

Check Point R80.10

PART OF CHECK POINT INFINITY

Security Management Architecture Overview

**WELCOME
TO THE FUTURE OF
CYBER SECURITY**

CLOUD • MOBILE • THREAT PREVENTION



Check Point
SOFTWARE TECHNOLOGIES LTD.

INFINITY

THE SECURITY MANAGEMENT ARCHITECTURE OF THE FUTURE

The exponential growth of data from various sources like network, cloud, mobile, and virtual systems constantly introduces new threats. Organizations need to quickly adapt and protect themselves, requiring more powerful, high-performance security.

Check Point Infinity, the first consolidated security across networks, cloud and mobile, provides the highest level of threat prevention against both known and unknown targeted attacks to keep you protected now and in the future.

Check Point R80.10 , part of Check Point Infinity, takes security management to new levels, merging security leadership with Ease of use –into a unified console for a full-spectrum visibility. The R80.10 provides organizations improved performance of powerful security features, with a new user-friendly front-end.


R80.10 SmartConsole is an integrated security management solution which includes policy, logging, monitoring, event correlation and reporting – all in a single system which enables administrators to easily identify security risks across the organization.

The R80.10 introduces greater efficiency, allowing multiple administrators to work simultaneously on the same management server, or even the same policy with no conflicts. It also enables delegation of routine tasks, so teams can better focus on monitoring and incident response.

A unified policy enables organizations to translate their security definitions into a simple set of rules, which then streamline policy administration and enforcement throughout the organization. Policy layers provide the ability to separate the policy into independent segments, which can be independently managed and automated.

The R80.10 improves overall productivity with state-of-the-art orchestration capabilities. Automation API orchestrates workflows, enabling security alignment with IT processes and systems.

MAIN COMPONENTS OF THE R80.10 SECURITY MANAGEMENT SOLUTION



SMARTCONSOLE

SmartConsole is the new unified application of Check Point R80.10 Security Management. The new SmartConsole provides a consolidated solution to manage the security of your organization:


- Security Policy Management
- Log Analysis
- System Health Monitoring
- Multi Domain Management

R80.10 SmartConsole provides numerous advantages. Security policy changes and logs of a modified rule are one click away. You can easily navigate from an item in a log to the relevant policy. Built-in support for multi-language and accessibility features, such as color configuration and keyboard navigation, make the new SmartConsole accessible to a wider audience. Efficient communication with the server loads views as you navigate.

Notes: The communication with the management server is based on web services on top of port 19009.

Some blades use components of the former SmartDashboard views. Those components communicate with the management server (FWM) using the CPMI API on port 18190.

R80.10 SECURITY MANAGEMENT PROCESSES



PROCESS	DESCRIPTION	INSTANCES	PORT	DEBUG OUTPUT PATH
CPM	New management server Serves requests from SmartConsole	One instance	9009 for localhost communication 19009 for remote communication	\$FWDIR/log/cpm.elg*
postgres	Database of new management server	Multiple (pool)	5432	\$FWDIR/log/postgres.elg
FWM	Part of server for specific requests from SmartConsole	One per Domain	18190 CPMI 6666 for connections of CPM to Domain Servers 6667 for connections of CPM to MDS	\$FWDIR/log/fwm.elg
API Server (jetty)	Automation server	One instance for all domains	Random high port. APIs exposed according to Gaia's web-UI port	\$MDS_FWDIR/log/api.elg
RFL	Logging core process	One instance	8218 for localhost communication	\$RTDIR/log/RFL.log* \$RTDIR/log/journal.csv*
SOLR	Logging database	One instance	8210 for localhost communication 8211 for remote communication	\$RTDIR/log/solr.log*
SMARTVIEW	Web UI for logs, reports and views	One instance	SmartView exposed according to Gaia's web-UI port	\$RTDIR/log/smartview.log* \$RTDIR/log/smartview-service.log*
INDEXER	Logging indexing process	One per Domain	Random high port	\$INDEXERDIR/log/log_indexer.elg*

PROCESS	DESCRIPTION	INSTANCES	PORT	DEBUG OUTPUT PATH
SMARTLOG_SERVER	SmartLog application server for logs	One per Domain	Random high port	\$SMARTLOGDIR/log/smartlog_server.elg*
SmartEvent (cpsemd)	SmartEvent application server for logs, reports and views	One instance	18266 for remote communication	\$RTDIR/log/cpsemd.elg*
SmartEvent CU (cpsead)	SmartEvent Correlation Unit server	One instance		\$RTDIR/log/cpsead.elg*
DBSync	Logging database synchronization process	One instance		Connected to R77 management: \$RTDIR/log/dbsync.log* Connected to R80.10 management: \$MDS_FWDIR/log/dbsync.elg*

For more details about ports used by Check Point processes , see [sk52421](#).

