opens.
7. Click **Save**.
   The Access Policy screen reopens.
8. On the rule branch, click the plus sign **(+)** between **Logon Page** and **Deny**.
9. Set up the appropriate authentication and client-side checks required for application access at your company, and click **Add Item**.
10. Change the Successful rule branch from **Deny** to **Allow** and click the **Save** button.
11. If needed, configure further actions on the successful and fallback rule branches of this access policy item, and save the changes.
12. At the top of the screen, click the **Apply Access Policy** link to apply and activate your changes to this access policy.
13. Click the **Close** button to close the visual policy editor.

Adding the access profile to the virtual server

Before you can perform this task, you need to create an access profile using Access Policy Manager™.

You associate the access profile with the virtual server created for the web application that Application Security Manager™ is protecting.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen displays a list of existing virtual servers.
2. Click the name of the virtual server that manages the network resources for the web application you are securing.
3. In the Access Policy area, from the **Access Profile** list, select the access profile.
4. Click **Update**.

Your access policy is now associated with the virtual server.

Setting up ASM session tracking with APM

You can use session tracking to track, enforce, and report on user sessions and IP addresses. To perform tracking, you enable session awareness and indicate how to associate the application user name with the session.

1. On the Main tab, click **Application Security** > **Sessions and Logins** > **Session Tracking**.
   The Session Tracking screen opens.
2. For **Session Awareness**, select the **Enabled** check box.
3. From the **Application Username** list, select **Use APM Usernames and Session ID**.
4. For **Track Violations and Perform Actions**, select the **Enabled** check box.
5. In the **Violation Detection Period** field, type the number of seconds that indicates the sliding time period to count violations for violation thresholds. The default is `900` seconds.
6. If you want the system to block all activity for a user, session, or IP address when the number of violations exceeds the threshold, specify one or more of the following settings on the Block All tab.

Note: *For the system to block requests, the security policy Enforcement Mode must be set to blocking (see* ***Policy*** > ***Blocking*** > ***Settings****) and some violations must be set to block.*

Option	Description
Blocked URLs	Specify which URLs to block after the number of violations exceeds the enabled thresholds. To block all URLs, select **Block all URLs**. To block authenticated URLs protected by login pages, select **Block Authenticated URLs**.
Username Threshold	Select **Enable** and specify the number of violations allowed before the system starts to block this user's activity.
Session Threshold	Select **Enable** and specify the number of violations allowed before the system starts to block activity for this HTTP session.
IP Address Threshold	Select **Enable** and specify the number of violations allowed before the system starts to block the activity of this IP address.
Block All Period	Specify how long to block users, sessions, or IP addresses if the number of violations exceeds the threshold. To block the user, session, or IP address indefinitely, click **Infinite**. Otherwise, click **User-defined** and type the number of seconds to block the traffic. The default is `600` seconds.

7. If you want the system to log activity when the number of user, session, or IP address violations exceeds the threshold during the violation detection period, specify one or more of the following settings on the Log All Requests tab.

Option	Description
Username Threshold	Select **Enable** and specify the number of violations allowed before the system starts logging this user's activity for the log all requests period.
Session Threshold	Select **Enable** and specify the number of violations allowed before the system starts logging activity for this HTTP session for the log all requests period.
IP Address Threshold	Select **Enable** and specify the number of violations allowed before the system starts logging the activity of this IP address for the log all requests period.
Log All Requests Period	Specify how long the system should log all requests when any of the enabled thresholds is reached. Type the number of seconds in the field.

8. If you want more tolerant blocking for selected violations, such as those prone to false positives, specify one or more of the following settings on the Delay Blocking tab.

