

Release Notes

Version 10.11.0

October 2021

This document applies to Apama 10.11.0 and to all subsequent releases.

Specifications contained herein are subject to change and these changes will be reported in subsequent release notes or new editions.

Copyright © 2013-2021 Software AG, Darmstadt, Germany and/or Software AG USA, Inc., Reston, VA, USA, and/or its subsidiaries and/or its affiliates and/or their licensors.

The name Software AG and all Software AG product names are either trademarks or registered trademarks of Software AG and/or Software AG USA Inc. and/or its subsidiaries and/or its affiliates and/or their licensors. Other company and product names mentioned herein may be trademarks of their respective owners.

Detailed information on trademarks and patents owned by Software AG and/or its subsidiaries is located at <https://softwareag.com/licenses/>.

Use of this software is subject to adherence to Software AG's licensing conditions and terms. These terms are part of the product documentation, located at <https://softwareag.com/licenses/> and/or in the root installation directory of the licensed product(s).

This software may include portions of third-party products. For third-party copyright notices, license terms, additional rights or restrictions, please refer to "License Texts, Copyright Notices and Disclaimers of Third Party Products". For certain specific third-party license restrictions, please refer to section E of the Legal Notices available under "License Terms and Conditions for Use of Software AG Products / Copyright and Trademark Notices of Software AG Products". These documents are part of the product documentation, located at <https://softwareag.com/licenses/> and/or in the root installation directory of the licensed product(s).

Document ID: PAM-RN-10110-20211015

Table of Contents

About this Guide.....	9
Documentation roadmap.....	10
Online Information and Support.....	11
Data Protection.....	12
 1 License Terms and Technical Restrictions.....	 13
 2 What's New in Apama 10.11.0.....	 15
Java upgrade in 10.11.0.....	16
Compiler upgrade in 10.11.0.....	16
Python upgrade in 10.11.0.....	17
PySys changes in 10.11.0.....	17
Apama enhancements in Software AG Designer in 10.11.0.....	20
Dashboard enhancements in 10.11.0.....	20
Miscellaneous enhancements and changes in 10.11.0.....	21
Platform changes in 10.11.0.....	21
Miscellaneous changes in 10.11.0 affecting backwards compatibility.....	23
Removed and deprecated features in 10.11.0.....	24
 3 What's New In Apama 10.7.2.....	 25
Connectivity plug-ins enhancements in 10.7.2.....	26
EPL enhancements in 10.7.2.....	26
Miscellaneous changes in 10.7.2 affecting backwards compatibility.....	27
 4 What's New In Apama 10.7.1.....	 29
Connectivity plug-ins enhancements in 10.7.1.....	30
EPL code coverage enhancements in 10.7.1.....	31
AnyExtractor enhancements in 10.7.1.....	32
Miscellaneous enhancements and changes in 10.7.1.....	32
Miscellaneous changes in 10.7.1 affecting backwards compatibility.....	32
Removed and deprecated features in 10.7.1.....	33
 5 What's New In Apama 10.7.0.....	 35
PySys changes in 10.7.0.....	36
Connectivity plug-ins enhancements in 10.7.0.....	38
Miscellaneous enhancements and changes in 10.7.0.....	39
Platform changes in 10.7.0.....	40
Miscellaneous changes in 10.7.0 affecting backwards compatibility.....	40
Removed and deprecated features in 10.7.0.....	42
 6 What's New In Apama 10.5.4.....	 47
Connectivity plug-ins enhancements in 10.5.4.....	48

Miscellaneous changes in 10.5.4 affecting backwards compatibility.....	48
7 What's New In Apama 10.5.3.....	49
PySys changes in 10.5.3.....	50
Connectivity plug-ins enhancements in 10.5.3.....	50
Miscellaneous enhancements and changes in 10.5.3.....	51
8 What's New In Apama 10.5.2.....	53
Cumulocity API enhancements in 10.5.2.....	54
Miscellaneous enhancements and changes in 10.5.2.....	54
Miscellaneous changes in 10.5.2 affecting backwards compatibility.....	55
9 What's New In Apama 10.5.1.....	57
PySys enhancements in 10.5.1.....	58
Connectivity plug-ins enhancements in 10.5.1.....	58
EPL enhancements in 10.5.1.....	58
Miscellaneous enhancements and changes in 10.5.1.....	59
Platform changes in 10.5.1.....	59
Support for Red Hat Enterprise Linux (RHEL) 8.0 in 10.5.....	59
Miscellaneous changes in 10.5.1 affecting backwards compatibility.....	59
10 What's New In Apama 10.5.0.....	61
New Apama image on Docker Hub in 10.5.0.....	62
Change of Java Log4j logging library in 10.5.0.....	62
PySys changes in 10.5.0.....	63
Connectivity plug-ins enhancements in 10.5.0.....	65
Apama enhancements in Software AG Designer in 10.5.0.....	66
EPL enhancements in 10.5.0.....	67
Miscellaneous enhancements and changes in 10.5.0.....	67
Platform changes in 10.5.0.....	68
Miscellaneous changes in 10.5.0 affecting backwards compatibility.....	68
Removed and deprecated features in 10.5.0.....	69
Migrating the Cumulocity IoT transports to 10.5.....	73
11 What's New In Apama 10.3.1.....	75
New command-line tool for adding Apama projects in 10.3.1.....	76
New Base64 codec connectivity plug-in in 10.3.1.....	76
New Batch Accumulator codec connectivity plug-in in 10.3.1.....	76
New Message List codec connectivity plug-in in 10.3.1.....	77
Connectivity plug-ins enhancements in 10.3.1.....	77
Apama enhancements in Software AG Designer in 10.3.1.....	79
EPL enhancements in 10.3.1.....	79
Miscellaneous enhancements and changes in 10.3.1.....	80
Platform changes in 10.3.1.....	81
Miscellaneous changes in 10.3.1 affecting backwards compatibility.....	81
Removed and deprecated features in 10.3.1.....	82

12 What's New In Apama 10.3.0.....	83
New Cumulocity IoT transport connectivity plug-in in 10.3.0.....	84
New EPL plug-in for JSON support in 10.3.0.....	85
New Prometheus support in 10.3.0.....	85
New Python EPL plug-in support in 10.3.0.....	85
PySys changes in 10.3.0.....	85
Connectivity plug-ins enhancements in 10.3.0.....	86
Apama enhancements in Software AG Designer 10.3.0.....	87
EPL enhancements in 10.3.0.....	88
Correlator utility enhancements in 10.3.0.....	88
Command Central support in 10.3.0.....	89
Dashboard enhancements in 10.3.0.....	90
Miscellaneous enhancements and changes in 10.3.0.....	90
Platform changes in 10.3.0.....	91
Miscellaneous changes in 10.3.0 affecting backwards compatibility.....	92
Removed and deprecated features in 10.3.0.....	93
 13 What's New In Apama 10.2.....	 95
New support for using the TCStore driver for queries in 10.2.....	96
New Kubernetes support in 10.2.....	96
Connectivity plug-ins enhancements in 10.2.....	96
Apama enhancements in Software AG Designer 10.2.....	97
Query enhancements in 10.2.....	98
EPL enhancements in 10.2.....	98
Correlator utility enhancements in 10.2.....	99
Command Central support in 10.2.....	100
Miscellaneous enhancements and changes in 10.2.....	100
Platform changes in 10.2.....	101
Miscellaneous changes in 10.2 affecting backwards compatibility.....	102
Removed and deprecated features in 10.2.....	103
 14 What's New In Apama 10.1.....	 105
New any type in 10.1.....	106
New HTTP server transport connectivity plug-in in 10.1.....	106
New EPL plug-in for R in 10.1.....	107
New TCStore driver (for Terracotta DB) for the MemoryStore in 10.1.....	107
MemoryStore enhancements in 10.1.....	107
Connectivity plug-ins enhancements in 10.1.....	108
Apama enhancements in Software AG Designer 10.1.....	109
Query enhancements in 10.1.....	110
EPL enhancements in 10.1.....	110
Correlator utility enhancements in 10.1.....	110
Dashboard enhancements in 10.1.....	112
Miscellaneous enhancements and changes in 10.1.....	112
Platform changes in 10.1.....	112
Miscellaneous changes in 10.1 affecting backwards compatibility.....	113
Removed and deprecated features in 10.1.....	115

15 What's New In Apama 10.0.....	119
New HTTP client transport connectivity plug-in in 10.0.....	120
Apama integration with MQTT in 10.0.....	120
Apama integration with Kafka in 10.0.....	120
Connectivity plug-ins enhancements in 10.0.....	121
New API for writing EPL plug-ins in C++ in 10.0.....	124
Apama enhancements in Software AG Designer 10.0.....	124
Query enhancements in 10.0.....	125
EPL enhancements in 10.0.....	125
Correlator utility enhancements in 10.0.....	125
Dashboard enhancements in 10.0.....	126
Miscellaneous enhancements and changes in 10.0.....	126
Platform changes in 10.0.....	127
Windows compiler update in 10.0.....	128
Miscellaneous changes in 10.0 affecting backwards compatibility.....	128
Removed and deprecated features in 10.0.....	130
Migrating the C and C++ EPL plug-ins written in 9.12 or earlier.....	131
 16 What's New In Apama 9.12.....	 133
Apama integration with Software AG Digital Event Services in 9.12.....	134
Apama integration with Universal Messaging in 9.12.....	134
New optional type in 9.12.....	135
New Microsoft Edge support in 9.12.....	135
Connectivity plug-ins enhancements in 9.12.....	135
Connectivity API changes in 9.12.0.1.....	137
Apama enhancements in Software AG Designer 9.12.....	139
Query enhancements in 9.12.....	140
EPL enhancements in 9.12.....	140
Correlator utility enhancements in 9.12.....	141
Dashboard enhancements in 9.12.....	142
PySys enhancements in 9.12.....	142
Command Central support in 9.12.....	143
Miscellaneous enhancements and changes in 9.12.....	143
Platform changes in 9.12.....	144
Miscellaneous changes in 9.12 affecting backwards compatibility.....	144
Removed and deprecated features in 9.12.....	145
 17 What's New In Apama 9.10.....	 149
New connectivity plug-ins feature in 9.10.....	150
New query and DataView metadata feature in 9.10.....	150
New EPL memory profiler feature in 9.10.....	151
New EPL code coverage feature in 9.10.....	151
New PySys version in 9.10.....	151
New dashboard server HTTP REST support in 9.10.....	152
New JSON-based REST support in 9.10.....	153
Command Central support in 9.10.....	153
Apama enhancements in Software AG Designer 9.10.....	154

EPL enhancements in 9.10.....	154
Correlator utility enhancements in 9.10.....	155
Dashboard enhancements in 9.10.....	156
Miscellaneous enhancements and changes in 9.10.....	157
Platform changes in 9.10.....	158
Supported JDK version in 9.10.....	159
Miscellaneous changes in 9.10 affecting backwards compatibility.....	159
Removed and deprecated features in 9.10.....	161
 18 What's New In Apama 9.9.....	163
Installation changes in 9.9.....	164
Query enhancements in 9.9.....	165
Apama enhancements in Software AG Designer 9.9.....	165
EPL enhancements in 9.9.....	166
Correlator utility enhancements in 9.9.....	166
Dashboard enhancements in 9.9.....	167
Miscellaneous enhancements and changes in 9.9.....	167
New samples in 9.9 showing how to containerize Apama for Docker (Linux only).....	168
Platform changes in 9.9.....	168
Renaming of Apama client libraries in 9.9.....	169
Miscellaneous changes in 9.9 affecting backwards compatibility.....	170
Removed and deprecated features in 9.9.....	170
 19 What's New In Apama 5.3.....	173
New Apama queries capability.....	174
Apama Studio enhancements in 5.3.....	175
Log rotation enhancements in 5.3.....	176
EPL enhancements in 5.3.....	177
Correlator-integrated messaging for JMS enhancements in 5.3.....	178
Dashboard enhancements in 5.3.....	179
Killing an object now requires specification of a monitor name.....	179
New sample project showing how to manage Universal Messaging DataGroups.....	179
C++ and C Correlator plug-in API enhancements.....	180
Miscellaneous enhancements and changes in 5.3.....	181
Platform changes in 5.3.....	182
JDK 7 now minimum JDK version supported.....	182
Removed and deprecated features in 5.3.....	183
Backwards incompatibility with persisted projects recovered to 5.3 from older versions...183	
 20 What's New in Apama 5.2.....	185
Apama Studio enhancements in 5.2.....	186
New feature for addressing contexts with channels in 5.2.....	187
New feature for using Universal Messaging to connect Apama components in 5.2.....	191
New and changed EPL statements in 5.2.....	192
Enhancements to Apama adapters in 5.2.....	193
Changes to correlator utilities in 5.2.....	193
Windows compiler and .NET version updates in 5.2.....	195
Linux compiler updates in 5.2.....	195

Upgrading to Apama-supplied JDBC and ODBC database drivers in 5.2.....	196
Removed and deprecated features in 5.2.....	197
Miscellaneous enhancements and changes in 5.2.....	199
 21 What's New in Apama 5.1.1.....	 201
 22 What's New in Apama 5.1.....	 203
New distributed MemoryStore.....	204
New EPL exception handling mechanism.....	204
New support for managing and monitoring with REST APIs.....	205
New support for creating correlator plug-ins using Java.....	206
New headless mode generation of ApamaDoc.....	206
New support for PySys test framework.....	206
Dashboard login redirect file name has changed.....	206
New EPL keywords.....	207
Miscellaneous enhancements in 5.1.....	207
Support removed from Apama 5.1.....	208
Migrating Apama applications.....	209

About this Guide

- Documentation roadmap 10
- Online Information and Support 11
- Data Protection 12

These *Release Notes* describe new features and changes introduced with the current Apama release as well as earlier releases.

Older versions covered in these *Release Notes* may no longer be supported. For an overview of the latest supported releases, see the Product Version Availability page at <https://empower.softwareag.com/Products/ProductAvailability/default.asp> (Empower login required).

Documentation roadmap

Apama provides documentation in the following formats:

- HTML (available from both the documentation website and the doc folder of the Apama installation)
- PDF (available from the documentation website)
- Eclipse help (accessible from Software AG Designer)

You can access the HTML documentation on your machine after Apama has been installed:

- **Windows.** Select **Start > All Programs > Software AG > Tools > Apama *n.n* > Apama Documentation *n.n***. Note that **Software AG** is the default group name that can be changed during the installation.
- **UNIX.** Display the `index.html` file, which is in the `doc/apama-onlinehelp` directory of your Apama installation directory.

The following guides are available:

Title	Description
<i>Release Notes</i>	Describes new features and changes introduced with the current Apama release as well as earlier releases.
<i>Installing Apama</i>	Summarizes all important installation information and is intended for use with other Software AG installation guides such as <i>Using Software AG Installer</i> .
<i>Introduction to Apama</i>	Provides a high-level overview of Apama, describes the Apama architecture, discusses Apama concepts and introduces Software AG Designer, which is the main development tool for Apama.
<i>Using Apama with Software AG Designer</i>	Explains how to develop Apama applications in Software AG Designer, which is an Eclipse-based integrated development environment.
<i>Developing Apama Applications</i>	Describes the different technologies for developing Apama applications: EPL monitors, Apama queries, and Java. You can use one or several of these technologies to implement a single Apama application. In addition, there are C++ and Java APIs for developing components that plug in to a correlator. You can use these components from EPL.

Title	Description
<i>Connecting Apama Applications to External Components</i>	Describes how to connect Apama applications to any event data source, database, messaging infrastructure, or application.
<i>Building and Using Apama Dashboards</i>	Describes how to build and use an Apama dashboard, which provides the ability to view and interact with DataViews. An Apama project typically uses one or more dashboards, which are created in the Dashboard Builder. The Dashboard Viewer provides the ability to use dashboards created in the Dashboard Builder. Dashboards can also be deployed as simple web pages. Deployed dashboards connect to one or more correlators by means of a dashboard data server or display server.
<i>Deploying and Managing Apama Applications</i>	Describes how to deploy components with Software AG Command Central, how to deploy and manage queries, and how to deploy Apama applications using Docker and Kubernetes. It also provides information for improving Apama application performance by using multiple correlators, for managing and monitoring Apama components over REST (Representational State Transfer), and for using correlator utilities and configuration files.

In addition to the above guides, Apama also provides the following API reference information:

- API Reference for EPL (ApamaDoc)
- API Reference for Java (Javadoc)
- API Reference for C++ (Doxygen)
- API Reference for .NET
- API Reference for Python
- API Reference for Component Management REST APIs

Online Information and Support

Software AG Documentation Website

You can find documentation on the Software AG Documentation website at <https://documentation.softwareag.com>.

Software AG Empower Product Support Website

If you do not yet have an account for Empower, send an email to empower@softwareag.com with your name, company, and company email address and request an account.

Once you have an account, you can open Support Incidents online via the eService section of Empower at <https://empower.softwareag.com/>.

You can find product information on the Software AG Empower Product Support website at <https://empower.softwareag.com>.

To submit feature/enhancement requests, get information about product availability, and download products, go to [Products](#).

To get information about fixes and to read early warnings, technical papers, and knowledge base articles, go to the [Knowledge Center](#).

If you have any questions, you can find a local or toll-free number for your country in our Global Support Contact Directory at https://empower.softwareag.com/public_directory.aspx and give us a call.

Software AG Tech Community

You can find documentation and other technical information on the Software AG Tech Community website at <https://techcommunity.softwareag.com>. You can:

- Access product documentation, if you have Tech Community credentials. If you do not, you will need to register and specify "Documentation" as an area of interest.
- Access articles, code samples, demos, and tutorials.
- Use the online discussion forums, moderated by Software AG professionals, to ask questions, discuss best practices, and learn how other customers are using Software AG technology.
- Link to external websites that discuss open standards and web technology.

Data Protection

Software AG products provide functionality with respect to processing of personal data according to the EU General Data Protection Regulation (GDPR). Where applicable, appropriate steps are documented in the respective administration documentation.

1 License Terms and Technical Restrictions

Using Apama 9.10 and higher without a valid license file means the product will function as Apama Community Edition. This binds the user and any related organization to the Freemium license terms in the Software AG license agreement (refer to the license agreement for further details). In addition, the following Apama Community Edition restrictions also apply:

- Support is not available from Software AG's Empower support portal.
- Projects cannot incorporate more than 10 correlators.
- Attribution is required in the end application's Help/About screen or relevant alternative (for example, "Streaming analytics by Software AG Apama").
- The correlator will execute EPL on a maximum of 4 CPU threads (and so performance may be restricted).
- The correlator does not permit processing Apama queries against distributed stores such as Terracotta's TCStore (that is, in-memory store only).
- Reliable messaging with connectivity plug-ins is not permitted.
- Correlator-integrated messaging for JMS only permits best-effort messaging (that is, reliable messaging is not permitted).
- The correlator log file includes notes highlighting it is running without a valid license file.
- The correlator will terminate if its resident memory usage rises above 1024MB.
- The correlator only permits 20 contexts to be created (and so performance may be restricted).
- The correlator only permits 5 EPL monitor types to be persistent.
- The correlator cannot read user-generated correlator deployment packages (CDPs).
- The correlator only handles 5 Apama query definitions.
- The correlator only handles 5 instances of each Apama query definition.
- The correlator only handles 50 unique partitions for an Apama query.

See also "Running Apama without a license file" in *Introduction to Apama*.

If you need to remove these Apama Community Edition restrictions, then please contact Software AG to purchase Apama.

2 What's New in Apama 10.11.0

■ Java upgrade in 10.11.0	16
■ Compiler upgrade in 10.11.0	16
■ Python upgrade in 10.11.0	17
■ PySys changes in 10.11.0	17
■ Apama enhancements in Software AG Designer in 10.11.0	20
■ Dashboard enhancements in 10.11.0	20
■ Miscellaneous enhancements and changes in 10.11.0	21
■ Platform changes in 10.11.0	21
■ Miscellaneous changes in 10.11.0 affecting backwards compatibility	23
■ Removed and deprecated features in 10.11.0	24

Apama 10.11.0 is the successor of Apama 10.7.2. There is no version 10.8, 10.9 or 10.10.

Apama 10.11.0 runs on the platforms listed in the *Supported Platforms* document for version 10.11. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>. Be sure to consult that document for details about supported versions of operating systems, compilers and other products.

Apama 10.11.0 includes new features, enhancements, and changes as described in the following topics.

Java upgrade in 10.11.0

This release of Apama comes with a Java 11 JDK instead of Java 8. All Java code executed within the correlator and IAF processes is now executed with Java 11, so you may need to make changes to your EPL plug-ins, connectivity plug-ins, IAF plug-ins, JMON applications, etc. that run inside the correlator or IAF.

If you build separate client processes that use the Apama APIs from `ap-client.jar`, a Java 8 runtime can still be used if you wish. However, this is deprecated and it is recommended that you migrate to using a Java 11 runtime when possible.

If you are using the "core components" version of the Apama Community Edition, then you must now provide a Java 11 runtime for use with Apama instead of Java 8.

There are plenty of great resources on the internet for finding out how to upgrade from Java 8 to Java 11, and lists of breaking changes between these versions. In practice, the main effort is likely to be upgrading any third-party libraries to Java 11-supporting versions, and adding additional libraries to replace some of the packages removed from the runtime itself (for example, JAXB). Java now uses a different default garbage collector (G1GC instead of the parallel garbage collector) which may result in different memory behavior.

The default locale data provider was changed (to the Unicode Consortium's CLDR), which may result in dates, times and numbers that used to parse successfully now being rejected. See the Java documentation for details of how to configure the Java locale data provider.

Compiler upgrade in 10.11.0

Apama 10.11.0 has upgraded the compilers that it uses:

- On Windows 10 and Windows Server 2019, Apama has upgraded from Microsoft Visual Studio 2015 Update 3 (Visual C++ v14.0) to Visual Studio 2019 (Visual C++ v14.2).
- On Linux, Apama builds on Red Hat Enterprise Linux (RHEL) 8 and Ubuntu 20.04 with GCC 8.4. Apama is certified for use on these platforms.
- On Linux on Arm (32-bit), Apama builds on Raspberry Pi OS 10 with GCC 8.3. Apama is also certified for use on this platform.

Consequently, you must rebuild the following components with the new compiler:

- C++ connectivity plug-ins

- C++ EPL plug-ins
- C and C++ IAF transport/codec plug-ins
- C and C++ code that uses the engine client API

Note:

Compiling connectivity or EPL plug-ins with Red Hat Enterprise Linux (RHEL) 8 no longer needs `-D_GLIBCXX_USE_CXX11_ABI=0` as described in “[Support for Red Hat Enterprise Linux \(RHEL\) 8.0 in 10.5](#)” on page 59.

In addition, when planning migration, consider that new compiler versions are more modern than the previous versions and less tolerant of ambiguous and potentially unsafe code patterns. It is likely that some time may be required to fix compiler errors after moving to the new versions.

Note:

Our recommendation is to enable “warnings as errors” in your projects where possible. This encourages good coding patterns and catches errors earlier in the development process. In particular, it helps you catch errors that might cause subtle, difficult to diagnose errors.

Python upgrade in 10.11.0

This release of Apama comes with Python 3.9.5, replacing Python 3.7.8. For a detailed overview of changes, including new features, deprecated and breaking changes, see <https://docs.python.org/3/whatsnew/3.9.html> and <https://docs.python.org/3.8/whatsnew/3.8.html>.

This change affects:

- All Python executed from within the Apama Command Prompt, for example, tests.
- EPL plug-ins written in Python.

If you are not using the version of Python distributed with the full Apama installation, you will have to update your Python installation to match.

PySys changes in 10.11.0

This release of Apama ships with a new version of PySys, version 2.0.

New PySys features

The 2.0 release contains many enhancements including:

- A new standard test structure that avoids the use of XML by allowing descriptor values such as the test title to be specified alongside your Python test class in a single `pysystest.py` file, instead of separate `run.py` and `pysystest.xml` files. You can mix and match the old and new styles within the same project. For new PySys projects, a simpler directory layout is now recommended in which the `self.input` directory is configured to be the main `testDir/` (which also contains the `pysystest.py` file) instead of having a separate `testDir/Input/` subdirectory for input files. This can make test contents easier to navigate.

- Some big extensions to the concept of “modes” that allow for more powerful configuration and use, including mode parameters for easier handling of multi-dimensional modes, and dynamic mode lists configured with a Python lambda expression.
- A new template-based implementation of `pysys make`, allowing easy configuration of how new tests are created - on a per-directory basis - and also automatic generation of test identifiers for new tests (when using numeric identifiers).
- Several improvements to the `pysys.mappers` API for more easily transforming text files during copy and grep operations, including support for multi-line exception stack traces.
- A large set of smaller additions, many based on end-user requests.

For more information on the above and also details on the many smaller enhancements in this PySys release, see the PySys Change Log (*PySys Test Framework > Change Log in the API Reference for Python*).

New Apama helper features

There are several improvements to the Apama helper classes for PySys in this release:

- The `ApamaPlugin` class (`self.apama`) has a new method `JoinEPLStackLines` which can be used in an `assertGrep(..., mappers=[...])` list to convert multi-line EPL stack dumps to a single line. This can make it easier to ignore expected errors while still checking for unanticipated errors. For example:

```
self.assertGrep('testCorrelator.log', '(ERROR|FATAL|Failed to parse) .*',
contains=False,
mappers=[self.apama.JoinEPLStackLines(),
pysys.mappers.JoinLines.JavaStackTrace()])
```

See *Apama Helper Classes for PySys* in the *API Reference for Python* for more information.

- The `CorrelatorHelper` and `IAFHelper` classes have a new method `waitForLogGrep` which (much like `waitForGrep`) can be used to wait for a log message to appear in the log file, but with automatic checking for errors and an early abort with a clear outcome reason if an error is detected. For example:

```
correlator.waitForLogGrep('INFO .* Test completed')
```

- The `ApamaPlugin` class (`self.apama`) has a new field `defaultLogIgnores` which holds a list of project-level regular expressions that should be ignored when checking for errors, warnings and other messages in Apama component logs. The default value contains a small set of `WARN` messages which should be ignored in testcases. Additional ignores can be added to this list with a multi-line list of regular expressions in your `pysysproject.xml` configuration:

```
<property name="apamaDefaultLogIgnores">
  WARN .*My warning to ignore
  ERROR .*My error to ignore
</property>
```

- It is recommended to add at least one maker-template to your `pysysproject.xml` file to provide a good starting point when creating new testcases using `pysys make`. A default Apama test template is provided with the product which includes a `test.mon` file and some typical `PySysTest`

commands. You should delete any aspects not needed in your testcase. To make this template available in existing projects, just add the following XML to your `pysysproject.xml` file:

```
<pysysproject>
  <pysysdirconfig>
    <maker-template name="apama-default-test"
      description="a test for an Apama application that runs in a local correlator"
      copy="{pysysTemplatesDir}/apama-default-test/*"
    />
  </pysysdirconfig>
</pysysproject>
```

Breaking changes

There are a several changes that could require changes to existing PySys projects. A full list is provided in the PySys Change Log, but the main ones are summarized below:

- When user-defined mappers= are used (for example, during `self.copy`), it is now an error for a mapper to strip off the trailing `\n` character at the end of each line, as failure to do so can have unintended consequences on later mappers. This requirement is also more clearly documented.
- Some mistakes in the `pysystest.xml` structure that were previously tolerated will now produce `stderr` warning messages (for example, incorrectly nesting `<modes>` inside `<groups>`) and others will produce a fatal error (for example, multiple occurrences of the same element). To find out if any tests need fixing up, just execute `pysys print` in your PySys project directory and act on any warning or error messages.
- The deprecated `supportMultipleModesPerRun=false` project property (only used in very old PySys projects) can no longer be used. Please change your tests to use the modern modes approach instead.
- On Windows, the `testDir` (and the input, output, and reference directories) no longer start with the `\\?\` long path prefix. Instead, this can be added for operations where it is needed using `pysys.utils.fileutils.toLongPathSafe` (as the standard PySys methods already do, for example `self.copy`). Where possible, it is recommended to avoid nesting tests and output directories so deeply that long path support is needed.

In addition, your tests may be subject to breaking changes from the upgrade to Python 3.9. See also [“Python upgrade in 10.11.0” on page 17](#).

Deprecations

- It is strongly recommended to use the new `pysys.constants.PREFERRED_ENCODING` constant instead of Python's built-in `locale.getpreferredencoding()` function, to avoid thread-safety issues in your tests.
- If you have a custom `pysys.utils.perfreporter.CSVPerformanceReporter` subclass, the signatures for `pysys.utils.perfreporter.CSVPerformanceReporter.getRunDetails` and `pysys.utils.perfreporter.CSVPerformanceReporter.getRunHeader` have changed to include a `testobj` parameter.

Although this should not immediately break existing applications, to avoid future breaking changes, you should update the signature of those methods if you override them to accept `testobj` and also any arbitrary `**kwargs` that may be added in future.

- The `pysys.xml` module is deprecated. Use `pysys.config` instead.
- The `pysys.utils.fileunzip` module is deprecated. Use `BaseTest.unpackArchive` instead. For example, replace `unzip(gzfilename, binary=True)` with `self.unpackArchive(gzfilename, gzfilename[:-3])`.
- The (undocumented) `DEFAULT_DESCRIPTOR` constant is now deprecated and should not be used.
- The old `<mode>` elements are deprecated in favor of the new Python lambda syntax (support for these will not be removed any time soon, but are discouraged for new tests).
- The `pysys.utils.pycompat` module is now deprecated.
- The `ConsoleMakeTestHelper` class is now deprecated in favor of `pysys.launcher.console_make.DefaultTestMaker`.
- If your `pysysproject.xml` contains an `APAMA_COMMON_JRE` property with its default value set to `${APAMA_HOME}/../jvm/jvm/jre`, you should either remove the property or change the default value to `${APAMA_HOME}/../jvm/jvm` as the `/jre` directory has been removed in this release of Apama.

A quick way to check for the removed and deprecated items in your test project is to use the regular expression as shown in the following `grep` command:

```
grep -r "\(\supportMultipleModesPerRun.*false\|DescriptorLoader
\|pysys.utils.loader\|_stringToUnicode\|pysys[.]xml\|pysys.utils.fileunzip
\|[\^_@]DEFAULT_DESCRIPTOR\|pysys.utils.pycompat\|PY2\|string_types
\|binary_type\|isstring[(]\|quotestring[(]\|openfile[(]\|ConsoleMakeTestHelper
\|def getRunDetails\|def getRunHeader\|locale.getpreferredencoding\|addResource
\|CommonProcessWrapper\|TEST_TEMPLATE\|DESCRIPTOR_TEMPLATE\|
\|/jvm/jvm/jre\)" .
```

Apama enhancements in Software AG Designer in 10.11.0

Apama 10.11.0 includes the following enhancements in Software AG Designer:

- You can now export and import the Apama-specific preferences from one workspace to another in Software AG Designer. For more information, see "Exporting the Apama-specific preferences" and "Importing the Apama-specific preferences" in *Using Apama with Software AG Designer*.

Dashboard enhancements in 10.11.0

Apama 10.11.0 includes the following dashboard enhancements:

- The Dashboard Builder now supports two sets of object palettes using different styles: classic and flat. The classic style was available in previous releases. The default is now to use the flat style. See "Setting Builder options" in *Building and Using Apama Dashboards* for more information.

Miscellaneous enhancements and changes in 10.11.0

Apama 10.11.0 includes the following miscellaneous enhancements and changes:

- Apama 10.11.0 incorporates the ICU (International Components for Unicode) time-zone data update 2021a, which is the most recent update at the time of release. This will update time-zone data used by the correlator and Time Format EPL plug-in.
- A new key for monitoring the peak physical memory usage of the correlator, `peakPhysicalMemoryMB`, has been added to the `/correlator/status` REST endpoint. This is useful for checking whether the correlator is running on an adequately provisioned machine. For more information, see "List of correlator status statistics" and "Managing and Monitoring over REST", both in *Deploying and Managing Apama Applications*.
- Apama now provides a `--shutdownTimeoutSecs` correlator option. This option is used to specify the maximum time allowed for the correlator to shut down after all persistence activities have been completed. When this time has elapsed, the correlator shuts down forcibly, ignoring any transport or plug-in shutdown activities. This option is especially useful to prevent indefinite hangs caused by plug-ins. If this option is not provided, a default value of 90 seconds is used. See "Starting the correlator" in *Deploying and Managing Apama Applications*.

A corresponding YAML element with the name `shutdownTimeoutSecs` is also provided. See "YAML configuration file for the correlator" in *Deploying and Managing Apama Applications*.

Platform changes in 10.11.0

Apama 10.11.0 runs on the platforms listed in the *Supported Platforms* document for version 10.11. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>.

Check the above mentioned *Supported Platforms* document for details about the newer platform versions supported by Apama 10.11.0. The platforms listed below are no longer supported.

Operating systems

- CentOS Linux (no replacement).
- Microsoft Windows 8.1 (replaced by support for a more recent version).
- Microsoft Windows Server 2012 R2 and 2016 (replaced by support for a more recent version).
- Red Hat Enterprise Linux 7 (replaced by support for a more recent version).
- SUSE Linux Enterprise Server (no replacement).

Compilers

- CentOS Linux (no replacement).
- Microsoft Windows 8.1 (replaced by support for a more recent version).

- Microsoft Windows Server 2012 R2 and 2016 (replaced by support for a more recent version).
- Red Hat Enterprise Linux 7 (replaced by support for a more recent version).
- SUSE Linux Enterprise Server (no replacement).

Application servers

- IBM WebSphere Application Server (no replacement).
- Oracle WebLogic Server (no replacement).
- Software AG Common Runtime 10.7 (replaced by support for a more recent version).

Databases

- Microsoft SQL Server 2016 (later versions are still supported).
- MySQL Community Edition 5.7 (replaced by support for a more recent version).
- Oracle 12c (a later version is still supported).

JMS providers

- IBM MQ 9.0 (replaced by support for a more recent version).
- IBM WebSphere Application Server (no replacement).
- Oracle WebLogic Server (no replacement).
- Software AG Universal Messaging 10.7 (replaced by support for a more recent version).

Integration platforms

- Software AG Universal Messaging Client 10.7 (replaced by support for a more recent version).

Web browsers

No changes.

Virtualization software

No changes.

Distributed memory store

- Software AG BigMemory Max 4.3.9 (replaced by support for a more recent version).
- Software AG Terracotta Store 10.7 (replaced by support for a more recent version).

Miscellaneous

- Apache Ant 1.10.5 (replaced by support for a more recent version).
- Java 8 (replaced by support for a more recent version).

Note:

In some situations, Java 8 can still be used. See [“Java upgrade in 10.11.0” on page 16](#) for more information.

- MathWorks MATLAB 2016b (replaced by support for a more recent version).
- Microsoft .NET Framework 4.5.1 (replaced by support for a more recent version).
- PySys 1.6.1 (replaced by support for a more recent version).
- Python 3.7 (replaced by support for a more recent version).
- Software AG Cumulocity IoT 10.7 (replaced by support for a more recent version).
- Software AG Designer 10.7 (replaced by support for a more recent version).

Miscellaneous changes in 10.11.0 affecting backwards compatibility

The following changes in Apama 10.11.0 affect backwards compatibility to previous Apama versions:

- Software AG Designer has been upgraded to a new Eclipse version. It is recommended that you back up your workspaces.
 - Software AG Designer now uses Eclipse 4.19. Eclipse 4.19 workspaces are upward-compatible, and workspaces created with older Eclipse versions are automatically migrated to 4.19 workspaces. Downward-compatibility of workspaces, however, is not supported. This means that a workspace which has been created or opened with Eclipse 4.19 cannot be used with an older Eclipse version. Therefore, it is recommended that you back up your workspaces before opening them with Eclipse 4.19.
 - If you have installed bundles by dropping them into the `plugins` or `dropins` directory, they may no longer resolve when you upgrade to a new Eclipse Platform version. In each new version of the Eclipse Platform, there are new versions of bundles included in the platform, and often a small number of removed bundles. This may cause your previously dropped-in bundles to no longer resolve and run. It is always recommended that you install software via the **Help > Install New Software** mechanism so you are made aware of any install-time failure to resolve dependencies.
- The Jackson Java serialization library is no longer present on the default classpath of the correlator, or of Apama JAR files such as `ap-util.jar`. Although the presence of Jackson was never documented, it is possible that some customers are depending on it. If you have Java classes (such as connectivity plug-ins) that depend on any `com.fasterxml.jackson.*` classes, you should explicitly add the Jackson JAR files to your classpath.

Note:

There are also other features that affect backwards compatibility. See the following topics for more information: [“Java upgrade in 10.11.0” on page 16](#), [“Compiler upgrade in 10.11.0” on page 16](#), [“Python upgrade in 10.11.0” on page 17](#), and [“PySys changes in 10.11.0” on page 17](#).

Removed and deprecated features in 10.11.0

The following Apama features are now deprecated or have been removed in Apama 10.11.0:

- Data groups are deprecated by Universal Messaging in this release. Therefore, the samples in the `samples\universal_messaging\datagroups` subdirectory have been removed from the Apama installation.
- With the new PySys release, several PySys items are now deprecated or have been removed. For detailed information, see [“PySys changes in 10.11.0” on page 17](#).

3 What's New In Apama 10.7.2

■ Connectivity plug-ins enhancements in 10.7.2	26
■ EPL enhancements in 10.7.2	26
■ Miscellaneous changes in 10.7.2 affecting backwards compatibility	27

Apama 10.7.2 runs on the platforms listed in the *Supported Platforms* document for version 10.7. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>. Be sure to consult that document for details about supported versions of operating systems, compilers and other products.

Apama 10.7.2 includes new features, enhancements, and changes as described in the following topics.

Connectivity plug-ins enhancements in 10.7.2

Apama 10.7.2 includes the following connectivity plug-ins enhancements:

HTTP server

- You can now change the timeout value for a connection using the new `keepAliveTimeSecs` configuration option. See also "Configuring the HTTP server transport" in *Connecting Apama Applications to External Components*.

Cumulocity IoT

- SendSMS events now always pick the latest defaults for the sender name and sender address from the tenant options. It is no longer required to restart the correlator. See also "Configuring the Cumulocity IoT transport" in *Connecting Apama Applications to External Components*.

JSON codec

- The `filterOnContentType` configuration option for JSON codec now accepts parameters in the `metadata.contentType`. Therefore, the filter pattern has been updated from `"^application/([^\s/]+)?json$" to "^application/([^\s/]*+)?json(;[^\s;]=([^\s]*|"[^"]*"|'['']*')*)*$". See also "The JSON codec connectivity plug-in" in Connecting Apama Applications to External Components.`

EPL enhancements in 10.7.2

Apama 10.7.2 includes the following EPL enhancements:

- It is now easier to identify and extract dictionary and sequence types contained within an any. For this purpose, the following new actions are now available for the any type:
 - `isDictionary`
 - `isEvent`
 - `isSequence`
 - `getGenericDictionary`
 - `getDeclaredElementType`
 - `getDeclaredKeyType`

See the *API Reference for EPL (ApamaDoc)* for more information.

Miscellaneous changes in 10.7.2 affecting backwards compatibility

The following changes in Apama 10.7.2 affect backwards compatibility to previous Apama versions:

- The `GenericResponse.body` member is now deprecated and will be removed in a future release. It is recommended that you now use the `GenericResponse.getBody` action instead. See the `com.apama.cumulocity` package in the *API Reference for EPL (ApamaDoc)* for more information on the `GenericResponse` event.

4 What's New In Apama 10.7.1

■ Connectivity plug-ins enhancements in 10.7.1	30
■ EPL code coverage enhancements in 10.7.1	31
■ AnyExtractor enhancements in 10.7.1	32
■ Miscellaneous enhancements and changes in 10.7.1	32
■ Miscellaneous changes in 10.7.1 affecting backwards compatibility	32
■ Removed and deprecated features in 10.7.1	33

Apama 10.7.1 runs on the platforms listed in the *Supported Platforms* document for version 10.7. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>. Be sure to consult that document for details about supported versions of operating systems, compilers and other products.

Apama 10.7.1 includes new features, enhancements, and changes as described in the following topics.

Connectivity plug-ins enhancements in 10.7.1

Apama 10.7.1 includes the following connectivity plug-ins enhancements:

HTTP client

- The HTTP client's generic JSON chain now allows you to invoke an HTTP service with a payload encoded to either `multipart/form-data` or `application/x-www-form-urlencoded` media types using the new predefined `FormRequest` event definition. See also "Handling HTML form encoding using a predefined generic event definition" in *Connecting Apama Applications to External Components*.

HTTP server

- The HTTP server transport now allows concurrent processing using multiple chains for multiple connections. The default is `false`, but you can turn it on by setting `concurrentChains` to `true` in the YAML configuration file. See also "Configuring the HTTP server transport" in *Connecting Apama Applications to External Components*. This enhancement was also delivered with Apama 10.5.4.
- The new monitoring metric `numChains` is now available for the HTTP server transport. It indicates the number of active connections into the HTTP server instance. See also "Monitoring status for the HTTP server" in *Connecting Apama Applications to External Components*. This enhancement was also delivered with Apama 10.5.4.
- The HTTP server transport now uses the OpenAPI/Swagger 3.0 specification. As of Apama 10.7.1, you should use the OpenAPI 3.0 tooling (<https://swagger.io/tools/swagger-codegen/>) to generate client code from the definitions. See also "About the HTTP server transport" in *Connecting Apama Applications to External Components*.

Cumulocity IoT

- When the HTTP client's generic JSON bundle has been added to the project configuration, defaults for the sender name and sender address are now retrieved from the tenant options in Cumulocity IoT if they are not specified in the `SendsMS` event. If the tenant options cannot be retrieved from Cumulocity IoT, you can now configure the defaults by adding `CUMULOCITY_SMS_SENDER_NAME` and `CUMULOCITY_SMS_SENDER_ADDRESS` to the `.properties` configuration file of the Cumulocity IoT transport. See also "Configuring the Cumulocity IoT transport" in *Connecting Apama Applications to External Components*.

- A geofence helper utility has been added to the **Utilities for Cumulocity** EPL bundle. See also "About the Cumulocity IoT transport" in *Connecting Apama Applications to External Components*. In addition, the following new events are now available in the `com.apama.cumulocity` package: `GeoFenceContainer` and `GeoFencePoint`. See the *API Reference for EPL (ApamaDoc)* for more information on these events.
- The `com.apama.cumulocity` events `Alarm`, `Event`, `ManagedObject` and `Operation` now have `isCreate()` and `isUpdate()` actions which can be used in event listeners to distinguish between events from Cumulocity IoT which are a result of a creation or an update. See the *API Reference for EPL (ApamaDoc)* for more information on these events.
- The `com.apama.cumulocity.Error` event now has a `reqId` field which corresponds to the identifier of the request from which the error event occurred. See the *API Reference for EPL (ApamaDoc)* for more information.
- The Cumulocity API now supports a new `withChannelResponse` action for the predefined types (`ManagedObject`, `Alarm`, `Event`, `Measurement`, `MeasurementFragment` and `Operation`). This allows your application to receive a response on the `type.SUBSCRIBE_CHANNEL` channel when one of these object types is created or updated.

It is recommended that you now use the new `withChannelResponse` action instead of the `withResponse` action which is now deprecated. See also "[Removed and deprecated features in 10.7.1](#)" on page 33.

See the following topics in *Connecting Apama Applications to External Components* for more information:

- Updating a managed object
- Creating a new alarm
- Creating a new event
- Creating a new measurement
- Creating measurement fragments
- Creating a new operation

EPL code coverage enhancements in 10.7.1

Apama 10.7.1 includes the following EPL code coverage enhancements:

- In addition to the existing support for generating code coverage reports in HTML and CSV format, the `epl_coverage` tool now also generates coverage reports in a widely-used XML format (the format that is used and was popularized by Cobertura, which is a code coverage utility for Java; see also <https://cobertura.github.io/cobertura/>). This allows coverage from EPL to be checked and visualized by third-party tools. No changes are required to get the XML report. An `epl_coverage.xml` file is generated automatically whenever the `epl_coverage` tool is executed. See also "Generating code coverage information about EPL files" in *Deploying and Managing Apama Applications*.

- Correlator EPL code coverage reports now include the coverage data of deleted monitors, with the exception of those deleted with `engine_delete --all` or when the file has been reinjected. For more details, see "Recording code coverage information" in *Deploying and Managing Apama Applications*.