R80.10 SECURITY MANAGEMENT DATABASE


The server is designed to simplify the daily work of your security teams.

- Database sessions - Enables the collaboration of multiple admins without conflicts.
- Database revisions - Allows you to see history of changes and improves the performance of many operations, such as Policy Installation and High Availability.
- Database domains - A solution used both in Security Management and Multi Domain Management that improves the performance of global policy assignment, Threat Prevention, and Application Control updates.

Database Sessions

With R80.10, multiple administrators can work on SmartConsole on the same domain, with the same policies, and at the same time. To avoid configuration conflicts, all work is done in sessions, as illustrated in the diagram below.

Sessions



Every session is private. Changes cannot be seen by other administrators until the changes are published. To keep the sessions private, objects are locked when an administrator is modifying it. Other administrators will only see that the object is locked; they will not be able to change it. When an object is locked, the name of the administrator working on that object is displayed. This helps administrators coordinate work on shared resources.

When all the modifications have been completed, the administrator publishes the session. Only then will the changes become public and visible to all other administrators. Only public data can be installed on gateways.

All modifications are saved instantly in the management server database. If there is an accidental disconnection, no work is lost. Administrators can discard changes during a session, and they can open a new session.

If required, administrators with the proper permissions can take over other administrator's sessions.

Database Domains


In R80.10, the management configuration is stored in the PostgreSQL database. This data is segmented into multiple database domains.

Consider the difference between the following terms:

- *Database Domains* - The segments in the postgres database in which data is stored, both in Security Management and Multi Domain Management.
- *Multi Domain Domains* – Logical Domains created by the user in Multi Domain Management (previously known as Provider-1 Customers) and used to manage different parts of the organization network.

These Domains are implemented with database domains.

Database Domain Types



The User Domain stores the configuration of information which is modified by the user, such as network objects and security policies.

The new R80.10 system comes with these Data Domains:

- *Default Data Domain* - Contains default network objects and services.
- *Threat Prevention Domain* - Stores Threat Prevention updates generated by Check Point.
- *Application Control Domain* - Stores Application Control updates generated by Check Point.

The *Data Domains* are modified by updates generated by Check Point only.

The *System Domain* stores the configuration of administrators, permission profiles and management settings.

The *Log Domain* stores the configuration for logs - data which is generated automatically and queries requested by the user.

The *Global Domain* stores the configuration of Global Policies and Global Objects (In use only in Multi Domain Management).

Peer Domains

Security policy configuration requires the use of data from numerous domains and some domains must recognize and share their data with other domains. These domains are *Peer Domains* to each other.

In a Multi Domain environment, Global Domain is peered to the User Domains. The peered domain structure eliminates the need to copy all of the Global objects to the User Domain. This results in improved performance and scalability.

Domains in Multi Domain Environments

In Multi-Domain, each Domain (formerly known as Customer) is implemented by a separate database domain of type User Domain. Other database domains - System Domain, Global Domain, Threat Prevention data, Application Control data and default services - each have their own domains in both Multi-Domain and Security Management environments.

The partitioning in the same database provides multiple benefits:

- *Separated User Domains.* Providing a database domain for each Domain ensures full separation from other Domains. User Domains are not peered to each other. Therefore, no data can be shared between them.
- *Enhanced global policy assignment.* Before R80.10, the global policy assignment copied all the global objects to the domain's database (\$FWDIR/ conf of the domain). In R80.10, the assignment operation sets the User Domain to detect the newer revision of the Global Domain.
- *Enhanced Threat Prevention updates in Multi Domain Environments.* When the user updates the Threat Prevention domain, it is updated only on the domain that it is peered with. The user has the choice in which domain the update revision should be applied. If reverting Threat Prevention, the user can revert just this domain and not the whole database.
- *Multi-Domain Management Server level data is stored in the System Domain.* It stores administrators, permissions, trusted clients and also the configuration of servers, Domains and Domain Servers, and allows managing this data on all MDS servers in parallel.

Database Revisions

In pre-R80.10 versions, revisions were saved on the server as a backup, to restore in case of disaster. Each revision replicates the management database completely.