Note: *For the system to block requests, the security policy Enforcement Mode must be set to blocking (see* ***Policy*** *>* ***Blocking*** *>* ***Settings****) and the specified violations must be set to block.*

Option	Description
Username Threshold	Select **Enable** and specify the number of violations a user must cause before the system begins blocking this user for the delay blocking period.
Session Threshold	Select **Enable** and specify the number of violations users must cause (during the violation detection period) before the system begins blocking this HTTP session for the delay blocking period.
IP Address Threshold	Select **Enable** and specify the number of violations allowed before the system begins blocking this IP address for the delay blocking period.
Delay Blocking Period	Type the number of seconds that the system should block the user, session, or IP address when any of the enabled thresholds is reached.
Associated Violations	Move the violations for which you want delay blocking from the **Available** list into the **Selected** list. If the selected violations occur, the system does not block traffic until one of the enabled thresholds is reached. At that point, the system blocks traffic causing those violations for the user, session, or IP address, but allows other transactions to pass.

9. Click **Save**.

After you set up session tracking, if any enabled threshold exceeds the number of violations during the detection period, the system starts the configured actions for block all, log all requests, and delay blocking.

Test that you can log in to the web application through the Access Policy Manager™ logon page. You can also test that the security policy works by generating violations and reviewing the application security logs.

Monitoring user and session information

To monitor user and session information, you first need to set up session tracking for the security policy.

You can use the reporting tools in Application Security Manager™ to monitor user and session details, especially when you need to investigate suspicious activity that is occurring with certain users, sessions, or IP addresses.

1. On the Main tab, click **Application Security Reporting**.
   The Requests screen opens and shows all illegal requests that have occurred for this security policy.
2. In the Requests List, click anywhere on a request.
   The screen displays details about the request including any violations associated with the request and other details, such as the source IP address, user name, and session ID.
3. In the General Details area, next to the **Username**, **Source IP Address**, or **Session ID**, click the **Show Session Awareness details** link.
   The screen displays the session awareness action flags that you can set.
4. Update the settings for your selections, as appropriate.

Option	Description
Log All Requests	When set to **Enabled**, the system immediately begins to log activity for the user, session, or IP address and continues for the log activity period (600 seconds by default).
Delay Blocking	When set to **Enabled**, the system is immediately more tolerant of blocking selected violations (configured using **Policy** > **Session Awareness**. The delay lasts for the delay blocking period (600 seconds by default).
Block All	When set to **Enabled**, the system blocks all activity for this user, session, or IP address until further notice.

5. On the menu bar, click **Session Tracking Status**.
   You can see the list of action flags that you previously set. You can also add or release action flags from the Session Awareness screen.
6. To see a graphical view of the violations, from the Charts menu, choose Charts.
   The Charts screen opens where you can view pie charts and bar charts.
7. In the Charts area, next to **View by**, click the viewing criteria for the report you want to see.

   For example, you can view information about illegal requests by user name, session ID, or IP address. Then you can filter the Requests list by the top violator and examine request details for the user, session, or IP address.
8. Examine the charts and review the data you need. Click **Export** to create a PDF of any charts you want to save.

After you set up session tracking, you can monitor the specific requests that cause violations by examining each request and reviewing graphical charts. From the Requests list, you can also set up logging, delay blocking, or block all requests for a specific user, session, or IP address.

Chapter 9

Automatically Creating Security Policies for AJAX Applications

Topics:

- *Application security for applications that use AJAX*
- *Overview: Creating a security policy for applications that use AJAX*
- *Implementation result*

Application security for applications that use AJAX

Application Security Manager™ can protect AJAX applications including those that use JSON or XML for data transfer between the client and the server. If the AJAX application uses XML for data transfer, the security policy requires that an XML profile be associated with a URL or parameter. If the AJAX application uses JSON for data transfer, the security policy requires that a JSON profile be associated with a URL or parameter. If the AJAX application uses HTTP for data transfer, no profile is needed.

You can also set up AJAX blocking response behavior for applications so that if a violation occurs during AJAX-generated traffic, the system displays a message or redirects the application user to another location.

Overview: Creating a security policy for applications that use AJAX

AJAX (Asynchronous JavaScript and XML) applications make requests to the server and send responses to the client formatted using XML or JavaScript Object Notation (JSON). You can create a security policy automatically for applications that use AJAX.