AnyExtractor enhancements in 10.7.1

Apama 10.7.1 includes the following enhancements:

- The `com.apama.util.AnyExtractor` event for extracting data from an any type now has actions `getAnyOr`, `getStringOr`, `getIntegerOr`, `getFloatOr`, `getBooleanOr`, `getSequenceOr`, and `getDictionaryOr`. These actions return a user-specified default value if the path to the nested value does not exist, or if the path exists but the value is not of the required type. These offer non-throwing alternatives to the existing getter actions in the `AnyExtractor` event. See the *API Reference for EPL (ApamaDoc)* for more information.
- `AnyExtractor` now supports extracting values from some additional types: events, `sequence<ANY_TYPE>`, `dictionary<string, ANY_TYPE>`, `dictionary<any, ANY_TYPE>` and location objects. The "." and "[]" operators are also interchangeable, so either can be used for extracting values from any of these types.

Miscellaneous enhancements and changes in 10.7.1

Apama 10.7.1 includes the following miscellaneous enhancements and changes:

- For correlators that are configured for integrated messaging, the JMS receiver-specific status lines in the correlator log file now show a new `maxMsgKB` field which indicates the maximum size in kilobytes of JMS messages that have been received so far. See also "Logging correlator-integrated messaging for JMS status" in *Connecting Apama Applications to External Components*.

Miscellaneous changes in 10.7.1 affecting backwards compatibility

The following changes in Apama 10.7.1 affect backwards compatibility to previous Apama versions:

- While sending a request, the HTTP client transport no longer sets a `content-type` header from `metadata.contentType` if the body is empty. If you want to send the `content-type` even if the body is empty, you now have to set `metadata.http.headers.content-type` explicitly in the request. See also "Handling HTTP headers" in *Connecting Apama Applications to External Components*.
- The YAML file that was provided with the HTTP client's generic JSON transport for using predefined generic event definitions is now used as a linked resource. Therefore, the bundle instance file `HTTP_Client_bundle_instance.yaml` is no longer created for the "generic" option (see also "Adding the HTTP client connectivity plug-in to a project" in *Using Apama with Software AG Designer*). To work with the newer version of this bundle, you have to call the `engine_deploy` tool again on your project.

Existing bundle instances for the “generic” option will be discarded. In case you modified the bundle and still want to use it, the chain definition name must be explicitly mentioned when creating a transport:

```
HttpTransport.getOrCreateWithConfigurations(host, port,  
{HttpTransport.CONFIG_DYNAMIC_CHAIN_DEF_NAME: "modified-chain-definition-name"});
```

Also, the chain name has been changed from `HTTPClientGenericChain` to `HTTPClientGenericJSONChain`. If you are using this chain, you have to update any EPL where you refer to `HTTPClientGenericChain` in the `chainDefnName` argument of a `createDynamicChain()` call to reflect the changed behavior. See also “Creating dynamic chains from EPL” in *Connecting Apama Applications to External Components*.

Removed and deprecated features in 10.7.1

The following Apama features are now deprecated or have been removed in Apama 10.7.1:

- The `withResponse` action, which was added to the Cumulocity API in Apama 10.5.2, is now deprecated for all of the predefined types (`ManagedObject`, `Alarm`, `Event`, `Measurement`, `MeasurementFragment` and `Operation`) and will be removed in a future release. It is recommended that you now use the new `withChannelResponse` action. See also [“Connectivity plug-ins enhancements in 10.7.1” on page 30](#).

5 What's New In Apama 10.7.0

■	PySys changes in 10.7.0	36
■	Connectivity plug-ins enhancements in 10.7.0	38
■	Miscellaneous enhancements and changes in 10.7.0	39
■	Platform changes in 10.7.0	40
■	Miscellaneous changes in 10.7.0 affecting backwards compatibility	40
■	Removed and deprecated features in 10.7.0	42

Apama 10.7.0 is the successor of Apama 10.5.3. There is no version 10.6.

Apama 10.7.0 runs on the platforms listed in the *Supported Platforms* document for version 10.7. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>. Be sure to consult that document for details about supported versions of operating systems, compilers and other products.

Apama 10.7.0 includes new features, enhancements, and changes as described in the following topics.

PySys changes in 10.7.0

This release of Apama ships with a new version of PySys, version 1.6.1. The 1.6.0 and 1.6.1 releases contain many enhancements including:

- Several new writers for recording test results, including GitHub Actions support and a writer that produces .zip archives of test output directories.

You may wish to copy some of the new writers from the `APAMA_HOME/samples/pysys/pysysproject.xml` file into your own projects to benefit from these, or see the PySys documentation in the *API Reference for Python* for more information.

- A new `BaseTest.assertThatGrep()` method that extracts a value using a grep expression and then allows asserting its value is as expected. For extract-and-assert use cases, this approach gives much clearer messages when the assert fails than using `assertGrep`.
- Numerous minor usability improvements to the `pysys.py` command line, project configuration file, and to many PySys helper methods.
- Other enhancements and changes as described below.

New plug-in API

PySys 1.6.0 introduces a new plug-in API that provides a simpler and more modular way to share application-specific utility methods across multiple tests than the traditional approach of creating custom `BaseTest` and `BaseRunner` classes. Apama makes use of this to provide a new `ApamaPlugin` that provides easier access to Apama functionality. To use it, just add this to your `pysysproject.xml`:

```
<test-plugin classname="apama.testplugin.ApamaPlugin" alias="apama"/>
```

You can then access the plug-in's methods using `self.apama`, for example, to start a correlator:

```
correlator = self.apama.startCorrelator('testCorrelator', ...)
```

New way of configuring EPL code coverage reporting

There is a new and easier way to configure EPL code coverage reporting. Just remove the references to `ApamaRunner` (now deprecated) from your `pysysproject.xml` and replace them with:

```
<writer classname="apama.coverage.EPLCoverageWriter">
  <property name="destDir" value="__coverage_epl.${outDirName}"/>

  <!-- The following (optional) properties take a comma-separated list,
       for example, "**/foo/Bar*.mon, **/baz/*.mon".
```

```

    See the EPLCoverageWriter API documentation for details on these and
    other available configuration parameters.
-->
<property name="srcIncludes" value=""/>
<property name="srcExcludes" value=""/>
</writer>

```

See the `EPLCoverageWriter` class in the *API Reference for Python* for more details.

New way of validating Apama events sent by your application

Your PySys tests can now validate the events that are sent out of your Apama application in a new way. The previous approach of validating the output of the `engine_receive` tool using `assertDiff` or `assertGrep/assertOrderedGrep` regular expressions has several disadvantages including uninformative error messages when the match fails, and no easy way to extract the information of interest from complex Apama events. Instead we now introduce a new way to check the events: instead of `CorrelatorHelper.receive(evt, channels= [...])`, call `CorrelatorHelper.injectTestEventLogger(channels=[...])` to request a message to be written to the correlator log when an event is sent to the specified channels, and then in your `validate()` method, use `ApamaPlugin.extractEventLoggerOutput(logfile)` to get a list of events as a standard Python dictionary. For example:

```

sensor1_temps = [
    # Extract only the field value(s) we care about
    # (allows us to ignore unimportant information, timestamps, etc):
    (evt['temperature'])
    for evt in self.apama.extractEventLoggerOutput('testCorrelator.log')

    # Filter to include the desired subset of events:
    if evt['.eventType']=='apamax.myapp.Alert' and evt['sensorId']=='TempSensor001'
]
self.assertThat('sensor1_temps == expected',
    sensor1_temps=sensor1_temps, expected=[
        111.0,
        120,
        145.2,
    ])

```

For a concrete example, see `Apama_cor_001` in the `samples/pysys` directory of your Apama installation.

Property settings affecting compatibility with existing projects

In previous versions of the Apama PySys helper classes, the following properties were set on the project object:

- `APAMA_BIN_DIR`
- `APAMA_COMMON_JRE`
- `APAMA_LIBRARY_VERSION`

As of this version, these properties are no longer set on the project object, except if they are set in the `pysysproject.xml` file. All uses of `APAMA_BIN_DIR` and `APAMA_LIBRARY_VERSION` have been removed from the helper classes. `APAMA_COMMON_JRE` is still read from the `pysysproject.xml` file, but is not set on the project otherwise.

If the above three properties are not set in the `pysysproject.xml` file, a relative directory to `APAMA_HOME` is now used by default. If you have any PySys tests which explicitly use the above properties, for example, `self.project.APAMA_BIN_DIR`, then we recommend you rewrite them to use a path relative to `self.project.APAMA_HOME` instead. As a workaround, you can replicate the original behavior by providing this property in your `pysysproject.xml` file:

```
<property name="APAMA_BIN_DIR" value="${env.APAMA_HOME}/bin"/>
```

Other changes affecting compatibility with existing projects

The new PySys release includes some changes in behavior that may affect existing projects. For example:

- Stricter checking of project configuration files so that mistakes are easier to spot.
- Checks to prevent dangerous practices like editing `os.environ` during test execution.
- Removal of support for `descriptor.xml/.pysystest.xml` (all tests should now be defined with a `pysystest.xml` file).
- Changes to several project properties (such as `defaultAbortOnError`) to match best practice and avoid the need for new projects to copy a long list of properties just to get the recommended behavior.

As a result, some existing PySys projects will need changes to work with the new PySys release. However, in most cases this will just require copying in some properties to `pysysproject.xml` (to maintain the previous behavior) and perhaps adding `default="..."` values for any optional properties, but significant changes to test cases are unlikely to be needed. For full details of how to migrate an existing test project to the new release, including the project properties that you may need to copy in to retain existing behavior, see *PySys Test Framework > Change Log* in the *API Reference for Python*.

Deprecated APIs

See [“Removed and deprecated features in 10.7.0” on page 42](#) for a list of all PySys APIs that are deprecated in this release.

Connectivity plug-ins enhancements in 10.7.0

Apama 10.7.0 includes the following connectivity plug-ins enhancements:

- The custom **HTTP Client** connectivity bundles and also the **HTTP Client Generic Events** EPL bundle now support sending requests via a proxy server, if configured. See also "Configuring the HTTP client transport" in *Connecting Apama Applications to External Components*.
- The Cumulocity IoT transport and the Cumulocity IoT REST API now support sending requests via a proxy server, if configured. See also "Configuring the Cumulocity IoT transport" in *Connecting Apama Applications to External Components*.
- The names of the Cumulocity IoT bundles that you can add using Software AG Designer no longer include a version number. These are the following bundles:
 - Cumulocity Client

- Event Definitions for Cumulocity
- Utilities for Cumulocity

This has the advantage that with future versions, you can seamlessly upgrade to the latest Cumulocity IoT bundles. See also "Adding the Cumulocity IoT connectivity plug-in to a project" in *Using Apama with Software AG Designer*.

- The Cumulocity IoT transport now allows you to set the following properties in a `.properties` configuration file:
 - `CUMULOCITY_ALLOW_UNAUTHORIZED_SERVER`
 - `CUMULOCITY_FORCE_INITIAL_HOST`

In addition, you can now also specify the username configuration option in the form of `tenant/username`.

For more information, see "Configuring the Cumulocity IoT transport" in *Connecting Apama Applications to External Components*.

- The Cumulocity IoT transport now enables you to distinguish between "create" and "update" operations for notifications that are received for managed objects, events, alarms and operations. The new constants `NOTIFICATION_CREATED` and `NOTIFICATION_UPDATED` are available for this purpose. For more information, see "Receiving update notifications" in *Connecting Apama Applications to External Components*. This enhancement was also delivered with Apama 10.5.4.
- The Apache Kafka library which is used by the Kafka transport has been updated to version 2.5.0. This includes support for TLS 1.3 and fixes several vulnerability issues. Some configurations may have changed with the new version. For more details, see the Kafka 2.5.0 release notes at https://downloads.apache.org/kafka/2.5.0/RELEASE_NOTES.html.

Miscellaneous enhancements and changes in 10.7.0

Apama 10.7.0 includes the following miscellaneous enhancements and changes:

- The `apama_project` tool now allows you to add multiple bundles to a project using a single command. Instead of using a separate `add bundle bundlename` command as in previous versions, you can now add multiple bundles using the `add bundle bundlename bundlename ...` command. See also "Creating and managing an Apama project from the command line" in *Deploying and Managing Apama Applications*.
- The EPL memory profiler has been enhanced to include the memory used in any streams networks in your application (for example, from `a in all A() within 10.0 select mean(a.i)`). This includes the memory used by stream data structures such as aggregate functions and `retain` and `within` clauses. This is useful for understanding memory consumption in any application that uses streams (see also "Working with Streams and Stream Queries" in *Developing Apama Applications*).

For information about how to use the EPL memory profiler, see "Using the EPL memory profiler" in *Deploying and Managing Apama Applications*.

- The jemalloc memory allocator library for Linux platforms has now been upgraded to version 5.2.1. See also "Using jemalloc to optimize memory usage" in *Deploying and Managing Apama Applications*.
- A small change has been made to the implementation of listeners using the `at` keyword which changes the behavior of some patterns when times repeat (for example, DST changes). Before, there were some types of pattern which would only execute during the first of repeated sections of time, and others which would execute during both. Now all patterns which match a time in a repeated section of time will execute both in the first and the second repeat of the time. For more details, see "Triggering event listeners at specific times" in *Developing Apama Applications*.

Platform changes in 10.7.0

Apama 10.7.0 runs on the platforms listed in the *Supported Platforms* document for version 10.7. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>.

Check the above mentioned *Supported Platforms* document for details about the newer platform versions supported by Apama 10.7.0. The following platforms are no longer supported:

- Software AG Common Runtime 10.5 (replaced by support for a more recent version).
- Oracle 11g as a JDBC database driver (a later version is still supported).
- IBM DB2 as a JDBC database driver.
- Software AG Universal Messaging 10.5 (replaced by support for a more recent version).
- IBM MQ 8.0 as a JMS provider (a later version is still supported).
- Software AG Universal Messaging Client 10.5 (replaced by support for a more recent version).
- Apache Kafka 1.0.0 (replaced by support for a more recent version).
- SOAP 1.2.
- Software AG Designer 10.5 (replaced by support for a more recent version).
- Software AG BigMemory Max 4.3.8 (replaced by support for a more recent version).
- Software AG Terracotta Store 10.5 (replaced by support for a more recent version).
- Software AG Cumulocity IoT 10.5 (replaced by support for a more recent version).

Miscellaneous changes in 10.7.0 affecting backwards compatibility

The following changes in Apama 10.7.0 affect backwards compatibility to previous Apama versions:

- Software AG Designer now uses Eclipse 4.15. Eclipse 4.15 workspaces are upward-compatible, and workspaces created with older Eclipse versions are automatically migrated to 4.15 workspaces. Downward-compatibility of workspaces, however, is not supported. This means

that a workspace which has been created or opened with Eclipse 4.15 can not be used with an older Eclipse version. Therefore, it is recommended that you back up your workspaces before opening them with Eclipse 4.15.

- As of Apama 10.7.0, the Cumulocity IoT bundles are no longer versioned. Therefore, you have to replace the Cumulocity IoT bundles that you are using. Proceed as follows:

1. Add the new Cumulocity IoT bundles to your Apama projects.
2. Copy any changes you applied to the `CumulocityIoT.yaml` and `CumulocityIoT.properties` files over from the old, versioned bundles into the new bundles.
3. Remove the old, versioned bundles from your Apama projects.

See also [“Connectivity plug-ins enhancements in 10.7.0” on page 38](#).

- Generic requests that are sent using the Cumulocity IoT transport now return the response body with exactly what is returned from Cumulocity IoT. In previous versions, this was put into a dictionary with an empty key. So if you have existing code such as the following

```
AnyExtractor dict :=
AnyExtractor(AnyExtractor(response.body).getDictionary("")[""]);
```

you have to replace this with:

```
AnyExtractor dict := AnyExtractor(response.body);
```

See also "Invoking other parts of the Cumulocity IoT REST API" in *Connecting Apama Applications to External Components*.

- To be more compliant with the HTTP RFC and IANA standards, when automatically constructing the content-type header for HTTP requests, the HTTP client transport now only appends the charset parameter from the `metadata.charset` field if the content type is in the `text/*` domain. If you need to add a charset for any other content type, you can explicitly add it by setting the `metadata.http.headers.content-type` field to include the charset modifier. See "Handling HTTP headers" in *Connecting Apama Applications to External Components* for more details.
- In 10.7, the defaults of the two configuration options `cookieJar` and `followRedirects`, which are used by the HTTP client transport, have changed. Since Apama 10.3.1, Software AG Designer projects created with the **HTTP Client** connectivity bundle have had these set to true in the bundle configuration, and for those projects nothing will change whether you had the default or changed the configuration in your bundle instance (see also the backwards compatibility note in [“Connectivity plug-ins enhancements in 10.3.1” on page 77](#)). For older projects or those using the HTTP client transport in some other fashion, the default is now true. To disable these options, set them to false in the transport configuration. See "Configuring the HTTP client transport" in *Connecting Apama Applications to External Components* for more details.
- The JSON codec connectivity plug-in now safely handles empty payloads and now converts "" to {} instead of throwing an error in the towards-host direction. See also "The JSON codec connectivity plug-in" in *Connecting Apama Applications to External Components*.
- If the `engine_deploy` tool is executed with an `--include` or `--exclude` option with an argument containing an unsupported pattern (such as `./`), the tool will now exit with a non-zero exit

code, abort execution and print an error message. With previous versions, a warning was printed in this case. See also "Deploying a correlator" in *Deploying and Managing Apama Applications*.

- To avoid confusion, creating user status items through either EPL or from plug-ins now strips any leading or trailing whitespace from key names. This avoids having two distinct keys with visually identical key names. This means that two keys which previously differed only in whitespace will now cause a conflict for having the same name when using the `createUserStatusItem` interface. When using the `setUserStatus` interface, they will now update the same value. Keys which previously had trailing whitespace that caused them not to appear in the Prometheus statistics will now appear. Status values in the correlator status REST API will now appear without whitespace if they had them before. In addition, empty keys, or keys only consisting of whitespace, are now illegal and will throw an exception. See also "Using the Management interface" in *Developing Apama Applications* and "User-defined status reporting from connectivity plug-ins" in *Connecting Apama Applications to External Components*.
- The Apama images that are available on Docker Hub (`apama-correlator`, `apama-builder` and `apama-cumulocity-jre`) have been changed to use Red Hat UBI 8 minimal as their base operating system. This may affect users who are extending those images. In addition to the upgrade from 7 to 8, the switch to a minimal image results in much smaller images, but as a result some packages that a user relies on may be missing. In particular, this includes the `yum` tool for installing packages. Instead, the `microdnf` tool must be used. We have taken the opportunity to also ensure that all of the images will work with non-UTC time zones out of the box. However, the only locale which is shipped is the new default `C.UTF-8` locale (previously, it was `en_US.UTF-8`). Any other locales which are required must be added to the image. See also "Deploying Apama Applications with Docker" in *Deploying and Managing Apama Applications*.
- In Apama 10.7.0, the JDBC DataDirect driver version 6.0.0 for Microsoft SQL Server has been moved from `APAMA_HOME/lib/datadirect-jdbc/6.0.0/eyssqlserver.jar` to `APAMA_HOME/lib/eyssqlserver.jar`. Software AG Designer projects using the old location for the driver classpath will need to update the configuration of their JDBC adapter to point to `APAMA_HOME/lib/eyssqlserver.jar`.
- The new PySys release includes some changes in behavior that may affect existing projects. See ["PySys changes in 10.7.0" on page 36](#) for more information.

Note:

There may also be removed features that affect backwards compatibility. See ["Removed and deprecated features in 10.7.0" on page 42](#) for more information.

Removed and deprecated features in 10.7.0

The following Apama features are now deprecated or have been removed in Apama 10.7.0:

- With the new PySys release, several PySys APIs are now deprecated. This includes the following:
 - Instead of `-Xep1coverage`, use `-XcodeCoverage` to enable EPL code coverage. If you need to enable EPL coverage without enabling any other types of code coverage in your project, this can be achieved with `-Xep1Coverage` (note the capitalization of "C").

- `ApamaRunner` is deprecated. Remove it from your `pysysproject.xml` and instead configure EPL code coverage as described in [“PySys changes in 10.7.0” on page 36](#).
- `ApamaBaseTest` is not currently deprecated, but its use is no longer encouraged. New tests subclass `pysys.basetest.BaseTest` rather than `ApamaBaseTest`.
- The `PySys ThreadFilter` class is deprecated and will soon be removed. Usually, it is not recommended to suppress log output and better alternatives are available, for example, the `quiet=True` option for `BaseTest.startProcess`, and the `BaseTest.disableLogging()` method.
- The `PySys` method `pysys.basetest.BaseTest.addResource` is deprecated and will be removed in a future release. Therefore, change your tests to stop using it. Use `pysys.basetest.BaseTest.addCleanupFunction` instead.

See also [“PySys changes in 10.7.0” on page 36](#).

- In previous releases, the correlator and IAF status key names `numSnapshots` and `virtualMemorySize` were deprecated and replaced with the keys `persistenceNumSnapshots` and `virtualMemoryMB`. The deprecated keys have now been removed and will no longer appear in the status output. This applies to the following:
 - the REST API (see also "Managing and Monitoring over REST" in *Deploying and Managing Apama Applications*),
 - the `engine_management`, `component_management` and `iaf_management` tools (see also "Shutting down and managing components" in *Deploying and Managing Apama Applications*), and
 - the C++, .NET and Java client APIs (see also "The Client Software Development Kits" in *Connecting Apama Applications to External Components*).
- The deprecated SOAP-based Web Services Client IAF adapter has been removed. It is recommended that you now use a REST-based HTTP server instead and connect to it with the HTTP client transport, if possible (see "The HTTP Client Transport Connectivity Plug-in" in *Connecting Apama Applications to External Components*). However, if you really need to communicate with a SOAP-based server, use the webMethods Integration Server and connect to it from Apama using the Digital Event Services transport (see "The Digital Event Services Transport Connectivity Plug-in" in *Connecting Apama Applications to External Components*).
- Support for Software AG's Event-Driven Architecture (EDA) has been removed. This was deprecated in Apama 10.1. Instead of using EDA events, it is recommended that your Apama applications use digital event types, as provided by Software AG Digital Event Services. For further information, see "The Digital Event Services Transport Connectivity Plug-in" in *Connecting Apama Applications to External Components*.
- Several minor features which have been deprecated over the past few years have been removed. The following table lists all of the removed APIs or options and their recommended replacements.

Removed item	Replacement
<code>float.nextafter</code>	<code>float.nextAfter</code>

Removed item	Replacement
<code>decimal.nextafter</code>	<code>decimal.nextAfter</code>
<code>dictionary.getDefault</code>	<code>dictionary.getOr</code>
<code>aggregate prior</code>	<code>aggregate nth</code>
<code>HttpTransport. CONFIG_TLS_ACCEPT_UNRECOGNIZED_CERTS</code>	<code>HttpTransport. CONFIG_TLS_ACCEPT_UNRECOGNIZED_CERTS</code>
<code>correlator --period</code>	<code>correlator --frequency</code>
<code>correlator -Xconfig</code> or <code>correlator --configFile</code>	<code>correlator --config</code>
<code>engine_receive -Xconfig</code> or <code>engine_receive --configFile</code>	No replacement.
<code>Component.getComponentLogicalId</code>	<code>Component.getComponentPhysicalId</code>
<code>engine_connect --persistent</code>	No replacement, all connections are persistent now.
<code>ADBC: Discovery.findAvailableServers</code>	<code>Discovery.findAvailableDataSources</code>
<code>ADBC: Discovery.findAvailableServersFull</code>	<code>Discovery.findAvailableDataSourcesFull</code>
<code>ADBC: Discovery.reopen</code>	<code>Discovery.reopenWithAck</code>
<code>ADBC: Discovery.getQueryResultToken</code>	No replacement.
<code>ADBC: Discovery.setRunUntilTime</code>	No replacement.
<code>ADBC: Discovery.getBatchDoneToken</code>	No replacement.
<code>ADBC: DBUtil.handleStatus</code>	No replacement.
<code>ADBC: DBReconnectPolicy.RECONNECT</code> and <code>DBReconnectPolicy.RECONNECT_AND_CLEAR_REQUEST_QUEUE</code>	<code>DO_NOT_RECONNECT</code> or <code>RECONNECT_AND_RETRY_LAST_REQUEST</code>
<code>ADBC: DBUtil.subscribeConnectionStatus</code>	<code>DBUtil.subscribeStatus</code>
C++ and Java connectivity APIs: <code>MapHelper</code>	<code>MapExtractor</code>
C++ EPL plug-in: <code>CorrelatorInterface::sendEvent</code>	<code>CorrelatorInterface::sendEventTo</code> To replicate <code>sendEvent</code> , call <code>sendEventTo</code> with the empty string for the channel.
Java EPL plug-in: <code>Correlator.enqueue</code> and <code>Correlator.enqueueTo</code>	<code>Correlator.sendTo</code> To replicate <code>enqueue</code> , call <code>sendTo</code> with the empty string for the channel. <code>sendTo</code> is a direct compatible replacement for <code>enqueueTo</code> .

Removed item	Replacement
<code>Chain.createChain</code>	<code>Chain.createDynamicChain</code> To replace <code>createChain</code> with <code>createDynamicChain</code> , pass the empty string as the new final argument.
PySys: <code>CorrelatorHelper.waitForCorrelatorUp</code> and <code>IAFHelper.waitForIAFUp</code>	<code>waitForComponentUp</code> This is available on the base class for all helper objects.
<code>apama-macros.xml</code> Ant script: <code>parallel</code> attribute for the <code>engine-management</code> and <code>iaf-management</code> macros	No replacement.
<code>apama-macros.xml</code> Ant script: <code>engine-inject-filelistid</code> macro	<code>inject-project-init-list</code>

6 What's New In Apama 10.5.4

- Connectivity plug-ins enhancements in 10.5.4 48
- Miscellaneous changes in 10.5.4 affecting backwards compatibility 48

Apama 10.5.4 runs on the platforms listed in the *Supported Platforms* document for version 10.5. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>. Be sure to consult that document for details about supported versions of operating systems, compilers and other products.

Apama 10.5.4 includes new features, enhancements, and changes as described in the following topics.

Connectivity plug-ins enhancements in 10.5.4

Apama 10.5.4 includes the following connectivity plug-ins enhancements:

- The HTTP server now allows concurrent processing using multiple chains for multiple connections. The default is `false`, but you can turn it on by setting `concurrentChains` to `true` in the YAML configuration file. See also "Configuring the HTTP server transport" in *Connecting Apama Applications to External Components*.
- The new monitoring metric `numChains` is now available for the HTTP server transport. It indicates the number of active connections into the HTTP server instance. See also "Monitoring status for the HTTP server" in *Connecting Apama Applications to External Components*.
- The Cumulocity IoT transport now enables you to distinguish between "create" and "update" operations for notifications that are received for managed objects, events, alarms and operations. The new constants `NOTIFICATION_CREATED` and `NOTIFICATION_UPDATED` are available for this purpose. For more information, see "Receiving update notifications" in *Connecting Apama Applications to External Components*.

Miscellaneous changes in 10.5.4 affecting backwards compatibility

The following changes in Apama 10.5.4 affect backwards compatibility to previous Apama versions:

- The JSON codec connectivity plug-in now safely handles empty payloads and now converts `""` to `{}` instead of throwing an error in the towards-host direction. See also "The JSON codec connectivity plug-in" in *Connecting Apama Applications to External Components*.

7 What's New In Apama 10.5.3

■ PySys changes in 10.5.3	50
■ Connectivity plug-ins enhancements in 10.5.3	50
■ Miscellaneous enhancements and changes in 10.5.3	51

Apama 10.5.3 runs on the platforms listed in the *Supported Platforms* document for version 10.5. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>. Be sure to consult that document for details about supported versions of operating systems, compilers and other products.

Apama 10.5.3 includes new features, enhancements, and changes as described in the following topics.

PySys changes in 10.5.3

Apama 10.5.3 now ships a newer version of the PySys test framework, version 1.5.1.

If you are already familiar with PySys, it is strongly recommended to read the full Change Log in detail to familiarize yourself with the new features and new best practices for improving how you write tests. There is more detail on many PySys features in the User Guide. You can find the Change Log and User Guide in the *API Reference for Python*, which you can also access via the `doc/pydoc/index.html` file of your Apama installation.

Key features for the new PySys 1.5.1 release include:

- The documentation is vastly expanded and easier to navigate. It is recommended that you bookmark the BaseTest class (`/doc/pydoc/pysys-doc/BaseTest.html`) as this is the main page that you will want to refer to regularly when writing PySys tests.
- `assertThat`, `assertGrep` and `waitForGrep` (formerly known as `waitForSignal`) have lots of new functionality.
- The project file has a new `<project-help>` element that allows you to customize the `pysys run --help` output.
- Many bug fixes.

A new Apama helper class for PySys called `ProjectHelper` is now available in the `apama.project` package. It provides methods for creating, manipulating and deploying Software AG Designer-compatible projects from your PySys tests and thus provides access to the `apama_project` tool; see "Creating and managing an Apama project from the command line" in *Deploying and Managing Apama Applications*.

Connectivity plug-ins enhancements in 10.5.3

Apama 10.5.3 includes the following connectivity plug-ins enhancements:

- Using the new `set` rule of the Mapper codec, you can now set the value of a field to a specific value, regardless of whether that field already exists in the event. For more information, see "The Mapper codec connectivity plug-in" in *Connecting Apama Applications to External Components*.
- The HTTP client generic event definitions now allow you to invoke services with non-JSON string payloads. For more information, see "Using predefined generic event definitions to invoke HTTP services with JSON and string payloads" in *Connecting Apama Applications to External Components*.

Miscellaneous enhancements and changes in 10.5.3

Apama 10.5.3 includes the following miscellaneous enhancements and changes:

- Apama 10.5.3 incorporates the ICU (International Components for Unicode) time-zone data update 2020a, which is the most recent update at the time of release. This will update time-zone data used by the correlator and Time Format EPL plug-in.
- You can now enable verbose garbage collection logging and application event logging by using the `--loglevel` option of the `correlator` executable or by using a YAML configuration file on correlator startup. In previous releases, this was only possible for a running correlator via the `-r` option of the `engine_management` tool. As a result, these messages are now logged with a `DEBUG` prefix when enabled, rather than `INFO`. There is also a small change to the format of the messages logged. See also "Viewing garbage collection and application events information" in *Deploying and Managing Apama Applications*.
- It is now possible to monitor the status for the `MemoryStore` via the user status mechanism. See also "Monitoring status for the `MemoryStore`" in *Developing Apama Applications*.
- The REST API and Prometheus now provide statistics information about the number of threads in use by the JVM: `jvmNumThreads` (REST API) and `sag_apama_correlator_jvm_num_threads` (Prometheus). These statistics will exist when the Java Virtual Machine (JVM) which is embedded in the correlator has been enabled. See also "List of correlator status statistics" in *Deploying and Managing Apama Applications*.
- The `@since` tag can now be used in ApamaDoc comments on event members and constants as well as for events and actions. See also "Inserting ApamaDoc tags" in *Developing Apama Applications*.

8 What's New In Apama 10.5.2

■ Cumulocity API enhancements in 10.5.2	54
■ Miscellaneous enhancements and changes in 10.5.2	54
■ Miscellaneous changes in 10.5.2 affecting backwards compatibility	55

Apama 10.5.2 runs on the platforms listed in the *Supported Platforms* document for version 10.5. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>. Be sure to consult that document for details about supported versions of operating systems, compilers and other products.

Apama 10.5.2 includes new features, enhancements, and changes as described in the following topics.

Cumulocity API enhancements in 10.5.2

Some events have been added to the Cumulocity API to allow easy retrieval of information in some predefined areas.

- You can get the values for all the tenant options in the tenant, or search for specific tenant options.
- You can look up the details and roles of the user you connected as, or the microservice service user. Alternatively, you can forward on the `Authorization` header or OAuth cookies of an incoming request to validate the roles of that user.
- Measurement events can be split into individual fragments and series to allow better performance when filtering on incoming events.

The Cumulocity API now allows the ability to retrieve a response for the creation or update of the predefined types (`Alarm`, `Operation`, `Event`, `Measurement`, `MeasurementFragment` and `ManagedObject`) by using the newly added `withResponse` action. Headers can now be set on the request, for example, to determine the processing mode within Cumulocity IoT.

For detailed information, see "The Cumulocity IoT Transport Connectivity Plug-in" in *Connecting Apama Applications to External Components*.

Miscellaneous enhancements and changes in 10.5.2

Apama 10.5.2 includes the following miscellaneous enhancements and changes:

- Using the new `--configLog file` option of the `correlator` executable, you can now specify the path to the configuration log file. The configuration log file contains the correlator's configuration, that is, the contents of all YAML configuration files and properties files as well as the correlator startup arguments and environment. The configuration log is a good way to capture useful diagnostic information from the correlator's startup in performance-critical situations where it is not acceptable to enable the input log. This option was already added with 10.5.1. However, the documentation has only been updated with 10.5.2. See also "Starting the correlator" in *Deploying and Managing Apama Applications*.
- The results of a `eplMemoryProfileMonitorInstanceDetail` request are now returned using the following sort order:
 1. monitor name (ascending)
 2. monitor instance ID (ascending)
 3. EPL type (ascending)

See also "Using the EPL memory profiler" in *Deploying and Managing Apama Applications*.

- The startup log stanza is now repeated even when logging is going into `stdout/stderr` via the `--rotateLogs` management request while using the `engine_management` tool (see also "Shutting down and managing components" in *Deploying and Managing Apama Applications*) or by using `SIGHUP` in Linux.

Miscellaneous changes in 10.5.2 affecting backwards compatibility

The following changes in Apama 10.5.2 affect backwards compatibility to previous Apama versions:

- Because it is not safe to iterate over a `list_t` or `map_t` container type while modifying it, these iterators have been modified to assert (or throw an exception) if accessed after the parent container has been modified. These will fail with a message saying that the iterator has been invalidated by modification of either a map or a list. Any code which now throws an exception will have been unsafe previously (even if it did not crash). The new check simply exposes the existing bug in the user's code. See also "C++ data types" in *Connecting Apama Applications to External Components*.
- If you are compiling either EPL or connectivity plug-ins in C++, and you build libraries that are used by multiple plug-ins and those libraries have an API using Apama's C++ data types, you may need to change how you compile your plug-ins. If you have multiple shared objects with such an API, you must ensure that all of the shared objects are compiled with the same service pack version of the Apama header files. This does not affect you if you are compiling independent plug-ins. For more details, see "Compiling C++ plug-ins" in *Developing Apama Applications* and "Building plug-ins" in *Connecting Apama Applications to External Components*.

9 What's New In Apama 10.5.1

■ PySys enhancements in 10.5.1	58
■ Connectivity plug-ins enhancements in 10.5.1	58
■ EPL enhancements in 10.5.1	58
■ Miscellaneous enhancements and changes in 10.5.1	59
■ Platform changes in 10.5.1	59
■ Support for Red Hat Enterprise Linux (RHEL) 8.0 in 10.5	59
■ Miscellaneous changes in 10.5.1 affecting backwards compatibility	59

Apama 10.5.1 runs on the platforms listed in the *Supported Platforms* document for version 10.5. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>. Be sure to consult that document for details about supported versions of operating systems, compilers and other products.

Apama 10.5.1 includes new features, enhancements, and changes as described in the following topics.

PySys enhancements in 10.5.1

The PySys `apama.correlator` and `apama.iaf` modules have been updated to accept arguments such as `logfile` and `verbosity` via the constructor as an alternative to specifying them in the `start()` method. Any individual arguments specified in the `start(...)` method will take precedence over the constructor arguments. See the *API Reference for Python* for more information.

Connectivity plug-ins enhancements in 10.5.1

Apama 10.5.1 includes the following connectivity plug-ins enhancements:

- Additional metrics are now available for the HTTP client and HTTP server transports. For details, see "Monitoring status for the HTTP client" and "Monitoring status for the HTTP server", both in *Connecting Apama Applications to External Components*.
- The Cumulocity IoT transport now provides status values via the user status mechanism and allows you to tune the batching and status reporting by setting the properties in a `.properties` configuration file. For more information, see "Monitoring status for Cumulocity IoT" and "Configuring the Cumulocity IoT transport", both in *Connecting Apama Applications to External Components*.

EPL enhancements in 10.5.1

Apama 10.5.1 includes the following EPL enhancements:

- The `ondie` action can now have either no argument or the arguments `optional<com.apama.exceptions.Exception>` and `optional<string>`. See also "About executing `ondie()` actions" in *Developing Apama Applications*.
- In a sequence of sequences, in a dictionary of sequences or in cast operations, it is no longer required to add a space between two closing angle brackets (`>`). A syntax error no longer occurs when the space is missing. For example, instead of

```
dictionary<decimal, <sometype> >
```

you can now write

```
dictionary<decimal, <sometype>>
```

Both are valid EPL code.

Miscellaneous enhancements and changes in 10.5.1

Apama 10.5.1 includes the following miscellaneous enhancements and changes:

- Apama 10.5.1 incorporates the ICU (International Components for Unicode) time-zone data update 2019c, which is the most recent update at the time of release. This will update time-zone data used by the correlator and Time Format EPL plug-in.
- The `/request` REST API on the correlator can now be called with a JSON payload as an alternative to the existing XML payloads. Responses can also be received in JSON. For more details, see the *API Reference for Component Management REST APIs*.
- The REST interfaces to all components for retrieving the environment variables of the process now replace the values of any environment variables with the names containing `PASSWORD`, `PASSWD` or `CREDENTIALS` with the string “<redacted>” for security.
- Apama now provides an `apamadoc` command-line tool which can be used to generate ApamaDoc. For more details, see "Generating ApamaDoc from the command line" in *Developing Apama Applications*.

Platform changes in 10.5.1

Apama 10.5.1 runs on the platforms listed in the *Supported Platforms* document for version 10.5. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>.

Due to these upgrades, the following versions are no longer supported:

- Oracle Linux 7 (x86) 64-bit (support has also been removed from 10.5.0).
- Otherwise, no changes as compared to “Platform changes in 10.5.0” on page 68.

Support for Red Hat Enterprise Linux (RHEL) 8.0 in 10.5

Apama 10.5.0 and 10.5.1 support running on RHEL 8.0. If you build your own C++ EPL or connectivity plug-ins on RHEL 8.0, you will have to add the following flags to your compiler command-line options:

```
-D_GLIBCXX_USE_CXX11_ABI=0 -Wno-parentheses -Wno-class-memaccess -Wno-aligned-new
```

Miscellaneous changes in 10.5.1 affecting backwards compatibility

The following changes in Apama 10.5.1 affect backwards compatibility to previous Apama versions:

- If the correlator or the IAF is configured to log to a file, then the log file being used is now output to `stderr`. In the case of the correlator, this is the full path and filename, whereas the IAF only logs the filename.

- The shutdown order of the connectivity chain components has changed. In previous versions, the codecs were shut down first, followed by the transport and host. As of this version, the host is shut down first (that is, calls into the host component of the chain are unblocked), followed by the codecs and the transport. See also "Lifetime of connectivity plug-ins" in *Connecting Apama Applications to External Components*.

When you develop your own connectivity plug-ins, you may have to update your plug-ins to reflect the changed behavior. Make sure that your testing includes shutting down the correlator (for example, using the PySys correlator.shutdown() method), since it is possible that this change in behavior makes shutdown issues in user-defined plug-ins more likely to show up.

10 What's New In Apama 10.5.0

■ New Apama image on Docker Hub in 10.5.0	62
■ Change of Java Log4j logging library in 10.5.0	62
■ PySys changes in 10.5.0	63
■ Connectivity plug-ins enhancements in 10.5.0	65
■ Apama enhancements in Software AG Designer in 10.5.0	66
■ EPL enhancements in 10.5.0	67
■ Miscellaneous enhancements and changes in 10.5.0	67
■ Platform changes in 10.5.0	68
■ Miscellaneous changes in 10.5.0 affecting backwards compatibility	68
■ Removed and deprecated features in 10.5.0	69
■ Migrating the Cumulocity IoT transports to 10.5	73

Apama 10.5.0 is the successor of Apama 10.3.1. There is no version 10.4.

Apama 10.5.0 runs on the platforms listed in the *Supported Platforms* document for version 10.5. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>. Be sure to consult that document for details about supported versions of operating systems, compilers and other products.

Apama 10.5.0 includes new features, enhancements, and changes as described in the following topics.

New Apama image on Docker Hub in 10.5.0

A new Apama base image called “apama-cumulocity-jre” is available on Docker Hub. It can be used as a smaller image for Cumulocity IoT microservices. It only contains a correlator and the Cumulocity IoT connectivity pieces. There is a JRE but no JDK, and the image does not come with Python support by default. See also “Apama image for Cumulocity IoT microservices” in *Deploying and Managing Apama Applications*.

For more details, see the installation instructions for this image on Docker Hub.

Note:

This image is also available for Apama version 10.3.1.

Change of Java Log4j logging library in 10.5.0

The correlator, the IAF, the dashboard servers, and the `com.apama.util.Logger` class used by the Java client library now use Log4j 2 rather than Log4j 1.

In Software AG Designer, Apama projects with a Java nature now use Log4j 2. If you are still using Log4j 1.2, you must now upgrade your applications to use Log4j 2 or manually add the Log4j-1-to-2 bridge or the Log4j-1-to-SLF4J bridge to your classpath from `common/lib/ext/log4j`.

In Log4j 2, the configuration file format, the system property that specifies the configuration file location, and the Java package are all completely different to Log4j 1. So if your application directly uses the Log4j 1 API (`org.apache.log4j.*`) or you have any custom Log4j configuration files (for example, a `log4j.properties` file), you will need to change them when upgrading to this Apama release. See the Log4j website at <https://logging.apache.org/log4j/2.x/> for detailed information on Log4j 1.x migration.

The Log4j 1.2 configuration files are no longer used. Thus, the `log4jConfigFile` configuration option for controlling the verbosity, which was previously used by the Kafka and Digital Event Services transports, is no longer supported in an Apama project. If you are using the `log4jConfigFile` configuration option, you should now set the log levels in the YAML configuration file for the correlator instead, using the `correlatorLogging` entry. See also “Setting correlator and plug-in log files and log levels in a YAML configuration file” in *Deploying and Managing Apama Applications*.

If you are using `com.apama.util.Logger` within the correlator or IAF processes, no changes should be required. However, if you use `com.apama.util.Logger` in a standalone Java application, perhaps using the Apama Engine Client or Scenario Service APIs, you will need to provide a Log4j 2 configuration file even if previously you did not, as the `com.apama.util.Logger` class no longer

performs automatic configuration of Log4j. Again, see the Log4j website for sample configuration files and detailed information on the configuration file format.

You may also need to take action if you use a third-party library inside the correlator that itself performs logging, for example, a connectivity plug-in, EPL plug-in, or JMS client library. The correlator's root classloader now contains Log4j 2, the widely-used SLF4J logging facade, and the Log4j 2 SLF4J implementation binding. So any libraries using SLF4J or Log4j 2 will correctly write to the correlator log file automatically, and their log levels can be customized using `correlatorLogging` in the YAML configuration file (as described in "Setting correlator and plug-in log files and log levels in a YAML configuration file" in *Deploying and Managing Apama Applications*). The root classloader also contains the Log4j 1.x bridge (`log4j-1.2-api.jar`) which redirects calls to most of the Log4j 1 API through to Log4j 2, so provided you do not use internal features such as programmatic configuration, Log4j 1 libraries should work fine. However, we recommend that you transition away from Log4j 1 since it has reached end-of-life status. If some other logging library is used which does not support SLF4J directly, you should add a jar that bridges from that library's API to SLF4J or Log4j2 if possible.

Previous Apama versions included the jars for Log4j 1 and SLF4J in the `lib` directory of your Apama installation. In this release, the SLF4J and Log4j logging jars are now located in the Software AG installation directory under `common/lib/ext` and `common/lib/ext/log4j`. Any build scripts (for example, Apache Ant `build.xml` scripts) that need these jars at compilation time should be updated to point to the new location.

The IAF-based Web Services Client adapter now writes some of its logging output to standard output rather than to a hardcoded file called `iaf.log` as in the previous version, though the messages written to standard output are likely to be useful only for debugging purposes. This behavior can be customized in `APAMA_HOME\adapters\config\log4j-iaf.xml` if needed. If you have an existing configuration, you should edit the `WebServicesClientAdapter.xml` file in your project to use the Log4j 2 configuration file and system property instead of the old `log4j.configuration` property:

```
<property name="log4j.configurationFile"
value="file:///@APAMA_HOME@/adapters/config/log4j-iaf.xml"/>
```

PySys changes in 10.5.0

Apama 10.5.0 now ships a newer version of the PySys system testing framework, version 1.5.0.

PySys 1.4.0 includes improvements to the writer API such as integrated support for running tests using Travis CI. New test projects can be created with the new `pysys.py makeproject` command. Among other additions, it includes support for creating background threads from test cases, a new `skipTest()` method, an `IS_WINDOWS` constant, and a new `stdouterr` argument for `ProcessUser.startProcess()` which makes it easier to specify to which files process output is to be written. There are also some important bug fixes in areas such as international (I18N) character support, logging, assertion error handling, cleanliness of the `startProcess` environment, and process monitoring.