The new R80.10 architecture has built-in revisions. Every time you publish, an additional revision is created and automatically stored. Each revision establishes a new baseline in the database and contains only the changes from the previous revision. Thus, it consumes less disk space and is created faster than before.

This solution improves performance and many of the new management capabilities are based on it.

- Faster policy verification using the difference between installed revisions.
- More efficient Management High Availability based on the changes made since the previous synchronization.

The following diagram illustrates the database revisions over time:


Revisions



Revisions and Peer Domains

Database revisions cross peer domains.

Revisions and Peer Domains



In this example:

The Assign Global Policy operation created revision 4 on the User Domain and peered it to revision 1 of the Global Domain.

The publish which created revision 2 on the Global Domain was not visible on the User Domain until the Reassign Global Policy. This reassign operation updated the User Domain to look at the latest revision (2) of the Global Domain.

The same technique is done also for Threat Prevention updates and Application Control updates. Each update creates a new revision in the relevant Data Domain (Threat Prevention Domain or Application Control Domain).

The User Domain is peered to a revision of the Data Domain.

In Multi-Domain Management, each User Domain can be peered to a different revision of the Global Domain or Data Domain.

Managing crisis cases using revisions:

Case:	Connectivity/security problem after changes on the policy are performed and installed.
Solution:	1. Go to Security Policies > Installation History. 2. In Policy Installation History, choose the last known good version and click Install specific version. 3. After a gateway is safely re-installed, query the audits and changes made in each revision on the management.
Case:	Network problem after downloading Threat Prevention update and installing it on gateways.
Solution:	From Threat Prevention > Updates , choose an update known to be good.

More cases:

- Full environment restore to a certain point in time is needed.
The recommended practice is to use **Restore from Backup**, but this will cause a loss of all the work done after that point.
For more details, see the R80.10 Installation and Upgrade Guide.
- **Revert Policy** is available to revert back to a previous state. It will revert the rulebase structure, but not the objects used in the rulebase.

Administrator Permissions

Check Point R80.10 Security Management offers Role Based Access Control for tasks across the security team members. You define which permissions to grant your team members.

The pre-defined administrator roles are:

- Access Control Administrator - Manages only the Access security (and not Threat Prevention security).
- Threat Prevention Administrator - Manages only the Threat Prevention security (and not Access Control security).
- Layer-specific Administrator - Delegates according to the responsibility of each team member.
- Auditor - Read-only connection.
- Device Management Administrator - Manages device deployment, licensing and provisioning.
- Super User - Manages all of the above aspects.

You can also assign granular permissions, to define a customized role.

Management High Availability Deployment

Compared to R77, R80.10 provides an enhanced solution for High Availability.

Utilizing the built-in revisions technology, R80.10 allows the High Availability process to synchronize only the changes done since the previous synchronization.

This provides:

- Real-time updates between management peers
- Minimal impact on the management server resources.

Install Policy

In R80.10, Policy Verification underwent a major improvement.

From our investigations, we found that in large rulebases, a significant part of the install policy time is taken by the verification phase. Using the R80.10 built-in revision technology, we were able to significantly improve the install policy time for large rulebases. The new solution verifies only the rules changed since the last policy installation.

For example: On tests we performed on a certain large rulebase, on-going policy installations of R80.10 were six times faster than a similar policy installation, with similar changes, in R77.30.

Note: First policy install after upgrading to R80.10 might still take more time, because it creates the first baseline. Consequent installations will benefit from the mentioned improvement.

Starting with R80.10, on R80.10 gateways, the Access Control policy can be installed separately from the Threat Prevention policy.

Benefits include:

- Improved performance when installing Access/Threat only.
- Role segregation support for organizations in which Threat is managed separately from Access. One administrator will not affect the other team's policy.

AUTOMATION

R80.10 provides a command line interface and REST APIs, which provide new ways to interact with Check Point management servers. For admins with expertise, the new APIs offer higher productivity and easier daily management.

The automation server is installed as an integral part of the R80.10 management server. By default, the API server is active on management servers with 4 GB of RAM (or more) and on stand-alone servers with 8 GB of RAM (or more).


If your management server has less RAM, and as a result, the API server is not activated, you can activate it through SmartConsole.

Open **Manage & Settings > Blades > Management API > Advanced Settings**. Select the option you want:

Automatic start – Select to start the automation server after the management server restarts.

Access settings – Enter which IP addresses are allowed to send API commands to the automation server. By default, the automation server will not accept requests unless they originate from the management server itself.