Task Summary

Creating a security policy automatically
Reviewing security policy status

Creating a security policy automatically

Before you can create a security policy, you must perform the minimal system configuration tasks including defining a VLAN, a self IP address, and other tasks required according to the needs of your networking environment.

Application Security Manager™ can automatically create a security policy that is tailored to secure your web application.

1. On the Main tab, click **Application Security** > **Security Policies**.
   The Active Policies screen opens.
2. Click the **Create** button.
   The Deployment wizard opens to the Select Local Traffic Deployment Scenario screen.
3. For the **Local Traffic Deployment Scenario** setting, specify a virtual server to use for the security policy.
   - Select **Existing Virtual Server** and click **Next** to use an existing virtual server (as long as it does not have an HTTP Class profile associated with it).
   - Select **New Virtual Server** and click **Next** to create a new virtual server and pool with basic configuration settings.

   The virtual server represents the web application you want to protect. The system automatically creates an HTTP Class with the same name as the virtual server.

   The Configure Local Traffic Settings screen opens.
4. Configure the new or existing virtual server, and click **Next**.

The Select Deployment Scenario screen opens.

5. For **Deployment Scenario**, select **Create a policy automatically** and click **Next**.
The Configure Security Policy Properties screen opens.
6. From the **Application Language** list, select the language encoding of the application, or select **Auto detect** and let the system detect the language.

> ***Important:*** *You cannot change this setting after you have created the security policy.*

7. If the application is not case-sensitive, clear the **Security Policy is case sensitive** check box. Otherwise, leave it selected.

> ***Important:*** *You cannot change this setting after you have created the security policy.*

8. Click **Next**.
The Configure Attack Signatures screen opens.
9. To configure attack signatures, move the systems used by your web application from the **Available Systems** list into the **Assigned Systems** list.
The system adds the attack signatures needed to protect the selected systems.
10. For the **Signature Staging** setting, verify that the default option **Enabled** is selected.

> ***Note:*** *Because the Real Traffic Policy Builder® begins building the security policy in Blocking mode, it is a good idea to keep signature staging enabled to make sure that false positives do not occur.*

New and updated attack signatures remain in staging for 7 days, and are not enforced (according to the learn, alarm, and block flags) during that time.

11. Click **Next**.
The Configure Automatic Policy Building screen opens.
12. For **Policy Type**, select an option to determine the security features to include in the policy.

Options	Description
Fundamental	Creates a security policy enforcing HTTP request protocol compliance, evasion techniques, allowed file types (including length checks), attack signatures, the violation Request Length Exceeds Defined Buffer Size, and host names.
Enhanced	Creates a security policy with all the elements of the **Fundamental** policy type; also checks for global parameters (including length checks), cookies, and allowed methods to the security policy.
Comprehensive	Creates a security policy with all the elements of the **Enhanced** policy type; also checks for allowed URLs, meta characters on URLs, meta characters on parameters, URL parameters (instead of global parameters), and dynamic parameters.

A bulleted list on the screen describes which security features are included in each type.

13. For **Rules**, move the slider to set the Policy Builder learning speed.

Option	Description
Fast	Use for a small number of requests from a small number of sessions; for example, useful for web sites with less traffic. However, there is a greater chance of adding false entities to the security policy.

Option	Description
Medium	Use for a medium number of requests, or if you are not sure about the amount of traffic on the application web site. This is the default setting.
Slow	Use for a large number of requests from many sessions; for example, useful for web sites with lots of traffic. This option creates the most accurate security policy, but takes Policy Builder longer to collect the statistics.

Based on the option you select, the system sets greater or lesser values for the number of different user sessions, different IP addresses, and length of time before it adds and enforces elements in the security policy.