PySys 1.5.0 brings some significant new features, including support for running a test in multiple modes and `pysysdirconfig.xml` files that allow you to specify defaults that apply to all test cases under a particular directory - such as groups, modes, a prefix to add to the start of each test identifier, and a numeric hint to help define the execution order of your tests. There are also new features in the `pysys run` and `pysys print` command lines, and a host of small additions to the

API to make test creation easier, for example, `assertEval`, `copy` (with filtering of each copied line) and `write_text` (for easy programmatic creation of files in the output directory).

If you are already familiar with PySys, it is strongly recommended to read the full PySys change log in detail to familiarize yourself with the new features and new best practices for improving how you write tests. You can find the PySys change log (`PYSYS_CHANGELOG.html`) in the `doc/pydoc` directory of your Apama installation. There is more detail on many PySys features in the user guide (`PYSYS_USERGUIDE.html`) available in the same location.

PySys 1.4.0 and 1.5.0 include some changes that might require changes to existing tests, including:

- Stricter checking of unknown or invalid keyword arguments passed to `assert*` methods. Incorrect arguments that were silently ignored in previous releases will now result in a test failure to alert you to the fact that they need fixing.
- The environment variables passed to processes started by `startProcess` no longer include environment variables of the parent PySys process by default. Note that this does not affect correlator or IAF processes which continue to be started with the same environment variables as PySys itself.
- Changes to the API and the default columns of the memory/CPU process monitoring framework.
- It is now mandatory to provide a `pysysproject.xml` file in the root directory containing your tests.
- Errors and typos in `pysystest.xml` XML descriptors will now prevent any tests from running (previously, they were just logged). Since an invalid descriptor prevents the associated test case from reporting a result, the new behavior ensures such mistakes will be spotted and fixed promptly. If you have any non-PySys files under your PySys project root directory with names such as `descriptor.xml` which PySys would normally recognize as test cases, you can avoid errors by adding a `.pysysignore` file to prevent PySys looking into that part of the directory tree.
- `ProcessUser.mkdir` now returns the absolute path (including the output directory) instead of just the relative path that was passed in. This makes it easier to use in-line while performing operations such as creating a file in the new directory. Code that relied on the old behavior of returning the path that was passed in may need to be updated to avoid having the output directory specified twice. If you are using `os.path.join`, then no change will be required.
- If you have a custom `BaseRunner` subclass, note that the `self.output` variable in `BaseRunner` is no longer set to the current directory, but instead to a `pysys-runner-OUTDIR` subdirectory of the test root (or to `OUTDIR/pysys-runner` if `OUTDIR` is an absolute path). This directory is not created (or cleaned) automatically, so if you need to write to `self.output`, you should add code to do that to your runner.
- The `PROJECT` variable in the constants module is deprecated. Use `self.project` instead (which is defined on classes such as `BaseTest` and `BaseRunner`).

If you have existing PySys tests, we recommend copying some or all of the properties from `APAMA_HOME/samples/pysys/pysysproject.xml` into your own `pysysproject.xml`, especially `supportMultipleModesPerRun`, `defaultAbortOnError`, `defaultIgnoreExitStatus`, and the `default-file-encodings` section. However, this will change the behavior of existing tests. See the PySys documentation and change log for further details.

We also strongly recommend you to update the `requires-pysys` and `requires-python` elements in your `pysysproject.xml`, to ensure that anyone attempting to run them using an older version than they were designed for gets a clear error message rather than just a subtle difference in behavior.

See the *Upgrade guide and compatibility* section in the above mentioned PySys change log for further details.

The documentation has been updated with information on how to get started with PySys. See "Testing Apama Applications Using PySys" in *Developing Apama Applications*.

Connectivity plug-ins enhancements in 10.5.0

Apama 10.5.0 includes the following connectivity plug-ins enhancements:

MQTT transport

- The MQTT transport has been changed to use a single connection rather than one per channel. This means that there is a single client identifier, which you can now also configure. The new configuration options `mqttClientId` and `cleanSession` are now available for this purpose. See also "Configuring the connection to MQTT" in *Connecting Apama Applications to External Components*.

HTTP transports

- The HTTP server transport now only accepts requests after an application has called the `onApplicationInitialized()` method in the correlator. As long as this method is not called, the server returns a "503 Host Not Ready" HTTP response. See also "HTTP response codes" in *Connecting Apama Applications to External Components*.
- Both the HTTP server and HTTP client transports now support gzip and deflate encoding (decoding is already supported since version 10.3.1):
 - For the HTTP server transport, the new configuration option `responseCompression` is available for the chain on the HTTP server. See also "Configuring the HTTP server transport" in *Connecting Apama Applications to External Components*.
 - For the HTTP client transport, the new field `metadata.http.headers.content-encoding` can now be added to the client request. See also "Mapping events between EPL and HTTP client requests" and "Content-encoding example", both in *Connecting Apama Applications to External Components*.

Cumulocity IoT transport

- A new version (10.5) of the Cumulocity IoT transport connectivity plug-in is now available. See also "[Migrating the Cumulocity IoT transports to 10.5](#)" on [page 73](#).
- The Cumulocity IoT transport now forwards delete notifications for managed objects, measurements and events to Apama using the following EPL events:

- `com.apama.cumulocity.ManagedObjectDeleted`
- `com.apama.cumulocity.MeasurementDeleted`
- `com.apama.cumulocity.EventDeleted`

See the *API Reference for EPL (ApamaDoc)* for more information.

- A new request/response pair has been added to the Cumulocity IoT transport event API. This allows you to make arbitrary requests to the Cumulocity API to support request types which are not provided as built-in events in the API. See also "Invoking other parts of the Cumulocity IoT REST API" in *Connecting Apama Applications to External Components*.
- A new `CumulocityRequestInterface` event has been added to the Cumulocity IoT transport which allows you to invoke other microservices within Cumulocity IoT from an EPL application. This can be used from an Apama instance outside of Cumulocity IoT and within Cumulocity IoT, either as an EPL application or a custom microservice. See also "Invoking microservices" in *Connecting Apama Applications to External Components*.
- The `subscribeToOperations` configuration option of the Cumulocity IoT transport is set to `false` by default when it is not explicitly specified in the YAML configuration file. However, for your convenience, when you add a new **Cumulocity Client** connectivity bundle in Software AG Designer or by using the `apama_project` tool, this option is now already set to `true` in the resulting YAML configuration file.

JSON codec

- The JSON codec now supports any content type that matches the pattern `“^application/([^\/*+]?)?json$”` when the configuration option `filterOnContentType` is set to `true`. See also "The JSON codec connectivity plug-in" in *Connecting Apama Applications to External Components*.

Apama enhancements in Software AG Designer in 10.5.0

Apama 10.5.0 includes the following enhancements in Software AG Designer:

- In Software AG Designer, some of the commands, control names and window titles containing the old term "MonitorScript" have been replaced with new names. The following table lists the new names:

Old name	New name
MonitorScript Build Path	Apama Build Path
MonitorScript File Selection	Add Apama External Dependency
MonitorScript Path Variables	Apama Build Path Variables
MonitorScript Search	Apama Search
Export as MonitorScript	Scenario EPL

Old name	New name
Hide non-MonitorScript projects	Hide non-Apama projects
Apama MonitorScript Editor	Apama EPL Editor
Launch/Inspect MonitorScript Application	Launch/Inspect Apama Application
Edit MonitorScript path	Edit EPL Path
MonitorScript Export Wizard	Scenario EPL Export Wizard

Also, the directory `samples/monitorscript` has been replaced with `samples/epl`.

EPL enhancements in 10.5.0

Apama 10.5.0 includes the following EPL enhancements:

- Expressions similar to `10.rand()` are now valid. Previously, this resulted in a compiler error as the compiler used to treat “10.” as a float. The workaround in previous versions was to either add a space after the integer (for example, `10 .rand()`) or to enclose the integer in parentheses (for example, `(10).rand()`). This is no longer required.

Miscellaneous enhancements and changes in 10.5.0

Apama 10.5.0 includes the following miscellaneous enhancements and changes:

- Apama 10.5.0 incorporates the ICU (International Components for Unicode) time-zone data update 2019b, which is the most recent update at the time of release. This will update time-zone data used by the correlator and Time Format EPL plug-in.
- The correlator now adheres to control group (cgroup) limitations on Linux platforms. REST and Prometheus APIs now provide `physicalCores` as the number of cores available to the correlator, instead of the number that exists on the host machine. A new metric `availableMemoryMB` has been added to represent the maximum available memory in MB that the correlator has access to. See also "Restricting Correlator Resource Usage with Control Groups" in *Deploying and Managing Apama Applications*.
- Correlator logging from user-supplied Java plug-ins and internal logging from the correlator itself can now be configured from a YAML configuration file using the `correlatorLogging` section, supporting the ability to change the log level and file used for individual categories. The categories in the `correlatorLogging` section are treated hierarchically. For example, setting the level for `foo.bar` will also set the level for `foo.bar.baz` (unless the log level for `foo.bar.baz` is also specified). This provides the same capabilities as the existing `eplLogging` section in a YAML configuration file. See "Setting correlator and plug-in log files and log levels in a YAML configuration file" in *Deploying and Managing Apama Applications* for more details.
- Apama no longer ships its own version of Apache Ant. It now uses the Ant version from the common library of your Software AG installation. That is, Apache Ant is now present at `SAG_HOME/common/lib/ant` instead of `APAMA_HOME/third_party/apache_ant`. In addition, Apache Ant has been upgraded from version 1.9.7 to version 1.10.5.

Platform changes in 10.5.0

Apama 10.5.0 runs on the platforms listed in the *Supported Platforms* document for version 10.5. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>.

Check the above mentioned *Supported Platforms* document for details about the newer platform versions supported by Apama 10.5.0. The following platforms are no longer supported:

- Pivotal Cloud Foundry Ubuntu.
- Software AG Common Runtime 10.3 and 10.4 (replaced by support for a more recent version).
- MySQL Community Edition 5.6 (a later version is still supported).
- Software AG Universal Messaging 10.3 and 10.4 (replaced by support for a more recent version).
- JBoss HornetQ.
- Pivotal Software RabbitMQ.
- Software AG Universal Messaging Client 10.3 and 10.4 (replaced by support for a more recent version).
- Azul Zulu Java 8.31 (a later version is still supported).
- Software AG Designer 10.3 and 10.4 (replaced by support for a more recent version).
- Python 2.7 (a later version is still supported).
- Apache Ant 1.9.7 (replaced by support for a more recent version).
- Software AG BigMemory Max 4.3.6 and 4.3.7 (replaced by support for a more recent version).
- Software AG Terracotta Store 10.3 and 10.4 (replaced by support for a more recent version).

Miscellaneous changes in 10.5.0 affecting backwards compatibility

The following changes in Apama 10.5.0 affect backwards compatibility to previous Apama versions:

- It is no longer possible to define a stream listener which executes a single statement without braces. Now you must always specify a full block. See also "Using output from streams" in *Developing Apama Applications*.
- Some expressions including calls to action variables did not correctly implement the documented left-to-right execution order within the expressions. In 10.5, we have changed the order of execution of such expressions to match the documentation and the behavior of other similar expressions.
- When logging from EPL, the default log level (if you do not specify a log level with `at`) has been changed from `CRIT` to `INFO`. See also "Specifying log statements" in *Developing Apama Applications*.

- Java logging in Apama now uses Log4j 2 rather than Log4j 1. For more information, see [“Change of Java Log4j logging library in 10.5.0” on page 62](#) and the announcement in [“Removed and deprecated features in 10.3.1” on page 82](#).
- The timeout for scenario operations in the Java Scenario Service API has been increased from 60 seconds to 5 minutes. Timeouts in instance creation, deletion and editing are now also logged at WARN level.
- Your applications must now call `onApplicationInitialized()` in the correlator before sending requests to the HTTP server. Any requests that are sent before calling this method will fail with a “503 Host Not Ready” response code. See “Sending and receiving events with connectivity plug-ins” in *Connecting Apama Applications to External Components* for more information on `onApplicationInitialized()`.
- With the Cumulocity IoT transport, the `url` configuration option is now mandatory. The configuration of a tenant identifier (tenant configuration option) is now optional. See also “Configuring the Cumulocity IoT transport” in *Connecting Apama Applications to External Components*.
- The Python interpreter supported by Apama has been upgraded from 3.6 to 3.7. For most users no changes will be necessary to your code, but new features may be available in 3.7. If you are not using the version of Python distributed with the full installation, you will have to update your Python installation to match.
- The `apama-correlator` and `apama-builder` images published on Docker Hub and other locations now use Red Hat's Universal Base Images for RHEL 7 as their base images, rather than CentOS which was used in previous releases. The `apama-cumulocity-jre` image still uses Ubuntu 18.04 (Bionic Beaver).
- The `apama-correlator` image published on Docker Hub now no longer contains a full JDK. Instead it only contains a JRE suitable for running Java applications in the correlator. If you need to build Java code at runtime, we recommend using a multi-stage build as documented in “Building Apama projects during the Docker build” in *Deploying and Managing Apama Applications*. If you must have the full JDK at runtime, then you should use the `apama-builder` image as your base.

Note:

There may also be removed features that affect backwards compatibility. See [“Removed and deprecated features in 10.5.0” on page 69](#) for more information.

Removed and deprecated features in 10.5.0

The following Apama features are now deprecated or have been removed in Apama 10.5.0:

- Apama's Event Modeler in Software AG Designer is no longer supported and has been removed. This includes any associated scenario files (`.sdf`), scenario block files (`.bdf`), scenario function files (`.fdf`), PySys test support (`CorrelatorHelper.injectScenario()` was removed), Ant macro definitions in `apama-macros.xml` (see also the list below), and `engine_deploy` support (the command-line option `--generateEventModelerConfig` was removed). See also the announcement in [“Removed and deprecated features in 9.12” on page 145](#). The Scenario Browser view and

the Scenario Service API are still fully supported, but they are now used only for DataViews and queries.

Important:

You have to re-export or edit all Ant scripts exported from Software AG Designer - even if they do not use scenarios - as several scenario-related macros and macro attributes that were present in the default `build.xml` file have now been removed.

To manually edit your `build.xml` file:

- remove the line that calls `<generate-event-modeler-config .../>`, and
- edit the line `<inject-project-init-list .../>` to remove the following attributes:
 - `scenario-config=`
 - `temp-location=`

If your project uses scenarios and/or you have customized your `build.xml` file, you have to make additional changes.

If you wish to clean up your `build.xml` file to remove redundant scenario-related items, the full list of removed items is:

- The following properties have been removed:
 - `catalogs.dir`
 - `scenarios.dir`
 - `event_modeler.config`
 - `event_modeler.xml`
- The following attributes have been removed from `inject-project-init-list`:
 - `scenario-config=`
 - `scenario-debug=`

In addition, `temp-location=` is no longer needed. So it is recommended that you remove this.

- The following macros have been removed:
 - `generate-event-modeler-config`
 - `start-event-modeler`
 - `run-event-modeler`
 - `generate-block`
 - `generate-scenario`
 - `generate-scenarios`

- `inject-scenario`
- `inject-scenarios`
- `delete-scenario`
- Python 2.7 is no longer supported. You should therefore port any PySys testcases written for Python 2.7 to Python 3.x when possible. Python 2.7 is no longer being tested by Software AG but may continue to work with PySys. You can choose to continue to use your own installation of Python 2.7 at your own risk. See also the announcement in [“Removed and deprecated features in 10.3.0” on page 93](#).
- The Cumulocity IoT transports for 9.0 and 9.8 have been removed. A new version (10.5) of the Cumulocity IoT transport connectivity plug-in is now available. See [“Migrating the Cumulocity IoT transports to 10.5” on page 73](#) for details.
- The `requestAllDevices` configuration option for configuring the Cumulocity IoT transport is now deprecated. You should now explicitly request for all available devices on startup using the `com.apama.cumulocity.FindManagedObject` API.

The `pageSize` and `useBatching` configuration options for configuring the Cumulocity IoT transport are now also deprecated and have been removed from the documentation. This is now handled automatically. Specifying values for either of these configuration options will be ignored.

See also "Configuring the Cumulocity IoT transport" in *Connecting Apama Applications to External Components*.

- The documentation for the old C and C++ EPL plug-in APIs that were supported until version 9.12 has been removed. These APIs are deprecated since 10.0 and will be removed in a future release (see also [“New API for writing EPL plug-ins in C++ in 10.0” on page 124](#)). For details on the current C++ EPL plug-in API, see "Writing EPL Plug-ins in C++" in *Developing Apama Applications*. For your convenience, the information on how to migrate the old plug-ins has been added to the Release Notes for 10.0; see [“Migrating the C and C++ EPL plug-ins written in 9.12 or earlier” on page 131](#).
- `com::apama::setLogFD` has been removed from the C++ API. If you are using this, it is recommended that you use `setLogFile` instead. See also the *API Reference for C++ (Doxygen)*.
- For the Java and C++ connectivity plug-ins, the `LOGGER` field and use of the legacy constructor signatures that do not use `TransportConstructorParameters` or `CodecConstructorParameters` have been removed.

For C++, the following macros have been removed:

- `SAG_DECLARE_CONNECTIVITY_CODEC`
- `SAG_DECLARE_CONNECTIVITY_TRANSPORT`

See also the announcement in [“Removed and deprecated features in 10.0” on page 130](#).

In addition, the previously deprecated `name` field was removed from the C++ connectivity plug-ins. You should now use `pluginName` instead.

- Calling a non-static action on a type is now deprecated and will be removed in a future release. Instead, you should now call the method directly on an instance.

For example, you should no longer call a non-static action on a type as follows:

```
integer range := 100;
integer.rand(range);
```

Use the following instead:

```
integer range := 100;
range.rand();
```

Another example - you should no longer use the following:

```
event A {
    action nonStaticAction() { log "non static" at INFO; }
    static action staticAction() { log "static" at INFO; }
}
A.nonStaticAction(new A); // Results in warning
```

Use the following instead:

```
A a = new A;
a.nonStatic(); // This is preferred way
```

- The enqueue statement (without a to) is now deprecated and will be removed in a future release. It is recommended that you use the send...to statement instead. For applications which require enqueueing a string rather than an event object, you can either use the any type (where it is an event but not statically known at compile time) or you can switch to a second correlator using the engine_connect tool and the emit statement, which still supports bare strings (for things like &TIME events).
- Use of Universal Messaging from the IAF is now deprecated and will be removed in a future release. It is recommended that you now change any IAF-based adapter configurations using Universal Messaging with a <universal-messaging> element in the configuration file to use an <apama> element to talk directly to the correlator. See "Apama correlator configuration" in *Connecting Apama Applications to External Components*.
- The SOAP-based Web Services Client IAF adapter is now deprecated and will be removed in a future release. It is recommended that you now use a REST-based HTTP server and connect to it with the HTTP client transport, if possible (see "The HTTP Client Transport Connectivity Plug-in" in *Connecting Apama Applications to External Components*). However, if you really need to communicate with a SOAP-based server, use the webMethods Integration Server and connect to it from Apama using the Digital Event Services transport (see "The Digital Event Services Transport Connectivity Plug-in" in *Connecting Apama Applications to External Components*).
- When using FindMeasurement from the Cumulocity API, use of fragmentType as a filter is now deprecated. You should instead use both of valueFragmentType and valueFragmentSeries together to filter measurements. Use of just one will be significantly slower. See also "Querying for measurements" in *Connecting Apama Applications to External Components*.

Migrating the Cumulocity IoT transports to 10.5

In order to migrate an existing Apama project to 10.5 which is using the removed 9.0+ and/or 9.8+ Cumulocity IoT connectivity bundles, you must proceed as described below. If you are using Cumulocity IoT with the built-in Apama EPL editor, then no migration steps should be needed.

To migrate the Cumulocity IoT transports to 10.5:

1. Back up the `CumulocityIoT.properties` files for 9.0+ and/or 9.8+. These are located in the `config/connectivity` folder of your project.
2. Remove the **Cumulocity Client 9.0+** and/or **Cumulocity Client 9.8+** connectivity bundles from your project.
3. If present, remove the **Cumulocity IoT > Event Definitions for Cumulocity 9.0+** and/or **Cumulocity IoT > Event Definitions for Cumulocity 9.8+** EPL bundles from your project.
4. If present, remove the **Cumulocity IoT > Utilities for Cumulocity 9.0+** and/or **Cumulocity IoT > Utilities for Cumulocity 9.8+** EPL bundles from your project.
5. Add the **Cumulocity IoT > Cumulocity Client 10.5+** connectivity bundle.
6. Add the **Cumulocity IoT > Event Definitions for Cumulocity 10.5+** EPL bundle.
7. Add the **Cumulocity IoT > Utilities for Cumulocity 10.5+** EPL bundle.
8. Populate the `CumulocityIoT.properties` file for 10.5 with the data from your backup.
9. Use the `com.apama.cumulocity.Util.generateReqId()` action to generate all request identifiers (`reqId`).

Important:

There should be no use of `integer.getUnique()` or `integer.incrementCounter()` to generate request identifiers.

For detailed information on this connectivity plug-in, see "The Cumulocity IoT Transport Connectivity Plug-in" in *Connecting Apama Applications to External Components*.

11 What's New In Apama 10.3.1

■ New command-line tool for adding Apama projects in 10.3.1	76
■ New Base64 codec connectivity plug-in in 10.3.1	76
■ New Batch Accumulator codec connectivity plug-in in 10.3.1	76
■ New Message List codec connectivity plug-in in 10.3.1	77
■ Connectivity plug-ins enhancements in 10.3.1	77
■ Apama enhancements in Software AG Designer in 10.3.1	79
■ EPL enhancements in 10.3.1	79
■ Miscellaneous enhancements and changes in 10.3.1	80
■ Platform changes in 10.3.1	81
■ Miscellaneous changes in 10.3.1 affecting backwards compatibility	81
■ Removed and deprecated features in 10.3.1	82

You can install Apama 10.3.1 by selecting "2019 April Innovation Release" in the Software AG Installer or by using the Software AG Update Manager or Command Central to upgrade an existing Apama 10.3.0 installation.

Important:

Installing Apama 10.3.1 as an "Innovation Release" from the Software AG Installer will limit its support to the end of support for the "2019 April Innovation Release". Only installation via an upgrade from 10.3.0 will result in support until the end of support for 10.3.0.

Apama 10.3.1 runs on the platforms listed in the *Supported Platforms* document for version 10.3.1. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>. Be sure to consult that document for details about supported versions of operating systems, compilers and other products.

Apama 10.3.1 includes new features, enhancements, and changes as described in the following topics.

New command-line tool for adding Apama projects in 10.3.1

A new command-line tool, the `apama_project` tool, is now available. It can be used to:

- create an Apama project,
- list all supported bundles that can be added to a project,
- list all bundles/instances that have already been added to a project,
- add connectivity, adapter and EPL bundles and their instances to a project,
- remove bundles/instances from a project.

Each Apama project generated by this tool is compatible with Software AG Designer and can seamlessly be imported into Software AG Designer as an Apama project.

For detailed information on the new tool, see "Creating and managing an Apama project from the command line" in *Deploying and Managing Apama Applications*.

New Base64 codec connectivity plug-in in 10.3.1

A new codec, the Base64 codec connectivity plug-in, is now available. The Base64 codec performs bidirectional translation between binary payloads and a string that is the Base64-encoding of that payload. The codec is able to operate on arbitrary fields of the message, or the entire payload. For detailed information, see "The Base64 codec connectivity plug-in" in *Connecting Apama Applications to External Components*.

New Batch Accumulator codec connectivity plug-in in 10.3.1

A new codec, the Batch Accumulator codec connectivity plug-in, is now available. The Batch Accumulator codec can be used to automatically batch messages from an unbatched transport such as the HTTP server transport when sending messages towards the correlator. This can be

used to increase the maximum throughput of the system. For detailed information, see "The Batch Accumulator codec connectivity plug-in" in *Connecting Apama Applications to External Components*.

New Message List codec connectivity plug-in in 10.3.1

A new codec, the Message List codec connectivity plug-in, is now available. The Message List codec can be used to automatically combine a batch of messages produced by the correlator into a single message containing the original batch as a list. This can be used with a service which supports a batch as a list to avoid the system being limited by request round-trip times. For detailed information, see "The Message List codec connectivity plug-in" in *Connecting Apama Applications to External Components*.

Connectivity plug-ins enhancements in 10.3.1

Apama 10.3.1 includes the following connectivity plug-ins enhancements:

HTTP server transport

- It is now possible for the response to an HTTP request to the HTTP server transport to be handled and created by an EPL application, instead of automatically returning an empty Accepted response. This allows the HTTP server to implement a request/response REST protocol and not just an event submission protocol. For detailed information, see "Handling responses in EPL" in *Connecting Apama Applications to External Components*.

The following new configuration options are available for handling responses in EPL: `automaticResponses` and `responseTimeoutMs`. For more information, see "Configuring the HTTP server transport" in *Connecting Apama Applications to External Components*.

A new `metadata.requestId` field is available for requests from the transport to EPL, and several new fields are available for EPL-supplied responses. For more information, see "Mapping events between EPL and HTTP server requests" in *Connecting Apama Applications to External Components*.

For an example, see "EPL-controlled responses" in *Connecting Apama Applications to External Components*.

- The HTTP server transport now supports gzip and deflate decoding. For more information, see "About the HTTP server transport" in *Connecting Apama Applications to External Components*.
- The HTTP server transport now supports HTML form decoding and can decode `multipart/form-data` or `application/x-www-form-urlencoded` media types to a dictionary payload. For more information, see "Handling HTML form decoding" in *Connecting Apama Applications to External Components*.

HTTP client transport

- The HTTP client transport now supports gzip and deflate decoding. HTTP client requests now set the `Accept-Encoding` header appropriately. For more information, see "About the HTTP client transport" in *Connecting Apama Applications to External Components*.

- The HTTP client transport now supports HTML form encoding and can encode a dictionary payload to either `multipart/form-data` or `application/x-www-form-urlencoded` media types. For more information, see "Handling HTML form encoding" in *Connecting Apama Applications to External Components*.
- The HTTP client transport now supports HTTP redirects transparently. If redirects are to be followed, set the new `followRedirects` configuration option to `true`. For more information, see "Configuring the HTTP client transport" in *Connecting Apama Applications to External Components*.
- The HTTP client transport can now transparently handle cookies for requests so that it is no longer required to handle them in EPL. If you want to enable this behavior, set the new `cookieJar` configuration option to `true`. The HTTP client then stores in memory all cookies received from the server and adds them to subsequent outgoing requests. For more information, see "Configuring the HTTP client transport" and "Dealing with cookies", both in *Connecting Apama Applications to External Components*.

Note:

For backwards compatibility, the new `followRedirects` and `cookieJar` configuration options are set to `false` by default when they are not explicitly specified in the YAML configuration file. However, for your convenience, when you add a new **HTTP Client** connectivity bundle in Software AG Designer, these options are already set to `true` in the resulting YAML configuration file.

Cumulocity IoT transport

- The Cumulocity IoT transport connectivity plug-in now honors the `pageSize` parameter while querying for resources. This applies to all versions of the Cumulocity bundles (9.0 and 9.8+). See also the following topics in *Connecting Apama Applications to External Components*:
 - Querying for managed objects
 - Querying for alarms
 - Querying for events
 - Querying for measurements
 - Querying for operations
- Starting with the Cumulocity 9.8+ bundles, the Cumulocity IoT transport connectivity plug-in honors the `fromDate` and `toDate` parameters while querying for operations. See also "Querying for operations" in *Connecting Apama Applications to External Components*.
- Starting with the Cumulocity 9.8+ bundles, the value of the Cumulocity IoT transport connectivity plug-in's `requestAllDevices` configuration option is set to `false`. See also "Configuring the Cumulocity IoT transport" in *Connecting Apama Applications to External Components*. You should now explicitly request for all available devices on startup using the `com.apama.cumulocity.FindManagedObject` API. For more information, see "Sample EPL" in *Connecting Apama Applications to External Components*.

String codec

- The String codec can now convert individual payload/metadata fields to a character set other than UTF-8. The following new configuration options are available for this purpose: `eventTypes`, `fields`, and `encoding`. For more information, see "The String codec connectivity plug-in" in *Connecting Apama Applications to External Components*.

Apama enhancements in Software AG Designer in 10.3.1

Apama 10.3.1 includes the following enhancements in Software AG Designer:

- The following connectivity bundle group has been added to Software AG Designer:
 - **Cumulocity Client 9.8+.** Software AG recommends that you use this new bundle over the Cumulocity Client 9.0+ bundle. See also "Adding the Cumulocity IoT connectivity plug-in to a project" in *Using Apama with Software AG Designer*.
- The following EPL bundles have been added to Software AG Designer:
 - **Event Definitions for Cumulocity 9.8+.** This EPL bundle defines all events that can be used for interacting with Cumulocity IoT. This includes definitions for events that you receive from Cumulocity IoT, events that you can send to Cumulocity IoT, and event APIs that you can use for requesting data from Cumulocity IoT. For more information, see the `com.apama.cumulocity` package in the *API Reference for EPL (ApamaDoc)*.
 - **Utilities for Cumulocity 9.8+.** This EPL bundle contains useful utilities for EPL code that is interacting with Cumulocity IoT. For more information, see the `com.apama.cumulocity.Util` event in the *API Reference for EPL (ApamaDoc)*.
- The default setting when adding an **MQTT** connectivity bundle to a project in Software AG Designer has been changed to use channels starting with `mqtt:` in Apama, rather than consume all channels in the whole system. This corresponds to the default setting of the `channelPrefix` configuration option as described in "Configuring the connection to MQTT" in *Connecting Apama Applications to External Components*.

EPL enhancements in 10.3.1

Apama 10.3.1 includes the following EPL enhancements:

- A new method `getActionParameterTypes()` has been added for the `any` type. It returns the required types as a sequence of strings, allowing reflection on action parameters, built-in methods, closures, and plug-in methods. See the *API Reference for EPL (ApamaDoc)* for more information.
- For consistency with our recommended case for action methods, we have replaced the `nextafter` method on `float` and `decimal` with `nextAfter`. The previous method name will still work for compatibility, but we recommend you to begin using `nextAfter` instead.
- The Management interface can now be used to fetch the list of configuration properties that the correlator was started with from EPL. This is available through the new `getConfigProperties`

method on the `com.apama.correlator.Component` event. See also "Using the Management interface" in *Developing Apama Applications* and the *API Reference for EPL (ApamaDoc)*.

- The Management interface now has another method for manipulating user-defined status values. `Component.incrementUserStatus` allows you to atomically increment or decrement a status value from multiple contexts. For more details, see the *API Reference for EPL (ApamaDoc)*.
- When exceptions are thrown from EPL plug-ins that are written in Python or Java, the resulting stack trace when caught or logged in EPL now includes both the stack elements from within the plug-in as well as the ones within EPL. Thus, stack traces are no longer logged separately for exceptions from EPL plug-ins written in Java.

Miscellaneous enhancements and changes in 10.3.1

Apama 10.3.1 includes the following miscellaneous enhancements and changes:

- The documentation has been updated with information on how to back up the persistence database using file-system snapshots while the correlator is still running. See "Backing up the persistence database while the correlator is running" in *Developing Apama Applications*.
- The `-XX:+HeapDumpOnOutOfMemoryError` argument has been added to the default JVM flags for the Java process within the correlator. This means that whenever a Java `OutOfMemoryError` is thrown, a `java_pidnumber.hprof` file is created containing the Java heap. If you want to disable this new behavior, start the correlator with `-J-XX:-HeapDumpOnOutOfMemoryError` (with a minus sign instead of a plus sign).
- In addition to using the `apama_env` script to set up your console environment to run Apama commands, you can also use it as a one-shot wrapper to run a single Apama command without changing your environment. For more details, see "Setting up the environment using the Apama Command Prompt" in *Deploying and Managing Apama Applications*.
- The Apama Docker images provided through Docker Hub and other container registries now include a default `HEALTHCHECK` instruction to monitor the status of the correlator. This will be monitored by Docker and other tooling automatically.
- The documentation has been updated with information on how to provide an Apama license file to the Docker image. See "Licensing Apama in Docker" in *Deploying and Managing Apama Applications*.
- The correlator now exposes non-numeric metadata status over Prometheus. Labels for the name of the slowest context and for the name of the slowest consumer/receiver of events are now given on the `sag_apama_correlator_slowest_input_queue_size_total` and `sag_apama_correlator_slowest_output_queue_size_total` metrics. See also "List of correlator status statistics" in *Deploying and Managing Apama Applications*. In addition, the following Prometheus metrics have been added where the labels give information about the status of the current license or about the configuration: `sag_apama_correlator_licensedata` and `sag_apama_correlator_metadata`. See also "Monitoring with Prometheus" in *Deploying and Managing Apama Applications*.
- The correlator now logs the user name of the user which it was started as in the start-up stanza of the correlator log file.

- The `engine_deploy` tool now always generates the `connectivity.yaml` file for a project which contains connectivity bundles. With previous versions, this file was generated only when a deployment file was provided to the `engine_deploy` tool. See also "Deploying a correlator" in *Deploying and Managing Apama Applications*.
- The file `client_status_names.hpp` has been added to the public includes to contain constants for all standard status values within the correlator. See also "List of correlator status statistics" in *Deploying and Managing Apama Applications*.
- As of Apama 10.3.1, it is possible to uninstall service packs and fixes. To do so, run the Software AG Update Manager or Command Central, select the service pack or fix to be uninstalled, and then follow the instructions in the user interface. Keep in mind that it is not possible to uninstall service packs and fixes for Apama versions before 10.3.1. For detailed information, see *Using Software AG Update Manager* and *Software AG Command Central Help*. The most up-to-date versions of these guides are always available at <http://documentation.softwareag.com/> (Empower login required).

Platform changes in 10.3.1

Apama 10.3.1 runs on the platforms listed in the *Supported Platforms* document for version 10.3.1. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>.

Check the above mentioned *Supported Platforms* document for details about the newer platform versions supported by Apama 10.3.1.

Due to these upgrades, the following versions are no longer supported:

- No changes as compared to "Platform changes in 10.3.0" on page 91.

Miscellaneous changes in 10.3.1 affecting backwards compatibility

The following changes in Apama 10.3.1 affect backwards compatibility to previous Apama versions:

- If you are using the unusual convention of `type.method(...)` to call a method, you will be passing an instance of `type` as the first parameter name. We have corrected an issue with the `any.getActionParameterNames()` method so that it now correctly returns `self` as the first parameter where there is no closure. This matches the behavior of the new `any.getActionParameterTypes()` method (see also "EPL enhancements in 10.3.1" on page 79).
- Correlator and IAF processes using the `--pidfile` option now take an exclusive lock on the pidfile. This means that multiple components using the same pidfile will be prevented from starting up. See also "Starting the correlator" in *Deploying and Managing Apama Applications* and "IAF command-line options" in *Connecting Apama Applications to External Components*.
- As of this version, the .NET Event Parser plug-in correctly interprets floating point separators on hosts that use a separator other than the period. Prior to Apama 10.3.1, the .NET Event Parser plug-in did not interpret `Location` events correctly. The issue occurred only with hosts set up to use a character other than the period (`.`) as the floating point separator. For example,

if the local host was set up to use a comma as the floating point separator, then a `Location([4.4,3.3,2.2,1.1])` was interpreted as `Location([44,33,22,11])`. Now this is correctly interpreted as `Location([4.4,3.3,2.2,1.1])`.

Removed and deprecated features in 10.3.1

The following Apama features are now deprecated or have been removed in Apama 10.3.1:

- The `MapHelper` class (`com.softwareag.connectivity.MapHelper` in the Java API and `com::softwareag::connectivity::MapHelper` in the C++ API) is now deprecated and will be removed in a future release. Use the `MapExtractor` class instead. This class provides an easy and type-safe way to extract values from maps such as those used for plug-in configuration. It has automatic support for generating user-friendly error messages if configuration items are missing or of the wrong type.

In the Java API, the `com.softwareag.connectivity.util.MapExtractor` class is available since version 10.0. See the *API Reference for Java (Javadoc)* for further information.

In the C++ API, the `com::softwareag::connectivity::MapExtractor` class is available since version 10.1. See the *API Reference for C++ (Doxygen)* for further information.

- The ADBC reconnection policy constants `RECONNECT` and `RECONNECT_AND_CLEAR_REQUEST_QUEUE`, which can be defined in the `ADBCHelper` API, are now deprecated and will be removed in a future release. It is recommended that you use any of the other policy constants as appropriate (that is, either `DO_NOT_RECONNECT` or `RECONNECT_AND_RETRY_LAST_REQUEST`). See also "Reconnection settings" in *Connecting Apama Applications to External Components*.
- Log4j 1.2 is now deprecated and support for it will be removed in a future release. For now, the Apama Java client libraries in `ap-util.jar` and `ap-client.jar` continue to support Log4j 1.2, but can now alternatively write to Log4j 2 if the Log4j-1-to-2 bridge `log4j-1.2-api-2.x.jar` is added to the classpath before those jars.

12 What's New In Apama 10.3.0

■ New Cumulocity IoT transport connectivity plug-in in 10.3.0	84
■ New EPL plug-in for JSON support in 10.3.0	85
■ New Prometheus support in 10.3.0	85
■ New Python EPL plug-in support in 10.3.0	85
■ PySys changes in 10.3.0	85
■ Connectivity plug-ins enhancements in 10.3.0	86
■ Apama enhancements in Software AG Designer 10.3.0	87
■ EPL enhancements in 10.3.0	88
■ Correlator utility enhancements in 10.3.0	88
■ Command Central support in 10.3.0	89
■ Dashboard enhancements in 10.3.0	90
■ Miscellaneous enhancements and changes in 10.3.0	90
■ Platform changes in 10.3.0	91
■ Miscellaneous changes in 10.3.0 affecting backwards compatibility	92
■ Removed and deprecated features in 10.3.0	93

Apama 10.3.0 runs on the platforms listed in the *Supported Platforms* document for version 10.3.0. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>. Be sure to consult that document for details about supported versions of operating systems, compilers and other products.

Apama 10.3.0 includes new features, enhancements, and changes as described in the following topics.

New Cumulocity IoT transport connectivity plug-in in 10.3.0

A new connectivity plug-in, the Cumulocity IoT transport, is now available. The Cumulocity IoT transport connectivity plug-in allows you to receive information from devices, measurements, events and alarms from Cumulocity IoT and send operations to Cumulocity IoT (and thus to the devices). For detailed information, see "The Cumulocity IoT Transport Connectivity Plug-in" in *Connecting Apama Applications to External Components*.

New EPL and connectivity bundles are also available which you can add to your projects using Software AG Designer. See also "[Apama enhancements in Software AG Designer 10.3.0](#)" on [page 87](#).

The Cumulocity IoT transport connectivity plug-in and the new bundles were also delivered with 10.2.0.1.

As of 10.3.0, and also with 10.2.0.3, the following enhancements are available:

- The Cumulocity IoT transport connectivity plug-in supports the following (for detailed information, see the corresponding topics in *Connecting Apama Applications to External Components*):
 - Updating a managed object.
 - Updating an existing event.
 - Updating an operation.
 - Querying for measurements.
 - Querying for operations.
- You can individually subscribe to (and unsubscribe from) alarms, events, managed objects, measurements, or operations.
- The Cumulocity IoT transport connectivity plug-in supports the following additional configuration options (for more information, see "Configuring the Cumulocity IoT transport" in *Connecting Apama Applications to External Components*):
 - `subscribeToOperations` for subscribing to all device operations.
 - `pageSize` for retrieving resources from Cumulocity IoT.
 - `useBatching` for creating bulk measurements.

New EPL plug-in for JSON support in 10.3.0

A new EPL plug-in, the JSON plug-in, is now available. You can use the JSON plug-in to convert EPL objects to JSON strings, and vice versa. The JSON plug-in requires the new **JSON Support** bundle for use with Software AG Designer. See also [“Apama enhancements in Software AG Designer 10.3.0” on page 87](#).

Equivalent functionality is available with the JSON codec connectivity plug-in. Use the JSON codec if the JSON is coming in as an entire message of a connectivity chain. Use the JSON EPL plug-in if only one field of an event is JSON, or if the events are going to or are coming from a connection that is not using a connectivity chain.

See "Using the JSON plug-in" in *Developing Apama Applications* for detailed information.

New Prometheus support in 10.3.0

The correlator now supports the Prometheus monitoring system for observing internal correlator metrics. See "Monitoring with Prometheus" in *Deploying and Managing Apama Applications* for detailed information.

New Python EPL plug-in support in 10.3.0

Apama supports a new framework for writing EPL plug-ins in Python. See "Writing EPL Plug-ins in Python" in *Developing Apama Applications* for detailed information.

EPL plug-ins written in Python support Python 3, which is shipped with Apama. They do not support Python 2.

Using the new `--python` option of the correlator executable, you can enable and disable Python support. See "Starting the correlator" in *Deploying and Managing Apama Applications*.

You can also enable and disable Python support using a YAML configuration file. The new `pythonSupport` element is available for this purpose. Using the new `eplPlugins` element in a YAML configuration file, you can configure your EPL plug-ins written in Python. See "Using Python plug-ins" in *Developing Apama Applications*.

Sample EPL plug-ins written in Python can be found in the `samples\correlator_plugin\python` directory of your Apama installation.

PySys changes in 10.3.0

Apama now ships PySys 1.3 instead of PySys 1.2. Release notes for PySys 1.3 can be found in `Apama\third_party\python\Lib\site-packages\pysys-release.txt` on Windows and in `Apama/third_party/python/lib/python3.6/site-packages/pysys-release.txt` on Linux.

The sample `pysysproject.xml` file has been updated to support the new features, details of which are explored further down below.

A new Apama PySys base test class for use in new PySys test cases has been added. Using this class means that PySys methods that read or write files (such as `waitForSignal` or `assertGrep`) will

use character strings rather than byte strings and UTF-8 encoding for certain file types: `.xml`, `.yaml`, `.json`, `.log`.

Key features for the new PySys 1.3 release include:

- PySys tests can now be written and executed using Python 3.
- Support was added for coloring console output to highlight passes, fails, warnings, and more. This is enabled by default in the example `pysysproject.xml` configuration file shipped with Apama. Coloring can also be enabled or disabled for a particular user and/or machine by setting the `PYSYS_COLOR` environment variable to `true` or `false`. Coloring works on any terminal that supports ANSI escape sequences (for example, most UNIX terminals). On Windows, coloring works using the `colorama` library that is installed into the Python included with Apama.
- A framework was added for recording performance results (such as throughput or latency) that are measured by PySys tests. This includes storage of results in a human-readable and/or machine-parseable CSV file, together with run-specific information such as the host where the test was executed. The CSV files can be aggregated across multiple test runs and/or cycles. See `BaseTest.reportPerformanceResult()` for more details.
- A `--progress` command line option was added (can also be switched on by setting the `PYSYS_PROGRESS` environment variable to `true`), which logs a summary of how many tests have executed, the outcomes, a list of most recent failure reasons, and a list of what other tests are currently executing. This provides helpful feedback to you while executing a long test run.
- A `defaultIgnoreExitStatus` project property was added which controls whether non-zero return codes from `startProcess()` result in test failures when the `ignoreExitStatus` flag is not explicitly specified. This property defaults to `True` if not explicitly configured, which retains behavior for existing tests.

However to promote best practice for new PySys projects, the example PySys project configuration file included with Apama sets `defaultIgnoreExitStatus` to `False`, which ensures that processes which return failure codes are not ignored unless explicitly intended by the author of the test case. We recommend changing existing projects to set `defaultIgnoreExitStatus` to `False` and also to set `abortOnError` to `True` in the project file so that unintentional process failures do not go unnoticed.

- With PySys 1.3, the `pysys.py run` command now returns a non-zero exit code if any tests failed (previously, this command returned 0 on any failed tests).

API documentation for PySys and the Apama PySys extensions can be found in the *API Reference for Python*.

The use of PySys 1.3 has introduced some backwards compatibility changes. See [“Miscellaneous changes in 10.3.0 affecting backwards compatibility” on page 92](#) for detailed information.

Connectivity plug-ins enhancements in 10.3.0

Apama 10.3.0 includes the following connectivity plug-ins enhancements:

- The MQTT transport connectivity plug-in now supports client-side TLS. For this purpose, the following new authentication options can be specified in the YAML configuration file: `certificateFile`, `certificatePassword` and `privateKeyFile`. For more information, see "Configuring the connection to MQTT" in *Connecting Apama Applications to External Components*.
- The Kafka transport connectivity plug-in now uses the Kafka default value for `max.block.ms` instead of the custom value (2147483647) that was used in previous versions. See also "Configuring the connection to Kafka (dynamicChainManagers)" in *Connecting Apama Applications to External Components*.
- The Classifier codec now supports regular expressions. For more information, see "The Classifier codec connectivity plug-in" in *Connecting Apama Applications to External Components*.