Automation Architecture



The Automation server communicates with the management server (CPM) in the same way SmartConsole does. This architecture ensures that:

- The same validation errors and warnings that SmartConsole shows will also be presented when using an automation session.
- The same audit logs generated when using SmartConsole will also be generated when using an automation session.
- The same permission profile that controls what an admin is allowed to view or modify when using SmartConsole will also be enforced when using an automation session.

In a Multi-Domain environment, there is just one automation server and it detects all the IP addresses of the Multi-Domain Management Server. To control the settings of the automation server using the GUI, launch SmartConsole and log in to the MDS Domain. These settings will apply to all domains. If you log in to a specific domain, call the “login” command, and enter the domain name as an argument.

The automation server listens to a random high port. Every time the management server restarts, the automation server port can change.

The Automation server receives its request from the Apache server process on the Gaia machine. If you change the Apache server’s port from 443 to a different port, this will affect both the Automation server and the Gaia web-UI interface.

When you call the login API using a REST API or command line, a new session is generated (multiple active sessions for the same admin are allowed). This ensures that when a script publishes its changes, the script is publishing its changes only, and not changes that were caused by other scripts that run at the same time.

mgmt_cli

The *mgmt_cli* tool is portable and can run on any Linux or Windows machine. A Linux version of the *mgmt_cli* command line tool is included in all R80.10 Gaia installations. A windows version (*mgmt_cli.exe*) is in the R80.10 SmartConsole installation.

The *mgmt_cli* tool transforms the arguments it receives to REST API calls. Any REST API request can be expressed as a *mgmt_cli* command, and any *mgmt_cli* command can be expressed as one or more REST API requests.

When logged in to a management server as a “root” user, you can run *mgmt_cli* commands without providing credentials, using a built-in user with “super-admin” rights. To use this option add `--root true` to your *mgmt_cli* command.

Every *mgmt_cli* command can accept its arguments by reading them from the command line or from a csv file. For example, these two options are equivalent:

Option 1:

```
mgmt_cli add host name host1 ip-address 192.168.0.1 color blue --root true
```

```
mgmt_cli add host name host2 ip-address 192.168.0.2 color red --root true
```

Option 2:

```
mgmt_cli add host --batch hosts_data.csv --root true
```

where **hosts.data.csv**:

```
name, ip-address, color  
  
host1, 192.168.0.1, blue  
  
host2, 192.168.0.2, red
```

You can control the output format of `mgmt_cli` using the “`-format`” option. By default, the tool will use a text format that follows YAML conventions, but if you need to parse the output of `mgmt_cli` and extract the value of a specific field, we recommend using JSON as the output format. JSON lets you use standard tools, such as ‘jq’ (which is built-in to Gaia) to parse the output in a way that is more efficient than using “standard” text parsing tools such as “grep”, “cut”, “awk”, etc.

Troubleshooting and diagnostics

To check the status of the automation server, in expert-mode on the management server, run:

```
api status
```

There is no need to turn-on a debug flag. The automation server saves debug data by default to this file: `$MDS_FWDIR/log/api.elg`

Automation FAQ

- Can I start make changes using a script (without publishing) and review the changes in SmartConsole?

Yes. Launch SmartConsole, go to **Manage and Settings > Sessions**, right-click the automation session you want to review, and select **Takeover**.

- What about my existing dbedit scripts?

dbedit (or GuiDbEdit) uses the cpmi protocol which is gradually being replaced by the new R80.10 automation architecture. cpmi clients are still supported in R80.10, but there are some functionalities that cannot be managed by cpmi anymore. For example, the Access and Threat policies do not have a cpmi representation. They can be managed only by the new `mgmt_cli` and not by cpmi clients. There are still many tables that have an inner cpmi representation (for example, network objects, services, servers, and global properties) and can still be managed using cpmi.

- Where can I find more information?

You can find the full API reference, samples and more information in <https://sc1.checkpoint.com/documents/R80/APIs/> and in the Exchange Point community > Developer Network section: <https://community.checkpoint.com/community/developers/developers-forum/>

- What do to if I still have questions?

Please post you question on <https://community.checkpoint.com> under the Developer Network forum.

LOGGING, EVENT MANAGEMENT, REPORTING, AND MONITORING

With R80.10, logging, event management, reporting, and monitoring are more tightly integrated than ever before. Security data and trends are easy to understand at a glance, with widgets and chart templates that optimize visual display. Logs are now integrated with the policy rules. To access logs associated with a specific rule, all you have to do is click that rule. Free-text search lets you enter specific search terms to retrieve results from millions of logs in seconds.