14. For **Trusted IP Addresses**, select which IP addresses to consider safe:

Options	Description
All	Specifies that the policy trusts all IP addresses. For example, if the traffic is in a corporate lab or preproduction environment where all of the traffic is trusted; the policy is created faster.
Address List	Specifies networks to consider safe. Fill in the **IP Address** and **Netmask** fields, then click **Add**. This option is typically used in a production environment where traffic could come from untrusted sources. The IP Address can be either an IPv4 or an IPv6 address.

If you leave the trusted IP address list empty, the system treats all traffic as untrusted. In general, it takes more untrusted traffic, from different IP addresses, over a longer period of time to build a security policy.

15. If you want the security policy to automatically detect JSON and XML protocols, select the **JSON/XML payload detection** check box.

    This option is available only for the **Enhanced** and **Fundamental** policy types.

    If requests contain legitimate XML or JSON data, the Policy Builder creates content profiles in the security policy according to the data it detects.
16. If you want to display a response page when an AJAX request does not meet the security policy, select the **AJAX blocking response behavior** check box.
17. Click **Next**.
    The Security Policy Configuration Summary opens where you can review the settings to be sure they are correct.
18. Click **Finish** to create the security policy.
    The Automatic Policy Building Status screen opens where you can view the current state of the security policy.

The Policy Builder starts and automatically begins building the security policy by examining the traffic to the web application. The system sets the enforcement mode of the security policy to Blocking, but it does not block requests until the Policy Builder processes sufficient traffic, adds elements to the security policy, and enforces the elements.

Tip: *This is a good point at which to test that you can access the application being protected by the security policy.*

Reviewing security policy status

You can monitor the general progress of the Real Traffic Policy Builder®, see what policy elements the system has learned, and view additional details on the Automatic Policy Building Status screen.

1. On the Main tab, click **Application Security** > **Policy Building** > **Automatic** > **Status**.
   The Automatic Policy Building Status screen opens.
2. In the **Current edited policy** list near the top of the screen, verify that the edited security policy is the one you want to work on.
3. Review any messages in the identification and messages area to learn what is currently happening on the system.

   For example, messages say when the Policy Builder is enabled, when the security policy was last updated, and the number of elements that were learned.
4. Review the status of the Real Traffic Policy Builder.

Options	Description
Enabled	The system is configured to automatically build a security policy, and the Policy Builder is processing traffic.
Disabled	The system is not processing traffic. Check the automatic policy building configuration.
Detecting Language	The system is still configuring the language after analyzing responses to identify the language of the web application. The Policy Builder is enabled, but it cannot add elements to the security policy until the language is set.

5. Examine the **General Progress** of the security policy.

   A progress bar indicates the stability level of the security policy. The progress bar reaches 100% when the policy is stable, no new policy elements need to be added, and time and traffic thresholds have been reached.
6. In the Policy Elements Learned table, review the number of elements that the Policy Builder has analyzed and added to the security policy, and the attributes that need to be updated.

   Tip: *Click the number in the Elements column to see the specific elements that were added.*
7. Optionally, in the Details tree view, click the expand button for any item to learn more about that security policy element, what the system has seen so far, and what it will take to stabilize the element.

When enough traffic from unique sessions occurs over a period of time, the system starts to enforce the file types and other elements in the security policy. When enforced as part of a stable policy, the files types and other elements are removed from the staging list.

Implementation result

The Real Traffic Policy Builder® creates a security policy that can protect applications that use AJAX with JSON or XML for data transfer between the client and the server. The system examines the traffic and creates an appropriate profile. If the application uses XML, the security policy includes one or more XML profiles associated with URLs or parameters. If the application uses JSON, the security policy includes one or more JSON profiles associated with URLs or parameters.

Chapter 10

Adding JSON Support to an Existing Security Policy

Topics:

- *Overview: Adding JSON support to existing security policies*
- *Implementation result*

Overview: Adding JSON support to existing security policies

This implementation describes how to add JSON (JavaScript® Object Notation) support to an existing security policy for an application that uses JSON for data transfer. You create a JSON profile to define what the security policy enforces and considers legal when it detects traffic that contains JSON data.

You can add JSON support to a security policy by completing these tasks.