In addition, the Classifier codec is no longer provided as a sample.

Apama enhancements in Software AG Designer 10.3.0

Apama 10.3.0 includes the following enhancements in Software AG Designer:

- The following connectivity bundle group has been added to Software AG Designer:
 - **Cumulocity IoT.** This was also delivered with 10.2.0.1 as a standalone bundle. See also "Adding the Cumulocity IoT connectivity plug-in to a project" in *Using Apama with Software AG Designer*.
- The following EPL bundles have been added to Software AG Designer:
 - **Event Definitions for Cumulocity 9.0+.** This EPL bundle defines all events that can be used for interacting with Cumulocity IoT. This includes definitions for events that you receive from Cumulocity IoT, events that you can send to Cumulocity IoT, and event APIs that you can use for requesting data from Cumulocity IoT. For more information, see the `com.apama.cumulocity` package in the *API Reference for EPL (ApamaDoc)*.
 - **Utilities for Cumulocity 9.0+.** This EPL bundle contains useful utilities for EPL code that is interacting with Cumulocity IoT. For more information, see the `com.apama.cumulocity.Util` event in the *API Reference for EPL (ApamaDoc)*.
 - **Automatic onApplicationInitialized.** For simple applications, you can add this EPL bundle to your project to ensure that `onApplicationInitialized` is called as soon as the entire application has been injected into the correlator. This is not appropriate in all cases: for more advanced applications where you need to wait for a `MemoryStore`, database or another resource to be prepared before your application is able to begin to process incoming messages, you should continue to write EPL code to call `onApplicationInitialized` explicitly instead of using the bundle. See also "Sending and receiving events with connectivity plug-ins" and "Using EPL to send and receive JMS messages" in *Connecting Apama Applications to External Components*.
 - **JSON Support.** The new JSON plug-in requires this bundle for use with Software AG Designer. See also "[New EPL plug-in for JSON support in 10.3.0](#)" on page 85.

These EPL bundles can be found in the **Standard bundles** node of the New Apama Project and Add Bundle dialogs. See also "Adding bundles to projects" in *Using Apama with Software AG Designer*.

The **Cumulocity IoT Event Definitions**, **Cumulocity Utilities** and **Automatic onApplicationInitialized** bundles were delivered with 10.2.0.1, but in that version they were placed in the **Connectivity bundles** node. In 10.3.0, they have been placed in the **Standard bundles** node. In addition in 10.3.0, the **Cumulocity IoT Event Definitions** bundle has been renamed to **Event Definitions for Cumulocity 9.0+**, and the **Cumulocity Utilities** bundle has been renamed to **Utilities for Cumulocity 9.0+**.

- The **Connectivity and Adapters** node and the **EPL Bundles** node are now created by default when you create an Apama project. You can use these nodes to add connectivity bundles, adapter bundles, and EPL bundles. In previous releases, these nodes were created only after adding a bundle. For more information, see "Adding connectivity and adapter bundles to projects" and "Adding EPL bundles to projects" in *Using Apama with Software AG Designer*. As part of this change, the following commands have been removed from the context menu which is shown when you right-click an Apama project:

- **Apama > Add Connectivity and Adapters**

- **Apama > Add Bundle**

- You can now add Docker support to your Apama project. The following new command is available for this purpose:

- **Apama > Add Docker Support**

For more information, see "Adding Docker support to Apama projects" in *Using Apama with Software AG Designer*.

EPL enhancements in 10.3.0

Apama 10.3.0 includes the following EPL enhancements:

- The new boolean action `empty()` is now available for the `any` type. It returns `true` if the `any` is empty, `false` if it wraps an object. See the *API Reference for EPL (ApamaDoc)* for more information on the `any` type.
- A new action `hash()` is now available on all built-in types which can be hashed. It returns an integer hash value in all cases. See the *API Reference for EPL (ApamaDoc)* for more information.
- Dictionary sizes can now be a larger than 2^{32} since the hash function used for keys has been expanded to 64bits.
- The LLVM back-end used for the compiled runtime has been upgraded in this release. In some benchmarks this can produce up to 30% improvement in throughput.

Correlator utility enhancements in 10.3.0

Apama 10.3.0 includes the following correlator utility enhancements:

- The `engine_send` macro in the `apama-macros.xml` Ant script now has an additional `channel` attribute which allows you send events directly to a specified delivery channel. See also "About deploying Apama applications with an Ant script" in *Deploying and Managing Apama Applications*.
- When starting the correlator using the `apama-macros.xml` Ant script, a pidfile (that is, a file containing the process identifier) is now created by default. The default location and name of this pidfile is `${log.dir}/${correlator_name}.pidfile`. You can override this location by modifying the corresponding entry in the `build.xml` file.
- The `correlator.properties` file, which is generated by the `engine_deploy` tool, now uses `${correlator.logsDir}/${correlator.name}.pidfile` as the default location and name of the pidfile (instead of `${correlator.dataDir}/correlator.pid`).
- As an alternative to specifying the `--outputDeployDir` action with the `engine_deploy` tool, you can now specify the shorter `-d` option. In addition, when the output directory or correlator deployment package already exists, it is now overwritten (rather than causing an error). See also "Deploying a correlator" in *Deploying and Managing Apama Applications*.
- The `-s` (or `--source`) option of the `epl_coverage` tool can now be specified more than once to allow multiple source directories to be searched. In addition, the performance of the `epl_coverage` tool has been improved when using source search directories to locate the source for `.mon` files that are not present in the location used during the inject, or are contained within `.cdp` files. See also "Creating code coverage reports" in *Deploying and Managing Apama Applications*.
- If the Java virtual machine (JVM) which is embedded in the correlator is enabled and when logging at `INFO` level, a new `jvm=n` field is now shown in the `Correlator Status:` line of the main correlator log file. Its value is the sum of all memory used by the JVM in megabytes (that is, the used heap memory, the used non-heap memory, and the used buffer pool memory). See also "Descriptions of correlator status log fields" in *Deploying and Managing Apama Applications*.
- Using the new `--python` option of the correlator executable, you can enable and disable Python support. See also "[New Python EPL plug-in support in 10.3.0](#)" on page 85.
- All configured log files now create their parent directory at startup if it does not exist.

Command Central support in 10.3.0

Apama 10.3.0 includes the following changes for its integration with Command Central:

- A new Command Central template has been added to the Command Central samples under the `Apama/samples/command_central/reference` directory. This template documents all the configuration options for the correlator, including how to deploy an Apama application in Command Central. It exists for informational purposes only; it is not a well-rounded sample application that you should build upon.

For more information, see "Deploying Apama Components with Command Central" in *Deploying and Managing Apama Applications*.

Dashboard enhancements in 10.3.0

Apama 10.3.0 includes the following dashboard enhancements:

- The HTML5 pie graph has been added in the **Graphs** tab of the Dashboard Builder's Object Palette. The HTML5 pie chart is an HTML implementation of a swing pie chart object. For more information, see "HTML5 Pie graph" in *Building and Using Apama Dashboards*.
- It is now possible to perform time-based rotation of the log files for the data server and display server. For more information, see "Rotating the log files of the data server and display server" in *Building and Using Apama Dashboards*.

Miscellaneous enhancements and changes in 10.3.0

Apama 10.3.0 includes the following miscellaneous enhancements and changes:

- Apama 10.3.0 incorporates the ICU (International Components for Unicode) time-zone data update 2018e, which is the most recent update at the time of release. This will update time-zone data used by the correlator and Time Format EPL plug-in.
- Apama now supports the jemalloc memory allocator as an optional alternative to its standard memory allocator on Linux platforms. For more information, see "Using jemalloc to optimize memory usage" in *Deploying and Managing Apama Applications*.
- The TCStore driver now supports TLS-secured connections to a secure Terracotta Server Array. See the [Terracotta DB documentation](#) for more information on these concepts. The new TCStore driver property `securityRootDirectory` is available for this purpose. For more information, see "TCStore (Terracotta DB) driver details" in *Developing Apama Applications*.
- The REST API now supports the following keys which provide statistics information when the Java virtual machine (JVM) which is embedded in the correlator has been enabled:
 - `jvmMemoryHeapUsedMB`
 - `jvmMemoryHeapFreeMB`
 - `jvmMemoryNonHeapUsedMB`
 - `jvmMemoryBufferPoolUsedMB`
 - `jvmMemoryAllUsedMB`

In addition, the REST API now also supports the following keys:

- `virtualMemoryMB` (new name for the deprecated `virtualMemorySize`)
- `physicalMemoryMB`
- `swapPagesRead`
- `swapPagesWrite`

For more information, see "List of correlator status statistics" in *Deploying and Managing Apama Applications*.

- The image for the Apama correlator in Docker Store now supports the Cumulocity IoT transport. See also "[New Cumulocity IoT transport connectivity plug-in in 10.3.0](#)" on page 84.
- A builder image which supports the `engine_deploy` tool is now available in Docker Store. If you are using the official Docker Store image, you can use this new builder image for project build steps, as part of a Docker multi-stage build. See also "Building Apama projects during the Docker build" in *Deploying and Managing Apama Applications*.
- Docker and Kubernetes offer methods of passing secrets to the correlator. This can be used to securely provide credentials to the correlator. New Docker and Kubernetes samples are now available which demonstrate how to use secrets to set variables in correlator configuration files. See also "Apama samples for Docker" and "Apama samples for Kubernetes" in *Deploying and Managing Apama Applications*.
- The documentation has been updated with information to help identifying how personal data can be protected when building applications on the Apama platform. See "Protecting Personal Data in Apama Applications" in *Developing Apama Applications* for detailed information.
- The EPL syntax highlighting files for Emacs and Vim have been updated with this release. You can find them in the `utilities` directory of your Apama installation.

Platform changes in 10.3.0

Apama 10.3.0 runs on the platforms listed in the *Supported Platforms* document for version 10.3.0. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>.

Check the above mentioned *Supported Platforms* document for details about the newer platform versions supported by Apama 10.3.0. Due to these upgrades, the following versions are no longer supported:

- Software AG Common Runtime 10.2 (replaced by support for a more recent version).
- Oracle WebLogic Server 12.1.3 as an application server (replaced by support for a more recent version).
- IBM WebSphere Application Server 8.5.5.8 (Liberty).
- Microsoft SQL Server 2014.
- Software AG Universal Messaging 10.2 as a JMS provider (replaced by support for a more recent version).
- Oracle WebLogic Server 10.3.5 and 12.1.3 as a JMS provider (replaced by support for a more recent version).
- Software AG Universal Messaging Client 10.2 (replaced by support for a more recent version).
- Oracle Java 8 (replaced by Java support from a different vendor).

- Software AG Designer 10.2 (replaced by support for a more recent version).
- Google Chrome 57 (replaced by support for a more recent version).
- Mozilla Firefox 52 (replaced by support for a more recent version).
- Apple Safari.
- Software AG Terracotta BigMemory Max 4.3.5 (replaced by support for a more recent version).
- Software AG Terracotta DB Store 10.2 (replaced by support for a more recent version).

Miscellaneous changes in 10.3.0 affecting backwards compatibility

The following changes in Apama 10.3.0 affect backwards compatibility to previous Apama versions:

- The `Row.set()` action of the `MemoryStore` can be used to set non-schema fields in distributed stores which support them. With previous versions, setting a field with an empty any has set a field with the value "empty". As of this version, setting a field with an empty any removes the non-schema field with the given name.
- The `engine_deploy` tool now fails to generate output when the input project contains errors. This changes the way in which the output is shown when Ant is run for an Ant-exported project. With previous versions, the Ant build was successful and the errors were shown afterwards while injecting the files. As of this version, the errors are already shown in the console when the Ant build fails.
- With previous versions, the correlator incorrectly converted a float of the form `-0.0` or a decimal of the form `-0.0d` to a `"0"` string (that is, the minus sign was missing). As of this version, this is correctly converted to a `"-0"` string. This change may lead to a backward compatibility issue for any software using Apama that is checking exclusively for `"0"` strings (without minus signs) and cannot handle the case of a string of the form `"-0"`.
- The location of the following files in your Apama installation has changed from `Apama/monitors/` to `Apama/monitors/cumulocity/9.0/`:
 - `Cumulocity_EventDefinitions.mon`
 - `Cumulocity_ServiceMonitor.mon`
 - `Cumulocity_Utills.mon`

These files were also delivered with 10.2.0.1. If you are explicitly using these files in a script or manual injection ordering, update the location of these files.

- The Docker images provided by Software AG are now configured to run processes as the `sagadmin` user by default within containers. For more information, see "Quick start to using an Apama image" in *Deploying and Managing Apama Applications*.
- Previous releases of Apama in Docker Store declared the `${APAMA_WORK}` directory to be a volume. This made it hard to use the work directory when creating your own derived images because of limitations on how it is handled. Therefore, the `VOLUME` declaration has been removed.

If you need to restore this functionality, you can add a `VOLUME` statement in your own derivative Dockerfile after populating the work directory.

- The MQTT transport connectivity plug-in now uses TLS version 1.2. This may affect existing deployments. See also "Configuring the connection to MQTT" in *Connecting Apama Applications to External Components*.
- Apama now ships Python 3.6.6 instead of Python 2.7. You may have existing PySys tests written in Python 2 that you may wish to migrate to Python 3. It is recommended to refer to the internet for the details of what this change from Python 2 to Python 3 entails, as well as recipes for upgrading test cases. It is likely that updates will be required, particularly around the handling of strings, file I/O and print statements. There are many tools, such as `2to3.py`, available to aid in this. Support for Python 2.7 is now deprecated; see also "[Removed and deprecated features in 10.3.0](#)" on page 93.

Although Python 2.7 is no longer shipped, you can continue to use existing PySys tests written for Python 2 with Apama 10.3.0. However, you have to manually copy the `pysys` and `apama` directories from the `Apama\third_party\python\Lib\site-packages` directory on Windows or the `Apama/third_party/python/lib/python2.7/site-packages` directory on Linux into your own Python 2.7.13 installation. On Windows, you must also ensure that the `pywin32` library is installed into your Python 2 installation (from <https://pypi.org/>).

- With PySys 1.3, the `pysys.py run` command now returns a non-zero exit code if any tests failed (previously, this command returned 0 on any failed tests). See also "[PySys changes in 10.3.0](#)" on page 85.
- The `docker` module contains classes to help create and orchestrate Docker and Kubernetes images from within PySys. This module has been renamed to `apama.docker`. In this release, the old name `docker` can still be used, but is deprecated. If you run Python with deprecation warnings enabled, you will see a warning if you are using the old name. It is recommended that you update your tests to use `apama.docker` because the `docker` module name will be removed in a future release.
- The `defaultValue` rule of the Mapper codec will no longer always be applied to metadata fields, and instead will only be applied in the case of an empty/null metadata field. See also "The Mapper codec connectivity plug-in" in *Connecting Apama Applications to External Components*.

Removed and deprecated features in 10.3.0

The following Apama features are now deprecated or have been removed in Apama 10.3.0:

- Support for Java Web Start has been removed in this version. Thus, it is no longer possible to deploy a dashboard as a Web Start application. Instead, you should now deploy it either as a simple, thin-client web page or as a locally installed desktop application (the Dashboard Viewer). See also "Deployment options" in *Building and Using Apama Dashboards*.
- In Software AG Designer, it is no longer possible to select the Web Start deployment type in the Dashboard Deployment configuration editor. Existing `dashboard_deploy.xml` files that specify this deployment type will automatically have this type removed from them when opened. It is then up to you to specify a different deployment type.

- As of this version, you can no longer use the CLASSPATH environment variable to configure the correlator's system classpath. See also the announcement in [“Removed and deprecated features in 10.1” on page 115](#).
- The REST API key `virtualMemorySize` is now deprecated and will be removed in a future release. Use the new key `virtualMemoryMB` instead. See also [“Miscellaneous enhancements and changes in 10.3.0” on page 90](#).
- Python 2.7 is no longer shipped with Apama. Support for Python 2.7 is now deprecated and will be removed in a future release. Until support is removed, PySys and the shipped Apama extensions can be used with a Python 2.7.13 installation if one is provided by the user. See [“Miscellaneous changes in 10.3.0 affecting backwards compatibility” on page 92](#) for details of how to do that. It is strongly recommended to use Python 3 for any new development.

13 What's New In Apama 10.2

■ New support for using the TCStore driver for queries in 10.2	96
■ New Kubernetes support in 10.2	96
■ Connectivity plug-ins enhancements in 10.2	96
■ Apama enhancements in Software AG Designer 10.2	97
■ Query enhancements in 10.2	98
■ EPL enhancements in 10.2	98
■ Correlator utility enhancements in 10.2	99
■ Command Central support in 10.2	100
■ Miscellaneous enhancements and changes in 10.2	100
■ Platform changes in 10.2	101
■ Miscellaneous changes in 10.2 affecting backwards compatibility	102
■ Removed and deprecated features in 10.2	103

Apama 10.2 runs on the platforms listed in the *Supported Platforms* document for version 10.2. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>. Be sure to consult that document for details about supported versions of operating systems, compilers and other products.

Apama 10.2 includes new features, enhancements, and changes as described in the following topics.

New support for using the TCStore driver for queries in 10.2

For Apama queries and applications that do *not* involve caching (such as store and "system of record" use cases), the TCStore driver is now the recommended driver for the distributed MemoryStore. See "TCStore (Terracotta DB) driver details" in *Developing Apama Applications*. See also "Configuring the TCStore driver" in *Deploying and Managing Apama Applications*.

Apama continues to support the BigMemory Max driver. This is still the recommended driver for caching use cases, but no longer for queries. See "BigMemory Max driver details" in *Developing Apama Applications*.

Support for using BigMemory Max for Apama queries is deprecated and will be removed in a future release. It is recommended that you now use Terracotta DB's TCStore for queries.

In the Distributed MemoryStore Configuration Wizard of Software AG Designer, when you select the third-party cache provider from the **Store provider** drop-down list box, you can no longer select the option **Apama Queries (using BigMemory)**. This option has been replaced by the **Apama Queries (using TCStore)** option. See also "Adding a distributed store to a project" in *Developing Apama Applications*.

If you still want to use BigMemory Max for queries in new projects, you must now select **BigMemory** as the store provider and specify "ApamaQueriesStore" as the store name. See also "Configuring the BigMemory Max driver" in *Deploying and Managing Apama Applications*. Existing deployments using BigMemory Max for queries are unaffected; this only covers developing new projects in Software AG Designer.

For detailed information on how to change a queries project that uses BigMemory Max to use TCStore instead, see "Migrating from BigMemory Max to TCStore" in *Developing Apama Applications*.

See also "Deploying and Managing Queries" in *Deploying and Managing Apama Applications*, which provides detailed information on how to use queries with the distributed MemoryStore.

New Kubernetes support in 10.2

Apama now supports Kubernetes. Kubernetes can be used for container orchestration, that is, for automating the deployment of your Apama applications. See "Deploying Apama Applications with Kubernetes" in *Deploying and Managing Apama Applications* for detailed information.

Connectivity plug-ins enhancements in 10.2

Apama 10.2 includes the following connectivity plug-ins enhancements:

- A new EPL sample application is now available in the `samples/connectivity_plugin/application/genericsendreceive` directory of your Apama installation. This simple sample application provides an easy way to get started with sending and receiving messages to or from any connectivity plug-in. For more information, see the `README.txt` file in the above directory and "Sending and receiving events with connectivity plug-ins" in *Connecting Apama Applications to External Components*.
- The Kafka transport is now available as a pre-built connectivity plug-in. See also "The Kafka Transport Connectivity Plug-in" in *Connecting Apama Applications to External Components*.
- Protection against brute force attacks on users and passwords is now available for the HTTP server. Using the new configuration options `cooldownSecs` and `maxAttempts`, a user-configurable cool-down and maximum number of login attempts now allows failed requests to be throttled. In addition, handling of unknown users now mirrors allowed users to avoid user information leakage. See also "Configuring the HTTP server transport" and "Authentication" in *Connecting Apama Applications to External Components*.
- The Unit Test Harness codec connectivity plug-in now orders all maps and not just top level maps. The ordering of nested maps in the output file may be different to previous versions. See also "The Unit Test Harness codec connectivity plug-in" in *Connecting Apama Applications to External Components*.
- The C++ connectivity plug-in API now also supports creating user-defined dynamic chain managers for transports, which was previously only possible in Java (see also "[Connectivity plug-ins enhancements in 10.0](#)" on page 121). The C++ `com::softwareag::connectivity::chainmanagers` namespace has been added for this purpose; this is equivalent to the Java `com.softwareag.connectivity.chainmanagers` package. See also "Requirements of a transport chain manager plug-in class" in *Connecting Apama Applications to External Components* and the *API Reference for C++ (Doxygen)*.

In addition, a new macro is available for C++:

`SAG_DECLARE_CONNECTIVITY_TRANSPORT_CHAIN_MANAGER_CLASS(class-name)`, which is used for exporting a chain manager class. See also "Exporting the required symbols from C++ plug-ins" in *Connecting Apama Applications to External Components*.

A skeleton sample for C++ is provided in the `samples/connectivity_plugin/cpp/skeleton_chainmanager` directory of your Apama installation. You can use this sample as a starting point to write your own C++ chain manager and transport.

- A new `to_string(...)` function is now available in the `com::softwareag::connectivity` namespace of the C++ API. This function operates on `data_t`, `list_t` and `map_t` classes and returns a string representation of the object. It is intended for logging and similar tasks. It is not tuned for performance. See also the *API Reference for C++ (Doxygen)*.

Apama enhancements in Software AG Designer 10.2

Apama 10.2 includes the following enhancements in Software AG Designer:

- You can now add the Kafka connectivity bundle to your projects. For further information, see "Adding the Kafka connectivity plug-in to a project" in *Using Apama with Software AG Designer*.

- Software AG Designer now offers several different options to choose from when adding the Universal Messaging connectivity plug-in to your project, including a simple configuration for getting started quickly using JSON messages and an advanced configuration showing a wide range of possible options for configuring the Universal Messaging transport and for customizing the way Universal Messaging messages are mapped to Apama events. See also "Adding the Universal Messaging connectivity plug-in to a project" in *Using Apama with Software AG Designer*.
- Software AG Designer now offers several different options to choose from when adding the HTTP client connectivity plug-in to your project: an option that uses application-specific event definitions, transports and mapping configurations, and a "generic" option that uses generic Apama event definitions and a chain configuration to invoke HTTP services with JSON payloads. For the "generic" option, the HTTP client adds a reference to `com.apama.util.AnyExtractor` which is useful for easily extracting data from an any type. See also "Adding the HTTP client connectivity plug-in to a project" in *Using Apama with Software AG Designer* and "Using predefined generic event definitions to invoke HTTP services with JSON and string payloads" in *Connecting Apama Applications to External Components*.

Query enhancements in 10.2

Apama 10.2 includes the following query enhancements:

- In the **Actions** pane of the Query Designer, you can now route event actions to the correlator. For more information, see "Adding query route event actions" in *Using Apama with Software AG Designer*.

EPL enhancements in 10.2

Apama 10.2 includes the following EPL enhancements:

- The reflection capabilities have been improved. For example, it is now possible for an EPL application to query what actions a type has. For this purpose, the following new methods are now available for the any type:
 - `getActionNames`
 - `getAction`
 - `getConstant`
 - `getConstantNames`
 - `getActionParameterNames`
 - `getActionParameters`
 - `getActionReturnTypeName`
 - `getGenericAction`

See also "Reflection on types" in *Developing Apama Applications*.

- When listening to events of all types, you can now provide the following optional values:

- `typeName` to specify which routable type you are listening for.
- `values` to name indexable fields in the event.

See also "Listening for events of all types" in *Developing Apama Applications*.

- ApamaDoc is now available for all of the built-in types and their methods. This includes the primitive types, the `string` type, the reference types, and the `monitor` pseudo-type. See the *API Reference for EPL (ApamaDoc)* for detailed descriptions of all of these types. You can find them under **All Types**, and under the headings **Built-in types** and **Aggregates**. Alternatively, you can find them in the following packages:

- `<Default Package>` (most of the built-in types can be found here)
- `com.apama` (includes the `Channel` type)
- `com.apama.aggregates` (includes all of the built-in aggregate functions such as `avg`)
- `com.apama.exceptions` (includes the `Exception` and `StackTraceElement` types)

As the type descriptions are now part of the ApamaDoc, the corresponding topics under "Types" have been removed from the EPL Reference.

- The new `@throws` tag can now be used with ApamaDoc. It documents exceptions which may be thrown from an action and the conditions which would cause the action to throw an exception. See also "Inserting ApamaDoc tags" in *Developing Apama Applications*.
- A new `getInfo()` action is now available from the `com.apama.correlator.Component` event of the Management interface, which provides a way for an EPL application to get access to a variety of correlator statistics and operational information. See also "Using the Management interface" in *Developing Apama Applications*.

Correlator utility enhancements in 10.2

Apama 10.2 includes the following correlator utility enhancements:

- Using the `--outputDeployDir` action of the `engine_deploy` tool, you can now alternatively specify the name of a Zip file that is to contain the contents of the correlator deployment directory. See also "Deploying a correlator" in *Deploying and Managing Apama Applications*.
- With previous versions, when a directory containing a Software AG Designer project was specified with the `engine_deploy` tool, the resulting correlator deployment directory would not configure any connectivity plug-ins or any other kind of configuration. This is no longer the case. The resulting correlator deployment directory now accurately reflects the entire launch configuration, unless more than one launch configuration exists in the project in which case you have to choose a launch configuration. See also "Deploying a correlator" in *Deploying and Managing Apama Applications*.
- The YAML configuration file for the correlator now supports many new elements which correspond to command-line options of the `engine_connect` tool. For more information, see

"Setting up connections between correlators in a YAML configuration file" in *Deploying and Managing Apama Applications*.

- To mitigate potential injection attacks against the correlator web interface, the address used to access the web interface is now, for certain URLs, validated against a list of permitted host names. By default, this corresponds to the configured IP addresses and host names of the machine on which the correlator is running and the configured correlator port. This means that for most applications this change will be transparent. If your correlator web interface is deployed behind another component or has otherwise non-standard configuration, you may need to set the new configuration option `server/permittedRootURIs` to include the addresses from which the correlator will be accessed. For more information, see "Using the correlator web interface with non-default addresses" in *Deploying and Managing Apama Applications*.

Note:

This only applies to the correlator web interface that is accessed via the correlator port (the default is 15903). It does not apply to the HTTP server transport connectivity plug-in.

- If persistence is enabled, a new `Persistence Status: log` line is now shown in the main correlator log file when logging at INFO level. See "Persistence status" in *Deploying and Managing Apama Applications* for detailed information on the log fields that are shown in this line.

REST API keys with similar names are also available now. See ["Miscellaneous enhancements and changes in 10.2" on page 100](#).

Command Central support in 10.2

Apama 10.2 includes the following changes for its integration with Command Central:

- The **Initialization configuration** and **Configuration path** configurations for the correlator have been removed from Command Central. Respectively, these are the `APAMA-CORRELATOR-INITIALIZE` configuration type and the `connectivityYamlPath` property of `APAMA-ARGS`. This was the mechanism by which you configured the EPL and connectivity configuration for the correlator to run, replaced by a new mechanism which is documented at "Deploying a Designer project to Command Central" in *Deploying and Managing Apama Applications*.
- You can now override the correlator configuration properties using the new `APAMA-PROP-OVERRIDES` configuration type. In practice, you would use the `APAMA-PROP-OVERRIDES` configuration type for configuration that is different between your development environment (in your Designer project) and your production environment (as managed by Command Central), like the hostname of a messaging system used by a connectivity plug-in. For more information, see "Overriding correlator configuration" in *Deploying and Managing Apama Applications*.

For more information, see "Deploying Apama Components with Command Central" in *Deploying and Managing Apama Applications*.

Miscellaneous enhancements and changes in 10.2

Apama 10.2 includes the following miscellaneous enhancements and changes:

- The REST API now supports the `enqueueQueueSize` key which indicates the number of enqueued events. The status log field `eq` with the same meaning, however, is no longer shown in the main correlator log file. See also "List of correlator status statistics" in *Deploying and Managing Apama Applications*.
- The REST API now supports the following keys which provide statistics information when the correlator is running in persistent mode:
 - `persistenceNumSnapshots`
 - `persistenceLastSnapshotTime`
 - `persistenceSnapshotWaitTimeEwmaMillis`
 - `persistenceCommitTimeEwmaMillis`
 - `persistenceLastSnapshotRowsChangedEwma`

For more information, see "List of correlator status statistics" in *Deploying and Managing Apama Applications*.

Status log fields with similar names are also available now. See "[Correlator utility enhancements in 10.2](#)" on page 99.

- Generated documentation for the REST API is now available in the `doc/rest` directory of your Apama installation. See the *API Reference for Component Management REST APIs*. The corresponding information has therefore been removed from "Managing and Monitoring over REST" in *Deploying and Managing Apama Applications*.
- The Docker samples have been updated to support a newer Docker version (see the *Supported Platforms* document for version 10.2 which is available from <http://documentation.softwareag.com/apama/index.htm>). This includes making the samples usable with both an image built from an installation using the sample and the published image in the Docker Store. The Docker samples now also come with Kubernetes support. See also "Deploying Apama Applications with Kubernetes" in *Deploying and Managing Apama Applications*.

In addition, documentation on how to use Docker with Apama is now available. See "Deploying Apama Applications with Docker" in *Deploying and Managing Apama Applications*.

- The correlator no longer refuses to start if a distributed `MemoryStore` is enabled and a license file cannot be found. See also "Using the distributed `MemoryStore`" in *Developing Apama Applications*.

Platform changes in 10.2

Apama 10.2 runs on the platforms listed in the *Supported Platforms* document for version 10.2. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>.

Check the above mentioned *Supported Platforms* document for details about the newer platform versions supported by Apama 10.2. Due to these upgrades, the following versions are no longer supported:

- Microsoft Windows 7.

- Microsoft Windows Server 2008.
- Software AG Common Runtime 10.1 (replaced by support for a more recent version).
- Microsoft SQL Server 2012.
- Software AG Universal Messaging 10.1 as a JMS provider (replaced by support for a more recent version).
- IBM WebSphere MQ 7.0.1 as a JMS provider (replaced by support for a more recent IBM MQ version).
- Software AG Universal Messaging Client 10.1 (replaced by support for a more recent version).
- Apache Kafka 0.10.1 (replaced by support for a more recent version).
- Oracle Java 8 update 141 (replaced by support for a more recent version).
- Software AG Designer 10.1 (replaced by support for a more recent version).
- Google Chrome 53 (replaced by support for a more recent version).
- Docker 1.11.0 (replaced by support for a more recent version).
- Docker Compose 1.7.0 (replaced by support for a more recent version).
- Software AG Terracotta BigMemory Max 4.3.4 (replaced by support for a more recent version).
- Software AG Terracotta DB Store 10.1 (replaced by support for a more recent version).

Miscellaneous changes in 10.2 affecting backwards compatibility

The following changes in Apama 10.2 affect backwards compatibility to previous Apama versions:

- Apama 10.2 incorporates the ICU (International Components for Unicode) time-zone data update 2017c, which is the most recent update at the time of release. This will update time-zone data used by the correlator and Time Format EPL plug-in.

In addition, Apama 10.2 contains an upgrade from ICU 58.1 to ICU 60.1 which contains a small behavioral difference which may affect your applications: When formatting time in UTC, the output for a “GMT+offset” time zone is now “UTC” rather than “GMT” (non-UTC time zones are unaffected).

- It is no longer legal to define two event types with the same name and package in a Software AG Designer project. In previous versions, two similar event declarations in a project were shown as a warning in Software AG Designer. As of this version, this is shown as an error.
- An `IllegalArgumentException` is now thrown at runtime if an EPL dictionary is initialized with duplicate keys. With previous versions, no exception was thrown; the dictionary contained the key only once and the value was undefined.
- The ADBC Event API has been changed such that a single `com.apama.database.Query` event can no longer be used to run multiple queries concurrently. A single query event was never

intended to run concurrent queries. It is recommended to instantiate multiple `com.apama.database.Query` events if you want to run queries concurrently. To reuse the same `com.apama.database.Query` event, you must now set up appropriate listeners on it again, by calling `setBatchDoneCallback` and `setSchemaCallback` before calling `start`. See also "The ADBC Event application programming interface" in *Connecting Apama Applications to External Components* and the description of the `Query` event in the *API Reference for EPL (ApamaDoc)*.

Removed and deprecated features in 10.2

The following Apama features are now deprecated or have been removed in Apama 10.2:

- It is no longer possible to configure more than one Universal Messaging dynamic chain manager per correlator. This means that a given correlator can connect to only one Universal Messaging cluster. See also "Configuring the connection to Universal Messaging (dynamicChainManagers)" in *Connecting Apama Applications to External Components*.
- The deprecated `subscribe()`, `unsubscribe()` and `sendToDES()` actions on the auto-generated EPL representations of Digital Event Services types have been removed. Instead, you should now use the static `CHANNEL` constant on the EPL type to access the channel name. With `CHANNEL`, you can send, subscribe and unsubscribe in the same way as any other Apama channel. For more information, see "Using Digital Event Services connectivity from EPL" in *Connecting Apama Applications to External Components*. See also the announcement in "[Removed and deprecated features in 10.0](#)" on page 130.
- Support for using BigMemory Max for queries is deprecated and will be removed in a future release. It is recommended that you now use Terracotta DB's TCStore for queries. See also "[New support for using the TCStore driver for queries in 10.2](#)" on page 96.
- Support for setting the `readVisibility` and `writeVisibility` properties in the TCStore driver configuration for distributed memory stores is deprecated and will be removed in a future release. It is recommended that you remove these properties from your Spring configuration files. Setting these properties does not have any effect in version 10.2, but will result in an error in future versions of Apama. See also "TCStore (Terracotta DB) driver details" in *Developing Apama Applications*.
- In Software AG Designer, the **Build adapter files for monitorscript changes** option has been removed from the EPL build path. See also "Project source files" section in *Using Apama with Software AG Designer*.
- The **Initialization configuration** and **Configuration path** configurations for the correlator have been removed from Command Central. Respectively, these are the `APAMA-CORRELATOR-INITIALIZE` configuration type and the `connectivityYamlPath` property of `APAMA-ARGS`. This was the mechanism by which you configured the EPL and connectivity configuration for the correlator to run, replaced by a new mechanism which is documented at "Deploying a Designer project to Command Central" in *Deploying and Managing Apama Applications*.

14 What's New In Apama 10.1

■ New any type in 10.1	106
■ New HTTP server transport connectivity plug-in in 10.1	106
■ New EPL plug-in for R in 10.1	107
■ New TCStore driver (for Terracotta DB) for the MemoryStore in 10.1	107
■ MemoryStore enhancements in 10.1	107
■ Connectivity plug-ins enhancements in 10.1	108
■ Apama enhancements in Software AG Designer 10.1	109
■ Query enhancements in 10.1	110
■ EPL enhancements in 10.1	110
■ Correlator utility enhancements in 10.1	110
■ Dashboard enhancements in 10.1	112
■ Miscellaneous enhancements and changes in 10.1	112
■ Platform changes in 10.1	112
■ Miscellaneous changes in 10.1 affecting backwards compatibility	113
■ Removed and deprecated features in 10.1	115

Apama 10.1 runs on the platforms listed in the *Supported Platforms* document for version 10.1. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>. Be sure to consult that document for details about supported versions of operating systems, compilers and other products.

Apama 10.1 includes new features, enhancements, and changes as described in the following topics.

New any type in 10.1

EPL now supports any types. Values of the any type can hold a value of a concrete EPL type (that is, a type other than the any type). Different values of any variables may have different types. This is useful when mapping to generic data formats such as JSON, or where the data in an event or being processed by an action may be one of a number of different types. An any value may also be empty and not contain a value (similar to the concept of `null` in other languages).

The new `switch` statement operates on an expression of the any type. See "Handling any values of different types with the switch statement" in *Developing Apama Applications* for more information.

EPL supports casting of the any type to a concrete target type and vice versa. See "Handling the any type" in *Developing Apama Applications* for more information.

The `MemoryStore` now supports the any type. A valid schema can be created from an event type using `schemaFromAny(event)`. New actions to get and set fields and rows from the any type and from `dictionary<string,any>` have been added. See "Description of row structures" in *Developing Apama Applications* for more information.

The connectivity plug-ins can now handle values of the any type. This means that an empty `data_t` or a `null` value in Java will translate to/from an empty type. See "Map contents used by the `apama.eventMap` host plug-in" in *Connecting Apama Applications to External Components*.

The EPL plug-ins written in C++ support the any type, converting it to a `data_t`. See "Method signatures" in *Developing Apama Applications*.

Note: any type support is currently not available for the EPL plug-ins written in Java.

The Apama event parser libraries for Java and .NET have been updated to support events with any fields. A new class `AnyFieldType` has been added to represent the any field type. To access it, you have to create new `AnyFieldType()`. For complete information about this class, refer to the description of

- `com.apama.event.parser` in the *API Reference for Java (Javadoc)*
- `Apama.Event.Parser` Namespace in the *API Reference for .NET*

The Apama event parser library for C has been updated to support events with any fields.

New HTTP server transport connectivity plug-in in 10.1

A new connectivity plug-in, the HTTP server transport, is now available. It can handle HTTP requests and deliver the contained events to the correlator, which will return acknowledgment of

the requests. In addition to this, it can serve static files. It also provides support for TLS (Transport Layer Security) alongside HTTP basic authentication. For detailed information, see "The HTTP Server Transport Connectivity Plug-in" in *Connecting Apama Applications to External Components*.

Using Software AG Designer, you can add the new HTTP server connectivity bundle to your projects. For further information, see "Adding the HTTP server connectivity plug-in to a project" in *Using Apama with Software AG Designer*.

New EPL plug-in for R in 10.1

A new EPL plug-in, the R plug-in, is now available. It enables an EPL application to use an external R server using Rserve. The R plug-in requires the new **R Support** bundle for use with Software AG Designer. See "Using the R plug-in" in *Developing Apama Applications* for detailed information.

New TCStore driver (for Terracotta DB) for the MemoryStore in 10.1

Apama 10.1 adds support for Terracotta DB. This is provided as a driver that can be used by any Apama application using the existing Apama MemoryStore API.

The new driver uses the TCStore API to provide access to Terracotta DB's storage capabilities. This allows Apama to read and write *records* in TCStore *datasets*, which may also be read and written by other Apama correlators and non-Apama components such as other Software AG products or custom clients written against the TCStore API.

The new TCStore driver can be added to projects and easily configured using Apama's Distributed MemoryStore editor in Software AG Designer (see "Configuring a distributed store" in *Developing Apama Applications*).

Apama's TCStore driver only supports clustered mode. Currently, there is no support for embedded TCStore, for performing searches from Apama, or for the Ehcache API.

Apama continues to support the MemoryStore driver for Terracotta BigMemory Max (see "BigMemory Max driver details" in *Developing Apama Applications*), which remains the recommended driver for both Apama queries and for caching use cases. For applications that do not involve queries or caching, TCStore is now the preferred MemoryStore driver. See "TCStore (Terracotta DB) driver details" in *Developing Apama Applications* for detailed information.

MemoryStore enhancements in 10.1

Apama 10.1 includes the following enhancements for the MemoryStore:

- The MemoryStore EPL API has several new actions that make use of the new any type to simplify operations. For example, a schema can be created from an event type using `schemaFromAny(event)` and there are actions that support getting and setting the contents of a Row using an Apama event type with fields matching the fields of the row. There are also new actions to get and set fields and rows making use of the any type. See "Description of row structures" in *Developing Apama Applications* for more information.

- The MemoryStore EPL API has a new `Row.toDictionary()` action that returns a `dictionary<string, any>` containing all fields in the row. This is useful for diagnostic purposes such as logging and testing, as well as for application development.
- The distributed MemoryStore driver API now allows drivers to provide access to *extra fields* that are present in individual rows but not in the table schema. The existing `Row.getXXX` and `Row.setXXX` actions in the EPL API can now be used to get and set such extra fields for stores that support this feature, such as TCStore. The new `Row.getKeys()` and `Row.toDictionary()` actions provide a way to find out what fields are present in a given row, including both fields from the schema and extra fields. See "Description of row structures" in *Developing Apama Applications* for more information about using this feature from EPL, and see "Creating a distributed MemoryStore driver" in *Developing Apama Applications* for more information about writing a driver that supports it.
- The `forceCommit()` action is now available for the `com.apama.memorystore.Row` event. This new action allows you to commit the local Row back to the table even if the row in the table has been modified after the Row event was obtained. See also "Using transactions to manipulate rows" in *Developing Apama Applications*.

Connectivity plug-ins enhancements in 10.1

Apama 10.1 includes the following connectivity plug-ins enhancements:

Reliable messaging

- Events can now be marked with a `MessageId` annotation. This is used for reliable receiving. The new `MessageId` annotation names a field that should contain the unique identifier of the connectivity plug-in message that the event came from. It is thus no longer necessary to make use of `copyFrom` to get the message identifier into a field. For more information, see "Adding predefined annotations" in *Developing Apama Applications*.

Dynamic chain managers API

- The new method `ConnectivityPlugins.createDynamicChain` is now available in EPL. It replaces the deprecated `ConnectivityPlugins.createChain` method. The new method allows you to specify an additional argument: `defaultChannelTowardsHost`. For more information, see "Creating dynamic chains from EPL" in *Connecting Apama Applications to External Components*.

Miscellaneous connectivity framework changes

- A new helper class called `com::softwareag::connectivity::MapExtractor` is available in the C++ API. This class provides an easy and type-safe way to extract values from `map_t` objects such as those used for plug-in configuration. It has automatic support for generating user-friendly error messages if configuration items are missing or of the wrong type. See the *API Reference for C++ (Doxygen)* for further information.

Note:

For Java, this functionality has already been added in 10.0. See also "[Connectivity plug-ins enhancements in 10.0](#)" on page 121.

Enhancements and additions to standard connectivity plug-ins

- The MQTT transport connectivity plug-in now supports TLS (Transport Layer Security). The following new configuration options are available for this purpose: `acceptUnrecognizedCertificates` and `certificateAuthorityFile`. For more information, see "Configuring the connection to MQTT" in *Connecting Apama Applications to External Components*.
- The HTTP client transport connectivity plug-in now requires TLS version 1.2 for connecting to secure https URLs. See also "About the HTTP client transport" in *Connecting Apama Applications to External Components*.
- The HTTP client transport connectivity plug-in now has an additional configuration option `basePath` which defines an optional path to be prefixed to the `metadata.http.path` for all messages sent to the transport. See also "Configuring the HTTP client transport" in *Connecting Apama Applications to External Components*.
- The Mapper codec now supports `forEach` to define mapping rules for elements nested in a list. See also "The Mapper codec connectivity plug-in" in *Connecting Apama Applications to External Components*.

Apama enhancements in Software AG Designer 10.1

Apama 10.1 includes the following enhancements in Software AG Designer:

- You can now export Apama event definitions as digital event types using the Digital Event Types editor. See also "Using the Digital Event Services connectivity bundle" in *Using Apama with Software AG Designer*.
- The following commands from the **File > New** menu (and from the **New** context menu) have been renamed:
 - **MonitorScript File** has been renamed to **EPL Monitor**.
 - **Event Definition** has been renamed to **EPL Event Definition**.

These name changes are also reflected in the wizards that are invoked by these commands.

For both commands, validations for the name have been implemented such that wizard now shows an error message if name contains special characters, except for the dollar (\$) and underscore (_) characters.

The EPL editor now provides a template with the specified monitor name or event definition name.

See the following topics in *Using Apama with Software AG Designer* for more information:

- "Creating new monitor files for EPL applications"
- "Creating event definitions by adding EPL code"
- "Creating event definitions from XML files"
- "Creating event definitions from XSD files"

Note:

Support for EDA events is deprecated. See also [“Removed and deprecated features in 10.1” on page 115](#).

Query enhancements in 10.1

Apama 10.1 includes the following query enhancements:

- It is now possible to use correlator-integrated JMS to ensure that queries do not lose messages if a node in a cluster is failing (for example, if the node is stopped for a planned maintenance). For more information, see "Handling node failure and failover" in *Developing Apama Applications*.
- The `route` statement is now allowed in queries, and is now the recommended mechanism for sending events between queries. See also "Communication between monitors and queries" in *Developing Apama Applications*.

EPL enhancements in 10.1

Apama 10.1 includes the following EPL enhancements:

- A new EPL method `integer.incrementCounter(string)` is now available. It can be used to generate a monotonically increasing integer for the corresponding string key passed as a function argument in the scope of the correlator. This is a type method as well as an instance method.
- A new EPL method `context.iqSize()` is now available. It can be used to return an integer that is the current number of elements in the input queue for the context.
- All string literals in EPL injections are now interned automatically. Interned strings are now also persisted if the correlator is persistent. This removes the need to inject monitors for the sole purpose of interning string constants.