One-click exploration makes it easy to move from a high-level overview, to specific event details, such as the type of attack, timeline, application type, and source. After you investigate an event, it is easy to act on it. Depending on the severity of the event, you can ignore it, act on it later, block it immediately, or toggle over to the rules associated with the event to refine your policy. You can send individualized reports to your manager or auditors that show only the content relevant to each reader.

In R80.10, SmartReporter and SmartEvent functionality is integrated into SmartConsole. With rich and customizable views and reports, R80.10 introduces a new experience for log and event monitoring.

The new views are available from the SmartConsole and from SmartView WebUI:

- **SmartConsole > Logs & Monitor**
- **SmartView Web Application:** browse to: <https://<Server IP>/smartview/>
where Server IP is IP address of the Security Management Server or SmartEvent server.

Logging Data Flow



LAYERS

R80.10 brings the Unified Rulebase and Layers as new ways to organize security policies. A policy can have one or more layers as its building blocks. You can use combinations of Ordered Layers, Inline Layers, and Domain Layers (in Multi-Domain environments).

In pre-R80.10 systems, each setting had its own rulebase. Now settings are unified. The Access Control Rulebase includes Firewall, Application Control, Data, and Mobile Access. The Threat Prevention Rulebase includes IPS, Anti-Bot, Anti-Virus, and Sandblast.

Using Layers gives you various benefits:

- Delegation - Assign different layers to different administrators.
- Modularity - Break a large policy into smaller chunks to improve orientation and simplify maintenance.
- Reuse - Mark layers as shared and reuse in multiple policies, instead of cloning.

Backward Compatibility

R80.10 SmartConsole automatically organizes pre-R80.10 Gateway policies with ordered layers.

The administrator will enjoy the new layers user experience immediately. On policy installation, the policy is converted to the pre-R80.10 format.

Examples:


Gateways that have the Firewall and Application Control blades enabled will have their policies split into two ordered layers: Network and Applications.

Gateways that have the IPS and Threat Emulation blades enabled will have their policies split into two ordered layers: IPS and Threat Prevention.

Enforcement

Ordered Layers are enforced by the gateway. When there is a match on the first layer, the gateway moves to evaluate the rules in the second layer, and continues on to each layer.

With R80.10 gateways, you can modify the default ordered layer structure. You can add or remove layers as is necessary.



✓ Ordered Layers - Best Practices

- Inspect additional content - First layer handles networking, the next layer handles apps and the third layer handles data.
 - Organize a high-level policy with must-drop rules before any additional inspection.
 - Apply drop filters as the last layers of your policy.
-

Inline Layers

Inline Layers are sub-policies. An Inline Layer contains a number of rules, and an action on the set of rules as a whole: Accept, Drop, or Evaluate.



Inline Layers Best Practices

- Inspect additional content - Add application and content awareness only after the parent rule is matched, or reuse the application control layer as the inline layer, to all Internet-facing rules.
- Segment-based policy - Parent rules can represent the scope of the organization (for example, "DMZ to Internet" and "Sales to Data Center"). The inline layer can be the actual policy for those scopes.



Be aware that a cleanup rule in the format of "any,any,ACTION" is automatically created at the bottom of every layer, and inline layer. You can control the action of the rule.

LAYERS IN MULTI-DOMAIN POLICIES

Multi-Domain environments in R80.10 use Layers to segregate different parts of the rulebase for different permission profiles. Inside a global policy layer, a placeholder for domain rules appears. It represents the place in which the domain rules will be applied upon global policy assignment. Global rules can be defined above and below the placeholder.

Upon global policy assignment, all of the domain's policies are updated with the global rules. The placeholder from the global domain is seen as a "parent rule for domain policy". Its action is "domain layer", and it has a "domain layer" inside with all of the local domain rules. The domain administrator can select a different domain layer or choose not to have any domain layer.

When the gateway evaluates the rules in the local policy, if there is no match for the global rules at the top of the rulebase, it starts to evaluate the rules from the domain layer. If there is still no match for those rules, the global rules that were created below the domain layer are evaluated.