Task Summary

Creating a JSON profile
Associating a JSON profile with a URL
Associating a JSON profile with a parameter

Creating a JSON profile

Before you can complete this task, you need to have already created a security policy for your application.

This task describes how to create a JSON profile that defines the properties that the security policy enforces for an application sending JSON payloads.

Note: *The system supports JSON in UTF-8 and UTF-16 encoding.*

1. On the Main tab, click **Application Security** > **Content Profiles** > **JSON Profiles**.
2. Click **Create**.
   The Create New JSON Profile screen opens.
3. Type the name of the profile.
4. Adjust the maximum values that define the JSON data for the AJAX application, or use the default values.
5. To change the security policy settings for specific attack signatures for this JSON profile, in the **Global Security Policy Settings** list, select the attack signatures and then move them into the **Overridden Security Policy Settings** list.

   Note: *If no attack signatures are listed in the* ***Global Security Policy Settings*** *list, create the profile, update the attack signatures, then edit the profile.*

6. In the **Overridden Security Policy Settings** list, enable or disable each attack signature as needed:

Options	Description
Enabled	Enforces the attack signature for this JSON profile, although the signature might be disabled in general. The system reports the violation `Attack Signature Detected` when the JSON in a request matches the attack signature.
Disabled	Disables the attack signature for this JSON profile, although the signature might be enabled in general.

7. To allow or disallow specific meta characters in JSON data (and thus override the global meta character settings), click **Value Meta Characters**.

- Select the **Check characters** check box, if it is not already selected.
- Move any meta characters that you want allow or disallow from the **Global Security Policy Settings** list into the **Overridden Security Policy Settings** list.
- In the **Overridden Security Policy Settings** list, change the meta character state to **Allow** or **Disallow**.

8. To mask sensitive JSON data (replacing it with asterisks), click **Sensitive Data Configuration**.
   - In the **Element Name** field, type the JSON element whose values you want the system to consider sensitive.
   - Click **Add**.

   Important: *If the JSON data causes violations and the system stops parsing the JSON part way through a transaction, the system masks only the sensitive data that was fully parsed.*

   Add any other elements that could contain sensitive data that you want to mask.
9. Click **Create**.
   The system creates the profile and displays it in the JSON Profiles list.

This creates a JSON profile which does not affect the security policy until you associate the profile with a URL or parameter.

Next, you need to associate the JSON profile with any URLs or parameters that might include JSON data.

Associating a JSON profile with a URL

Before you can associate a JSON profile with a URL, you need to have created a security policy with policy elements including application URLs, and the JSON profile.

You can associate a JSON profile with one or more explicit or wildcard URLs.

1. On the Main tab, click **Application Security** > **URLs**.
2. In the **Current edited policy** list near the top of the screen, verify that the edited security policy is the one you want to work on.
3. From the Allowed URLs List, click the name of a URL that might contain JSON data.
   The Allowed URL Properties screen opens.
4. Next to **Allowed URL Properties**, select **Advanced**.
5. For the **Header-Based Content Profiles** setting, in the **Request Header Name** field, type the explicit string or header name that defines when the request is treated as the **Parsed As** type; for example, `content-type`. This field is not case sensitive.

   Note: *If the URL always contains JSON data, just change the default header-based content profile to be **Parsed As JSON**, then you do not have to specify the header name and value.*
6. For the **Header-Based Content Profiles** setting, in the **Request Header Value** field, type the wildcard (including *, ?, or [chars]) for the header value that must be matched in the **Request Header Name** field; for example, `*json*`. This field is case sensitive.
7. From the **Parsed As** list, select **JSON**.
8. From the **Profile Name** list, select the JSON profile appropriate for this URL.
9. Click **Add**.
   Add as many header types as you need to secure this URL, clicking **Add** after specifying each one.
10. To override the global meta character settings for this URL, adjust the meta character policy settings:

- Select the **Check characters on this URL** check box, if it is not already selected.
- Move any meta characters that you want allow or disallow from the **Global Security Policy Settings** list into the **Overridden Security Policy Settings** list.
- In the **Overridden Security Policy Settings** list, change the meta character state to **Allow** or **Disallow**.