Correlator utility enhancements in 10.1

Apama 10.1 includes the following correlator utility enhancements:

- Using the new `-v category=level` (or `--loglevel category=level`) option of the correlator executable, it is now possible to configure the log level for messages logged by the correlator-internal systems. You can do this either for the whole correlator or for individual categories. This includes setting the log level for individual connectivity plug-ins and EPL plug-ins. For detailed information, see "Starting the correlator" and "Setting correlator and plug-in log files and log levels in a YAML configuration file" in *Deploying and Managing Apama Applications*.
- With the correlator executable, it is now possible to set distinct and individual log levels for the correlator and the application. The following options are now available: `--applicationLogLevel` and `--applicationLogFile`. If these two options are not set, then the current settings of the `-v` (or `--loglevel`) and `-f` (or `--logfile`) options are used as the default. For detailed information, see "Starting the correlator" in *Deploying and Managing Apama Applications*.

Note that less log levels are now allowed for the correlator log messages, that is, for the `-v` (or `--loglevel`) option. `OFF`, `CRIT` and `FATAL` are no longer supported. If specified, they are reset to `ERROR` and a warning message is printed at the top of the correlator log file. In addition, the use of `ERROR` and `WARN` is no longer recommended. If specified, a warning message is also printed at the top of the correlator log file.

- The YAML configuration file for the correlator now supports many new elements which correspond to command-line options of the `correlator` executable. For more information, see "YAML configuration file for the correlator" in *Deploying and Managing Apama Applications*.
- The `engine_deploy` tool now converts all of the options in a Software AG Designer launch configuration to the equivalent YAML configuration options, now that they exist. Previously, only a small number were honored. The new options are:
 - `javaApplicationSupport`
 - `jmsConfigDir`
 - `distMemStoreConfigDir`
 - `inputLogFile`
 - `logFile`
 - `logLevel`
 - `name`
 - and the JVM heap and offheap settings (supported via `jvmOptions`)

This may be a backwards-incompatible change if you had settings in your launch configuration that were previously being ignored. The generated files also contain support for some options not present in Software AG Designer, but which are often useful at deployment time: `pidFile`, `runtime` and `runtimeCacheDir`.

The generated YAML file also contains configuration to set an environment variable `APAMA_PROJECT_DIR`, which can be used to specify classpath dependencies for EPL plug-ins written in Java in the deployment descriptor file (see also "Specifying the classpath in deployment descriptor files" in *Developing Apama Applications*). For example:

```
<classpath>${env:APAMA_PROJECT_DIR}/lib/foo.jar</classpath>
```

The properties file now contains the properties `logsDir` and `dataDir` to allow easily changing the location where data (for example, runtime cache directory, pid file) and log files are written without needing to modify the properties for each file path individually.

The properties file also contains a property `startTimeoutSecs`. This value is not currently used by any Apama tools, but can be used to provide a hint to customer-developed deployment or testing tools about how long to wait for the component to start up.

Note that the `extraArgs` settings are copied into the properties file, but cannot be used by the YAML configuration file for the correlator. Therefore, they must be manually passed on the command line by whatever tool is responsible for starting the correlator.

For more information on the `engine_deploy` tool, see "Deploying a correlator" in *Deploying and Managing Apama Applications*.

- Instead of using the `engine_deploy` tool to generate a correlator deployment package (CDP), you can now also use Ant to do this. The `apama-macros.xml` file includes the new `project-to-cdp` macro for this purpose. See also "About deploying Apama applications with an Ant script" in *Deploying and Managing Apama Applications*.
- In previous versions, the correlator automatically and silently substituted a dot (.) for `${APAMA_HOME}` or `${APAMA_WORK}` if they were used within a YAML configuration file without being defined as an environment variable or as a property within the YAML configuration itself. This was hard to diagnose and incorrect. The correlator will now no longer do this automatic substitution and will fail to start up, as with any other undefined YAML configuration variable. Note that this does not affect correlators that are started through Software AG Designer or using the Apama Command Prompt, both of which correctly set these variables.

Dashboard enhancements in 10.1

Apama 10.1 includes the following dashboard enhancements:

- The HTML5 tables now support the tree grid view. For more information, see "HTML5 tables" in *Building and Using Apama Dashboards*.
- A new field **Hidden menu commands** has been added to the **Startup and Server** section. This field specifies the context menu commands that are to be hidden for the display server. See also "Startup and Server" in *Building and Using Apama Dashboards*.

Miscellaneous enhancements and changes in 10.1

Apama 10.1 includes the following miscellaneous enhancements and changes:

- Apama 10.1 incorporates the ICU (International Components for Unicode) time-zone data update 2017b, which is the most recent update at the time of release. This will update time-zone data used by the correlator and Time Format EPL plug-in.
- A new sample is available in the `samples/monitorscript/hot-redeploy/StoreState` directory of the Apama installation which shows how to transfer monitor state using the `MemoryStore`. For more information, see "Versioning and upgrading monitors" in *Developing Apama Applications*.
- The Apama event parser library for C has been updated to support events with `optional` fields.

Platform changes in 10.1

Apama 10.1 runs on the platforms listed in the *Supported Platforms* document for version 10.1. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>.

Check the above mentioned *Supported Platforms* document for details about the newer platform versions supported by Apama 10.1. Due to these upgrades, the following versions are no longer supported:

- Red Hat Enterprise Linux 6 (x86) 64-bit.
- Red Hat Enterprise Linux 7.0 (x86) 64-bit (replaced by support for a more recent version).
- CentOS Linux 6 (x86) 64-bit.
- CentOS Linux 7.0 (x86) 64-bit (replaced by support for a more recent version).
- Oracle Linux 6 (x86) 64-bit.
- Oracle Linux 7.0 (x86) 64-bit (replaced by support for a more recent version).
- Software AG Common Runtime 10.0 (replaced by support for a more recent version).
- Apache Tomcat 8.5.5 (replaced by support for a more recent version).
- Software AG Universal Messaging 10.0 as a JMS provider (replaced by support for a more recent version).
- IBM WebSphere Application Server 8.5.5.0 as a JMS provider (replaced by support for a more recent version which supports Java 8).

Note:

Support for IBM WebSphere Application Server 8.5.5.8 (Liberty) is now deprecated and will be removed in a future release.

- Software AG Universal Messaging Client 10.0 (replaced by support for a more recent version).
- Oracle Java 8 update 121 (replaced by support for a more recent version).
- Software AG Designer 10.0 (replaced by support for a more recent version).
- Google Chrome 48 (replaced by support for a more recent version).

Miscellaneous changes in 10.1 affecting backwards compatibility

The following changes in Apama 10.1 affect backwards compatibility to previous Apama versions:

- The evaluation order of some of the operations in the interpreted EPL runtime has been changed and is now in line with the compiled EPL runtime. Anyone who is reliant on the old evaluation order may see a change of behavior. More specifically, the evaluation order for assignment operations and the lookup/store operations for the `dictionary` and `sequence` types have been changed. For assignment operations, the left-hand side is evaluated before the right-hand side. The expression to get the `dictionary` or `sequence` is evaluated before the expression to get the key or index.
- For the sake of consistent naming, several SSL configuration options have been renamed in the standard connectivity plug-ins.

Renamed options in the Universal Messaging transport connectivity plug-in (see also "Configuring the connection to Universal Messaging (dynamicChainManagers)" in *Connecting Apama Applications to External Components*):

- `authentication/certificateAuthority` is now `certificateAuthorityFile` and no longer under `authentication`.
- `authentication/certificate` is now `authentication/certificateFile`.

Renamed options in the HTTP client transport connectivity plug-in (see also "Configuring the HTTP client transport" in *Connecting Apama Applications to External Components*):

- `acceptUnrecognizedCertificates` is now named `tlsAcceptUnrecognizedCertificates`.
- `certificateAuthorityPath` is now named `tlsCertificateAuthorityFile`.
- A new subsection `authentication` has been added which now contains the following options: `authenticationType`, `username` (which has been renamed from `user`), and `password`.

If you are using the above SSL configuration options, it is recommended that you change your YAML configuration files to use the new names.

- For the sake of consistent naming, the following elements in the YAML configuration file for the correlator have been renamed:
 - In the correlator section: `runtime-cache` has been renamed to `runtimeCacheDir`.
 - In the server section: `pidfile` has been renamed to `pidFile`.
 - In the correlator section, under the persistence entry: `snapshotInterval` has been renamed to `snapshotIntervalMillis`.

Note:

The `-PsnapshotInterval` option of the correlator executable has also been renamed to `-PsnapshotIntervalMillis`. You can still use the old option name, but it is deprecated, and it is recommended that you now use the new name. See also "Starting the correlator" in *Deploying and Managing Apama Applications*.

In addition, the elements that were previously located in the messaging section of the YAML configuration file have been moved to the server section. These are `maxOutstandingSecs` and `maxOutstandingKb`. Note that the description for `maxOutstandingKb` has been corrected (it was wrongly called `maxOutstandingBytes` in the documentation for the previous version). It takes the same units as the `-M` command-line option.

If you are using the above elements, it is recommended that you change your YAML configuration files to use the new names.

See also "YAML configuration file for the correlator" in *Deploying and Managing Apama Applications*.

- The `engine_deploy` tool now converts all of the options in a Software AG Designer launch configuration to the equivalent YAML configuration options. This may be a backwards-incompatible change if you had settings in your launch configuration that were previously being ignored. See also the information on `engine_deploy` in ["Correlator utility enhancements in 10.1" on page 110](#).
- Software AG Designer no longer sets the `CLASSPATH` environment variable when it launches the correlator. This brings Software AG Designer behavior in line with deployment behavior

(for example, Ant macros) which also does not set any classpath environment at the correlator level. EPL plug-ins which require classpaths to be set should do this via the `apama_java.xml` configuration in Software AG Designer (as described in "Specifying the classpath in deployment descriptor files" in *Developing Apama Applications*) or by starting the correlator with `-J-Djava.class.path=path`.

This also applies to the IAF adapters, which have the classpath defined in their configurations. Software AG Designer no longer sets the `CLASSPATH` environment variable when it launches the IAF. Therefore, any projects that relied upon that behavior should be changed to set the classpath in the IAF configuration files instead (as described in "The IAF configuration file" in *Connecting Apama Applications to External Components*).

- Any Software AG Designer project containing queries now automatically requires the **Correlator-Integrated JMS Application Support** bundle. This enhancement will only require an action on your part if your launch configuration uses the **Manual Ordering** option for injections, in which case you will need to take care that the new JMS files are injected *before* any query (`.qry`) files. See also "Injections" in *Using Apama with Software AG Designer*.
- A Universal Messaging cluster which has dropped below quorum of operating nodes is now considered completely down and is therefore disconnected, instead of remaining connected to a cluster which is non-functioning. In particular, this now means that it is not possible to start up a correlator if the cluster to which it is connecting is below quorum. If a cluster drops below quorum while the correlator is operating, the normal reconnect logic applies as it does in a single node or a cluster that has been entirely disconnected.
- The query statistics are now updated every 10 seconds by default. In previous versions, they were updated every second. By sending a `SetQueryStatisticsPeriod` event, you can now control the frequency of the statistics gathering or disable query statistics entirely. See also "Monitoring running queries" in *Deploying and Managing Apama Applications*.

Removed and deprecated features in 10.1

The following Apama features are now deprecated or have been removed in Apama 10.1:

- Software AG's Event-Driven Architecture (EDA) is deprecated. Support for EDA events will be removed in a future release. It is recommended that your Apama applications use digital event types instead, as provided by Software AG Digital Event Services. For further information, see "The Digital Event Services Transport Connectivity Plug-in" in *Connecting Apama Applications to External Components*.
- Use of the `CLASSPATH` environment variable to configure the correlator's system classpath is deprecated and will be removed in a future release. It is recommended that you set the classpath on a per-plug-in basis where possible, for example, in the descriptor file for an EPL or JMon plug-in (see "Specifying the classpath in deployment descriptor files" in *Developing Apama Applications*) or in the configuration file for a connectivity plug-in, JMS or distributed MemoryStore library. For cases where you really need to set the global system classpath, you can use the `-J` (or `--javaopt`) option of the correlator executable, along with `-Djava.class.path=path` (as described in "Starting the correlator" in *Deploying and Managing Apama Applications*).

- The `ConnectivityPlugins.createChain` method in EPL has been deprecated. Use `ConnectivityPlugins.createDynamicChain` instead. See also [“Connectivity plug-ins enhancements in 10.1” on page 108](#).
- The `--remoteDeploy` and `--remoteUndeploy` options of the `dashboard_management` tool (see “Managing and stopping the data server and display server” in *Building and Using Apama Dashboards*) have been removed in 10.1. These options relied on third-party technology that has not been officially supported since 5.1.
- As of this version, the deprecated EngineManagement API for C can no longer be used. It is recommended that you use the EngineManagement API for C++ instead. See “EngineManagement API” in *Connecting Apama Applications to External Components*. Any clients written using `engine_client_c.h` will need to be rewritten to use `engine_client_cpp.hpp` instead. This includes clients written in pure C and clients written using the C++ shims defined in the `engine_client_c.h` file. See also the announcement in [“Removed and deprecated features in 9.12” on page 145](#).
- As of this version, the deprecated native Universal Messaging integration can no longer be used. It is recommended that you use the Universal Messaging connectivity plug-in instead. See also [“Apama integration with Universal Messaging in 9.12” on page 134](#).

When using the connectivity plug-in to mimic the removed native integration, you should use a single `dynamicChains` definition:

```
dynamicChains:
  umString:
    - apama.eventString:
        suppressLoopback: true
        description: "@{um.rnames}"
        remoteAddress: "@{um.rnames}"
    - stringCodec
        nullTerminated: true
    - umTransport:
        channelPattern: ".*"
```

The configuration of the chain manager depends on your previous Universal Messaging configuration:

Removed configuration	New configuration
<pre>um.channels.prefix=UM_ um.channels.mode=precreate</pre>	<pre>managerConfig: channel: prefix: UM_ includePrefixOnUM: true missingChannelMode: ignore</pre>
<pre>um.channels.prefix=UM_ um.channels.mode=autocreate</pre>	<pre>managerConfig: channel: prefix: UM_ includePrefixOnUM: true missingChannelMode: create</pre>

Removed configuration	New configuration
<pre>um.channels.prefix=UM_ um.channels.mode=mixed</pre>	<pre>managerConfig: channel: prefix: UM_ includePrefixOnUM: true missingChannelMode: create</pre>
<pre>um.channels.mode=precreate</pre>	<pre>managerConfig: channel: prefix: "" missingChannelMode: ignore</pre>

In the removed configuration, it was possible to configure some channels (without a prefix) to be looked up on Universal Messaging but not created, while also specifying a prefix under which channels would be created. This specific use case is not possible to directly port to the new-style configuration.

Note:

The above configuration style which has been removed for the correlator is still used for the IAF (see "Defining Universal Messaging properties for the IAF" in *Connecting Apama Applications to External Components*). For communication with the correlator, however, the Universal Messaging transport connectivity plug-in is required, and this connectivity plug-in must be configured to be equivalent to the Universal Messaging properties you have configured for the IAF (see also "Communicating with the correlator over Universal Messaging" in *Connecting Apama Applications to External Components*).

15 What's New In Apama 10.0

■ New HTTP client transport connectivity plug-in in 10.0	120
■ Apama integration with MQTT in 10.0	120
■ Apama integration with Kafka in 10.0	120
■ Connectivity plug-ins enhancements in 10.0	121
■ New API for writing EPL plug-ins in C++ in 10.0	124
■ Apama enhancements in Software AG Designer 10.0	124
■ Query enhancements in 10.0	125
■ EPL enhancements in 10.0	125
■ Correlator utility enhancements in 10.0	125
■ Dashboard enhancements in 10.0	126
■ Miscellaneous enhancements and changes in 10.0	126
■ Platform changes in 10.0	127
■ Windows compiler update in 10.0	128
■ Miscellaneous changes in 10.0 affecting backwards compatibility	128
■ Removed and deprecated features in 10.0	130
■ Migrating the C and C++ EPL plug-ins written in 9.12 or earlier	131

Apama 10.0 is the successor of Apama 9.12.

Apama 10.0 runs on the platforms listed in the *Supported Platforms* document for version 10.0. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>. Be sure to consult that document for details about supported versions of operating systems, compilers, BigMemory, and Universal Messaging.

Apama 10.0 includes new features, enhancements, and changes as described in the following topics.

New HTTP client transport connectivity plug-in in 10.0

A new connectivity plug-in, the HTTP client transport, is now available. It can be used to connect to external services over HTTP/REST, perform requests on them and return the response as an event. For detailed information, see "The HTTP Client Transport Connectivity Plug-in" in *Connecting Apama Applications to External Components*.

Using Software AG Designer, you can add the new HTTP client connectivity bundle to your projects. For further information, see "Adding the HTTP client connectivity plug-in to a project" in *Using Apama with Software AG Designer*.

This HTTP client transport connectivity plug-in is also delivered with version 9.12.0.5. In addition to the functionality provided with that fix, version 10.0 provides the following enhancements:

- The HTTP client can now encode query parameters from a map provided in the metadata of requests.
- The HTTP client now uses nested maps in metadata for metadata keys containing a period (.). For example, `metadata.http` is now a map containing `statusCode`, `statusReason`, `path`, and so on. Applications using the Mapper codec to set or read these will continue to work without change. Custom codecs may need changes to support this.

Apama integration with MQTT in 10.0

A new standard connectivity plug-in, the MQTT transport, is now available. It can be used to communicate between the correlator and an MQTT broker. This implementation of the MQTT transport is based on the Eclipse Paho C library which supports the MQTT v3 protocol. For detailed information, see "The MQTT Transport Connectivity Plug-in" in *Connecting Apama Applications to External Components*.

Using Software AG Designer, you can add the new MQTT connectivity bundle to your projects. For further information, see "Adding the MQTT connectivity plug-in to a project" in *Using Apama with Software AG Designer*.

Apama integration with Kafka in 10.0

A new standard connectivity plug-in, the Kafka transport, is now available. It can be used to communicate with a Kafka distributed streaming platform. In this release, it is provided as a sample in the `samples/connectivity_plugin/java/KafkaTransport` directory of your Apama

installation. See also "The Kafka Transport Connectivity Plug-in" in *Connecting Apama Applications to External Components*.

Connectivity plug-ins enhancements in 10.0

Apama 10.0 includes the following connectivity plug-ins enhancements.

Reliable messaging

- Connectivity plug-ins can now perform reliable messaging. If you have done APP_CONTROLLED reliable messaging with the Java Message Service (JMS) before, then doing the same with connectivity plug-ins will look very similar. Your old way of doing acknowledgments and flushes in EPL will be quite easy to translate across to the new way of doing this with connectivity plug-ins. For more information, see "Using reliable transports" in *Connecting Apama Applications to External Components*.

If you want to write your own transport that supports reliable messaging, see "Developing reliable transports" in *Connecting Apama Applications to External Components*.

- The Digital Event Services transport connectivity plug-in now supports reliable sending and receiving of events. For more information, see "Reliable messaging with Digital Event Services" in *Connecting Apama Applications to External Components*.
- A new metadata value, `sag.messageId`, is now available. This is used for reliable receiving. For more information, see "Metadata values" in *Connecting Apama Applications to External Components*.

User-defined status reporting

Connectivity plug-ins can now add any number of user-defined status values which are reported as part of the correlator's status information from the REST API, the `engine_watch` tool, the EngineClient API, and from the EPL Management interface. If the status keys follow the conventions listed in "Monitoring the KPIs for EPL applications and connectivity plug-ins" in *Deploying and Managing Apama Applications*, the status and KPIs of your application's connectivity plug-ins can be displayed by Command Central. For more information about how to report status, see "User-defined status reporting from connectivity plug-ins" in *Connecting Apama Applications to External Components*. See also the `samples\connectivity_plugin\cpp\httpclient` and `samples\connectivity_plugin\java\HTTPServer` directories of your Apama installation for examples of how to report status information from a connectivity plug-in.

Message metadata support for non-string values

The message metadata can now also store non-string values. With previous versions, it was allowed to only store string keys and values.

In both C++ and Java Message implementations, the `getMetadataMap()` function has been added which returns a view on the metadata which can contain non-string values (`map_t` for C++ and `Map<String, Object>` for Java). The `getMetadata()` function still exists, but is deprecated.

While the changes for Java maintain backwards compatibility for users of the `getMetadata()` function, the changes for C++ do not. For C++, the `metadata_t` iterator's `key()` and `value()` functions

have been changed to return `std::string` instead of `const char*`. Only stringifiable values (bool, integer, double, decimal) can be obtained by calling `value()`, otherwise an exception is thrown. Moreover, for C++, the metadata operator `[]` has been changed such that it returns `std::string` instead of `const char*`. To access the full set of types in C++ metadata, use the new `getMetadataMap()` function instead.

For the existing connectivity plug-ins (such as the Mapper codec and Classifier codec), the assumption has been removed that only string keys and values can be stored into the metadata.

For more information, see "Metadata values" in *Connecting Apama Applications to External Components*.

New dynamic chain managers API

Connectivity chains can be created in three different ways:

- Statically using the `startChains` section of a YAML file.
- Dynamically by an explicit EPL call to `ConnectivityPlugins.createChain`.
- Dynamically by a connectivity plug-in. For example, a plug-in could create a chain in response to an EPL application subscribing or sending to a channel with a specific name or prefix.

In the previous release, it was not possible for user-defined connectivity plug-ins to dynamically manage chain creation, although some standard plug-ins shipped with the product (such as Universal Messaging and Digital Event Services) made use of this capability. In 10.0, there is a new "chain manager" API providing the ability for user-defined plug-ins written in Java to manage chains. In this release, it is not yet possible to do this from C++ plug-ins (it's only Java right now).

A transport plug-in that manages its own chains must provide a subclass of `com.softwareag.connectivity.chainmanagers.AbstractChainManager` which can create and destroy chains whenever it wishes, by instantiating any dynamic chain definition in the configuration files that contain the transport plug-in managed. A chain manager can register a listener that is notified whenever a new channel is used for the first time by the Apama EPL application, that is, when the application subscribes or sends to a channel, which it can use to create a new chain to send/receive using that channel as needed.

For more information about configuring transport connectivity plug-ins that have a dynamic chain manager for use in your project, see "Configuration file for connectivity plug-ins" in *Connecting Apama Applications to External Components*.

For more information about how to create a dynamic chain manager transport plug-in, see "Static and dynamic connectivity chains" in *Connecting Apama Applications to External Components*.

Miscellaneous connectivity framework changes

- For Java, it is now also possible to set the classpath by defining a semicolon-delimited list of strings in the configuration file. See also "Configuration file for connectivity plug-ins" in *Connecting Apama Applications to External Components*.
- A new helper class called `com.softwareag.connectivity.util.MapExtractor` is available in the Java API. This class provides an easy and type-safe way to extract values from maps such as those used for plug-in configuration. It has automatic support for generating user-friendly

error messages if configuration items are missing or of the wrong type. See the *API Reference for Java (Javadoc)* for further information.

Enhancements and additions to standard connectivity plug-ins

- A new standard connectivity plug-in, the String codec, is now available. It can be used to translate string payloads to binary payloads, and vice versa. For detailed information, see "The String codec connectivity plug-in" in *Connecting Apama Applications to External Components*.
- A C++ implementation of the JSON codec is now available. It behaves in the same way as the already existing Java implementation of that codec. For more information, see "The JSON codec connectivity plug-in" in *Connecting Apama Applications to External Components*.
- The JSON codec now has an additional configuration parameter `filterOnContentType` which allows it to ignore messages which are not encoded with the "application/json" content type. For more information, see "The JSON codec connectivity plug-in" in *Connecting Apama Applications to External Components*.
- The Mapper codec has been enhanced as follows:
 - With the `defaultValue` action, an empty or null value is now treated in the same way as an unset value.
 - It is now possible to set metadata or payload fields to lists or maps and top-level payloads to maps, in addition to setting them to strings.
 - A new `copyFrom` action is available. This new action is exactly the same as the `mapFrom` action except that the source field is not removed after the copy.

For more information, see "The Mapper codec connectivity plug-in" in *Connecting Apama Applications to External Components*.

- The Mapper and Classifier codecs can now contain a period (.) for both payload and metadata rules. These correspond to nested maps. Mapper rules for built-in codecs and transports will continue to work without changes, but a mapper which is setting a metadata field containing a period will now produce a nested map structure. Custom codecs and transports may need changes to support this. See also "The Mapper codec connectivity plug-in" and "The Classifier codec connectivity plug-in" in *Connecting Apama Applications to External Components*.
- When mapping `nEventProperties` to metadata, the Universal Messaging transport connectivity plug-in now uses the prefix `um.properties`. You can use the Mapper codec to enable the original behavior. In addition, the ability to set `nEventProperties` while also sending a binary payload has been added. To do this, set `um.properties.Key=Value` in the metadata. See also "Supported payloads" in *Connecting Apama Applications to External Components*.
- With previous versions, only the host that created a Universal Messaging channel was able to publish to the channel created by Apama. As of this version, it is possible to specify ACL permissions for channels that are automatically created by Apama. This can be done by specifying the new `createChannelPermissions` option in the YAML configuration file for the Universal Messaging connectivity plug-in. See "Configuring the connection to Universal Messaging (dynamicChainManagers)" in *Connecting Apama Applications to External Components*. This behavior has also been changed in the deprecated native Universal Messaging integration.

New API for writing EPL plug-ins in C++ in 10.0

In 10.0, the API for writing EPL plug-ins in C++ has been replaced. Support for the old C and C++ APIs defined in `correlator_plugin.hpp` and `correlator_plugin.h` is still retained in 10.0 for backwards compatibility but are deprecated and will be removed in a future release. Use of them will produce warnings both at compile time and when injecting monitors into the correlator which use a deprecated plug-in.

EPL plug-ins should now be written using the API in `epl_plugin.hpp` instead. See "Writing EPL Plug-ins in C++" in *Developing Apama Applications* for information on the new API.

For information on migrating plug-ins from the old APIs to the new one, see ["Migrating the C and C++ EPL plug-ins written in 9.12 or earlier" on page 131](#).

The new EPL plug-in API does not include a pure-C API. Instead, the C++ API has been written with a compiler-agnostic ABI. This removes the requirement to use the platform compiler version when creating your plug-ins and therefore the need for a separate C API. All non-Java plug-ins should now be written against the new C++ API.

Apama enhancements in Software AG Designer 10.0

Apama 10.0 includes the following enhancements in Software AG Designer:

- The color in which annotations are shown in the EPL editor can now be configured in the Apama preferences. See also "Editor Colors" in *Using Apama with Software AG Designer*.
- The new field **Flush between events** has been added to the Data Player query page. It controls how many times the correlator's queues are flushed between every event from the database. This gives correct behavior in multi-context applications at the cost of speed. The default value for new queries is 2, which flushes the correlator's queues twice. Existing queries are not affected by this enhancement. A blank value is used in this case, which means that the correlator queues are not flushed. If you want to change this, you have to set the required value on the Data Player query page. See also "Specifying Data Player playback queries" in *Using Apama with Software AG Designer*.
- If a query pattern uses the `or` operator (see also ["Query enhancements in 10.0" on page 125](#)), then the New Query Send Event Action dialog includes the new **Required events** section. It sets which events must be matched by the pattern in order for the send event action to send the event. See also "Adding query send event actions" in *Using Apama with Software AG Designer*.
- The Open Type dialog functionality has been added for all Apama perspectives. See "Opening EPL elements and Java classes using Open Type dialog" in *Using Apama with Software AG Designer*.
- The Correlator Deployment Package wizard now accepts a launch configuration instead of a project. The ability to select the files to be included has been removed from the wizard. Instead, if you want to create a correlator deployment package (CDP) containing a subset of project's files, you must now create a dedicated launch configuration to represent the required files before opening the Correlator Deployment Package wizard. See "Exporting correlator deployment packages" in *Using Apama with Software AG Designer*.

- When you add a connectivity bundle, the `standard-codecs.yaml` file is now created in the **Linked Resources** node. This file contains standard codecs defined by Apama. See also "Connectivity bundles" in *Using Apama with Software AG Designer*.

Query enhancements in 10.0

Apama 10.0 includes the following query enhancements:

- A query may now use a decimal type for a parameter. See also "Format of query definitions" in *Developing Apama Applications*.
- In a `find` statement, you can now specify the `or` operator in the event pattern. The events on one side or the other of the `or` operator must be matched for the query to fire. For detailed information, see "Query or operator" in *Developing Apama Applications*.
- Queries can now use the output of another query as input, thus allowing chaining of queries. For detailed information, see "Using the output of another query as query input" in *Developing Apama Applications*.
- To delete a query, you need to run `engine_delete -F` followed by a fully qualified query name. In previous releases, it was possible to delete a query by just using `engine_delete` if nothing else depended on it. Note that `engine_delete -F` deletes *everything* that depends on the query or the query output event. Therefore, caution is advised when deleting a query. See also "Deleting code from a correlator" in *Deploying and Managing Apama Applications*.

EPL enhancements in 10.0

Apama 10.0 includes the following EPL enhancements:

- The `throw` statement is now available. It causes an exception to be thrown. See also "The throw statement" in *Developing Apama Applications in EPL*.
- The behavior for disconnected out of band events has changed. In the `com.apama.oob` package, `componentName` and `address` have been added to `ReceiverDisconnected` and `SenderDisconnected` so that the behavior is now in line with `ReceiverConnected` and `SenderConnected`. A new dictionary<string, string> field called `extraParams` has been added to all the out of band events. Additionally in the logging, `componentName` has been changed to be `"name"`, unless no name is provided in which case the old format `"name (on port port_number)"` is used. See also "Out of band connection notifications" in *Developing Apama Applications in EPL*.
- It is now possible to invoke EPL plug-ins written in Java from persistent monitors provided that they do not use chunks, which means that the behavior is now the same as for EPL plug-ins written in C++.

Correlator utility enhancements in 10.0

Apama 10.0 includes the following correlator utility enhancements:

- Instead of specifying a list of configuration files on the command line with the `--config` option, it is now also possible to specify these files in the `includes` section of another configuration

file. For detailed information, see "Including YAML configuration files inside another YAML configuration file" in *Deploying and Managing Apama Applications*.

- It is now possible to specify the correlator port number in a YAML configuration file. For detailed information, see "Specifying the correlator port number" in *Deploying and Managing Apama Applications*.
- Using the new `engine_deploy` tool, it is now possible to generate an initialization file list, a correlator deployment package (CDP), or a correlator deployment directory from an Apama project that has been created with Software AG Designer. You can also use this tool to perform the initialization in a running correlator. If you are not using Software AG Designer, you can also use this tool with a directory of Apama files. For detailed information, see "Deploying a correlator" in *Deploying and Managing Apama Applications*.

In addition, you can now also use Ant to generate a correlator deployment directory. The `apama-macros.xml` file includes the new `generate-correlator-deploy-dir` macro for this purpose. See also "About deploying Apama applications with an Ant script" in *Deploying and Managing Apama Applications*.

- A new management request has been added: `startInternalClock`. This is available on correlators via the `--doRequest` option of the `engine_management` tool (see "Shutting down and managing components" in *Deploying and Managing Apama Applications*) and via the `apama-macros.xml` Ant script (see also "About deploying Apama applications with an Ant script" in *Deploying and Managing Apama Applications*). This request takes a previously externally clocked correlator and starts to clock it internally.

Dashboard enhancements in 10.0

Apama 10.0 includes the following dashboard enhancements:

- The REST API now also supports the `PUT /logLevel` request for dashboard data servers and display servers. See also "Managing and Monitoring over REST" in *Deploying and Managing Apama Applications*.

Miscellaneous enhancements and changes in 10.0

Apama 10.0 includes the following miscellaneous enhancements and changes:

- When using the IAF-based CSV codec plug-in (JCSVCodec), you can now enable Excel compatibility mode by setting the new `excelCompatible` property to `true` in the IAF configuration file. Double quotes are then used to match the behavior of Excel. By default, this new property is set to `false`. See also "The CSV codec IAF plug-in" in *Connecting Apama Applications to External Components*.

In addition, you can also enable Excel compatibility mode when sending a configuration event from the transport that is communicating with the CSV codec using the method described in "Multiple configurations and the CSV codec" in *Connecting Apama Applications to External Components*.

- The REST API now also supports the `/info/stats` URI for the correlator and the IAF. It provides access to the same information as the `/correlator/status` and `/iaf/status` URIs. Previously,

the `/info/stats` URI was only available for dashboard data servers and display servers. See also "Managing and Monitoring over REST" in *Deploying and Managing Apama Applications*.

- The new method `GenericComponentInfo.getDataValue()` has been added to the Apama API for Java. Using this method, you can easily look up a value from a `GenericComponentManagement` information category that has unique keys (that is, dictionary content) without having to iterate through every item. See the *API Reference for Java (Javadoc)* for further information.
- For applications that report status to Command Central, the recommended values for the `".status"` key are now `STARTING`, `ONLINE` or `FAILED` (instead of `ONLINE` and `OFFLINE`, although `OFFLINE` is still permitted for legacy applications). See also "Monitoring KPIs for EPL applications and connectivity plug-ins" in *Deploying and Managing Apama Applications*.
- A title can now be provided when generating ApamaDoc from an Ant script or from Software AG Designer. See also "Generating ApamaDoc from an Ant script" in *Developing Apama Applications* and "Exporting ApamaDoc" in *Using Apama with Software AG Designer*.
- The ATM Fraud queries sample has been moved to the demos folder of the Apama installation, and it is now available as a demo from the Welcome page. With previous versions, it was located in the `samples\queries\ATM_Fraud_Sample` folder.
- The demos and samples have been updated to use EPL instead of scenarios. Keep in mind that the Event Modeler is deprecated and that support for creating and deploying scenarios will be removed in a future release (see also "[Removed and deprecated features in 9.12](#)" on page 145).

Platform changes in 10.0

Apama 10.0 runs on the platforms listed in the *Supported Platforms* document for version 10.0. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>.

Check the above mentioned *Supported Platforms* document for details about the newer platform versions supported by Apama 10.0. Due to these upgrades, the following versions are no longer supported:

- Red Hat Enterprise Linux 6 update 1 (x86) 64-bit (replaced by support for a more recent version).
- CentOS Linux 6 update 1 (x86) 64-bit (replaced by support for a more recent version).
- Oracle Linux 6.5 (x86) 64-bit (replaced by support for a more recent version).
- Software AG Common Runtime 9.12 (replaced by support for a more recent version).
- Apache Tomcat 7.0.62 (replaced by support for a more recent version).
- Software AG Universal Messaging 9.12 (replaced by support for a more recent version).
- Software AG Universal Messaging Client 9.12 (replaced by support for a more recent version).
- Oracle Java 8 update 101 (replaced by support for a more recent version).
- Software AG Designer 9.12 (replaced by support for a more recent version).
- Google Chrome 44 (replaced by support for a more recent version).

- Firefox Mozilla 45 (replaced by support for a more recent version).
- Apple Safari (desktop) 9 (replaced by support for a more recent version).
- Apple Safari (mobile) iPad iOS 4 (replaced by support for a more recent version).
- VMWare ESXi 5.0 (replaced by support for a more recent version).
- Software AG Terracotta BigMemory Max 4.3.3 (replaced by support for a more recent version).
- MathWorks MATLAB 2014b (replaced by support for a more recent version).

Windows compiler update in 10.0

Apama has upgraded from Microsoft Visual Studio 2013 (Visual C++ v12) to Visual Studio 2015 Update 3 (Visual C++ v14). Consequently, you must rebuild the following components with the new compiler:

- C++ connectivity plug-ins
- C/C++ EPL plug-ins
- C/C++ IAF transport/codec plug-ins
- C/C++ code that uses the engine client API

In addition, when planning migration, consider that the C++ compiler in Visual Studio 2015 is more modern than the previous version and less tolerant of ambiguous and potentially unsafe code patterns. It is likely that some time may be required to fix compiler errors after moving to the new version.

For example, one of our samples previously included this macro definition:

```
#define snprintf _snprintf
```

This produces a compilation error with the new compiler. So if your code includes this line, you should simply remove it.

Note:

Our recommendation is to enable “warnings as errors” in your projects where possible. This encourages good coding patterns and catches errors earlier in the development process. In particular, it helps you catch errors that might cause subtle, difficult to diagnose errors.

Miscellaneous changes in 10.0 affecting backwards compatibility

The following changes in Apama 10.0 affect backwards compatibility to previous Apama versions:

- Apama 10.0 incorporates the ICU (International Components for Unicode) time-zone data update 2016j, which is the most recent update at the time of release. This will update time-zone data used by the correlator and Time Format EPL plug-in.

In addition, Apama 10.0 contains an upgrade from ICU 3.8.1 to ICU 58.1 which contains some small behavioral differences which may affect your applications:

- When formatting time in UTC, the output for a “GMT+offset” time zone is now “GMT” rather than “GMT+00:00” (non-UTC time zones are unaffected).
- The format/parse symbol `z` no longer uses time-zone short codes such as “EST” due to ambiguity. It now uses GMT+offset formats. The behavior of the `v`, `vv` and `vvv` format/parse symbols has also changed. If you wish to specify time zones by name, you should always use the location-based names, such as “America/New_York” using the format/parse symbol `vv`. For more information, see "Format specification for the TimeFormat functions" in *Developing Apama Applications*.
- When parsing time strings, trailing characters in the string after matching the whole pattern are ignored.
- When parsing time strings, separator characters (such as a hyphen (-) or others which do not have a specific meaning) will match any such string of those characters of any length (at least one). This means that the pattern “yyyy-MM-dd” will now match “2016//11/15”.
- Some locale settings have fixes that may change the returned value for the “week of year” (that is, `w` and `W`).
- The `chainInstanceId` argument of the `createChain` method must now specify a unique identifier. Thus, an error is now thrown if a chain with a duplicate ID is to be created. See also "Creating dynamic chains from EPL" in *Connecting Apama Applications to External Components*.
- There is a change in functionality in the case of YAML configuration files. With previous versions, it was possible to specify, for example, `${PARENT_DIR}foo` on Windows. This was a valid specification as `PARENT_DIR` was expanded with a trailing slash. As of Apama 10.0, this is no longer valid. It is now required that you also specify the slash:

```
${PARENT_DIR}/foo
```

This behavior is now in line with `PARENT_DIR` on Linux, and with `APAMA_HOME` and `APAMA_WORK` on Windows and Linux.

- The message metadata in connectivity plug-ins can now also store non-string values. For connectivity plug-ins in C++, these changes are not backwards compatible. See [“Connectivity plug-ins enhancements in 10.0” on page 121](#) for more information.
- The `nullTerminated` configuration option of the String codec now defaults to `false`. This means that the String codec no longer adds a null-terminator in messages towards the transport when converting to a `byte[]` (Java) or `buffer_t` (C++) type. This change has been made to be more compatible with the common use case. To enable the old behavior (such as when using the Universal Messaging transport connectivity plug-in to connect to a legacy Apama application), set `nullTerminated` to `true` in the String codec configuration. See also "The String codec connectivity plug-in" in *Connecting Apama Applications to External Components*.
- When mapping `nEventProperties` to metadata, the Universal Messaging transport connectivity plug-in now uses the prefix `um.properties`. You can use the Mapper codec to enable the original behavior. See also [“Connectivity plug-ins enhancements in 10.0” on page 121](#).

- The Mapper and Classifier codecs can now contain a period (.) for both payload and metadata rules. Custom codecs and transports may need changes to support this. See also [“Connectivity plug-ins enhancements in 10.0” on page 121](#).
- The `ConnectivityPlugins.mon` file now also requires the new `ConnectivityPluginsControl.mon` file from the same directory. This is important if you inject files manually. If you are using Software AG Designer to develop applications that use any connectivity plug-in bundles, this does not matter to you as Software AG Designer automatically upgrades the set of `.mon` files that you are using.
- As of 10.0, `ifpresent` and any are reserved words in EPL. Projects that are using either of these as an identifier will fail in 10.0.

Removed and deprecated features in 10.0

The following Apama features are now deprecated or have been removed in Apama 10.0:

- The `subscribe()`, `unsubscribe()` and `sendToDES()` actions on the auto-generated EPL representations of Digital Event Services types are now deprecated. The channel name is now accessed via the static `CHANNEL` constant on the EPL type. For more information, see "Using Digital Event Services connectivity from EPL" in *Connecting Apama Applications to External Components*.
- For the Java and C++ connectivity plug-ins, the `LOGGER` field and use of the legacy constructor signatures that do not use `TransportConstructorParameters` or `CodecConstructorParameters` are now deprecated.

For C++, the following macros are now deprecated:

- `SAG_DECLARE_CONNECTIVITY_CODEC`
- `SAG_DECLARE_CONNECTIVITY_TRANSPORT`

You should now use the new signature for the constructor, the new `logger` field, and the new macros as announced in [“Connectivity API changes in 9.12.0.1” on page 137](#). For more information, see "Requirements of a plug-in class" in *Connecting Apama Applications to External Components*.

- The `UMStringCodec` is now deprecated. Use the new generic `String` codec instead. For detailed information, see "The String codec connectivity plug-in" in *Connecting Apama Applications to External Components*.
- For connectivity plug-ins, the `getMetadata()` function is now deprecated. Use the new `getMetadataMap()` function instead. See [“Connectivity plug-ins enhancements in 10.0” on page 121](#) for more information.
- Support for the old APIs for writing EPL plug-ins in C and C++ is now deprecated and will be removed in a future release. EPL plug-ins should now be written using the new C++ API. See also [“New API for writing EPL plug-ins in C++ in 10.0” on page 124](#).
- Dashboard support for applet deployment is no longer supported. The available deployment options are therefore Web Start, Local, or display server.

- Compiling with Visual Studio 2013 is no longer supported. See [“Windows compiler update in 10.0” on page 128](#) for more information.
- The following information pertains to the Log File Manager plug-in (LoggingManager) which has been removed in 9.12, except for the associated bundle which was only removed in 10.0. If you still have the deprecated Logging Manager bundle in an Apama project, this results in an error indicating that the LoggingManager.bnd bundle cannot be found. To remove the Logging Manager bundle from an Apama project in Software AG Designer, proceed as follows:
 1. Select the project, invoke the context menu and choose **Properties**.
 2. In the properties dialog, select **Apama > MonitorScript Build Path**.
 3. Go to the **Bundles** tab.
 4. Select the **Cannot read bundle file** message and click **Remove**.
 5. Click **OK**.
 6. Click **Yes** to confirm the deletion of the bundle instance files.

Migrating the C and C++ EPL plug-ins written in 9.12 or earlier

To migrate an existing EPL plug-in that has been written in C or C++ and has been created with version 9.12 or earlier, you will need to follow the steps below. The two APIs are quite different, but you will find the implementation overall simpler and easier to manage.

1. Include `epl_plugin.hpp` instead of `correlator_plugin.hpp`.

```
#include <epl_plugin.hpp>
```

2. Import the `com::apama::epl` namespace.

```
using namespace com::apama::epl;
```

3. Define a class which inherits from `EPLPlugin`. This will hold all the methods which are exported to EPL. Previously, these could have been free functions or static member functions. Now they are non-static member functions.

```
class MyPlugin: public EPLPlugin<MyPlugin>
```

4. Add a zero-argument constructor.

```
MyPlugin(): base_plugin_t("MyPlugin") {}
```

5. Change all your exported methods to be members of your class and change the `void(const AP_Context &, const AP_TypeList &, AP_Type &, AP_TypeDiscriminator)` signature to be the actual C++ types that your method will take and return. You will no longer need to unpack the typelist argument or explicitly set the return value, you can just read the arguments and return directly.

```
void myMethod(int64_t arg1)
{
    data.push_back(arg1);
}
```

```
bool anotherMethod(double arg1, double arg2)
{
    return arg1 < arg2;
}
```

6. Change use of functions on the `AP_Context` argument to functions to use the `getCorrelator()` static method instead.

```
getCorrelator().sendEvent("MyEvent(42)");
```

7. Replace the `AP_Function` array, type signatures and version declaration from the `AP_INIT_FUNCTION_NAME` free function with an `initialize()` static function on your class.

```
static void initialize(base_plugin_t::method_data_t &md)
{
    md.registerMethod<decltype(&MyPlugin::myMethod),
        &MyPlugin::myMethod>("myMethod");
    md.registerMethod<decltype(&MyPlugin::anotherMethod),
        &MyPlugin::anotherMethod>("anotherMethod");
}
```

8. Other code from `AP_INIT_FUNCTION_NAME` goes in your constructor.
9. Code from `AP_SHUTDOWN_FUNCTION_NAME` goes in your destructor.

```
~MyPlugin()
{
    ...
}
```

10. Remove `AP_LIBRARY_VERSION_FUNCTION_NAME`.
11. Add a call to `APAMA_DECLARE_EPL_PLUGIN` for your class.

```
APAMA_DECLARE_EPL_PLUGIN(MyPlugin)
```

12. If you use chunks, you no longer need to inherit from a particular class, just declare methods which return `custom_t<T>`, where `T` is the type you want to store in EPL.

```
struct MyChunkType
{
    int64_t datum;
};
custom_t<MyChunkType> createChunk(int64_t datum)
{
    return custom_t<MyChunkType>(new MyChunkType{datum});
}
int64_t getChunkValue(const custom_t<MyChunkType> &chunk)
{
    return chunk->datum;
}
```

16 What's New In Apama 9.12

■ Apama integration with Software AG Digital Event Services in 9.12	134
■ Apama integration with Universal Messaging in 9.12	134
■ New optional type in 9.12	135
■ New Microsoft Edge support in 9.12	135
■ Connectivity plug-ins enhancements in 9.12	135
■ Connectivity API changes in 9.12.0.1	137
■ Apama enhancements in Software AG Designer 9.12	139
■ Query enhancements in 9.12	140
■ EPL enhancements in 9.12	140
■ Correlator utility enhancements in 9.12	141
■ Dashboard enhancements in 9.12	142
■ PySys enhancements in 9.12	142
■ Command Central support in 9.12	143
■ Miscellaneous enhancements and changes in 9.12	143
■ Platform changes in 9.12	144
■ Miscellaneous changes in 9.12 affecting backwards compatibility	144
■ Removed and deprecated features in 9.12	145

Apama 9.12 is the successor of Apama 9.10. There is no version 9.11.