This is how a global layer looks when it is defined at the Global Domain:

The screenshot shows the Check Point SmartConsole interface with the 'internal-gw' policy selected. The left sidebar shows the 'Access Control' section with 'Global-applications' highlighted. The main table lists five rules:

No.	Name	Source	Destination	VPN	Services & Applications	Action	Track	Install On
1	ask for purpose of browsing	Internal-users	* Any	* Any	Facebook Games, Facebook File Shari..., File Storage and Sh..., P2P File Sharing, VoIP, Supports VoIP	Ask, Ask-for-reason, Once a day, Per applicatio...	None	* Policy Targets
2	block high risk categories	* Any	* Any	* Any	Critical Risk, High Risk, Illegal / Questionable, Child Abuse, Sex	Drop, Website-not-a...	None	* Policy Targets
3	Placeholder for domain rules							
4	report suspicious categories	Negated content-checkers	* Any	* Any	Alcohol, Illegal Drugs, Sex Education	Accept	Network Log	* Policy Targets
5	cleanup rule	* Any	* Any	* Any	* Any	Accept	None	* Policy Targets

The bottom status bar shows 'No tasks in progress', 'TomerDomain', '172.23.56.68', and 'Discarded aa'.

And this is how it is defined in the user domain after the global policy assignment:

The screenshot shows the Check Point SmartConsole interface with the 'TomerDomain' policy selected. The left sidebar shows the 'Access Control' section with 'Global-applications' highlighted. The main table lists five rules, including a 'Parent rule for Domain's policy' (rule 3) which is expanded to show sub-rules:

No.	Name	Source	Destination	VPN	Services & Applications	Action	Track	Install On
1	ask for purpose of browsing	Internal-users	* Any	* Any	Facebook Games, Facebook File Shari..., File Storage and Sh..., P2P File Sharing, VoIP, Supports VoIP	Ask, Ask-for-reason, Once a day, Per applicatio...	None	* Poli
2	block high risk categories	* Any	* Any	* Any	Critical Risk, High Risk, Illegal / Questionable, Child Abuse, Sex	Drop, Website-not-a...	None	* Poli
3	Parent rule for Domain's policy							
3.1	blocked categories	* Any	* Any	* Any	Gambling, Pornography	Drop	Log	* Poli
3.2	limit access to high traffic websites	Brokers	* Any	* Any	News / Media, Streaming Media Pr..., Social Networking	Accept, Upload_10Mbps	Network Log	* Poli
4	report suspicious categories	Negated content-checkers	* Any	* Any	Alcohol, Illegal Drugs, Sex Education	Accept	Network Log	* Poli
5	cleanup rule	* Any	* Any	* Any	* Any	Accept	None	* Poli

The bottom status bar shows 'TomerDomain', '172.23.56.68', and '4 Draft'.

THREAT PREVENTION IN R80.10

In R80.10, the IPS Blade is managed by the Threat Prevention policy.

Note: Gateways from previous releases will still receive the IPS content as part of the Access Control Policy. Inspections made by the Firewall Blade will not appear in the list of IPS Protections in SmartConsole. Examples are “non-compliant HTTP” and “aggressive aging”. Instead of accessing it through IPS Protection, they can be accessed from the Inspection Settings page.

R80.10 introduced new concepts which improve the productivity of our customers:

- **Tags** – The IPS protections have dynamic tags. You can tune the Threat Prevention profiles to automatically activate or deactivate all protections that are tagged by relevant aspect such as protocols, affected software and file types.
- **Staging** – Automatically mark newly downloaded protections and control their activation mode (replaces the Follow-Up Flag from previous versions).
- **New default profiles** – R80.10 includes Basic, Optimized and Strict profiles, providing different levels of balance between security and performance.
- **Threat Prevention Exception** - With a single click from logs in the Threat Prevention views, add exceptions to your policy.

DEBUGGING

Debug with Database Modifications. Use `mgmt_cli`.

In Pre-R80.10 systems, debugging was done using `dbedit`, which can still be used for the objects still managed from FWM (such as gateways, VSX objects and QoS policy). Also in earlier systems, you could manually edit config files in `$FWDIR/conf`.

In R80.10, manual editing of files does not work. The configuration is stored in the PostgreSQL database, and not in the `$FWDIR/conf` files. R80.10 now uses these files to create a representation of the installed revision, on Install Policy.

R80.10 Logging Processes are from one CLI command set, rather than various tools of before.

```
SmartEventSetDebugLevel <component> <debug level>
```

Example:

RFL:

```
Start: SmartEventSetDebugLevel rfl debug
```

```
Stop: SmartEventSetDebugLevel rfl info
```

SOLR:

```
Start: SmartEventSetDebugLevel solr debug
```

```
Stop: SmartEventSetDebugLevel solr warn
```

SmartView:

```
Start: SmartEventSetDebugLevel smartview debug
```

```
Stop: SmartEventSetDebugLevel smartview info
```