11. Click **Update**.
12. To activate the updated security policy, on the top right of the screen, click **Apply Policy**, then click **OK** to confirm.

The JSON profile is associated with the URL.

Continue to associate JSON profiles with any URLs in the application that might contain JSON data.

Associating a JSON profile with a parameter

You need to have created a security policy with policy elements including parameters and a JSON profile.

You can associate a JSON profile with a parameter.

1. On the Main tab, click **Application Security** > **Parameters**.
2. In the **Current edited policy** list near the top of the screen, verify that the edited security policy is the one you want to work on.
3. From the Parameters List, click the name of a parameter to which to assign a JSON profile.
   The Parameter Properties screen opens.
4. For the **Parameter Value Type** setting, select **JSON value**.
5. From the **JSON Profile** list, select the JSON profile to use for this parameter.
6. Click **Update**.
   The system associates the JSON profile with the parameter.
7. To activate the updated security policy, on the top right of the screen, click **Apply Policy**, then click **OK** to confirm.

Continue to associate JSON profiles with any parameters in the application that might contain JSON data.

Implementation result

You have manually added JSON support to the active security policy. The policy can now secure applications that use JSON for data transfer between the client and the server. If web application traffic includes JSON data, the system checks that it meets the requirements that you specified in the JSON profile.

Chapter 11

Adding AJAX Blocking Response Behavior to a Security Policy

Topics:

- *Overview: Adding AJAX blocking and login response behavior*
- *Configuring the blocking response for AJAX applications*

Overview: Adding AJAX blocking and login response behavior

Normal policy blocking and login response behavior could interfere with applications that use AJAX. If you want to display a message or redirect traffic without interfering with the user experience while browsing to an AJAX-featured web application, you need to enable AJAX blocking behavior (JavaScript injection). You can implement blocking and login response behavior for applications that use AJAX with JSON or XML for data transfer.

Important: *You can implement AJAX blocking behavior only for applications developed using one of the following frameworks:*

- *Microsoft® ASP.NET*
- *jQuery*
- *Prototype®*
- *MooTools*

By default, if you enable AJAX blocking behavior, when an AJAX request results in a violation that is set to **Block**, Application Security Manager performs the default AJAX response page action. The system presents a login response if the application user sends an AJAX request that attempts to directly access a URL that should only be accessed after logging in.

Note: *Enabling AJAX blocking behavior has performance implications.*

Configuring the blocking response for AJAX applications

Before you can complete this task, you need to have already created a security policy for your web application. The application needs to have been developed using ASP.NET, jQuery, Prototype®, or MooTools to use AJAX blocking behavior.

This task describes how to configure blocking response and login response pages for web applications that use AJAX.

1. On the Main tab, click **Application Security** > **Policy** > **Response Pages**.
2. In the **Current edited policy** list near the top of the screen, verify that the edited security policy is the one you want to work on.
3. Click **AJAX Response Page**.
4. Select the **Enable AJAX blocking behavior (JavaScript injection)** check box.
   The system displays the default blocking response and login response actions for AJAX.
5. For the **Default Response Page action**, select the type of response you want the application user to receive when they are blocked from the application:
   - **Custom Response** lets you specify HTML text or upload a file to use as a replacement for the frame or browser page that generated the AJAX request. Include the text, then click **Show** to preview the response.
   - **Popup message** displays text in a popup window (default text is included).

- **Redirect URL** redirects the user to the URL you specify. You can also include the support ID. For example:
  `http://www.example.com/blocking_page.php?support_id=<%TS.request.ID()%>`.

6. For the **Login Page Response action**, select the type of response (types are the same as for default response page in Step 5).
7. Click **Save**.
8. To activate the updated security policy, on the top right of the screen, click **Apply Policy**, then click **OK** to confirm.