Apama 9.12 runs on the platforms listed in the *Supported Platforms* document for version 9.12. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>. Be sure to consult that document for details about supported versions of operating systems, compilers, BigMemory, and Universal Messaging.

Apama 9.12 includes new features, enhancements, and changes as described in the following topics.

Apama integration with Software AG Digital Event Services in 9.12

Apama now supports Software AG Digital Event Services, which is a messaging system for communicating between different Software AG products using events. Digital Event Services itself has event definitions of its own, called digital event types, that all participating products can have in common. For more information, see "The Digital Event Services Transport Connectivity Plug-in" in *Connecting Apama Applications to External Components*.

Using Software AG Designer, you can now add a new connectivity bundle named **Digital Event Services** to your projects. You use it to automatically generate the Apama event types from the Software AG Digital Event Services type definitions.

The new connectivity bundle includes the Digital Event Services connectivity plug-in. You use this plug-in to send and receive the generated Apama events. All of the configuration for this plug-in is driven by Software AG Designer, and you can edit the configuration by just changing a few properties.

For further information, see "Using the Digital Event Services connectivity bundle" in *Using Apama with Software AG Designer*.

Apama integration with Universal Messaging in 9.12

The Apama support for Universal Messaging has changed. A new connectivity plug-in is now available for connecting to Universal Messaging realms and channels. This provides for new features in mapping of message types on Universal Messaging.

Using Software AG Designer, you can add a new connectivity bundle named **Universal Messaging** to your projects. See "Adding the Universal Messaging connectivity plug-in to a project" in *Using Apama with Software AG Designer*. This bundle includes the Universal Messaging connectivity plug-in. All of the configuration for this plug-in is driven by Software AG Designer, and you can edit the configuration by just changing a few properties. For more information, see "The Universal Messaging Transport Connectivity Plug-in" and especially the topic "Configuring the Universal Messaging connectivity plug-in" in *Connecting Apama Applications to External Components*.

The new Universal Messaging connectivity plug-in replaces the old native Universal Messaging integration, which will be removed in a future release. It is recommended that you use the new connectivity plug-in instead.

New optional type in 9.12

EPL now supports optional types. An `optional` is a value that may be a value of some other specified EPL type, or empty and thus have no value. This is useful for mapping to `null` values in other languages such as Java, or for results or fields which may be missing a value.

`optional` values can be extracted using the new `ifpresent` statement. This statement is used to check if one or more values are empty. It unpacks the values into new local variables and conditionally executes a block of code. See "Defining conditional logic with the `ifpresent` statement" and "The `ifpresent` statement" in *Developing Apama Applications* for more information.

Note:

The `ifpresent` keyword can no longer be used as a valid identifier.

The connectivity plug-ins can now handle values of `optional` type. This means that an empty `data_t` or a `null` value in Java will translate to/from an empty `optional`. See also "Map contents used by the `apama.eventMap` host plug-in" in *Connecting Apama Applications to External Components*.

The Apama event parser libraries for Java and .NET have been improved to support events with `optional` fields. A new class `OptionalFieldType` (accessible via `FieldTypes.Optional`) has been added to represent the `optional` field type. For complete information about this class, refer to the description of

- `com.apama.event.parser` in the *API Reference for Java (Javadoc)*
- `Apama.Event.Parser` Namespace in the *API Reference for .NET*

Note:

`optional` event field support will not be available in the C parser in `AP_EventParser` and `AP_EventWriter`.

New Microsoft Edge support in 9.12

Apama now supports the Microsoft Edge web browser.

Microsoft Edge does not support Java plug-ins and therefore does not support deployment of Apama dashboards as Java applets. If you need to deploy Apama dashboards as Java applets, use Microsoft Internet Explorer or Mozilla Firefox instead.

The supported versions of Microsoft Edge, Microsoft Internet Explorer and Mozilla Firefox are listed in the *Supported Platforms* document for version 9.12. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>.

Connectivity plug-ins enhancements in 9.12

Apama 9.12 includes the following connectivity plug-ins enhancements:

- The `plugins` stanza that is used in the configuration file for the connectivity plug-ins has been renamed to `connectivityPlugins`. It is recommended that you start using the new name in

your configuration files. See also "Overview of using connectivity plug-ins" in *Connecting Apama Applications to External Components*.

- The `--connectivityConfig` option which is available for the `correlator` executable has been renamed to `--config`. It is recommended that you start using the new name. In addition, it is now possible to specify multiple `.yaml` files containing connectivity plug-in configuration, which will be merged together. It is also possible now to specify a directory containing `.yaml` and `.properties` files. It is no longer possible to specify a YAML file that does not have the `.yaml` file extension; this will now result in an error. See also "Starting the correlator" in *Deploying and Managing Apama Applications*.
- A new `-D` option is available for the `correlator` executable, which specifies a value for a substitution property to be used by configuration files specified with the `--config` option of the `correlator` executable. See "Starting the correlator" in *Deploying and Managing Apama Applications*. Using this option is similar to specifying the substitution using a `.properties` file.
- Properties file substitution is now supported, allowing any user-supplied `${varname}` substitution variables to be replaced inside the configuration files for the connectivity plug-ins (`.yaml` files). See "Using properties files" in *Deploying and Managing Apama Applications*.
- Chains can now be defined under both the `startChains` and `dynamicChains` stanzas in the configuration file. Chains defined under `dynamicChains` are not instantiated at startup, but instead can be created dynamically from EPL. See also "Configuration file for connectivity plug-ins" and "Creating dynamic chains from EPL" in *Connecting Apama Applications to External Components*.

Dynamic chains are also used by the Digital Event Services transport and the Universal Messaging transport.

- A new predefined Unit Test Harness codec is provided. The Unit Test Harness codec and the new Null Transport transport make it easy to send test messages into a connectivity chain and/or to write out received messages to a text file, without the need to write any EPL in the correlator, which is very useful for writing unit tests for connectivity plug-ins, and especially codecs. For more information, see "The Unit Test Harness codec connectivity plug-in" in *Connecting Apama Applications to External Components*.
- A new predefined Diagnostic codec is provided, which can be used to diagnose issues with connectivity plug-ins. It logs the events that go through a connectivity chain in either direction. For more information, see "The Diagnostic codec connectivity plug-in" in *Connecting Apama Applications to External Components*.
- The behavior of the Mapper codec has changed as follows: A `mapFrom` rule where the source field does not exist now uses the default value if the `defaultValue` exists or if a subsequent `mapFrom` rule exists for the same destination field. An error is no longer raised in this case.

For example, let us assume that there is no field `b` in the metadata and the configuration looks as follows:

```
mapperCodec:
  SomeType:
    towardsHost:
      mapFrom:
        - payload.a: metadata.b
```



```
defaultValue:
- payload.a: "someDefaultValue"
```

In the previous version, an error was raised with the message that `b` does not exist. As of 9.12, the default value for the payload field `a` (that is, `someDefaultValue`) is used rather than raising an error. However, if the destination field (which is `payload.a` in the above example) is also not defined with `defaultValue`, then an error is raised.

In addition, a new boolean `allowMissing` is available. By default, this is set to `false`. When set to `true`, an error is not raised when a `defaultValue` has not been set and a source field is missing. `allowMissing` needs to be defined at the same level as the event types. For more information, see "The Mapper codec connectivity plug-in" in *Connecting Apama Applications to External Components*.

- The behavior of the `allowMissing` configuration property in the `apama.eventMap` host plug-in has changed. Previously only missing fields were set to their default values. As of this version, fields with empty values (`null` in Java) are also set to their default values. Similarly, empty values in nested events, elements in sequences and key/value pairs in dictionaries are also set to their default values. See also "Translating EPL events using the `apama.eventMap` host plug-in" in *Connecting Apama Applications to External Components*.
- New configuration properties `suppressLoopback`, `description` and `remoteAddress` are now available for the `apama.eventMap` and `apama.eventString` host plug-ins. For more information, see "Host plug-ins and configuration" in *Connecting Apama Applications to External Components*.
- You can now use the `jvmOptions` key in the configuration file to specify JVM options which the correlator is to pass to the embedded JVM. See also "Specifying JVM options" in *Deploying and Managing Apama Applications*.
- If you are building C++ plug-ins, you now need to link against the `apclient` library (instead of `apconnectivity`). See also "Building C++ plug-ins" in *Connecting Apama Applications to External Components*.

Connectivity API changes in 9.12.0.1

Version 9.12.0.1 introduces a new signature for the constructor of a C++ or Java connectivity plug-in, which passes in a `TransportConstructorParameters` or `CodecConstructorParameters` object instead of a list of individual parameters.

For Java, there is a new `logger` field that uses the Simple Logging Facade for Java (SLF4J) instead of the `Log4J` `LOGGER` available in earlier versions.

For C++, there is a new `logger` field to match the Java field, but the behavior is identical to `LOGGER` although we recommend to use the new name where possible for consistency.

In C++ only, there are the following new macros which must be used for classes with the new 9.12.0.1 constructor:

- `SAG_DECLARE_CONNECTIVITY_CODEC_CLASS` (instead of `SAG_DECLARE_CONNECTIVITY_CODEC`)
- `SAG_DECLARE_CONNECTIVITY_TRANSPORT_CLASS` (instead of `SAG_DECLARE_CONNECTIVITY_TRANSPORT`)

The previous legacy constructors, macros and logger are still available but will be deprecated in a future release. Therefore, we recommend migrating to the new ones when possible, especially for any newly created plug-ins.

To migrate from the old C++ transport constructor and macro to the new one, change the following:

```
class MyPlugin: public AbstractSimpleTransport
{
public:
    MyPlugin(const std::string &name, const std::string &chainId,
              const map_t &config)
        : AbstractSimpleTransport(name, chainId, config)
    {
        // do stuff with config here
        // use "LOGGER" for logging
    }
    ...
}
SAG_DECLARE_CONNECTIVITY_TRANSPORT(MyPlugin)
```

to:

```
class MyPlugin: public AbstractSimpleTransport
{
public:
    MyPlugin(const TransportConstructorParameters &params)
        : AbstractSimpleTransport(params)
    {
        // do stuff with params.getConfig() (or the config class member) here
        // recommended to use "logger" for logging instead of "LOGGER"
    }
    ...
}
// use new 9.12.0.1 macro for declaring the class
SAG_DECLARE_CONNECTIVITY_TRANSPORT_CLASS(MyPlugin)
```

The same applies to the C++ codecs, just replace “transport” with “codec”.

To migrate from the old Java transport constructor and logger to the new one, change the following:

```
public MyTransport(Map<String, Object> config, String chainId,
                  org.apache.log4j.Logger log4jLogger) throws Exception {
    super(config, chainId, log4jLogger);
    // do stuff with params.getConfig() (or the config class member) here
    LOGGER.info("Transport was created correctly with config: "+config);
    ...
}
```

to:

```
public MyTransport(org.slf4j.Logger logger,
                  TransportConstructorParameters params) throws Exception {
    super(logger, params);
    logger.info("Transport was created correctly with config: "
               +params.getConfig());
    ...
}
```

The same applies to the Java codecs, just replace “transport” with “codec”.

Be sure to change all uses of the legacy “LOGGER” to the new “logger” in your Java plug-ins. Most of the logger methods are identical (except that the new logger does not permit logging an exception without an accompanying string message).

See "Requirements of a plug-in class" in *Connecting Apama Applications to External Components* for more information about plug-in constructors.

Apama enhancements in Software AG Designer 9.12

Apama 9.12 includes the following enhancements in Software AG Designer:

- When you create an adapter, the instance name that Software AG Designer offers by default now
 - includes the adapter type,
 - no longer includes a space (thus, it is no longer required to escape the name in various places), and
 - the number that is provided with the name is incremented depending on the adapter type.

For example, when you create a File adapter, the default instance name is now “FileAdapter1”, and when you create a Web Services Client adapter next, the default instance name is now “WebServicesClientAdapter1”. The instance names are no longer “instance 1” for the first adapter and “instance 2” for the second adapter.

It is recommended that you specify your own instance names instead of just accepting the default names. An instance name should clearly identify the purpose of the adapter. For example, if you use the correlator-integrated adapter for JMS to connect to a customer tracking system, you should use a name such as “CustomerTracking” (and not just “JMS”).

When you specify your own instance name, it is no longer possible to specify spaces within the name.

- Software AG Designer now supports adding multiple User Connectivity configuration instances (each represented by a directory of configuration files) to your project. The instances can contain multiple .yaml and .properties files. You can selectively enable or disable individual User Connectivity instances from the correlator launch configuration dialog. See "Creating Apama projects" in *Using Apama with Software AG Designer* for more information about adding User Connectivity support to a project.
- The **Connectivity Plug-ins Support** option has been removed from under **Startup options** in the Correlator Configurations dialog, and the **Connectivity** tab has been added.
- In the dialog for configuring correlators in a launch configuration, the **Xconfig** text box has been removed and replaced with a text box called **Configuration**. The **Configuration** text box accepts the path of a configuration file in either YAML (preferred) or the deprecated TXT configuration file format. This configuration can be used for a variety of purposes such as configuring per-package EPL logging settings. See "Configuring the correlator" in *Deploying and Managing Apama Applications*.

Query enhancements in 9.12

Apama 9.12 includes the following query enhancements:

- A query may now also use the result of an action call on the event as the key for a query. For further information, see "Defining actions as query keys" in *Developing Apama Applications*.
- Coassignment in query files can now use the `as` operator. For example:

```
using com.apama.aggregates.last;
query Example {
    find Event as e select last(e.value) as lastValue without Stop as s {
    }
}
```

Software AG Designer has been updated to use the `as` operator when constructing queries graphically. The current `:` assignment operator is still valid for backwards compatibility. For further information, see "Defining event patterns" in *Developing Apama Applications*.

EPL enhancements in 9.12

Apama 9.12 includes the following EPL enhancements:

- The following new `string` methods are available to search and replace strings using regular expressions:
 - `matches(string regex)`
 - `search(string regex)`
 - `replace(string regex, string replacement)`

Apama currently uses ICU 3.8.1 to implement the regular expressions.

- Event listeners (on statements) and stream listeners (from statements) can now coassign matching events to an implicitly declared variable using the `as` operator. This variable is only in scope in the processing block of the listener and has no effect on clashing local or global variables. For example:

```
// Variable newTick is not yet defined
on StockTick(*,*) as newTick {
    // The variable newTick is only in scope in this block
}
```

The current `:` assignment operator is still available for backwards compatibility. See "Specifying the on statement" and "Using output from streams" in *Developing Apama Applications* for further information.

- Data views now support the `decimal` field type. Decimals can thus be displayed in the Scenario Browser and used in dashboards. An event with a `decimal` field type can now be sent from a dashboard.
- Decimals are now also supported in C, C++ and Java plug-ins, in the `MemoryStore` and in the distributed `MemoryStore`.

In C and C++, the decimal value is passed as an `AP_decimal` which is really a `uint64` in the IEEE 754-2008 decimal format. We suggest using the Intel Decimal Floating Point library v2.0U1 (see <https://software.intel.com/en-us/articles/intel-decimal-floating-point-math-library>) if the actual decimal value needs to be used.

In Java, both `BigDecimal` (for backwards compatibility) and `Number` are now supported types for decimal in a plug-in API. For `Number`, it is either a `BigDecimal` type for a value, or a `Double` type for NaN/infinity. NaN/infinity are still not supported when using the `BigDecimal` type, and will throw an exception that kills the monitor instance if not caught.

- It is now possible to generate ApamaDoc from an Ant script on UNIX. For further information, see "Generating ApamaDoc from an Ant script" in *Developing Apama Applications*.
- A new `OutOfBandConnections` event has been added to the `com.apama.oob` package, which provides a mechanism to synchronously get the currently connected senders and receivers. For further information, see "Out of band connection notifications" in *Developing Apama Applications*.
- It is now possible to declare static actions inside an event type. For further information, see "Defining static actions" in *Developing Apama Applications*.
- The `if` statement no longer requires the `then` keyword. It can still be specified optionally (for backwards compatibility), but you can now write a shorter form. For example:

```
if seq.size() > 0 { print "First element is " + seq[0]; }
```

See also "Defining conditional logic with the `if` statement" and "The `if` statement" in *Developing Apama Applications*.

Correlator utility enhancements in 9.12

Apama 9.12 includes the following correlator utility enhancements:

- It is now possible to configure persistence using a YAML configuration file that is specified with the new `--config` option when starting the correlator. For detailed information, see "Configuring persistence in a YAML configuration file" in *Deploying and Managing Apama Applications*.
- It is now possible to deploy Apama applications using YAML configuration files. This is useful for Docker containers or other minimal environments where only part of an Apama installation is available or it is not practical to run Java tools to perform the injections. For detailed information, see "Deploying Apama applications with a YAML configuration file" in *Deploying and Managing Apama Applications*.
- A new option `--pidfile file` is available for the `correlator`, `iaf`, `dashboard_server` and `display_server` tools. It specifies the name of the file that contains the process ID. For more information, see:
 - "Starting the correlator" in *Deploying and Managing Apama Applications*
 - "IAF command-line options" in *Connecting Apama Applications to External Components*

- "Command-line options for the data server and display server" in *Building and Using Apama Dashboards*
- A new request type `toStringQueues`, which can be issued with the `-r` (or `--dorequest`) option, is available for the `engine_management` tool. It outputs the current contents of all input and output queues within the running correlator. This can be helpful for identifying slow senders and receivers, and potential causes (such as very large events or excessive flow). See also "Shutting down and managing components" in *Deploying and Managing Apama Applications*.
- A new request type `getEventTypes`, which can be issued with the `-r` (or `--dorequest`) option, is available for the `iaf_management` tool. It returns a string representation of the event types known to the running IAF. The output is equivalent to that of the `iaf -e config.xml` command. The `getEventTypes` request type, however, does not require a configuration file as input. See also "Shutting down and managing components" in *Deploying and Managing Apama Applications*.

Dashboard enhancements in 9.12

Apama 9.12 includes the following dashboard enhancements:

- A new dashboard data table "Query instance table" is now available in the Dashboard Builder. Similar to scenario instance table, the new query instance table contains values of all input parameters for instances of an Apama query. You can define commands in the Dashboard Builder to create, edit and delete an instance of a query using the following commands:
 - Create query instance
 - Edit query instance
 - Delete query instance
 - Delete all instances of a query
- A new Tab Control object is now available on the **Controls** tab of the Dashboard Builder's Object Palette. This object provides an alternative to using a Tab Panel ini file.
- The REST API now supports the following new URI GET request, both for XML and JSON: `/info/stats`. This displays the statistics of instance/trend table queries to dashboard data servers and display servers. In addition, `/connections` now also shows the connected correlators as well as any configured named dashboard data servers and display servers. See also "Managing and Monitoring over REST" in *Deploying and Managing Apama Applications*.

PySys enhancements in 9.12

Apama 9.12 includes the following PySys enhancements:

- Apama 9.12 ships a new release of PySys, version 1.2.0. This new release includes some usability enhancements such as more informative messages when an assertion fails. For full details, see the PySys release notes.
- The PySys Apama Correlator class has a new method called `initialize()` which can be used to initialize the correlator with all the files making up an Apama application (such as `.mon`, `.qry` and `.evt`) by specifying just the path of the Software AG Designer project's `.deploy` file,

or a .txt file or directory, without the need to list all the files making up the application individually. See the Apama PySys samples and Pydoc for more information and an example of how to use this.

Command Central support in 9.12

Apama 9.12 now supports the following in Command Central:

- Installing Apama using Command Central web user interface.
- Creating and deleting instances using the web user interface.
- Configuring correlator engine connect options. You can use this configuration to connect a source correlator (the sender) to a target correlator (the receiver). The target correlator will receive events from the specified channels of the source correlator. For more information, see "Configuring pipelining with engine_connect" in *Deploying and Managing Apama Applications*.
- Configuring correlator initialization. You can use this configuration to set an Apama application to initialize the correlator on startup.
- Configuring and deploying projects using the Digital Event Services connectivity plug-in.
- Monitoring KPIs for EPL applications.
- Monitoring KPIs for dashboard servers.

For more information, see "Deploying Apama Components with Command Central" in *Deploying and Managing Apama Applications*.

Miscellaneous enhancements and changes in 9.12

Apama 9.12 includes the following miscellaneous enhancements and changes:

- Apama 9.12 incorporates the ICU (International Components for Unicode) time-zone data update 2016f, which is the most recent update at the time of release. This will update time-zone data used by the correlator and Time Format EPL plug-in.
- The BigMemory Max driver now supports supplying converters per table to handle non-Apama managed data. The following new properties, which can be set in the `spring.xml` configuration file, are available for this purpose:
 - `converterConfig.default`
 - `converterConfig.byTable`

For more information, see "BigMemory Max driver details" in *Developing Apama Applications*.

- In Apama versions 9.10 and prior, Universal Messaging channel names were escaped so that only a small subset of the ASCII characters were used in channel names in Universal Messaging. As of Apama 9.12, the new boolean configuration property `um.channels.escaped` is available which specifies whether channel names are escaped or not. To maintain backwards compatibility, this property is currently set to `true`, but this may change in future releases. If

you set `um.channels.escaped` to `false`, Apama passes channel names directly to Universal Messaging without escaping.

When the slash (/) and backslash (\) characters are not escaped, they can now be used to create nested channels. However, you must take care not to use both slash and backslash characters as path separators within the same application as this will result in undefined behavior. See also "Considerations for using Universal Messaging channels" in *Connecting Apama Applications to External Components*.

- The Docker samples have been updated to support the currently supported Docker and Docker Compose versions.

Platform changes in 9.12

Apama 9.12 runs on the platforms listed in the *Supported Platforms* document for version 9.12. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>.

Check the above mentioned *Supported Platforms* document for details about the newer platform versions supported by Apama 9.12. Due to these upgrades, the following versions are no longer supported:

- Software AG Terracotta BigMemory 4.3.2 (replaced by support for a more recent version).
- Software AG Universal Messaging 9.10 (replaced by support for a more recent version).
- Software AG Designer 9.10 (replaced by support for a more recent version).
- Software AG Common Runtime 9.10 (replaced by support for a more recent version).
- Oracle Java 8 update 71 (replaced by support for a more recent version).
- Google Chrome 40 (replaced by support for a more recent version).
- Firefox Mozilla 38 (replaced by support for a more recent version).
- SUSE Linux Enterprise Server 11 (replaced by support for a more recent version).
- Oracle WebLogic Server 12c 12.1.1 (replaced by support for a more recent version - only supported as a JMS provider).
- Apache Ant 1.7.1 (replaced by support for a more recent version).
- Docker 1.6.0 (replaced by support for a more recent version).
- Docker Compose 1.3.0 (replaced by support for a more recent version).

Miscellaneous changes in 9.12 affecting backwards compatibility

The following changes in Apama 9.12 affect backwards compatibility to previous Apama versions:

- Due to tightening of the `canParse` and `parse` methods on the EPL types of `integer`, `float` and `decimal`, some previously allowed behavior is now invalid. These EPL types now only allow

trailing and leading whitespace, excluding special characters on the `float` and `decimal` types such as "d".

As of this version, integer parsing will only accept integers and whitespace. Any other characters are invalid, which also means that whole numbers such as "1.0" and "2.0" will not parse. The following examples illustrate the new behavior for the `canParse` and `parse` methods:

Now false	Was previously parsed to
" "	0
"2.2"	2
"2e3"	2
"2 2"	2
"2garbage"	2

`float` and `decimal` parsing remains unchanged, except that hexadecimal integers will no longer parse. The following example illustrates the new behavior for the `canParse` and `parse` methods:

Now false	Was previously parsed to
0x1d	0

- In previous releases, any invalid keyword arguments passed to methods in the Apama PySys extension classes were silently ignored with no error. For example, the following did not produce an error despite being incorrect:

```
correlator.start(foobar='x')
```

As of this release, an exception will be thrown if an invalid argument is specified. Thus, test cases that have errors in them are now easy to notice.

Removed and deprecated features in 9.12

The following Apama features are now deprecated or have been removed in Apama 9.12:

- Event Modeler is now deprecated. Support for creating and deploying scenarios will be removed in a future release. It is recommended that you use EPL or queries to build new Apama applications.
- The EngineManagement API for C is now deprecated. It is recommended that you use the EngineManagement API for C++ instead.
- The old native Universal Messaging integration is now deprecated and will be removed in a future release. It is recommended that you use the new Universal Messaging connectivity plug-in instead. See also [“Apama integration with Universal Messaging in 9.12” on page 134](#).
- The `-r` (or `--rnames`) and `-UMconfig` (or `--umConfigFile`) options of the `correlator` executable, which pertain to the old native Universal Messaging integration, are now deprecated. Use the

--config option together with one or more .yaml files and (typically) also .properties files instead.

- The `um.install.dir` property, which is used by the old native Universal Messaging integration to set the location of the Universal Messaging libraries in the `UM-config.properties` file, is now deprecated. If Universal Messaging has been installed in the same Software AG installation directory as Apama, Apama will now locate the Universal Messaging installation automatically at `${APAMA_HOME}/../UniversalMessaging` and will load the Universal Messaging libraries from the appropriate subdirectory. See also "Setting up Universal Messaging for use by Apama" in *Connecting Apama Applications to External Components*.
- The `-Xconfig` option of the correlator executable, which specifies a special configuration file (see "Using the Apama component extended configuration file" in *Deploying and Managing Apama Applications*), is now deprecated in favor of YAML configuration via `--config`. For detailed information on configuration using YAML, see "Configuring the correlator" and the description of the `--config` option in "Starting the correlator", both in *Deploying and Managing Apama Applications*.
- As of this version, the deprecated `profiling` request can no longer be used with the `-r` option of the `engine_management` tool. Instead, you now have to use the `cpuProfile` request, which provides the same functionality. For further information, see "Using the CPU profiler" in *Deploying and Managing Apama Applications*. See also the announcement in ["Removed and deprecated features in 9.10" on page 161](#).
- When using the `engine_management` tool with the `-r` option, it is no longer possible to specify a single string, enclosed in quotation marks, which delimits multiple arguments by spaces. As of this version, it is only possible to specify the arguments as shown in the following example:

```
engine_management -r cpuProfile frequency
```

Make sure that multiple arguments are no longer enclosed in quotation marks. Quotation marks are now only used for arguments that contain spaces, such as a file name. See also the announcement in ["Correlator utility enhancements in 9.9" on page 166](#).

- As of this version, the deprecated Log File Manager plug-in (`LoggingManager`) can no longer be used. Instead, you now have to use the built-in logging functionality. For further information, see "Logging and printing" in *Developing Apama Applications*. See also the announcement in ["Removed and deprecated features in 9.9" on page 170](#).
- As of this version, the deprecated Enterprise Management and Monitoring (EMM) and sentinel agent features can no longer be used. It is recommended that you now use Command Central for all deployment tasks. For detailed information, see "Deploying Apama Components with Command Central" in *Deploying and Managing Apama Applications*. See also the announcement in ["Removed and deprecated features in 9.10" on page 161](#).
- The `apamadoc.xml` and `apamadoc.bat` files have been removed from the `utilities` directory of the Apama installation. To generate ApamaDoc without a user interface on both Windows and UNIX, you now have to use the `generate-apamadoc` macro from the `apama-macros.xml` Ant script, which is located in the `etc` directory of your Apama installation. For further information, see "Generating ApamaDoc from an Ant script" in *Developing Apama Applications*.

- As of this version, the Apama EPL Visualization tool for Software AG Designer is no longer installed with Apama. The tool is now available for download from the Apama area of the Software AG TECHcommunity website at <http://techcommunity.softwareag.com>.

17 What's New In Apama 9.10

■ New connectivity plug-ins feature in 9.10	150
■ New query and DataView metadata feature in 9.10	150
■ New EPL memory profiler feature in 9.10	151
■ New EPL code coverage feature in 9.10	151
■ New PySys version in 9.10	151
■ New dashboard server HTTP REST support in 9.10	152
■ New JSON-based REST support in 9.10	153
■ Command Central support in 9.10	153
■ Apama enhancements in Software AG Designer 9.10	154
■ EPL enhancements in 9.10	154
■ Correlator utility enhancements in 9.10	155
■ Dashboard enhancements in 9.10	156
■ Miscellaneous enhancements and changes in 9.10	157
■ Platform changes in 9.10	158
■ Supported JDK version in 9.10	159
■ Miscellaneous changes in 9.10 affecting backwards compatibility	159
■ Removed and deprecated features in 9.10	161

Apama 9.10 runs on the platforms listed in the *Supported Platforms* document for version 9.10. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>. Be sure to consult that document for details about supported versions of operating systems, compilers, BigMemory, and Universal Messaging.

Apama 9.10 includes new features, enhancements, and changes as described in the following topics.

New connectivity plug-ins feature in 9.10

Apama now provides a new way for developing adapters called *connectivity plug-ins*.

Connectivity plug-ins can be written in Java or C++, and run inside the correlator process to allow messages to be sent and received to/from external systems. Individual plug-ins are combined together to form *chains* that define the path of a message, with the correlator host process at one end and an external system or library at the other, and with an optional sequence of message mapping transformations between them. A configuration file describes both the chains and the plug-ins making up each chain. The configuration file is written using the YAML markup language, and can express structured configuration (maps, lists and simple values) for plug-ins. For detailed information, see "Using Connectivity Plug-ins" in *Connecting Apama Applications to External Components*.

A new option, `--connectivityConfig`, is available for the correlator executable, which specifies the path to the above mentioned configuration file. See "Starting the correlator" in *Deploying and Managing Apama Applications*.

A new predefined annotation, `com.softwareag.connectivity.ExtraFieldsDict`, is available, which can be used to place map keys that do not name fields into a dictionary. See "Adding predefined annotations" in *Developing Apama Applications*.

New query and DataView metadata feature in 9.10

You can now record information about a query in the metadata section. This can be, for example, the recording author, the version number, or the last modified date of a query. See "Defining metadata in a query" in *Developing Apama Applications* for more detailed information.

Metadata properties can be specified for a DataView by adding keys with the prefix `DataViewDefinition.EXTRA_PARAMS_METADATA_PREFIX` to the `extraParams` dictionary of `DataViewAddDefinition` when adding the new DataView definition.

A new constant, `EXTRA_PARAMS_METADATA_PREFIX`, has been added to the `ScenarioService` API. Any metadata properties that were set as part of a query or DataView definition will appear in the `extraParams` with this prefix. See the `IScenarioDefinition` interface in the *API Reference for Java (Javadoc)* or *API Reference for .NET* for more information.

Once defined, metadata information about a query can be viewed in the Scenario Browser. The Scenario Browser shows metadata properties for queries and DataViews. Metadata properties are shown on the main panel when the query node is selected.

See also "[Miscellaneous enhancements and changes in 9.10](#)" on page 157 for information on irrelevant sections in the Scenario Browser that are now hidden from the main panel.

See also “[Dashboard enhancements in 9.10](#)” on [page 156](#) for information on the new “Definition extra params table” which can be used to display metadata (and other `extraParams` information) in a dashboard.

New EPL memory profiler feature in 9.10

An EPL memory profiler is now provided that you can use to display information on monitors and monitor instances. This is helpful for checking the performance and for finding memory leaks.

To make use of the EPL memory profiler, you use the `-r` (or `--dorequest`) option of the `engine_management` tool with one of the following request types:

- `eplMemoryProfileOverview`
- `eplMemoryProfileMonitorInstanceDetail`
- `eplMemoryProfileMonitorDetail`

The information that is returned for a request can be written to a comma-separated values (CSV) file which can easily be viewed in tabular form using a tool such as Microsoft Excel. For detailed information, see “Using the EPL memory profiler” in *Deploying and Managing Apama Applications*.

New EPL code coverage feature in 9.10

The correlator can now generate “code coverage” information about EPL files, indicating which lines have been executed. This is useful for measuring the quality of test cases, discovering lines of EPL code which are not being exercised by any tests, as well as for helping diagnose bugs or understand complex interactions in the EPL.

The recording of code coverage information can be enabled and written to disk using management requests, or using an environment variable that automatically writes out a coverage file when the correlator is shut down or when code is deleted from the correlator.

The new `epl_coverage` tool can then be used to merge together the coverage files that have been produced by the correlator and produce summary statistics about how much of each source file is covered, as well as an HTML report where each source line is shown annotated with different colors to indicate which lines are not being covered.

The Apama extensions to the PySys test framework can now enable code coverage recording, and automatically run a coverage report at the end of test execution, which will help users to create better test cases and to find code paths in their EPL applications that do not have adequate test coverage.

For detailed information, see “Generating code coverage information about EPL files” in *Deploying and Managing Apama Applications*.

New PySys version in 9.10

Apama 9.10 ships a new release of PySys, version 1.1.0. This new release includes many useful fixes and enhancements. For full details, see the PySys release notes.

One significant addition in PySys 1.1 is an optional `abortOnError` capability that produces clearer and more timely “outcome reason” messages to explain why a testcase has failed when something goes wrong, such as an injection failure or an expected signal not being found in a file. The `abortOnError` functionality can be enabled by setting the following property in your project file:

```
<property name="defaultAbortOnError" value="true"/>
```

See the `pysysproject.xml` file in the `samples/pysys` directory of your Apama installation for examples of all the settings we recommend for all new test projects. Existing projects may wish to continue with the older settings to avoid the need to fix up any failing tests.

Support has been added for generating EPL code coverage reports from PySys test executions. See "Using EPL code coverage with PySys tests" in *Deploying and Managing Apama Applications* for more details.

Correlator and IAF components are now shut down cleanly by default at the end of each test execution, which is a prerequisite for getting EPL code coverage information. See also [“Miscellaneous changes in 9.10 affecting backwards compatibility” on page 159](#) for information on how to disable this new behavior.

A new `antBuild` method has been added which can be used to invoke Apache Ant builds with the Apama environment variables defined. This can be used, for example, to execute an exported `build.xml` file which references the `apama-macros.xml` file.

Several new methods have been added to the Apama PySys correlator object:

- `inspect`
- `sendEventStrings`
- `inputLog` argument to `startCorrelator`

On Linux, PySys Docker extensions are now available in the `third_party/python/lib/python2.7/site-packages/docker` directory of your Apama installation.

New dashboard server HTTP REST support in 9.10

You can now monitor Apama dashboard data server and display server components with the same Apama Representational State Transfer (REST) HTTP protocol that is already supported by the correlator and IAF. This provides monitoring capabilities to third-party managing and monitoring tools or to any application that supports sending and receiving XML or JSON documents over the HTTP protocol.

The HTTP protocol uses the management port of the dashboard servers. As a result, in addition to the existing `dashboard_management` tool, it is now possible to use the `component_management` tool, the Java `GenericComponentManagement` API or any HTTP REST client supporting XML or JSON to perform tasks such as pinging the server, listing the version number, memory usage, environment variables and start-up arguments, shutting down the server, and sending other dashboard-specific requests. See "Managing and Monitoring over REST" in *Deploying and Managing Apama Applications* for more information about using REST with Apama.

New JSON-based REST support in 9.10

Support has been added for the JSON equivalent to most XML URI GET requests as per those listed below:

- `/correlator/info`
- `/correlator/status`
- `/correlator/appLogging`
- `/correlator/types`
- `/iaf/status`
- `/info`
- `/logLevel`
- `/connections`

For detailed information, see "Managing and Monitoring over REST" in *Deploying and Managing Apama Applications*.

The correlator types are now also accessible from the web browser using `http://localhost:15903/correlator/types`. Note that there is no security around these requests. Anyone who has access to the host and port will be able to make these requests. See "Security Requirements for Apama" in *Deploying and Managing Apama Applications* for more information.

Command Central support in 9.10

A **Platform Manager Plug-in** node is now available in the product selection tree of the Software AG Installer. You can see it when you expand the **Apama** node. This plug-in allows you to manage the server components of Apama using Command Central.

For more information, see the Command Central documentation at <http://documentation.softwareag.com/> (Empower login required).

Apama integration with Command Central supports the following features:

- **Inventory actions.** Creating and deleting an Apama instance representing the configuration needed to start a correlator, IAF, dashboard data server or display server process.
- **Configuration Manager.** Update instance configuration.
- **Lifecycle operations.** Start, stop, and restart an Apama instance (correlator, IAF, dashboard data server and dashboard display server instances).
- **Monitoring components** Monitoring correlator, IAF, dashboard data server and dashboard display server instances.

For more information, see "About deploying components with Command Central" in *Deploying and Managing Apama Applications*.

Apama enhancements in Software AG Designer 9.10

Apama 9.10 includes the following enhancements in Software AG Designer:

- The times for doing incremental rebuilds of large Apama projects have been significantly reduced. Tests have shown that this has been speeded up by more than a factor of two on a typical desktop.
- When configuring a JDBC adapter, you can now specify the timeout for queries in the **General Properties** section. The new property **Query timeout** is available for this purpose. Keep in mind that different database vendors define a query timeout differently; see the documentation for these databases for more information. See also "Configuring an ADBC adapter" in *Connecting Apama Applications to External Components*.
- You can now enable or disable diagnostic logging for a query file. This feature was available only for scenarios and was labelled as **Generate Debug Code**. In this release, the label **Generate Debug Code** is renamed to **Diagnostic Logging**.
- The top-level page of the Apama preferences now informs you in which directory you can find the Apama log file for Software AG Designer, and a link is now provided for accessing that directory directly. See also the description of the "Apama" preferences page in *Using Apama with Software AG Designer*.

EPL enhancements in 9.10

Apama 9.10 includes the following EPL enhancements:

- A new method `add()` has been added to the `Table` class of the `MemoryStore`. `add()` does the same as `get()`, except that it does not check if the row that is to be added already exists in the table until `commit()` is called and it therefore never throws an exception. If you are sure that the row does not yet exist, you can use `add()` as this is faster than `get()`. See also "Creating and removing rows" in the description of the `MemoryStore` in *Developing Apama Applications*.
- All comparable types can now be used with the comparison operators (`<`, `>`, `<=`, `>=`, `=` and `!=`). This includes sequences and dictionaries where all members are of a comparable type. See also "Comparable types" in *Developing Apama Applications*.
- The following new built-in aggregate functions are now available:
 - `countUnique(value)`
 - `percentile(value,rank)`

The built-in positional aggregate functions `first(value)`, `last(value)` and `nth(value,index)` now also support `integer`, `string`, `boolean` and `location` types. In addition, `nth(value,index)` now also supports negative indices to get items from the end of the window (`-1` means the last item, `-2` means the second last item, and so on).

The built-in aggregate functions `min(value)` and `max(value)` now also support `integer` types.

Correlator utility enhancements in 9.10

Apama 9.10 includes the following correlator utility enhancements:

- Using the compiled runtime now requires that the `binutils` package is installed on the Linux system. For the Linux platforms supported by Apama, this is part of the default installation.

Injection times of applications using the compiled runtime have been improved in this release. There is a performance improvement in initial injections for codebases with a large number of actions.

The correlator can also be configured to use a cache of the compiled artifacts to speed up subsequent re-injection of the same source files. For this purpose, the following new option is available for the correlator executable:

```
--runtime-cache dir
```

For more information, see "Starting the correlator" and "Injection time of compiled runtime" in *Deploying and Managing Apama Applications*.

- Using the new optional `${PID}` tag, you can now add the process ID to the name of a log file. This is helpful if you want to run multiple correlators with the same arguments but with separate log files. For more information, see "Specifying log filenames" in *Deploying and Managing Apama Applications*.
- New options for log handling which were previously only available as request types of the `-r` option are now available with the `engine_management`, `component_management` and `iaf_management` tools.

The following options are available for all of the above-mentioned tools:

```
--setLogFile path
```

```
--reopenLog
```

The following options are only available for the `engine_management` tool:

```
--rotateLogs
```

```
--setApplicationLogLevel [package=]level
```

```
--setApplicationLogFile [package=]path
```

```
--unsetApplicationLogLevel package
```

```
--unsetApplicationLogFile package
```

```
--unsetRootApplicationLogLevel
```

```
--unsetRootApplicationLogFile
```

```
--getApplicationLogLevel package
```

```
--getApplicationLogFile package
```

```
--getRootApplicationLogLevel
```

```
--getRootApplicationLogFile
```

For more information on these options, see "Shutting down and managing components" in *Deploying and Managing Apama Applications*.

- The following new options are now available for the `engine_management`, `component_management` and `iaf_management` tools, which can be used to get the uptime, and the virtual and physical memory usage of the component:

```
-M | --getuptime
```

```
-Vm | --getvmemory
```

```
-Pm | --getpmemory
```

For more information on these options, see "Shutting down and managing components" in *Deploying and Managing Apama Applications*.

Corresponding methods are also available in the Java `EngineClient` API.

Dashboard enhancements in 9.10

Apama 9.10 includes the following dashboard enhancements:

- The Google map object is now available on the **Graphs** tab of the Dashboard Builder's Object Palette. This object is available in Thin Client deployments only.
- The **Table (HTML5)** table object is now available on the **Tables** tab of the Dashboard Builder's Object Palette. This is an advanced HTML implementation of the table object which provides features such as enhanced column filtering, column re-ordering and column locking. This object is available in Thin Client deployments only.
- A new tabular data attachment type "Definition extra params table" is now available. This new table shows the metadata of the `extraParams` member of a `definition` (scenario, `DataView` or query). An entry in `extraParams` is considered metadata if its key name starts with the `Metadata:` prefix. This prefix will be automatically added for the query properties.
- A new command line argument `--dashboardDir folder` is now available. The display server can be run from any directory by using this option to specify an arbitrary *folder* where deployed dashboards can be found. If this option is not used, then `APAMA_WORK/dashboards` is assumed.
- A new command line argument `--dashboardExtraJars jarFiles` is now available. Custom function, custom command and/or any extra jar files referenced by the dashboard can now be specified by using this `--dashboardExtraJars` option. *jarFiles* contains a list of semicolon-separated JAR files to be added to the classloader at runtime. If this option is not used, the `APAMA_DASHBOARD_CLASSPATH` environment variable must be set accordingly for the JAR files to be accessible by dashboard processes.