When the enforcement mode of the security policy is set to blocking and a request triggers a violation (that is set to block), the system displays the AJAX blocking response according to the action set. If a login violation occurs when requesting the login URL, the system sends a login response page, or redirects the user.

Index

A

access policy
    configuring 65
access profile
    adding to virtual server 65
    creating 63
AJAX applications
    and login response page 80
    configuring blocking policy 80
    configuring blocking response 80
    configuring login response 80
    overview 70
    securing JSON data 70
    using JSON 76
anonymous proxy, disallowing 48
application security class, See HTTP class
application security synchronization
    considerations 15, 23, 31
Automatic Sync
    enabling 18, 34
awareness, user and session 52, 58

B

blocking policy
    configuring for AJAX 80
blocking response page
    configuring for AJAX 80

C

config sync address
    specifying 15, 23, 32
configuration synchronization
    configuring for ASM 14, 22, 31
    performing basic networking 15, 23, 31
    result for synchronizing multiple ASM systems 20
    result for two ASM systems 27
    result of synchronizing ASM devices 36
    synchronizing ASM devices 22, 30
    syncing to group 17, 25, 33

D

device discovery
    for device trust 16, 24, 32
device groups
    creating 17, 18, 25, 33, 34
    enabling ASM 19, 35
    enabling ASM synchronization 26
    synchronizing ASM-enabled 19, 27, 35
device management
    considerations 15, 23, 31
    setting up application security synchronization 14, 22, 30, 31
device trust
    establishing 16, 24, 32
disaster recovery
    result of synchronizing ASM devices 36
    synchronizing ASM devices 30

E

enforcement, login 53, 54

F

failover IP addresses
    specifying 18, 26, 34

G

geolocation
    enforcing 48
    enforcing from Requests list 49
    overview 48

H

HTTP class
    creating 60
HTTPS traffic
    creating virtual servers for 60

I

IP address exceptions
    creating 44
    deleting 45
    overview 44
    updating 45
IP address intelligence
    categories 42
    downloading the database 38
    enabling 38
    logging information 40
    rejecting bad requests 41
IP address intelligence blocking
    overview 38
    reviewing statistics 40
    setting up 39
IP intelligence database 42
iprep.autoupdate command 38
IP reputation blocking
    overview 38

J

JSON data
    masking 76
JSON profile
    creating 76

L

local traffic pools
    creating 59
local trust domain
    and device groups 17, 18, 25, 33, 34
local trust domains
    defined 16, 24, 32
login enforcement 53
login pages
    creating 52
    enforcing 53
login response page
    configuring for AJAX 80
log IP address intelligence information 40

M

malicious IP address 39

N

network failover
    configuring 17, 25, 33

P

parameters
    associating with JSON profiles 78
policy, See security policy.
pools
    creating local traffic 59
profiles
    associating parameters with JSON 78
    associating URLs with JSON 77
    creating JSON 76

S

security policy
    creating automatically 61, 70
    reviewing status 72
    synchronizing ASM 19, 27, 35
security policy synchronization
    configuring 14, 22, 30, 31
self IP addresses
    and VLANs 59
    creating 59
sensitive data
    masking JSON 76
session awareness
    enabling 54
    enabling with APM 66
sessions, monitoring 55, 67
session tracking
    enabling 54
    enabling with APM 66
    overview 52, 58
    prerequisites 58
Sync-Failover device groups
    creating 17, 25, 33
synchronization
    configuring for ASM 22, 30
    enabling for ASM 26
    result for configuring multiple ASM systems 20
    result for configuring two ASM systems 27
Sync-Only device groups
    creating 18, 34
    enabling ASM 19, 35

T

trust domains, See local trust domains

U

URLs
    associating with JSON profiles 77
user awareness
    overview 52, 58
user information, monitoring 55, 67

V

virtual servers
    creating for HTTPS traffic 60
VLANs
    and self IP addresses 59
    creating 58

X

x509 certificates
    for device trust 16, 24, 32