Miscellaneous enhancements and changes in 9.10

- It is now possible to pass the username and password on the JDBC connection URL, rather than having to set it in EPL. To do so, you must set the special value `ADBC_NULL` on both the username and password parameter on the EPL side, for example, via the `Connection.openDatabaseFull()` action in `ADBCEvents.mon`.
- The new `IAFStatusDataViewService` provides support for publishing the status of an IAF adapter as a `DataView` item. It can be used to easily monitor the adapter status using Apama's Scenario Browser or to visualize the adapter status using an Apama dashboard. For more information, see "DataView support" in *Connecting Apama Applications to External Components*.
- ApamaDoc documentation is now provided for most of Apama's public event APIs.
- By default, the ApamaDoc export wizard does not generate documentation for inner fields of a monitor. If you want to include inner fields of a monitor, run ApamaDoc in headless mode using the `--includeMonitorMembers` option.
- ApamaDoc now supports `<code/>` tags inside ApamaDoc comments.
- The correlator now logs `INFO` log lines if an external sender sending events in to the correlator is blocking, giving a reason as to why the queue is full. For example:

```
Blocked adding 23 events to context main (id=1)'s queue; context is blocked
waiting on correlator (on port 62156) which is flagged as a slow receiver
Blocked adding 2 events to context main (id=1)'s queue; context is blocked
waiting on ctx1 (id=2) which is running a plug-in
```

- The Java APIs now implement the `AutoCloseable` interface for simple cleanup using the Java `try-with-resources` statement, where applicable. See the *API Reference for Java (Javadoc)* for more information.
- The Apama work directory (`APAMA_WORK`) now has a `lib` directory which is on the `PATH/LD_LIBRARY_PATH` set up by the Apama Command Prompt (`apama_env` script). This `lib` directory can be used for user-supplied C++ correlator and IAF plug-ins.
- Client applications can now add user-defined status values which can be appended to the standard status messages of the correlator:
 - `get`, `set` and `delete` actions for the user-defined status values are now available from the EPL Management interface. See the corresponding section in "Using the Management interface" in *Developing Apama Applications*.
 - When the new `-a` (or `--all`) option is specified for the `engine_watch` tool, any user-defined status values are now appended to the standard status messages. See "Watching correlator runtime status" in *Deploying and Managing Apama Applications*.
 - User-defined status values are now visible over the REST API. See "Managing and Monitoring over REST" in *Deploying and Managing Apama Applications*.
 - User-defined status values can now be retrieved from the `EngineClient` API. See "EngineClient API" in *Connecting Apama Applications to External Components*.

Note that the correlator status statements that appear in the log files will not have the user-defined status values, and will remain unaffected.

- The `getMostBackedUpInput` and `getMostBackedUpQueueSize` methods of the `EngineStatus` interface have changed meaning. They now give the most backed up, or slowest, context name and queue size across all contexts, not just public contexts, as non-public channels can subscribe to specific channels and thus receive external events.
- The Scenario Browser now hides the irrelevant sections from the main panel.
 - For Scenario, the metadata section is hidden.
 - For Dataview, the input section is hidden.
 - For Query, the output section is hidden.
- On Linux, a Universal Messaging Docker sample is now available in the `samples/docker/applications/UniversalMessaging` directory of your Apama installation.
- With the previous Apama version, the correlator automatically shut down after 30 minutes when a license file could not be found. As of this version, when a license file cannot not found, Apama can be used without a time limit, but with reduced capabilities. See "Running Apama without a license file" in *Introduction to Apama*.

Platform changes in 9.10

Apama 9.10 runs on the platforms listed in the *Supported Platforms* document for version 9.10. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>.

Support for the following has been removed:

- Google Chrome no longer supports Java applets. Therefore, applet-deployed dashboards will not work on Chrome browsers. WebStart and display server deployment will continue to work on all supported browsers.

Check the above mentioned *Supported Platforms* document for details about the newer platform versions supported by Apama 9.10. Due to these upgrades, the following versions are no longer supported:

- Software AG Terracotta BigMemory 4.3.1 (replaced by support for a more recent version).
- Software AG Universal Messaging 9.9 (replaced by support for a more recent version).
- Software AG Designer 9.9 (replaced by support for a more recent version).
- Software AG Common Runtime 9.9 (replaced by support for a more recent version).
- Oracle Java 8 update 51 (replaced by support for a more recent version).
- Google Chrome 36 (replaced by support for a more recent version).
- Firefox Mozilla 31 (replaced by support for a more recent version).
- Microsoft Internet Explorer 9 (replaced by support for a more recent version).

- Apple Safari (desktop) 5 (replaced by support for a more recent version).

Due to changes to the supported JDK version (see [“Supported JDK version in 9.10” on page 159](#)), support for the following application servers has been removed in order to support application server versions that would run using a supported JRE version:

- Oracle WebLogic Server
 - 11g 10.3.5 (still supported as a JMS provider)
 - 12c 12.1.1 (still supported as a JMS provider)
- IBM WebSphere Application Server 8.5.5.0 (still supported as a JMS provider)

Supported JDK version in 9.10

Apama is now built against Oracle JDK 8, and installs with Oracle JRE 8. JDK 7 is no longer supported.

To develop, build, and test an Apama 9.10 application, you must use Oracle JDK 8, which is the version that Software AG Designer installs. Use of any JRE other than the one that the product ships with is discouraged.

However, for applet and/or Webstart dashboards to run properly in 32-bit browsers, the latest 32-bit JREs are recommended on the client machines.

Miscellaneous changes in 9.10 affecting backwards compatibility

- With previous versions, the EPL coassignment mechanism allowed coassignments from different triggerings of a listener instance to leak into each other. For example, consider the following code:

```
A a := A(0);
B b := B(0);
on all (A():a or B():b) {
    print a.toString() + " " + b.toString();
}
```

Previously, sending A(1) and B(1) printed the following, where the coassignment of A(1) carried over into the triggering of the listener caused by B(1):

```
A(1) B(0)
A(1) B(1)
```

This has been corrected in this version so that the above code will now print the following:

```
A(1) B(0)
A(0) B(1)
```

This change may also cause a slight reduction in memory usage and an increase in performance of persistent applications.

- The `isExternal()` method on events in EPL has changed behavior. While in previous releases, the `route..to` or `enqueue..to` statements would change an event to always be internal, in 9.10, these do not modify the `isExternal()` property. For example, an external event that is forwarded to a second context will still appear to be external. The `clone()` method always returns an internal event, so if an internal event is required, then `enqueue evt.clone()` to target will give the same behavior as previous releases.
- Apama now ships with the files `QueryServices.cdp` and `SetSingleContext.cdp` instead of `QueryServices.mon` and `SetSingleContext.mon`. That is, correlator deployment packages are now provided instead of EPL files.
- The Apama extensions to the PySys testing framework now perform a clean shutdown of correlators and IAFs before the test ends, rather than terminating the process without warning. This makes it easier to detect problems associated with the shutdown, and it also allows code coverage information to be retrieved. If you want to disable this new behavior, add the boolean property `shutdownApamaComponentsAfterTest` to the project file.
- When a launch configuration is created for a project in Software AG Designer, the name “defaultCorrelator” is now used. Backwards compatibility is provided for the Apama projects that you have created with previous versions: the name “Default Correlator” is automatically considered as “defaultCorrelator”.

However, if you still have Apama projects that have been created before version 5.2 and if these projects have adapter configurations that use the `${Default Correlator.port}` and `${Default Correlator.hostname}` properties and do not have any launch configurations, the Ant export of these projects will not work seamlessly. As a workaround, you have to apply any of the following changes to these projects:

1. When a launch configuration has been created in Software AG Designer which uses the name “defaultCorrelator”, change this back to the old “Default Correlator” name.
 2. Use the new properties in your adapter configuration. That is, change the old `${Default Correlator.port}` and `${Default Correlator.hostname}` properties to `${defaultCorrelator.port}` and `${defaultCorrelator.hostname}`.
 3. Modify the Ant exported `environment.properties` file and add the value `Default Correlator.port=portValue`.
- The random number generator in the Apama correlator has been updated to provide higher resolution and improved performance. This change will result in different sequences of random numbers being produced and as such any code or tests which rely on specific sequences of random numbers via seed values (provided, for example, via the `-XsetRandomSeed` option of the correlator executable) or which use the “replay log” feature will need to be updated.
 - Performance improvements in this version have eliminated the need for different correlator scheduler types. The `--scheduler` option has therefore been removed from the correlator executable. Note that the correlator will fail to start up if you continue to use this option.
 - The location of Apache Ant which is shipped with Apama has been changed and it can now be found in the `third_party\apache_ant` directory of your Apama installation. If you have customized any scripts to rely on the previous location, we recommend that you test your scripts carefully to ensure that they work as expected.

- The correlator now uses separate Java child classloaders for loading the libraries required for distributed MemoryStore and JMS support, which substantially reduces the number of libraries present in the top-level system classloader that all correlator, connectivity, JMS and distributed MemoryStore plug-ins share, and therefore reduces the potential for library conflicts between them and customer plug-ins. In previous versions it was possible (though not recommended) to start the correlator with distributed MemoryStore and JMS configuration files in the same directory; this is no longer possible and will result in "Failed to initialize correlator: [ADAPTER_NAME] configuration is invalid: Line XX in XML document from URL..." errors.

To avoid "class not found" problems, ensure you are specifying classpaths using the settings in the relevant XML configuration files for JMS and distributed MemoryStore rather than using command-line options that change the correlator's system classpath such as

`-Djava.class.path`. The vast majority of applications will be doing this already and will require no modification. In some situations, if your plug-in or a library it depends upon uses the context (thread-specific) classloader to locate classes, it may also be necessary to set that explicitly using the following:

```
Thread.currentThread().setContextClassLoader(
    YourClassHere.class.getClassLoader() )
```

Removed and deprecated features in 9.10

- Apama's Enterprise Management and Monitoring (EMM) and sentinel agent features are deprecated in this release. Apama recommends that you use Command Central instead of EMM for all deployment tasks.
- Actions without parameters that are declared as `action action-name {...}` are deprecated as of this version. If you still use such a declaration in your code, make sure to change this to `action action-name() {...}` (with additional parentheses after the action name).
- The following built-in aggregate functions are deprecated. It is recommended that you now use the alternative functions mentioned below.
 - `count(value)`. Use the alternative predicate aggregate function `count(not value.isNaN())` instead. Note that only the aggregate function is deprecated which uses the decimal and float values. The aggregate functions `count()` (without an argument) and `count(predicate)` (with a boolean argument) are still supported.
 - `prior(value, index)`. Use the alternative aggregate function `nth(value, index)` with a negative index instead.
- The profiling request of the `engine_management` tool has been deprecated. A new `cpuProfile` request, which provides the same functionality, is available instead. You can still use the deprecated profiling request, however, it is recommended that you now use the new `cpuProfile` request.

18 What's New In Apama 9.9

■ Installation changes in 9.9	164
■ Query enhancements in 9.9	165
■ Apama enhancements in Software AG Designer 9.9	165
■ EPL enhancements in 9.9	166
■ Correlator utility enhancements in 9.9	166
■ Dashboard enhancements in 9.9	167
■ Miscellaneous enhancements and changes in 9.9	167
■ New samples in 9.9 showing how to containerize Apama for Docker (Linux only)	168
■ Platform changes in 9.9	168
■ Renaming of Apama client libraries in 9.9	169
■ Miscellaneous changes in 9.9 affecting backwards compatibility	170
■ Removed and deprecated features in 9.9	170

Apama 9.9 is the successor of Apama 5.3.

Apama 9.9 runs on the platforms listed in the *Supported Platforms* document for version 9.9. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>. Be sure to consult that document for details about supported versions of operating systems, compilers, BigMemory, and Universal Messaging.

Apama 9.9 includes new features, enhancements, and changes as described in the following topics.

Installation changes in 9.9

As of version 9.9, Apama is installed using the Software AG Installer, and official fixes are installed using the Software AG Update Manager. See the new *Installing Apama* guide for more information.

Previous Apama releases (such as 5.3) allowed you to select one of three installation types: Server, User and Developer. In the Software AG Installer, these correspond to selecting the following options from the product selection tree:

- **Apama > Server.** Installs the server components of Apama, including the correlator and associated command line tools.

Note:

On UNIX, this is the only installation option. All other options listed here are only available on Windows.

- **Apama > Dashboard Viewer.** Installs the viewer tool that allows you to view and interact with dashboards that are receiving live data from an Apama dashboard server.

Note:

Unlike the old User installation type, the Software AG Installer only installs the viewer tool, and does not include any other Java/C++ libraries.

- **Designer > Apama > Apama Development Environment.** Installs Software AG Designer which provides graphical tooling for creating Apama applications.

All Apama documentation is now installed regardless of which installation type you select.

You can quickly select all Apama components (that is, the full functionality of Apama) by clicking **Typical Installations** which is shown next to the product selection tree and then selecting **Apama**.

Due to the above-mentioned changes, all Windows start menu shortcuts have changed.

The Eclipse-based development environment previously known as Apama Studio is replaced by integration into Software AG Designer. You can open Software AG Designer, which is also based on Eclipse, by choosing the following from the Windows Start menu: **All Programs > Software AG > Tools > Software AG Designer n.n**. Note that **Software AG** is the default group name that you can change during the installation.

On Windows, the Apama work directory (APAMA_WORK) now defaults to the all-users %PUBLIC% directory, for example, C:\Users\Public\SoftwareAG\ApamaWork_n.n (with previous versions, it was in C:\Users\username).

On UNIX, the default for the Apama work directory is now `~/softwareag/apamawork_n.n`.

With previous versions, the Eclipse workspace was stored under `APAMA_WORK` by default. Now it is in the location that you select when you start Software AG Designer for the first time. By default, the workspace is located in the `C:\Users\username` directory.

It is now also possible to install Apama using Command Central. You can then install Apama to any machine on which the Platform Manager is running. See the new *Installing Apama* guide for more information.

Query enhancements in 9.9

Apama 9.9 includes the following query enhancements:

- Queries can be configured to use the source timestamp in an event, by defining an action that returns the source time from the event's fields (performing any required transformation or parsing of the time). See "Using source time stamps of events" in *Developing Apama Applications*. Events can be marked with an annotation as not expected to arrive in order, and queries can specify an event type to use as a heartbeat to indicate that events have been delivered and are not delayed.
- Queries can contain a `with unique` clause which discards all but the latest event when many events have the same value for a specified field. See "Matching only the latest event for a given field" in *Developing Apama Applications*.
- When a number of correlators are operating in a cluster, failure of one of the correlators will now result in any timers which are active on that node being re-distributed to the other members of the cluster, so that timers are not lost. Note that Apama queries do not support reliable event delivery, so there may still be a small proportion of message loss as a result of a node failing.
- The Pysys Apama extensions now support injecting queries as part of tests.

Apama enhancements in Software AG Designer 9.9

As of version 9.9, Apama's main tool for implementing Apama applications runs in Software AG Designer, which is based on Eclipse, and is no longer called "Apama Studio".

New Apama features in Software AG Designer include:

Queries

- **Completion proposals** - completion proposals are now enabled in several parts of the Query Designer editor. This includes completion proposals in places where text editing is possible in dialogs of Filter (Where), Time From (Within), Exclusion (Without), Calculation (Select) and Filter (Having). The completion proposals are ordered by groups (variables, actions, etc.) and the items within the groups are sorted in alphabetical order.
- **Send Event action** - the channel name and the field values of the events can now be edited in a rich-text-enabled format. The string values can be edited using a mix of text and expression. The field values can now be edited in a dialog. It is now possible to add variables for expressions using a context menu.

- **Input section** - you can now configure timestamps in the Input section dialog. In the **Design** tab, the Input section table has been enhanced (additional columns for new features in this version, icon-based representation and improved tooltips), and validation markers are now shown in the corresponding parts of the dialog.
- **Error markers in table rows** - in addition to showing errors at the section level, the error markers are now shown at the row level.
- **Configure dialog** - the configure dialog allows you to manage the package (namespace) of the query and its name.
- **Dashboards for queries** - the Dashboard Generation editor can now take query files to generate the dashboard pages.

Other

- The location of the folder from which the .ste files are picked up has been changed to `software_ag_install_dir/Designer/extensions`.
- The launch configuration dialog for a correlator has an additional entry for referencing the Xconfig file. This supports variables in paths.
- The ANT export action no longer copies the `apama-macros.xml` file. The `build.xml` file, which is generated from the ANT export action, now references the `apama-macros.xml` file in `APAMA_HOME`.

EPL enhancements in 9.9

Apama 9.9 includes the following EPL enhancements:

- Some predefined annotations are now supported in EPL; see "Adding predefined annotations" in *Developing Apama Applications*. These can be used to specify default configurations for source timestamps (see "Using source time stamps of events" in *Developing Apama Applications*).
- The Time Format plug-in now has an event wrapper, which should be used in preference to directly calling the plug-in. For detailed information, see "Using the TimeFormat Event Library" in *Developing Apama Applications* and the information on the `TimeFormat` event in the *API Reference for EPL (ApamaDoc)*.

Correlator utility enhancements in 9.9

Apama 9.9 includes the following correlator utility enhancements:

- `engine_management`

When using the `-r` option, you should now specify distinct arguments: the request type and a list of strings. For example:

```
engine_management -r profiling frequency
```

In previous versions, you had to specify a single string that delimited the arguments by spaces:

```
engine_management -r "profiling frequency"
```

When specifying your management requests now, make sure that the arguments are no longer enclosed in quotation marks. Quotation marks are now only used for arguments that contain spaces, such as a file name. For example:

```
engine_management -r setLogFile "my new log file.log"
```

For more information on `engine_management`, see "Shutting down and managing components" in *Deploying and Managing Apama Applications*.

Equivalent changes have been made to the `doRequest` methods in the Java client API and the `engine-management-request` macro in the Ant macros. Applications making use of either must be updated accordingly.

Dashboard enhancements in 9.9

Apama 9.9 includes the following dashboard enhancements:

- A new **Trends HTML5** tab is now available in the object palette. The trend charts on this tab support HTML5. The `webChartFlag` property is set to `true` by default.
- A new property `webLabelFlag` is now available for all objects on the **Labels** and **General** tabs of the object palette. It is also available for the checkbox object on the **Controls** tab. The `webLabelFlag` property is set to `true`. With this property, the display server client will render the objects from the above-mentioned tabs using HTML5 (on supported browsers) rather than in the image generated by display server. This enhancement can improve system performance in some cases and will also allow users to copy text strings from the objects to the clipboard. This property will be ignored in `dashboard_builder`, `dashboard_viewer`, `Applet` and `WebStart` clients.
- A new scalar function, `Replace All`, is now available. It replaces all occurrences of a given string which matches the pattern of the regular expression with another string.

Miscellaneous enhancements and changes in 9.9

- The Apama installation includes an Oracle Java(TM) SE Development Kit (JDK). To write and compile Java applications, you need a Java compiler such as `javac`, and the `jar` utility. Software AG recommends that you use Oracle JDK 8 included in the Apama installation to develop, build, and test your applications. The minimum JDK version you can use is JDK 7.
- Apama 9.9 incorporates ICU (International Components for Unicode) time-zone data update 2015e, which is the most recent update at the time of release. This will update time-zone data used by the correlator and `TimeFormat` correlator plug-in.
- The correlator will now search `APAMA_WORK/license/ApamaServerLicense.xml` to locate a license file if none is specified on the command line. If the license file is not found at that location, the correlator will also search `Apama/etc/ApamaServerLicense.xml` in the installation directory. If the found license file is expired, the correlator will fail to start up. In order to run the correlator in evaluation mode, any expired license file must be removed.

Note:

The default name for the license file is now `ApamaServerLicense.xml` and no longer `license.txt`.

- If the license cannot be found, remote connections are now allowed. You are no longer restricted to communicating only with processes on the same machine.
- The correlator log will now include the type and size of the largest container (sequence, dictionary, listener, stream window) in each `MThread` when the logging of garbage collection events has been enabled using the `verbosegc` command. This is to allow easier diagnostics of applications that are using unusual amounts of memory.
- The file `jms-mapping-spring.xml`, which stores the rules for mapping between JMS messages and Apama events, is now generated directly into the `bundle_instance_files\Correlator-Integrated_JMS` folder of an Apama project. With previous versions, this file was generated into the `bundle_instance_files\Correlator-Integrated_JMS\Generated` folder. The `Generated` folder is no longer created and used by Apama.
- The connection event now has a new reopen action: `reopenWithACK` which requires a callback. This is useful to know when a reopen succeeds or fails. We recommend that you use this method over the deprecated `reopen()`.
- The `shutdownConnectionOnError` action will now be based on the reconnect policy and will therefore no longer remove outstanding queries from the queue if, for example, you have a reconnect policy of `RECONNECT_AND_RETRY_LAST`. In this case, outstanding queries will remain on the queue to be retried on a successful reconnection.

New samples in 9.9 showing how to containerize Apama for Docker (Linux only)

The Docker sample directory contains configuration and samples to help you containerize and run Apama components and applications on the Docker platform, including

- Turning an Apama install into a Docker image
- Running a Docker image containing an Apama correlator, IAF adapter, Dashboard server, etc.
- Using Docker Compose to orchestrate an application running multiple connected containers

For more information about the sample, see `README.txt` in `APAMA_HOME/samples/docker`.

Platform changes in 9.9

Apama 9.9 runs on the platforms listed in the *Supported Platforms* document for version 9.9. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>.

Support for the following has been removed:

- Installing Apama into a standalone Eclipse (the Apama development environment is now available only through Software AG Designer).

- 32-bit Windows platform (no longer supported even for development tasks, following its removal as a production platform in the previous release).

Check the above mentioned *Supported Platforms* document for details about the newer platform versions supported by Apama 9.9. Due to these upgrades, the following versions are no longer supported:

- Software AG Terracotta BigMemory 4.3.0 (replaced by support for a more recent version).
- Software AG Universal Messaging 9.8 (replaced by support for a more recent version).
- Software AG Common Runtime 9.8 (replaced by support for a more recent version).
- Oracle Java 8 update 25 (replaced by support for a more recent version).
- Microsoft SQL Server 2008 R2 (replaced by support for a more recent version).
- Apache Tomcat 7.0.53 (replaced by support for a more recent version).
- Google Chrome below version 36 (replaced by support for a more recent version).

Renaming of Apama client libraries in 9.9

Several libraries have been renamed. In most cases, the version number has been removed from the file names and an “ap” prefix has been added. Due to this change, you may need to update your build scripts and automated tests. Software AG Designer will automatically upgrade any Apama project with the “Apama Java” nature to have the renamed jar files on its build path.

The following table lists the most important name changes:

Old name	New name
Windows:	
apcommon5.3.lib	apcommon.lib
engine_client5.3.dll	apclient.dll
engine_client5.3.lib	apclient.lib
engine_client_dotnet5.3.dll	apclientdotnet.dll
iafcore5.3.lib	apiaf.lib
Linux:	
libengine_client.so	libapclient.so (-lapclient)
Jar files:	
dashboard_client5.3.jar	ap-dashboard-client.jar
correlator_extension_api5.3.jar	ap-correlator-extension-api.jar

Old name	New name
engine_client5.3.jar	ap-client.jar
jplugin_public5.3.jar	ap-iaf-extension-api.jar
util5.3.jar	ap-util.jar
distmemstore-bigmemory.jar	ap-distmemstore-bigmemory.jar
distmemstore5.3.jar	ap-distmemstore.jar

Note:

You should no longer reference `dashboard_studion.n.jar`. All required classes are now available in `ap-dashboard-client.jar`.

Miscellaneous changes in 9.9 affecting backwards compatibility

- Correlator persistence can now only be enabled with a valid license file. Otherwise, it will not be possible to run the correlator in persistent mode.
- Apama supports various licensing models, one of which is licensing per physical CPU core. In previous releases, any restrictions on the number of licensed cores were not enforced. In this release, the number of threads that can be used for running Apama applications will be capped at the number of cores specified in the license, if the correlator is run on a machine that has more cores than the license allows. The correlator will log a `WARN` message to make it clear when this has happened. If you have previously been running with an insufficient license, then you may experience a corresponding reduction in performance after upgrading.
- The Ant macro `<start-correlator>` (generated by the **Ant Export** feature) checks whether there is already a correlator running. However, previously that was ignored allowing `<start-correlator>` to succeed even if a new correlator with the specified parameters was not started. Now it will log a warning and (by default) fail the task if a correlator is already running. An Ant build script exported from a launch configuration with **Reuse Correlator** selected will therefore now fail. To return to the previous behavior, edit the `build.xml` to set the `failoncorrelatorrunning=false` attribute on the `<start-correlator>` task, or set the global property `correlator.failonrunning`, for example, by specifying `-Dcorrelator.failonrunning=false` on the command line when executing Ant.
- When you upgrade to a new major version of the product, you must regenerate any CDP that was created by exporting query files or scenario definitions from the previous version.

Removed and deprecated features in 9.9

- The Log File Manager plug-in (`LoggingManager`) is being deprecated and should not be used any more, as it will be removed in a later release. In the future, you should just use the EPL `log` statement to write log messages, and configure logging using the Management interface (see "Using the Management interface" in *Developing Apama Applications*). New actions have been added to the `com.apama.correlator.Logging` event, such as `setApplicationLogFile`, `setLogFile` and `setApplicationLogLevel`. These actions are the equivalent of using the

`engine_management` requests to configure logging (see also "Shutting down and managing components" in *Deploying and Managing Apama Applications*). In previous releases `com.apama.correlator.Logging` had only one action on it, namely `rotateLogs()`.

- Apama no longer provides ODBC drivers. Please use your own ODBC drivers if you need to use ODBC. Please note that Software AG recommends using JDBC rather than ODBC with Apama.
- The ADBCHelper actions `findAvailableServers` and `findAvailableServersFull` have been deprecated, and replaced with more appropriately named `findAvailableDataSources` and `findAvailableDataSourcesFull` respectively. These are identical in behavior.
- The `reopen` action on the connection event of `ADBCEvents.mon` has been deprecated. Please use the action `reopenWithAck()` instead.

19 What's New In Apama 5.3

■ New Apama queries capability	174
■ Apama Studio enhancements in 5.3	175
■ Log rotation enhancements in 5.3	176
■ EPL enhancements in 5.3	177
■ Correlator-integrated messaging for JMS enhancements in 5.3	178
■ Dashboard enhancements in 5.3	179
■ Killing an object now requires specification of a monitor name	179
■ New sample project showing how to manage Universal Messaging DataGroups	179
■ C++ and C Correlator plug-in API enhancements	180
■ Miscellaneous enhancements and changes in 5.3	181
■ Platform changes in 5.3	182
■ JDK 7 now minimum JDK version supported	182
■ Removed and deprecated features in 5.3	183
■ Backwards incompatibility with persisted projects recovered to 5.3 from older versions .	183

Apama 5.3 runs on the platforms listed in the *Supported Platforms* document for version 5.3. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>. Be sure to consult that document for details about supported versions of operating systems, compilers, BigMemory, and Universal Messaging.

Apama 5.3 includes the following new features, enhancements, and changes.

New Apama queries capability

The new Apama queries capability lets business analysts and developers create scalable applications that process events originating from very large populations of real-world entities. Scaling, both vertically (same machine) and horizontally (across multiple machines), is inherent in Apama query applications. Scaled deployments on multiple machines use distributed cache technology to maintain and share application state. This makes it easy to deploy across multiple servers, and keep the application running even if some servers are taken down for maintenance or fail. Apama queries can be used alongside EPL monitors in the same correlator process, interacting by sending events between them.

Incoming events that queries process are partitioned by, for example, customer account numbers, car license plate numbers, devices or some other entity. In a query application, the correlator processes the events in each partition independently of other partitions.

Advantages of Apama queries over Apama monitors:

- When used in conjunction with BigMemory, queries provides active-active availability. That is, queries can be run in a cluster, where every node in the cluster contributes processing resources. The number of nodes can be changed dynamically without losing state.
- Queries can be run on correlator clusters.

Disadvantages of Apama queries compared to Apama monitors:

- Higher latency than monitors. Latency is of the order of milliseconds to seconds rather than microseconds to milliseconds. Exact values depend on the deployment and the types of events being processed.
- Apama monitors allow you to write custom and more powerful EPL applications that do not have the declarative and structural bounds that queries have.

To take advantage of the scalability and availability that the queries platform offers, the problem your application needs to solve should meet one or more of the following requirements:

- Different partitions for a given query must be completely independent. However, different queries can use different partition keys for the same event types. For example, one query may partition ATM withdrawals by `cardNumber`, and another by `atmId`.
- The average number of events in each event window in a partition is low. The recommendation is less than 50 events. For example, if ATM withdrawals are partitioned by `cardNumber` then a window that retains withdrawals for a three-day period is fine because the typical number of withdrawals per card is likely to be low. While it is possible to have hundreds of withdrawals for a single card number, that would be an exceptional case and probably indicative of suspicious behavior.

- Other than the history of events, no state is required. Queries do not provide for state to be stored. However, it is possible to mix monitors and queries in the same deployment.
- The time between events destined for the same partition is typically long, that is, more than a few seconds between events.
- The exact ordering between events is not critical. A query may treat two events for the same partition that occur close in time as having occurred in an order that is different from the order in which they were sent.

Apama queries are designed to be easy to develop for both the business analyst and the application developer. Graphical tools to specify the application design and full round-trip engineering allow both the business analyst and the developer to work on the same queries. At the developer level, an Apama query is defined using the Apama event processing language, EPL. A new channel, `com.apama.queries`, is provided for sending events to queries. To implement queries, Apama provides:

- The new Query Designer editor in Apama Studio. This graphical user interface lets business analysts define queries without the need to write code.

See "Adding query files to projects" in *Using the Apama Studio Development Environment* (renamed to *Using Apama with Software AG Designer* in version 9.9).

- New EPL constructs for application developers to write code that defines queries. A query is defined in a `.qry` file. A query cannot contain a monitor and a monitor cannot contain a query. A query can use many, but not all EPL constructs.

To use EPL to define a query, see "Defining Queries" in *Developing Apama Applications*.

- Information and instructions for configuring, deploying and managing queries. To scale query deployment across multiple machines, you use Apama's distributed `MemoryStore` with a JMS bus such as Universal Messaging.

See "Deploying and Managing Queries" in *Deploying and Managing Apama Applications*.

Apama Studio enhancements in 5.3

The main enhancements to Apama Studio support the implementation of Apama query applications. New Apama Studio features include:

- Query Designer editor offers two views for defining and updating query definitions:
 - **Design** — Business analysts define and update query definitions in this view. The **Design** tab provides an event palette and a canvas for defining an event pattern of interest. Toolbars, tooltips, and dialogs make it easy to define the events of interest to the query as well as any required parameters, conditions for finding matches, aggregate calculations and actions to take when a match set is found.
 - **Source** — Application developers can define and update query definition source code in the EPL code editor available from the **Source** tab.
- Scenario Browser has been enhanced so that you can now use it to do the following:

- Create new instances of queries by specifying query parameter values.
- Change the parameter values for a query that is running.
- Monitor query activity in the local correlator.
- EPL Objects view: A new view that appears beside the outline view in the Apama developer perspective that shows logical view of Apama objects of an Apama project in a tree structure. The view presents monitors, events and aggregates available in the Apama project. It also has options to filter the tree content and group by.
- The Apama Workbench perspective default behavior is that it shows query and event files in addition to dashboard and scenario files.
- Exporting an Ant deployment script includes all files needed to deploy queries on a correlator host machine.
- Exporting a project to a correlator deployment package (CDP) can include any query files in your project.
- A default filter with value `com.apama.*` has been added to the Event Chooser dialog to automatically remove events in the `com.apama` namespace from the event list.
- All events starting with `__` (two underscore characters) will be treated as internal events and automatically filtered from the Event Chooser, EPL Query, and EPL Objects View dialogs.

Log rotation enhancements in 5.3

When you specify the name of any of the following log files:

- Main correlator log
- Correlator input log
- Any EPL log files

You can now use one, two or all three of the following string substitutions in any order:

`${START_TIME}`

`${ROTATION_TIME}`

`${ID}`

In the previous release, string substitutions were allowed for only the name of the input log file. Also, this release adds the `${ROTATION_TIME}` string substitution. For example, to specify an EPL log filename for messages generated in `com.example.mypackage`, you could specify the log filename as follows:

```
mypackage_${ID}_${ROTATION_TIME}.log
```

The format is as follows:

```
file[one_or_more_substitutions_in_any_order].log
```


<i>file</i>	<p>Replace <i>file</i> with the name of the file that you want to be the log file.</p> <p>If you also specify <code>\${START_TIME}</code> and/or <code>\${ROTATION_TIME}</code> and/or <code>\${ID}</code>, the correlator prefixes the name you specify to the time the correlator was started and/or the time the log file was rotated (logging to a new file began) and/or an ID, beginning with 001.</p>
<code>\${START_TIME}</code>	Tag that indicates that you want the correlator to insert the date and time that it started into the log filename.
<code>\${ROTATION_TIME}</code>	Tag that indicates that you want the correlator to insert the date and time that it starts sending messages to a new log file into the log filename.
<code>\${ID}</code>	Tag that indicates that you want the correlator to insert a three-digit ID into the filename of the log file. The ID that the correlator inserts first is 001. The log ID increment is related only to rotation of log files. The ID allows you to create a sequence of log files. Each time the log file is rotated, the correlator increments the ID.

If you plan to rotate log files then be sure to specify `${ROTATION_TIME}` or `${ID}`. You can also specify both.

Apama 5.3 provides two new interfaces for rotating all Apama log files. Log file rotation means that the log file in use is closed and a new log file begins to be used. Each of the following new interfaces rotates the main correlator log, correlator input log if the correlator is generating one, and any EPL log files.

- The `engine_management` utility now takes the following option:

```
engine_management -r rotateLogs
```

- In an EPL monitor that uses the Management interface correlator plug-in, you can now call the `rotateLogs()` action. The `rotateLogs()` action is defined in the `com.apama.correlator.Logging` event.

For details, see "Rotating correlator log files" and "Rotating all correlator log files" in *Deploying and Managing Apama Applications*.

EPL enhancements in 5.3

Apama 5.3 includes the following EPL enhancements:

- You can now use dictionary and sequence literals anywhere in an EPL program and not just in variable initializations as in previous releases. For example, the following code fragments are now valid EPL:

```
setMapping({"SOW":"Software AG", "MSFT":"Microsoft"});
route FooBar(true, [100.0, 1000.0, 100000.0]);
```

- A new channel is provided. To send an event to all running Apama queries, send the event to the `com.apama.queries` channel.

Correlator-integrated messaging for JMS enhancements in 5.3

Enhancements to correlator-integrated messaging for JMS include the following:

- APIs for implementing custom mapper

Apama now provides APIs for implementing your own Java class for mapping between Apama event strings and JMS message objects. If the mapping tools provided with Apama do not meet your needs you can use these APIs to create a custom mapper. A custom mapper can handle some event types and delegate handling of other event types to the mapping tools provided with Apama or to other custom mappers. Typically, you will want to use the `SimpleAbstractJmsMessageMapper` class, which is in the `com.apama.correlator.jms.config.api.mapper` package.

- Apama mapping tools support custom resolvers and functions

The mapping tools for JMS that are provided with Apama have been enhanced to include support for custom resolvers.

- New APP_CONTROLLED reliability mode

Apama now supports JMS APP_CONTROLLED reliability mode for receivers. This reliability mode allows applications to implement no-loss messaging without using correlator persistence.

- New support for applications to send messages reliably without using correlator persistence. An application can request BEST_EFFORT senders to flush messages to the JMS broker. Applications are notified after all messages already sent to the JMS sender channel have been sent to the broker. The application defines a strategy for

- Preserving state that allows re-generation of messages in case of a correlator failure.
- Cleaning up preserved state after receiving an acknowledgement that all messages sent to a JMS receiver channel have been flushed to the JMS broker.

- JUEL mapping expressions that can be used to get or set elements of a JMS message now include resolver expressions for obtaining sender, receiver, and connection IDs.

Reference information for the APIs is provided in Javadoc format. Introductory and general information is in *Connecting Apama Applications to External Components* in the following topics:

- "Implementing a custom Java mapper"
- "Using custom EL mapping extensions"
- "Receiving messages with APP_CONTROLLED acknowledgements"
- "Sending messages reliably with application flushing notifications"
- "JUEL mapping expressions reference for JMS"

Dashboard enhancements in 5.3

A dashboard trend graph object can now be rendered by using HTML5, which provides added functionality to the chart object without requiring a Flash Player or other browser plug-in. To enable this feature, you select the `webChartFlag` option from the list of trend graph properties.

This feature applies only to display server dashboard deployment.

When `webChartFlag` is selected some minor trend graph properties are not supported, there are two new properties, and there are some differences in behavior. For details about using trend charts when `webChartFlag` is selected, see "Rendering trend charts in HTML5" in *Building and Using Apama Dashboards*.

Killing an object now requires specification of a monitor name

In all client APIs (Java, C++, C and .NET), killing an object now requires specification of a monitor name. In previous releases, specification of any name (for example, package name, event name, aggregate name) was allowed but the result was unpredictable. Apama 5.3 makes the behavior consistent and predictable. The APIs that now require the name of a monitor are:

Java	<code>DeleteOperationsInterface.killName()</code> , <code>killNames()</code> and <code>killNamesFromFile()</code>
C++	<code>EngineManagement.killName()</code>
C	<code>AP_EngineManagement_Functions.killName()</code>
.NET	<code>ICorrelatorManagement.KillName()</code> , <code>KillNames()</code> and <code>KillNamesFromFile()</code>

Execution of the `engine_delete` utility with the `-k` option, which kills the specified object, also now requires specification of the name of at least one monitor. For example:

```
engine_delete -k com.mycompany.MyOrderMonitor
```

Execution of `engine_delete` with the `-k` option kills every instance of the specified monitor(s) whether or not the instance is processing anything. Any `ondie()` and `onunload()` actions defined in killed monitors are not executed.

New sample project showing how to manage Universal Messaging DataGroups

A new Apama sample project shows how Apama can

- Publish to Universal Messaging (UM) DataGroups
- Dynamically manage the membership of UM DataGroups as clients, known as DataStreams, connect and disconnect to the UM server.

UM DataGroups use a publish/subscribe messaging technology, which has a lightweight grouping structure that allows a remote process to manage the data received by each client.

Note:

UM DataGroups use a completely different mechanism than the one used by JMS topics or UM channels, which can be accessed through Apama's correlator-integrated messaging for JMS or through UM native integration.

The new sample has two parts:

- A suggested EPL event API to expose the DataGroup publisher and manager functionality. It is implemented using a correlator plug-in that is written in Java. Complete source code is provided.
- A sample application that shows simple and advanced ways to use this API for publishing and/or managing DataGroups. This demonstrates the powerful combination of the Apama platform and the UM DataGroups feature.

For more information about the sample, see `README.txt` in `APAMA_HOME/samples/universal_messaging/datagroups`. For more information about UM DataGroups, refer to the [Universal Messaging documentation](#).

C++ and C Correlator plug-in API enhancements

Correlator plug-ins written in C++ and C can now return sequences. The APIs for developing correlator plug-ins in C++ and C have been enhanced as follows:

- The C++ API class `AP_Type` now defines the following functions:

- To set the length of an existing sequence:

```
void setSequenceLength(size_t length)
```

- To create a sequence that contains empty items of the type you specify:

```
void createSequence(AP_TypeDescriptor inner)
```

Also, `AP_Type` objects can now be directly assigned from one to another. For example, to return the first argument, regardless of type:

```
rval = args[0];
```

- The C API class `AP_PluginType_Functions` now defines the following comparable functions:

- `void (AP_PLUGIN_CALL *setSequenceLength)(const AP_PluginType *obj, AP_uint32 len)`
- `void (AP_PLUGIN_CALL *createSequence)(const AP_PluginType *obj, AP_TypeDescriptor inner)`
- `void (AP_PLUGIN_CALL *copyFrom)(const AP_PluginType *obj, const AP_PluginType *other);`

- For both C++ and C APIs, a sequence you create can contain items of type integer, float, boolean, string, or chunk.

Miscellaneous enhancements and changes in 5.3

- It is now possible to specify the Java classpath individually for each JMon application or Java plug-in, using a new `classpath` element in the jar file's deployment descriptor XML or `@Application` annotation. See "Specifying the classpath in deployment descriptor files" in *Developing Apama Applications*.
- When you start a correlator and specify the `-J` option you can now specify `-J-Djava.class.path=path` and the path you specify will be appended to all internal Apama JAR files. In previous releases, Apama used only the setting of the `CLASSPATH` environment variable. If you specify both the `CLASSPATH` environment variable and a classpath on the command line then the classpath specified on the command line takes precedence.
- The MATLAB plug-in has been moved from the Apama Capital Markets Foundation into Apama 5.3. This plug-in includes the MATLAB bundle, libraries and executables that let Apama applications use MATLAB products. Documentation for using MATLAB products in an Apama application is now in *Developing Apama Applications*, *Developing Apama Applications in EPL*, "Using Correlator Plug-ins in EPL", "Using MATLAB products".
- The Apama extensions for the Pysys testing framework now include `injectCDP()` method to allow a CDP to be injected into a correlator from within a test.
- The event parse error logged by the correlator now contains the character offset of the error and the field member name that was being parsed. The expected event format logged is now recursive so nested event fields are also fully logged.
- Apama 5.3 incorporates ICU (International Components for Unicode) time-zone data update 2014j_44, which is the most recent update at the time of release. This will update time-zone data used by the correlator and TimeFormat correlator plug-in.
- The behavior of escaping in mapping literal expressions in the correlator JMS and Web Services Client adapter has been changed. Previously, literal expressions in mapping were evaluated as is without any escaping. Now a slash (`\`) is used as the escape character. Now a double quote (`"`) is escaped as `\`", a single quote (`'`) is escaped as `\`' and a slash (`\`) is escaped as `\\`. Quotes only need to be escaped if a string value is enclosed in the same type of quote. So the literal expression "a quote `\`' is string" was previously evaluated as "a quote `\`' is string", but it is now evaluated as "a quote `'` is string".
- Several performance improvements have been made to the `Parse()` and `Text` (formatting) implementation of the .NET event parser. This has resulted in a small change of behavior in `Event.FieldsMap`, which previously returned a `SortedDictionary` with a deterministic key iteration order (based on the natural ordering of the keys) but now returns an unordered dictionary with non-deterministic iteration order. For rare cases where deterministic ordering is required, users should wrap the `FieldsMap` result in a new `SortedDictionary` instance before iterating over it.
- The default consistency mode for `BigMemory` configurations is now `EVENTUAL` (the `BigMemory` default) rather than `STRONG`, and `synchronousWrites` has been disabled. Existing configurations will not be updated automatically, and users are still able to enable `STRONG` consistency if needed, but due to its superior performance, `EVENTUAL` is recommended when using Apama with the latest `BigMemory` release.

Platform changes in 5.3

Apama 5.3 runs on the platforms listed in the *Supported Platforms* document for version 5.3. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>.

ADBC databases are now supported only using the drivers included with Apama (except for MySQL, which is supported using MySQL's driver). Use of the JDBC drivers is strongly recommended in all cases, and the use of ODBC drivers (those supplied with Apama and the MySQL ODBC driver) is deprecated in this release.

The 32-bit Windows platform is now supported only for development tasks. Production use of Apama is not supported on this platform.

Check the Apama 5.3 Supported Platforms document for details about the newer platform versions supported by Apama 5.3. Due to these upgrades, the following versions are no longer supported:

- Software AG Terracotta BigMemory 4.1.4.
- Software AG Universal Messaging 9.7 is no longer supported, for either JMS or native Apama messaging.
- Oracle Java 7 (update 25 and higher) is supported for only client-side (client API applications, dashboards) use. Oracle Java 6 is no longer supported.
- MySQL 5.1.35
- IBM DB2 9.5.
- HornetQ 2.3.0.
- Eclipse 4.3.
- Mozilla Firefox below version 31.
- Microsoft Internet Explorer below version 9.
- Apple Safari desktop below version 5. (NB: the supported Safari version on the iPad is unchanged).
- WebSphere Application Server 7.0.0.13.

JDK 7 now minimum JDK version supported

Apama is now built against JDK 7, and installs with JRE 8.

You can no longer use JDK 6 (or older) to develop Apama applications nor can you use JRE 6 (or older) to deploy Apama applications. While it should not be necessary to rebuild applications previously built with an older Java version, the recommendation is that you build your application with the newer JDK for two important reasons:

- To ensure that you detect compatibility problems before you want to deploy the application
- To take advantage of performance improvements in newer Java compiler versions

To develop, build, and test an Apama 5.3 application, the recommendation is that you use Oracle JDK 8. The minimum you can use to build your application with Apama 5.3 libraries is JDK 7.

To deploy an Apama 5.3 application, the recommendation is that you use Oracle JRE 8, which is the version that Apama installs. Use of any JRE other than the one that Apama ships with is discouraged.

With Apama 5.3, web application servers and client web browsers must be upgraded to JRE 7 or later to access WebStart or Applet dashboards (JRE 8 is recommended). This requirement is due to the Java compiler changes in Apama 5.3.

Removed and deprecated features in 5.3

The following requests have been removed from the `engine_management` utility:

- `-r reopenLog`
- `-r rotateInputLog`
- `-r rotateReplayLog`

Instead, specify the `-r rotateLogs` request. See [“Log rotation enhancements in 5.3” on page 176](#).

The ODBC drivers included with Apama are deprecated in this release. Use the provided JDBC drivers.

The JDBC-ODBC bridge has been removed due to the Java 8 upgrade. Use the provided JDBC drivers.

Dashboards no longer support ODBC as a database source, due to the JDBC-ODBC bridge removal. The "Use ODBC Driver" option is enabled by default, this must now be disabled and the JDBC driver used.

The Visual Event Mapper is no longer available for ODBC data sources, due to the JDBC-ODBC bridge removal. The ODBC adapter Editor will no longer show the mapping tab to map events and generate supportive monitor script.

The event action name `getSourceTime` is now reserved for future use.

Backwards incompatibility with persisted projects recovered to 5.3 from older versions

To support the new Queries feature in Apama 5.3, some internal (non-public) event definitions have been changed in `ScenarioService.mon`. The change will have no effect on most applications as this API is not officially supported or documented. However it may affect customers using a persistent correlator to maintain application state across an upgrade from Apama 5.2 (or earlier) to 5.3 (or later), if they also make use of scenarios, DataViews or the MemoryStore plug-in.

This is because the EPL that is recovered from the persistent store after the upgrade will use the older (pre-5.3) event definitions and these are not compatible with post-5.3 Scenario Service clients and dashboards, or with the MemoryStore plugin. As a result users may see event parsing failures logged by the correlator, and/or Scenario Service clients and dashboards reporting errors such as

ScenarioService cannot parse event due to mismatch between the EPL event definition and the ScenarioService client.

In most cases it should still be possible to recover from a pre-5.3 persistent store without losing application state, provided the following actions are taken immediately after starting the correlator post-upgrade:

- If using DataViews and the `DataViewService_Impl_Dict.mon` monitor, delete it using `engine_delete com.apama.dataview.DataViewService_Impl_Dict`
- If using MemoryStore and the `MemoryStoreScenarioImpl.mon` monitor, delete it using `engine_delete com.apama.memorystore.MemoryStoreScenarioImpl`
- Delete any application Scenario definitions, using `engine_delete Scenario_SCENARIO_NAME`

Note that Scenario instances are never persistent so the state of any running Scenario instances would not be preserved across a correlator restart anyway.

- Delete the Scenario Service EPL using `engine_delete com.apama.Scenario`
- Re-inject the latest version of the Scenario Service from `APAMA_HOME/monitors/ScenarioService.mon`
- If using `MemoryStoreScenarioImpl` (see above), reinject the latest version from `APAMA_HOME/monitors/data_storage/MemoryStoreScenarioImpl.mon`
- If using `DataViewService_Impl_Dict` (see above), reinject the latest version from `APAMA_HOME/monitors/DataViewService_Impl_Dict.mon`
- Re-inject any Scenario definitions that are part of your application
- Send any events that are required to reinstantiate Scenario instances

20 What's New in Apama 5.2

■ Apama Studio enhancements in 5.2	186
■ New feature for addressing contexts with channels in 5.2	187
■ New feature for using Universal Messaging to connect Apama components in 5.2	191
■ New and changed EPL statements in 5.2	192
■ Enhancements to Apama adapters in 5.2	193
■ Changes to correlator utilities in 5.2	193
■ Windows compiler and .NET version updates in 5.2	195
■ Linux compiler updates in 5.2	195
■ Upgrading to Apama-supplied JDBC and ODBC database drivers in 5.2	196
■ Removed and deprecated features in 5.2	197
■ Miscellaneous enhancements and changes in 5.2	199

Apama 5.2 runs on the platforms listed in the *Supported Platforms* document for version 5.2. This is available from the following web page: <http://documentation.softwareag.com/apama/index.htm>. Be sure to consult that document for details about supported versions of operating systems, compilers, BigMemory, and Universal Messaging.

Apama 5.2 includes the following new features, enhancements, and changes.

Apama Studio enhancements in 5.2

Apama 5.2 adds the following enhancements to Apama Studio.

Apama Studio now uses Eclipse 4.3.

When Apama Studio starts it now collects any files in Apama's new `studio\extensions` folder and uses those files to add resources to the Apama Studio environment. This is useful when you want to import projects that have dependencies such as environment variables or catalogs of blocks, bundles or functions. In previous releases, when you imported projects that had such dependencies you manually added each resource before you could build the project.

See "Setting up the environment before importing projects" in *Using the Apama Studio Development Environment* (renamed to *Using Apama with Software AG Designer* in version 9.9).

There are several enhancements for creating launch configurations:

- In the **Run Configurations** dialog, when you select the **Common** tab, the default selection for **Save as** is now **Shared file** with a default path of `project_name\config\launch`. In previous releases, **Local file** was selected by default.

When **Shared file** is selected and you click **Apply** to create a new launch configuration Apama Studio now creates two files and places them in the directory specified by **Shared file**. The names of these files have the following format and the launch configuration information is split between them:

- `launch_config_name.deploy`
- `launch_config_name.launch`

In previous releases, creating a launch configuration generated only a `.launch` file, which contained Apama information but was managed by Eclipse. Now Apama manages the content of the `.deploy` file.

Any changes you make to the launch configuration will be reflected in the `.deploy` file.

- When you export an Ant launch configuration and you select **Generate initialization list during launch**, Apama Studio uses the `.launch` and `.deploy` files to create the deployment script. The Ant deployment script then uses `.deploy` file at runtime to read the injection order details.
- In the **Run Configurations** dialog, in the **Components** tab, a new **Connections** button is available. After you add two or more correlators, click **Connections** to specify connections between correlators. This has the same effect as using the `engine_connect` correlator utility to connect correlators.

The Ant export facility uses this connection information to create the required `engine_connect` tasks in the deployment script. Any `engine_connect` options that you specify in the **Connections** tab are saved in the `.deploy` file.

- In the **Correlator Configuration** dialog, the **Initialization** tab has been rearranged into two tabs as follows:
 - The **Injections** tab lists the files to be injected except for event files. The files will be injected in the order in which they appear in the **Injections** tab. You can accept the default **Automatic Ordering** calculated by Apama Studio or you can select **Manual Ordering** to change the default order. If you change the default order then it is up to you to ensure that you specify a correct injection order. You can also de-select a file so that it is not injected in a particular launch configuration. The file remains in the project.
 - The **Event Files** tab lists the `.evt` files to be injected. You can de-select a file so it is not injected in a particular launch configuration. You can select a `.evt` file and move it up or down in the list. The order of the files in the list on the **Event Files** tab is the order in which they will be injected. The default behavior is that Apama Studio lists the files in lexicographical order.

The information in the `.deploy` file accommodates any injection order changes you make in the **Injections** tab or **Event Files** tab of the **Correlator Configuration** dialog.

In Apama 5.2, EPL files are injected first and then `.evt` files are injected. In previous releases, EPL files and `.evt` files were injected in the order in which they appeared in the **Initialization** tab, which meant that they might be interspersed. This might cause a backward incompatibility for launch configurations created with previous releases.

New feature for addressing contexts with channels in 5.2

Apama 5.2 extends the ability to send events to named channels by providing a more powerful and general feature for using channels. In your application, each context, external receiver, correlator plug-in, or Apama client application can subscribe to one or more channels to receive the events sent on that channel. Adapter configurations can also now specify the channels on which the adapter sends events to correlators.

When a monitor instance subscribes to a channel it receives events delivered to that channel from adapters and other contexts. All monitor instances in the same context as the subscribing monitor instance receive the events delivered to the subscribed channel. Within a context, there can be subscriptions to different channels. The delivery of events from adapters, over multiple channels, to multiple contexts, is parallel from adapter to processing context, and does not impose any serial bottlenecks.

To achieve the best performance, the way in which data is organized for delivery from adapters to multiple channels must be determined. Typically, this organization matches the way you distribute the work of your application into contexts. Data sent to a particular channel should have no ordering dependency on data sent to any other channel.

The new channel features include the following enhancements and changes as well as additions to Apama Studio that support the use of channels.

EPL

- The new `send...to` statement sends the specified event to the specified channel. Use the `send...to` statement in place of the `enqueue...to` or `emit...to` statement. The `enqueue...to` and `emit...to` statements will be removed in a future release.
- The new `monitor.subscribe()` and `monitor.unsubscribe()` statements subscribe/unsubscribe a context to the specified channel. The monitor that contains the statement and any other monitors in the same context are subscribed/unsubscribed. Note that this is a special use of the `monitor` keyword.
- The new `com.apama.Channel` type holds a string or a context. You can send events to `com.apama.Channel` objects. If the `Channel` object contains a string then the event is sent to the channel of that name. If the `Channel` object contains a context then the event is sent to that context.

Correlator utilities

- The `engine_send` utility and `engine_receive` utilities can send/receive a file that associates events with channels. For details, see "Event association with a channel" in *Deploying and Managing Apama Applications*.
- The `engine_connect` utility has a new `parallel` mode (`-m parallel` or `--mode parallel`) that uses a connection per channel between correlators and delivers events to the target correlator on the channel the events were sent from in the source correlator. Previous behavior of `engine_connect` is preserved in `legacy` mode (`-m legacy` or `--mode legacy`), which uses one connection between two correlators and delivers all events to the default channel.

All connections to the correlator are now persistent. Consequently, the `engine_connect` utility no longer needs to provide the `--persistent` option and this option has been removed.

- The `engine_inspect` utility now provides information about channels being used and about receivers.
- When the correlator sends status to its log it now provides information for the following new status indicators:
 - `srn` — Slowest receiver name is the name of the receiver whose queue has the largest number of entries. If no receivers have queue entries then this value is "<none>".
 - `srq` — Slowest receiver queue. For the receiver identified by `srn`, the slowest receiver, this is the number of entries on its queue.

Apama client APIs

In C, C++, Java and .NET, there is a new enumeration called `ConnectMode` with values `CONNECT_LEGACY` or `CONNECT_PARALLEL`. This is used by new overloads for the existing methods `attachAsEventConsumerTo()`, `detachAsEventConsumerFrom()`, `attachAsConsumerOfEngine()` and `detachAsConsumerOfEngine()`.

There is now a separate connection for each EngineManagement object (EngineClientBean object, EventService object, or ScenarioService object) that you create in a client application. In previous releases, multiple EngineManagement objects shared a connection to the same Apama component.

For events that have a channel attribute set, the value of that attribute is now used when the event is sent to a correlator. In previous releases, the value of a channel attribute was ignored. For events that do not have a channel attribute set, the behavior is unchanged. That is, the event is delivered on the default channel (the empty string) to all public contexts.

The setEngineParams() method no longer accepts the LogicalID argument.

Correlator plug-in APIs

- Correlator plug-ins must be rebuilt against the new header files.
- When writing a correlator plug-in in C++, the following new methods are defined by AP_CorrelatorInterface:

```
virtual void sendEventTo(const char* event, AP_uint64 targetContext,
    AP_uint64 sourceContext)
void subscribe(const AP_EventHandlerInterface::ptr_t &handler,
    std::initializer_list<const char *> channels)

template<typename ITER>
    void subscribe(const AP_EventHandlerInterface::ptr_t &handler,
        const ITER &start, const ITER &end)

void subscribe(const AP_EventHandlerInterface::ptr_t &handler,
    const T &channel)

void unsubscribe(const AP_EventHandlerInterface::ptr_t &handler,
    std::initializer_list<const char *> channels)

template<typename ITER>
    void unsubscribe(const AP_EventHandlerInterface::ptr_t &handler,
        const ITER &start, const ITER &end)

void unsubscribe(const AP_EventHandlerInterface::ptr_t &handler,
    const T &channel)

virtual void unsubscribe(const AP_EventHandlerInterface::ptr_t &handler)
```

- When writing a correlator plug-in in Java, the following new methods are available on com.apama.epl.plugin.Correlator

```
public static native void sendTo(String evt, String chan);
public static native void sendTo(String evt, Context ctx);
public static native void sendTo(String evt, Context[] ctxs);

public static native void subscribe(EventHandler handler,
    String[] channels);
public static native void unsubscribe(EventHandler handler,
    String[] channels);
public static native void unsubscribe(EventHandler handler);
```

Adapter configurations

- When specifying connections to correlators, you can specify the optional `parallelConnectionLimit` attribute in a `<sinks>` element. Normally, you do not need to specify this attribute. The default behavior is that the IAF limits itself to an internally set number of connections with each specified sink. This number scales according to the number of CPUs that the IAF detects on the host that is running the IAF. While this number is usually sufficient, there are some situations in which you might want to change it. For example, if you are trying to conserve resources you might want to limit the number of connections to 1, or if you want to prevent multiple threads from sharing a connection you might allow a higher number of connections than the default allows.
- When specifying the mapping between an Apama correlator event type and a kind of external message, you can now specify the attributes `transportChannel` and/or `presetChannel` to support getting and setting the channel of each normalized event that passes through the IAF adapter.

For details, see "Event mappings configuration" in *Connecting Apama Applications to External Components*.

Application upgrade considerations arising from channels improvements

After you install Apama 5.2, you will need to make some modifications to your applications to take advantage of the new channel features. Before you make these modifications, consider the following:

- Does your application depend on event ordering? Applications that use multiple channels might re-order events.
- How many channels do you need? Which events would use which channels? What is your naming convention for channels? Typically, the way your application organizes work into contexts is a good indicator for how to use channels.
- In multi-correlator deployments, can you use fewer correlators? By using channels, a correlator provides greater parallel processing, which might render some additional correlators unnecessary. Using fewer correlators reduces the complexity of your deployment and can improve performance.
- Can you use Universal Messaging rather than the `engine_connect` utility? This also reduces the complexity of your applications and allows you to benefit from other UM features. See [“New feature for using Universal Messaging to connect Apama components” on page 191](#).

To start using channels:

- In EPL monitors, `subscribe` to receive events delivered to particular channels. Also, use `send` statements to send events to particular channels.
- In adapter configurations, when specifying the mapping between an Apama correlator event type and a kind of external message, specify new channel attributes in `<event>` and `<unmapped>` elements. See [“New feature for addressing contexts with channels in 5.2” on page 187](#) for details.
- In correlator plug-ins, use the new methods for subscribing to receive events sent on particular channels and for sending events to particular channels. These methods are provided by:

- C++: `AP_CorrelatorInterface`
- Java: `com.apama.epl.plugin.Correlator`

To upgrade multi-correlator deployments, the additional recommended steps are:

- In `engine_connect` command lines, specify the `-m parallel` or `--mode parallel` option. In parallel mode, for each specified channel, Apama will create a separate connection between the sending and target correlators.
- In event files to be sent by the `engine_send` utility, specify channels. For details, see "Event association with a channel" in *Deploying and Managing Apama Applications*.
- In the parts of your application that receive output from the `engine_receive` utility, make any changes needed to accommodate the new channel specification associated with events.

New feature for using Universal Messaging to connect Apama components in 5.2

Universal Messaging (UM) is Software AG's middleware service that delivers data across different networks. It provides messaging functionality without the use of a web server or modifications to firewall policy.

Previous Apama releases provided access to UM only through the correlator-integrated adapter for JMS. With Apama 5.2, there is support for Apama applications to directly use UM to transport events between Apama components. This can make deployments more flexible and can simplify configuration.

In an Apama application, correlators and adapters can connect to UM realms or clusters. A correlator or adapter connected to a UM realm or cluster uses UM as a message bus for sending Apama events between Apama components. Connecting a correlator or adapter to UM is an alternative to

- Defining connections between an adapter and particular correlators in the `<apama>` element of an adapter configuration file
- Specifying a connection between two correlators by executing the `engine_connect` correlator utility

Using UM can simplify an Apama application configuration. Instead of specifying many point-to-point connections you specify only the address (or addresses) of the UM realm or cluster. Apama components connected to the same UM realm use UM channels to send and receive events. (UM channels are similar to JMS topics.)

When an Apama application uses UM a correlator can be configured to automatically connect to the required UM channels. There is no need to explicitly connect UM channels to individual correlators. A correlator automatically receives events on UM channels that monitors subscribe to and automatically sends events to UM channels.

You can use UM only to send and receive events. You cannot use UM for EPL injections, delete requests, engine send, receive, watch or inspection utilities, or `engine_management -r` requests.

Only UM channels can be used with Apama. UM queues and datagroups are not supported in this Apama release.

In Apama Studio, you can add UM support to a project. After you do that, you can choose whether to specify UM messaging in correlator launch configurations and in adapter configurations.

In an IAF adapter configuration file, the new `enabled` attribute in an `<apama>` element and in the new `<universal-messaging>` element indicates whether the adapter uses UM. The default behavior is that an adapter does not use UM. To configure an adapter to use UM, set `enabled` to `true` in the `<universal-messaging>` element. Also, set `enabled` to `false` in the `<apama>` element, if there is one. This indicates that the configuration specified in the `<apama>` element should be ignored and the configuration specified in the `<universal-messaging>` element should be used.

When you are upgrading an application to use Apama 5.2 and you are configuring it to use UM keep in mind that events must be sent to named channels. You might need to add channel specifications to events. For adapters, you can use the `transportChannel`, `presetChannel`, and `defaultChannel` attributes to set channel names. If you do not specify values for any of these attributes, but you use Apama Studio to add UM support to your project then a default channel name will be in place for adapters that use UM. If you do not use Apama Studio then you must at least specify a value for the `defaultChannel` attribute in IAF adapter configurations. For details, see "Event file format" in *Deploying and Managing Apama Applications* and "Configuring IAF adapters to use Universal Messaging" in *Connecting Apama Applications to External Components*.

Apama 5.2 supports only the 9.7 release of Universal Messaging. The *Supported Platforms* document for version 5.2 in the [Apama documentation](#) lists supported releases for all Apama components.

Details for using Universal Messaging in an Apama application are in *Connecting Apama Applications to External Components*.

New and changed EPL statements in 5.2

Apama 5.2 includes the following enhancements to EPL:

- The new `send...to` statement supersedes the `enqueue...to` and `emit...to` statements. The `send...to` statement sends the specified event to the specified channel.

You can no longer use `send` as an identifier since it is now a keyword.

The `enqueue...to` and `emit...to` statements will be deprecated in a future release. Use the `send...to` statement instead.

- The new `monitor.subscribe()` and `monitor.unsubscribe()` statements subscribe/unsubscribe a context to the specified channel. The monitor that contains the statement and any other monitors in the same context are subscribed/unsubscribed. Note that this is a special use of the `monitor` keyword.
- The new type `com.apama.Channel` represents a channel in the correlator. See [“New feature for addressing contexts with channels in 5.2” on page 187](#).

Enhancements to Apama adapters in 5.2

Apama's Web Services Client adapter and Correlator-Integrated adapter for JMS now provide interfaces for registering and using custom mapping extensions that use Java Unified Expression Language (EL) resolvers and methods. In previous releases, while Apama provided EL resolvers and EL methods for you to use in mapping expressions, you could not add your own.

Changes to correlator utilities in 5.2

Apama 5.2 enhances and changes correlator utilities as follows:

■ engine_connect

The `engine_connect` utility has a new parallel mode (`-m parallel` or `--mode parallel`) that uses a connection per channel between correlators and delivers events to the target correlator on the channel the events were sent from in the source correlator. Previous behavior of `engine_connect` is preserved in legacy mode (`-m legacy` or `--mode legacy`), which uses one connection between two correlators and delivers all events to the default channel.

All connections to the correlator are now persistent. Consequently, the `engine_connect` utility no longer needs to provide the `--persistent` option and this option has been removed.

The change in behavior of `engine_connect` switching to persistent connections and other changes in Apama messaging means that two connections for the same component and channel but using different hostnames for the connections (for example, some combination of full hostname, short hostname or IP address, or multiple different addresses or hostnames) will create multiple connections, potentially duplicating events. Apama recommends ensuring all correlators are named consistently. In default (not parallel) mode, multiple connections using the same hostname will be flattened to a single connection, avoiding duplicate events.

There is a change in behavior when disconnecting. Previously, when you specified the `-x` option without also specifying the `-c` option the two correlators remained connected even though the source correlator stopped sending events to the target correlator. In this release, specification of the `-x` option without also specifying the `-c` option disconnects the two correlators.

■ engine_debug

When you plan to run this utility you must specify the `-g` (or `--nooptimize`) option when you start the correlator. This option disables optimizations that hinder debugging. In previous releases, you could run `engine_debug` even if `-g` was not specified when the correlator was started. Apama Studio automatically uses the `-g` option when it starts a correlator from a debug launch configuration. However, if you connect Apama Studio to an externally started correlator and you want to debug in that correlator then you must ensure that `-g` was specified when the externally-started correlator was started.

■ engine_inspect

The new `-P` option displays the name of any plug-in receivers, what channels each plug-in receiver is subscribed to, and the number of entries on the input queue of each plug-in receiver. A plug-in receiver is a correlator plug-in that is subscribed to at least one channel.

The new `-R` option displays the names of any external receivers, their addresses, what channels each external receiver is subscribed to, and the number of entries in the output queue of each external receiver.

The `-x` option now also lists the channels each context is subscribed to.

■ `engine_management`

When you specify the `-xr` or `-xs` option, you must now specify the component ID as *physical_ID/logical_ID*.

■ `engine_receive`

The new `-C` or `--logChannels` option causes the `engine_receive` utility to include the channel on which an event is received in `engine_receive` output. If you then use `engine_receive` output as input to the `engine_send` utility the events will be delivered back to their channels.

■ `engine_send`

It is now possible to prefix an event with a quoted string that denotes the channel that event is to be sent on. The prefix string follows the same escaping rules as string in events. For example, an entry in a `.evt` file might be: `"MyChannel",Tick("SOW", 35)`.

The new `-c channel` option identifies the delivery channel for events for which a delivery channel is not specified.

■ `engine_watch`

This utility now provides the following information:

Events on input context queues	The total number of events on the input queues of all public contexts in the correlator.
Most backed up input queue	The input queue that has the most events waiting to be processed.
Most backed up queue size	The number of events on the input queue that has the most events waiting to be processed.
Slowest receiver	The receiver with the largest number of incoming events waiting to be processed.
Slowest receiver queue size	For the receiver with the largest number of incoming events waiting to be processed, this is the number of events that are waiting.

■ `extract_replay_log`

When you run this script on a correlator input log you can specify the `-c` or `--correlator` option. This option specifies that the script that `extract_replay_log` generates should include the command line for starting a correlator. When you run the generated script, the correlator will be started with all of the command-line options needed to replay the input log. While this option is not new in this release, it was previously undocumented.

Support for replay logs has been removed. The `extract_replay_log` utility is now for extracting only input logs.

Windows compiler and .NET version updates in 5.2

Apama 5.2 has upgraded from Microsoft Visual Studio 2008 (Visual C++ v9) to Visual Studio 2013 Update 2 (Visual C++ v12). Consequently, you must rebuild the following components with the new compiler:

- C and C++ correlator plug-ins
- C and C++ IAF transport/codec plug-ins
- C and C++ code that uses the engine client API

For .NET clients, Apama now supports .NET 4.5.1 instead of 2.0. It is no longer possible to build .NET 2.0 or 3.5 clients with the Apama libraries.

In addition, you should consider the following:

- Visual Studio 2013 uses a different format for its C++ project files, which now have the `.vcxproj` extension rather than `.vcproj`. Visual Studio's migration wizard can usually convert the project files automatically. Alternatively, it is simple to create new projects in Visual Studio 2013 (from scratch or based on Apama samples) and re-add existing source files and project customization to them. Creating new projects might result in cleaner configuration files.
- The Apama C and C++ API samples all have new Visual Studio 2013 project files. Some non-standard preprocessor macro definitions that were present in older versions of the samples have been removed, for example, `_APBUILD_WIN32_ALL_` and `_AMD64_`. The samples now consistently use Visual Studio's standard macros instead (`_WIN32` and `_M_AMD64`). If you are incorporating existing code that depends on the old macros you might want to re-add the old macros to your project properties. To do this, in the Visual Studio, select your project and then select **C/C++ > Preprocessor > Preprocessor Definitions**.
- When planning migration, consider that the C++ compiler in Visual Studio 2013 is considerably more modern than the 2008 version, and also less tolerant of ambiguous and potentially unsafe code patterns. It is likely that some time may be required to fix compiler errors after moving to the new version. The recommendation is to enable **'warnings as errors'** in your projects where possible. This encourages good coding patterns and catches errors earlier in the development process. In particular, it helps you catch errors that might cause subtle, difficult to diagnose errors.
- .NET code that uses the engine client API may need to be rebuilt. The `.csproj` project files will need to be upgraded if you are using the new compiler.

Linux compiler updates in 5.2

Apama 5.2 builds on Red Hat Enterprise Linux (RHEL) 6 Update 1 with GCC 4.4.5. Apama is certified for use on this platform.

You should rebuild the following components with the new compiler:

- C and C++ correlator plug-ins
- C and C++ IAF transport/codec plug-ins
- C and C++ code that uses the engine client API

If you are using SUSE Linux Enterprise Server (SLES) 11 Update 3 with the default shipped GCC 4.3 compiler, note that while Apama 5.2 is certified to work on this platform, you will see a **no version information available** warning when building plug-ins against Apama. This is because the compiler used for building Apama was GCC 4.4.5. At runtime, this should not present any problems.

Upgrading to Apama-supplied JDBC and ODBC database drivers in 5.2

As in past releases, Apama 5.2 provides JDBC and ODBC database drivers. For Apama 5.2, these drivers have been upgraded and configured slightly differently than in past releases.

The drivers provided in previous releases are incompatible with Apama 5.2. The drivers provided in Apama 5.2 are incompatible with previous Apama releases. If you have pre-Apama 5.2 projects that use the previously-supplied drivers and you want to run those projects with Apama 5.2 then you must manually convert those projects to use the Apama 5.2 drivers.

To upgrade your project to use Apama 5.2 ODBC drivers, change your application so that it uses an ODBC driver supplied with Apama 5.2. See "Registering your ODBC database DSN on Windows" in *Connecting Apama Applications to External Components*.

To upgrade your project to use Apama 5.2 JDBC drivers, change your application to use the new JDBC URLs and class names.

MSSQL

Old path:	<code>apama_install_dir\lib\pgsqlserver.jar</code>
New path:	<code>apama_install_dir\lib\eysqlserver.jar</code>
Old URL:	<code>jdbc:progress:sqlserver://HOSTNAME::PORT;databaseName=DATABASENAME</code>
New URL:	<code>jdbc:sag:sqlserver://HOSTNAME::PORT;databaseName=DATABASENAME</code>
Old class name:	<code>com.prqs.jdbc.sqlserver.SQLServerDriver</code>
New class name:	<code>com.apama.jdbc.sqlserver.SQLServerDriver</code>

Oracle

Old path:	<code>apama_install_dir\lib\pgoracle.jar</code>
New path:	<code>apama_install_dir\lib\eyoracle.jar</code>

Old URL: `jdbc:progress:oracle://HOSTNAME::PORT;databaseName=DATABASENAME`
 New URL: `jdbc:sag:oracle://HOSTNAME::PORT;databaseName=DATABASENAME`
 Old class name: `com.prgs.jdbc.oracle.OracleDriver`
 New class name: `com.apama.jdbc.oracle.OracleDriver`

DB2

Old path: `apama_install_dir\lib\pgdb2.jar`
 New path: `apama_install_dir\lib\eydb2.jar`
 Old URL: `jdbc:progress:db2://HOSTNAME::PORT;databaseName=DATABASENAME`
 New URL: `jdbc:sag:db2://HOSTNAME::PORT;databaseName=DATABASENAME`
 Old class name: `com.prgs.jdbc.db2.DB2Driver`
 New class name: `com.apama.jdbc.db2.DB2Driver`

Removed and deprecated features in 5.2

The following features are removed or deprecated in this release:

- Support for the correlator replay log has been removed. You can now replay correlator behavior only with a correlator input log.

The following options were needed when starting a correlator that would capture or run a replay log. These options are no longer needed and they have been removed:

```
--replayLog
--replayMode
```

With the removal of the replay log, the `extract_replay_log` utility now extracts data from input logs only.

When using Apama Studio you can specify an input log but not a replay log in the **Correlator Arguments** tab of the **Correlator Configuration** dialog. The **Replay log mode** field has been removed. from the **Correlator Arguments** tab.

Replay log support has also been removed from Apama Studio Ant features. Trying to run an Ant deployment script with a replay log will result in an error when trying to start the correlator.

- The following correlator start-up options have been removed:

```
--highThroughput
--lowLatency
```

When using EMM, you can no longer select these options in the correlator **Configurations** tab.

These options were useful when running Apama on single-core machines. With the recommendation to always run Apama on multiple core machines, these options are no longer needed.

- Now that all connections to the correlator are persistent, the `engine_connect --persistent` option is no longer needed and it has been removed.
- The following methods in the `MemoryStore RowValue` API class have been deprecated and will be removed in a future release:

```
public final void setString(int typedIndex, int maxSizeHint, String v)
public final void setString(int typedIndex, int maxSizeHint,
    byte[] utf8ByteString)
public final void setInteger(int typedIndex, int maxSizeHint, long v)
public final void setFloat(int typedIndex, int maxSizeHint, double v)
public final void setBoolean(int typedIndex, int maxSizeHint, boolean v)
```

In place of these methods, use the `set*()` methods without the `maxSizeHint` argument.

- The Payload Extraction correlator plug-in has been deprecated as of Apama 5.0. While it has not yet been removed because of internal requirements, it will be removed in a future release and its use is strongly discouraged.

The dictionary-format payload is considerably more efficient than the string-format payload in almost all cases. Third-party adapters should now be using dictionary-format payloads.

- The following Event Modeler functions are deprecated and will be removed in a future release:

`ADD_EXTRAPARAM`

`GET_EXTRAPARAM`

`HAS_EXTRAPARAM`

In place of these deprecated functions, use the following functions:

`DICTIONARY_GET`

`DICTIONARY_GETORDEFAULT`

`DICTIONARY_HASKEY`

`DICTIONARY_SET`

- The following overloading of the C++ plug-in API method for sending an event has been deprecated:

```
sendEventTo(const char* event, AP_uint64 targetContext,
    AP_uint64 sourceContext) = 0;
```

It will be removed in a future release. In place of that overloading, use the following method:

```
sendEventTo(const char *event, AP_uint64 targetContext,
    const AP_Context &source) = 0;
```

- With the addition of the `send...to` statement, `send` is now a keyword and you can no longer use it as an identifier unless you escape it with a hash symbol, for example, `#send`.

While the following EPL statements are not deprecated in this release, they will be deprecated in a future release. Use the EPL `send...to` statement instead.

- `emit` and `emit...to`
- `enqueue...to`
- The `DataViewService_Impl_Dict` monitor now uses the built-in EPL `log` statement instead of having its own custom logger, and the `com.apama.dataview.SetLogLevel` event has been removed. The log level for this monitor can now be changed as described in "Setting EPL log files and log levels dynamically" in *Deploying and Managing Apama Applications*.

Miscellaneous enhancements and changes in 5.2

Apama 5.2 includes the following enhancements:

- In an IAF adapter configuration file, the new default behavior is that all sinks defined for an Apama correlator configuration now receive the events generated by the IAF Semantic Mapper. To change the default behavior, add the `sendEvents="false"` attribute to each `<sink>` element that represents a component that should not receive events. In previous releases, the default behavior was that only the first sink would receive events.
- Several changes have been made to the `EngineClient`, `EventService` and `ScenarioService` layers in Java and .NET to make it easier to use event and property change listeners safely without risking thread deadlock situations. As a result of these changes, Apama client developers might notice that engine client property change listeners now fire asynchronously (on a separate dedicated thread), which is a slight change in behavior. See the release notes for details.
- When using EDA events in Apama applications, non-EDA namespaces are now allowed. Also, namespaces that contain colons (`:`) can now be converted to Apama package names.
- In Dashboard Builder, when you define the **Send event** command, you can optionally specify the channel on which to send the event.
- The new **Correlator Time Format** dashboard function lets you convert a correlator timestamp to epoch time or to a date/time format you specify.
- The BigMemory `MemoryStore` driver now sets the Java type of each search attribute, which makes it easier to use the search attributes from other products. For example, Software AG's Presto can extract the `ehcache.xml` configuration from a running BigMemory instance.
- The distributed `MemoryStore RowValue` API class has a number of usability improvements in its API, including easier to use set and get methods. (The old-style set methods have been deprecated.) There is also a new `RowValueHelper` class that allows more efficient allocation of `RowValue` objects, and simplifies getting, setting and iterating over the contents of a `RowValue` object.

- Apama 5.2 incorporates ICU (International Components for Unicode) time-zone data update 2014e, which is the most recent update at the time of release. This will update time-zone data used by the correlator and TimeFormat correlator plug-in.
- C++ correlator plug-ins can be declared as nonblocking to prevent creation of unneeded processing threads. If a method in a nonblocking plug-in might block you can override the nonblocking designation for that method.
- The C++ plug-in API method for sending an event has a new overloading, which you should now use:

```
sendEventTo(const char *event, AP_uint64 targetContext,  
            const AP_Context &source) = 0;
```

The following is deprecated and will be removed in a future release:

```
sendEventTo(const char* event, AP_uint64 targetContext,  
            AP_uint64 sourceContext) = 0;
```


21 What's New in Apama 5.1.1

Apama 5.1.1 includes the following enhancements:

- The correlator now parses incoming events from different connections in parallel. This improves performance when many adapters or clients are sending events to the correlator.
- Apama dashboards now support IBM WebSphere Application Server 8.5.5. Apama's correlator-integrated messaging for JMS feature also now supports WAS 8.5.5.
- Apama Studio can now export an Ant launch configuration that accommodates changes to your application. This means you need to export the launch configuration only once even if you subsequently add, remove or edit a file in your application. In previous releases, to ensure that your exported scripts were in sync with your application you had to re-export the launch configuration each time you changed your application.

To turn on the new export-once behavior, select the new **Generate initialization list during launch** option when exporting an Ant launch configuration from Apama Studio. The exported launch configuration contains the location of your project directory. When the exported deployment script is executed, the file initialization list is generated from the project directory.

If you select **Generate initialization list during launch** then any scenario files are always converted to monitor files at the time of injection to the correlator.

For details, see "Exporting to a deployment script" in *Using the Apama Studio Development Environment* (renamed to *Using Apama with Software AG Designer* in version 9.9).

- Search attributes can now be exposed for Apama MemoryStore application tables that are stored in Terracotta BigMemory caches. Turning this configurable capability on enables BigMemory clients to query data stored by Apama applications in BigMemory caches. The default is that search attributes are not exposed.

For more information, see "Using the distributed MemoryStore" in *Developing Apama Applications*.

- Apama applications now support the use of Software AG Event-Driven Architecture (EDA) event types. In Apama Studio, you can generate Apama event type definitions from EDA event definition schemas and you can map these Apama event definitions to JMS messages that communicate with EDA applications.

22 What's New in Apama 5.1

■ New distributed MemoryStore	204
■ New EPL exception handling mechanism	204
■ New support for managing and monitoring with REST APIs	205
■ New support for creating correlator plug-ins using Java	206
■ New headless mode generation of ApamaDoc	206
■ New support for PySys test framework	206
■ Dashboard login redirect file name has changed	206
■ New EPL keywords	207
■ Miscellaneous enhancements in 5.1	207
■ Support removed from Apama 5.1	208
■ Migrating Apama applications	209

What's New in Apama 5.1 summarizes the new, enhanced, and changed features in Apama 5.1.

If you are upgrading from a version of Apama prior to 5.0 then consult the appropriate What's New and Migration information available with the intermediate versions. For further assistance, contact customer support.

New distributed MemoryStore

The Apama MemoryStore now provides the ability to create distributed stores in which data can be shared by applications running in multiple correlators. Distributed stores are supported by distributed caching software from a variety of third-party vendors. Apama provides a driver for integrating the distributed MemoryStore with the Terracotta BigMemory Max distributed caching software. Apama also provides a Service Provider Interface (SPI) for creating drivers to use with other third-party distributed cache providers.

In Apama Studio, you can configure a correlator to use the distributed MemoryStore by selecting **Distributed MemoryStore Support** in the **Correlator Configuration** dialog. See "Correlator arguments" in *Using the Apama Studio Development Environment* (renamed to *Using Apama with Software AG Designer* in version 9.9).

For information on using the distributed MemoryStore and for creating drivers for third-party distributed caching software, see "Using the MemoryStore" in *Developing Apama Applications*.

New EPL exception handling mechanism

EPL now supports the `try ... catch` exception handling structure for processing runtime errors. The statements in each block must be enclosed in curly braces. For example:

```
using com.apama.exceptions.Exception;
...
action getExchangeRate(
    dictionary<string, string> prices, string fxPair) returns float {
    try {
        return float.parse(prices[fxPair]);
    } catch(Exception e) {
        return 1.0;
    }
}
```

An exception that occurs in `try block1` causes execution of `catch block2`. Two new types have been added to support exception handling:

- `com.apama.exceptions.Exception` — A variable of this type contains an exception message and an exception type. It also contains a sequence of `com.apama.exceptions.StackTraceElement` objects. The sequence represents the stack trace for when the exception was first thrown. Exception objects have methods for accessing the exception message, the exception type, and the stack trace.
- `com.apama.exceptions.StackTraceElement` — A variable of this type contains information for one stack trace entry. A `StackTraceElement` object has methods for accessing the details for the stack trace entry it represents.

try, catch and throw are now keywords. If you want to use one of these words as an identifier you must prefix it with a hash symbol (#), otherwise it is an error.

For details, see "Exception handling" in *Developing Apama Applications*.

New support for managing and monitoring with REST APIs

You can now monitor Apama components with the Apama REpresentational State Transfer (REST) HTTP API. This provides monitoring capabilities to third-party managing and monitoring tools or to any application that supports sending and receiving XML documents over the HTTP protocol.

Apama components expose several URIs which can be used to either monitor or manage different parts of the system. Generic management URIs are exposed by most Apama components, while other URIs are exposed only by specific types of components. Most URIs are purely for informational purposes and will only respond to HTTP GET requests and interacting with them will not change the state of the component. However, some URIs allow the state of the correlator to be modified. For example, to set a component's log level, the /logLevel URI accepts an HTTP PUT request containing an XML document that specifies the log level.

For example: GET `http://localhost:15903/correlator/status` returns an XML document that contains:

```
<?xml version="1.0"?>
<?xml-stylesheet type="text/xsl" href="/resources/transform.xslt"?>
<map name="apama-response">
  <prop name="numConsumers">0</prop>
  <prop name="numOutEventsQueued">0</prop>
  <prop name="numOutEventsUnAked">0</prop>
  <prop name="numOutEventsSent">0</prop>
  <prop name="uptime">67657</prop>
  <prop name="numMonitors">0</prop>
  <prop name="numProcesses">0</prop>
  <prop name="numJavaApplications">0</prop>
  <prop name="numListeners">0</prop>
  <prop name="numEventTypes">0</prop>
  <prop name="numQueuedFastTrack">0</prop>
  <prop name="numQueuedInput">0</prop>
  <prop name="numReceived">0</prop>
  <prop name="numFastTracked">0</prop>
  <prop name="numEmits">0</prop>
  <prop name="numProcessed">0</prop>
  <prop name="numSubListeners">0</prop>
  <prop name="numContexts">1</prop>
  <prop name="virtualMemorySize">262540</prop>
  <prop name="numSnapshots">0</prop>
  <prop name="numInputQueuedInput">0</prop>
  <prop name="mostBackedUpInputContext"><none></prop>
  <prop name="mostBackedUpICQueueSize">0</prop>
  <prop name="mostBackedUpICLatency">0.0</prop>
</map>
```

For more information on using the REST API, see "Managing and Monitoring over REST" in *Deploying and Managing Apama Applications*.

New support for creating correlator plug-ins using Java

Custom correlator plug-ins can now be written in Java as well as in C and C++. When a Java plug-in is injected to the correlator, it is analyzed and any suitable method in the plug-in can be called from applications running in the correlator. Correlator plug-ins in Java are deployed using JMon applications and are packaged in a .jar file.

Correlator plug-ins that are written in Java can be created by using Apama Studio or from scratch.

- To use Apama Studio, see "Adding an EPL Plug-in written in Java" in *Using the Apama Studio Development Environment* (renamed to *Using Apama with Software AG Designer* in version 9.9).
- To write a correlator plug-in in Java from scratch, see "Writing EPL Plug-ins in Java" in *Developing Apama Applications*.

New headless mode generation of ApamaDoc

On Windows platforms, you can now generate ApamaDoc in headless mode. As a standalone operation from the command line, run the `apamadoc.bat` script, which is in the `%APAMA_HOME%\bin` folder. The `apamadoc.bat` file uses the `%APAMA_HOME%\utilities\apamadoc.xml` ant script to generate ApamaDoc.

For more information, see "Generating ApamaDoc from an Ant script" in *Developing Apama Applications*.

New support for PySys test framework

PySys, an open source automated system testing framework, is recommended as the best way to test Apama applications developed with EPL, Event Modeler, or JMon. The open source PySys project operates on generic elements such as starting and stopping processes, and searching text files for correct or incorrect output. In addition, Apama includes extensions to PySys that let you test correlators, IAF components, EPL injection, sending and receiving events, and similar operations. The PySys framework can start adapters, which enables you to perform end-to-end testing of an entire system.

Samples for how to use PySys are in the `samples\pysys` folder of your Apama installation. Reference documentation for PySys and Apama extensions is in the `doc\pydoc` directory of your Apama installation.

Dashboard login redirect file name has changed

The `loginRedirect.html` file has been renamed to `loginRedirect.jsp`.

In previous releases, the `loginRedirect.html` file appeared if you tried to display a dashboard before you logged in or after your session expired. With this release, the new name of this file is `loginRedirect.jsp`.

The default behavior of `loginRedirect.jsp` is to redirect the user to the dashboard's form authentication page. If you have implemented a custom login, you must create a custom version of `loginRedirect.jsp`. For example:

```
<meta HTTP-EQUIV="REFRESH" content="0; url=../userLogin">
<html>
  <head>
    <title>Apama: Login Redirect from Dashboard</title>
    <meta http-equiv="pragma" content="no-cache">
    <meta http-equiv="cache-control" content="no-cache">
    <meta http-equiv="expires" content="0">
  </head>
  <body/>
</html>
```

In this sample custom `loginRedirect.jsp` file, `../userLogin` is a custom login page.

To enable your custom `loginRedirect.jsp` file, replace the provided version of it in the following locations:

- `APAMA_INSTALL\etc\template.war`
- `APAMA_INSTALL\etc\dispTemplate.war`

After your custom `loginRedirect.jsp` file is in place, you must re-deploy your dashboards.

New EPL keywords

New EPL keywords are listed below. If you want to use one of these words as an identifier you must prefix it with a hash symbol (`#`), otherwise it is an error.

- `catch`
- `optional`
- `static`
- `throw`
- `try`

Miscellaneous enhancements in 5.1

Apama 5.1 includes the following enhancements:

- Correlator Deployment Packages (CDPs) can now contain `.jar` files.
- In Apama Studio, you can now export a project's Web Services Client adapter configuration to an archive file. You can then import the archive file into any Apama project.
- The Apama correlator-integrated adapter for JMS now supports Software AG's Universal Messaging 9.5.2.

- In previous releases, when using the Apama correlator-integrated messaging adapter for JMS, you had to explicitly obtain and install the `javax.jms.jar` file. This file is now provided with Apama and installed as part of Apama installation. Manual steps are no longer required.
- In the **Event Mappings** tab for the correlator-integrated messaging adapter for JMS and for the Web Services Client adapter, Apama Studio now displays error markers if events that were previously mapped are missing. You must correct the event mappings before you can run the adapter.
- The `com.apama.jmon.Unloadable` interface has been added to JMon. For any monitor that implements this interface, the `Unloadable.onUnload()` method will be called when the application is being unloaded, for example, due to an `engine_delete` request.
- Dashboards now support the following functions. Details for using these functions are in the *Dashboard Function Reference*.
 - **Init Local Variable** initializes a local variable to a specified value.
 - **Quick Set Sub** sets a substitution string to a specified value.
- For dashboard processes, the `filterInstance` attribute in the `apama-macros.xml` file now defaults to `true`. This ensures that users can see only the instances that are owned by them. To change this behavior, edit the `apama-macros.xml` file or create the `DataView` instances with "*" as the owner.
- As of Apama 5.0, Apama applications running in the correlator can make use of Apama *out of band notifications*. Out of band notifications are events that are automatically sent to all public contexts in a correlator whenever any component (an IAF adapter, dashboard, another correlator, or a client built using the Apama SDKs) connects or disconnects from the correlator. (This was not previously listed in the *Release Notes*.)

Support removed from Apama 5.1

Support for the following features has been removed from Apama:

- IAF-based adapter for JMS
- Routers - A router is a specialized correlator that is optimized to partition events so they go to different correlators. With this release, any correlator can be used as a router and the documentation includes information and instructions for configuring this.
- The engine client simple/bean receiver in the engine client API for Java and for .NET. Applications still using this deprecated event receiving mechanism (that is, using the `setReceiveEnabled()` Java method or `ReceiveEnabled` .NET property) should be changed to use the `addConsumer()` method instead, which will require moving event handling code out of the property changed listener into a new `IEventListener` implementation.
- Third-party products
 - Actional adapter
 - Control Tower

- Corticon
- OpenEdge
- Savvion adapter
- Sonic ESB service

Migrating Apama applications

For up to date information about migrating to Apama 5.1, refer to the Apama Knowledge Base articles on Empower (empower.softwareag.com).

